Title
An Exploration of Meso-scale Plasticity

Author
Silversmith, Daniel B

Publication Date
2018

Peer reviewed|Thesis/dissertation

An Exploration of Meso-scale Plasticity

by

Daniel Silversmith

DISSERTATION

Submitted in partial satisfaction of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

Bioengineering

in the

GRADUATE DIVISION

of the

UNIVERSITY OF CALIFORNIA, SAN FRANCISCO

AND

UNIVERSITY OF CALIFORNIA, BERKELEY

[redacted]

.....

Committee in Charge

Dedications & Acknowledgements

Just about five-and-half years ago I entered grad school as a young, excited, confident (*perhaps a little cocky*) student ready to embark on an exciting journey. The long and winding road of grad school quickly put me in my place and never ceased to surprise and challenge me. Overcoming these challenges has been an extraordinarily demanding, yet astonishingly common journey—one of perseverance and resilience, but more importantly one of collaboration and support from mentors, peers, friends, and family. The best part of grad school, by far, has been meeting a tremendous group of enthusiastic, brilliant people who are all incredibly ambitious and inspiring but also fun, collegial, and thoughtful.

First I'd like to acknowledge all of the mentors that I have had during my time in grad school. These individuals have collectively advised and supported me throughout this process. My thesis committee (*this is real, not at all a bribe to let me graduate*) has really helped to shape my work by asking insightful questions and pushing me to do better work. But they have also asked me about my career goals and made sure that my work aligns with those goals. Thank you, Rich Ivry. Rich has a knack for asking me some of the most seemingly ridiculous, but in the end, really insightful, assumption-rejecting questions. I spent a few months in his lab and that time shaped my initial grad school trajectory by getting me excited about the psychophysics of sensory and motor learning but also left me with a burning need to always question assumptions. Thank you, Jeanne Paz. I have only met Jeanne recently, which is a shame. In the short time that we have known each other Jeanne has never failed to provide feedback about my science, advised me about what it

would take to elevate my work to the next level, and shown a much appreciated interest in my professional goals. Karunesh Ganguly, the head of my thesis committee, but also my advisor and mentor, thank you for stepping up and helping me during a transition period in my graduate school career. For that I will forever be grateful. Karunesh has a bottomless repository of energy and enthusiasm. After every meeting with him, I am left exhausted but eager and excited to try and achieve our next goals. His breadth of knowledge and ability to keep up with ALL of the literature, past and present, is nothing short of remarkable. In my first couple years as a grad student I learned so much, and I have to credit Philip Sabes (aka, Flip) for imparting most of that wisdom. Flip was always able to use his intuitive understanding of modeling and statistics to both teach and inspire different analyses. I learned so much from his first-principles approach to neuroscience.

But most importantly, Karunesh and Flip both brought together and inspired wonderful groups of bright, talented researchers. I will always appreciate the kindness that the Ganguly Grad Gang showed me by organizing a group hike to welcome my transition into the lab. Since then, I have come to have interesting, speculative conversations with students and post-docs about all of our burning questions… What on earth is sleep for? Are any of these signals even real? What's next in life? Is the Wi-Fi working? How can scientific institutions be improved? In Flip's lab I met many talented, smart people with a diverse set of attitudes and opinions. I learned a ton during my formative grad school years from all of the members of this lab. I would like to especially thank Azadeh Yazdan who acted as a tremendous mentor and became a great friend in and out of the lab. Azadeh has an infectious optimism that was a necessary counterpoint on our team. She has been a role model, demonstrating the ideals of determination, goal setting, mentorship, and teamwork.

Perhaps most importantly, she added a dose of silliness to our daily routine of working and was willing to begrudgingly debate me about anything and everything.

Over the last few years I have come to a new level of appreciation for the activities and the fantastic people who have had fun with me outside the lab. Soccer, running, and basketball with friends have provided countless hours of entertainment, shielded me from a descent into the lab-sleep-lab cycle, and kept me mentally and emotionally sane throughout school. Thank you to the cooking and baking crew who has never failed to impress me and introduced me to all sorts of tasty new cuisines. Thank you to my roommates who minimized their eye rolling during our political debates and spurred on our intense Catan showdowns. All of these friends have helped me keep all aspects of life in perspective during grad school, all while maintaining their incredible, brilliant, and inspiring outlooks. I am grateful to have met you all during grad school.

I would also like to thank the long-time friends and family who have provided endless support and encouragement over video chats and during cross-country visits. I am simultaneously amazed yet completely unsurprised that we have kept in touch over the years despite the long distance. This gives me a resounding confidence that we will never lose touch, which I am so grateful for. Finally, I would like to give a special thanks to my family. Thank you for everything—for raising me, for teaching me, for playing with me, for believing in me. During this journey my family has never waivered in their confidence in me (*or at least they didn't show that to me, which I am so, so grateful for*). They have also never stopped asking when I'm going to graduate. Eventually, I learned to accept that they would never stop asking that question until now when I can finally tell them, "I'm done!"

Beyond the people who have taught, mentored, and supported me, there are

a few others that I would like to acknowledge. The monkeys that I worked with: Dmitri, Feivel, Gypsy Tom, Jalapeno, and Penfield, were incredible subjects. This work, and many, many other scientific contributions would not have been possible without them; their contribution to the scientific community is much appreciated. Lastly, financial support from the U.S. government via the National Science Foundation, the National Institutes of Health, the Department of Defense, and the Department of Veterans' Affairs made this work possible. I thank these funding sources and, of course, the taxpayers that funded them. I hope that the work in this thesis and my training has been a worthwhile investment for the American public.

Publications & Presentations Resulting from Thesis

Chapter 2

Yazdan-Shahmorad*, A., **Silversmith, D. B.***, Kharazia, V., & Sabes, P. N. (2018). Targeted cortical reorganization using optogenetics in non-human primates. *Elife*, *7*. http://doi.org/10.7554/eLife.31034

Silversmith, D. B., Yazdan-Shahmorad, A., Kharazia, V., & Sabes, P. N. (2017). Optogenetic stimulation leads to connectivity changes across sensorimotor cortex in nonhuman primates. Society for Neuroscience Annual Meeting, Washington D.C. (poster)

Silversmith, D. B., Yazdan-Shahmorad, A., & Sabes, P. N. (2016). Targeted cortical reorganization using optogenetics in non-human primates. Society for Neuroscience Annual Meeting, San Diego, CA (poster)

Silversmith, D. B., Yazdan-Shahmorad, A., & Sabes, P. N. (2016). Optogenetic stimulation leads to connectivity changes across sensorimotor cortex in nonhuman primates. Advances in Motor Learning & Motor Control, San Diego, CA. (talk)

Chapter 3

Silversmith, D. B., Ramanathan, D., Lemke, S. M., & Ganguly, K. (2018). Changes in spindle induced spiking activity after motor learning. Society for Neuroscience Annual Meeting, San Diego, CA. (poster)

Abstract

An Exploration of Meso-scale Plasticity

by

Daniel Silversmith

Specialized sensorimotor systems allow us to perform dexterous, graceful movements in complicated, dynamic environments. One remarkable feature of the sensorimotor system is its ability to learn new skills. The neural underpinnings of sensorimotor learning have traditionally been studied at two levels: low-level changes between individual neurons, and high-level changes in behavior. Directly connecting these distant levels of study is challenging. In this thesis, I try to bridge this gap by focusing on plasticity at an intermediate level that considers how neuron populations change with respect to each other, i.e., *meso-scale* plasticity.

In Chapter 2, I study changes in meso-scale connectivity in response to neurostimulation. A large-scale optogenetic interface enabled us to simultaneously stimulate and record population activity across primary somatosensory (S1) and primary motor (M1) cortex. We tracked two measures of network connectivity—one based on responses to focal stimulation and the other based on spontaneous activity patterns. Within minutes of stimulation, the inter-area functional connectivity strengthened. At a finer scale, stimulation led to heterogeneous changes across the network, which reflected the correlations introduced by stimulation-evoked activity, consistent with Hebbian models of synaptic plasticity. This work extends Hebbian plasticity models to meso-scale

circuits.

In Chapter 3, I connect low-level changes in M1 spiking structure to two meso-scale oscillations—spindles and slow oscillations (SOs), which are thought to support sensorimotor learning and memory consolidation. During spindles, individual neurons fired at a preferred phase of spindle cycles and neuron pairs synchronized activity during spindle peaks, signifying an increase in pairwise correlations and local functional connectivity. We found a direct relationship between the temporal proximity of SO and spindles (which are thought to interact), and changes to the distribution of spike correlations; closer oscillations were associated with narrowing of the distribution of correlations, with a reduction in low- and high-correlation pairs. Such narrowing is consistent with exploration of novel neural states and may be a key mechanism through which the interaction of meso-scale oscillations can support sensorimotor consolidation.

Throughout this thesis I advocate for studying and modeling neuron populations. I provide insight into meso-scale plasticity through a direct study of meso-scale connectivity changes and by relating the meso-scale to coordinated spiking activity.

Table of Contents

List of Figures

List of Tables

Chapter 1
Introduction

1.1 Synaptic Plasticity

Historically, much of the progress that has been made studying connectivity and learning in the brain has circumvented the complexity of the nervous system by focusing on changes in synaptic efficacy. In these experiments, performed in slice and *in vivo*, researchers focused on a single neuron or even a single synapse at a time, exerting precise control over the structure of inputs and outputs at specific synapses. This approach muted much of the complexity that arises in an interconnected network during learning and enabled researchers to model precise, low-level learning rules. Despite their apparent simplicity, the plasticity rules that this body of work has discovered reflect important principles for understanding connectivity and learning at all levels of the nervous system, including the meso-scale. With the goal of extrapolating from synaptic plasticity to meso-scale plasticity in mind, I go through several seminal synaptic plasticity experiments and then draw broad conclusions from this work, which I will use as a framework to investigate and understand changes in meso-scale connectivity.

1.1.1 Measuring Synaptic Strength

Experimentalists have enlisted a few methods to determine synaptic efficacy or strength. The most widely used methods are direct measures of synaptic strength and require simultaneous recordings from presynaptic spiking and postsynaptic currents (e.g., Bi & Poo 1998; Song et al. 2000). These currents are often measured in short time windows following presynaptic spikes (spontaneous spikes or induced with electrical stimulation) and can be depolarizing or hyperpolarizing depending on the synapse type. Postsynaptic currents have consistent responses (G. Q. Bi & Poo, 1998), and researchers have

characterized these currents in terms of their size and timing after presynaptic spikes. Often the maximum amplitude is used as a proxy for synaptic strength, but slope—a measure that takes into account both amplitude and timing—is also used as a measure of synaptic strength (G. Bi & Poo, 2001; G. Q. Bi & Poo, 1998; S. Song et al., 2000).

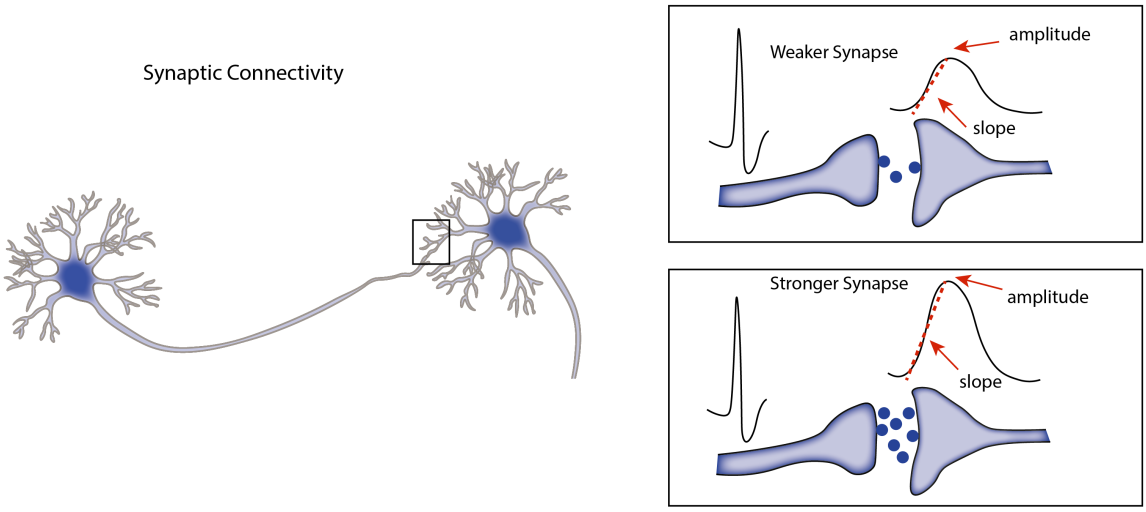

Figure 1.1 – Cartoon of Synaptic Strength Measurement. After presynaptic spikes, an excitatory postsynaptic current is characterized. The amplitude or slope of this current is used as a measure of synaptic strength.

1.1.2 Synaptic Plasticity Rules

Cellular plasticity has long thought to be driven, in part, by an associative process in which the consistent contribution of one neuron to another neuron's spiking leads to their strengthening (Hebb, 1949). This type of associative learning is often referred to as Hebbian learning or Hebbian plasticity, since it was first proposed by Hebb in 1949. Scientific investigation into Hebbian plasticity took hold several decades later; this exploration initially discovered that short, tetanic bursts of stimulation for a few minutes could strengthen synapses for hours or even days (T. V. P. Bliss & Lømo, 1973). Experimentalists later determined that tetanic input stimulation induced postsynaptic

spikes, and input-output spike pairing on the order of ~100 ms was an important factor in changing synaptic weights (Baranyi & Fehér, 1981; Gustafsson, Wigström, Abraham, Huang, & Wigstrom, 1987). A few studies also detected a consequence of spike order; these studies demonstrated that if presynaptic inputs preceded postsynaptic outputs then synapses were strengthened, but if presynaptic inputs followed postsynaptic outputs then synapses were weakened (D Debanne, Gähwiler, & Thompson, 1994; Dominique Debanne, Gähwiler, & Thompson, 1997; Levy & Steward, 1983; Markram, Lübke, Frotscher, & Sakmann, 1997). This work was later elaborated and characterized (G. Q. Bi & Poo, 1998) and ultimately termed spike timing dependent plasticity (STDP; Song et al. 2000; Bi & Poo 2001).

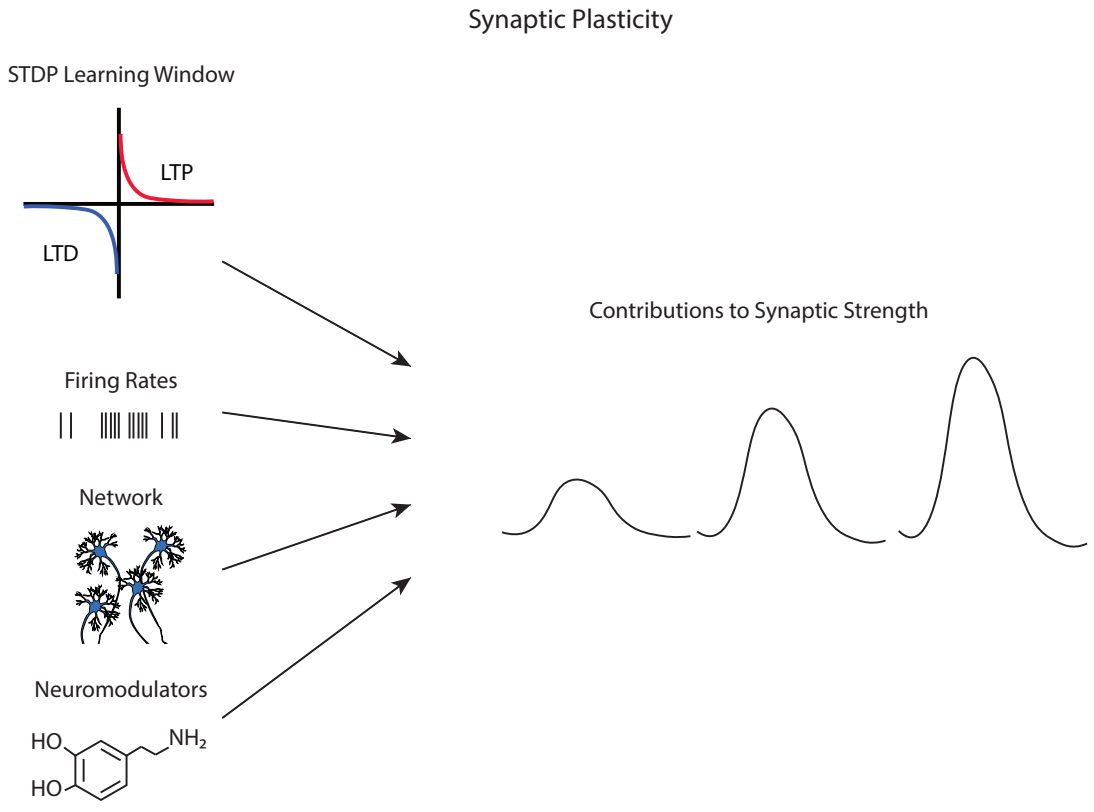

Figure 1.2 – Holistic View of Synaptic Plasticity. STDP, along with other factors contribute to modify the strength of synapses across networks of neurons.

However, even at single synapses, STDP is not so simple. Current models of STDP incorporate a spike timing dependent learning window as one of many factors that influence synaptic plasticity (Feldman, 2012). STDP learning windows are highly dependent on other factors such as firing rates (Markram et al., 1997; Sjöström, Turrigiano, & Nelson, 2001), the location of synapses on dendrites (Sjöström et al., 2001), and baseline synaptic weights (G. Q. Bi & Poo, 1998; Morrison, Diesmann, & Gerstner, 2008; Sjöström et al., 2001). Moreover, other network activity, neuromodulators, and brain states, all of which are constantly fluctuating, are known to modify STDP learning windows (Cassenaer & Laurent, 2012; González-Rueda, Pedrosa, Feord, Clopath, & Paulsen, 2018; Markram et al., 1997; McNaughton, Douglas, & Goddard, 1978; Pawlak & Kerr, 2008; Seol et al., 2007; Shen, Flajolet, Greengard, & Surmeier, 2008; Sjöström et al., 2001).

1.1.3 Functional Connectivity & Plasticity

Studying synaptic plasticity has traditionally relied on fine control of the inputs and outputs to an individual neuron or even an individual synapse. This high degree of control has enabled precise estimates of STDP learning windows, but these experiments are limited in some important ways. The main downside of this paradigm is that maintaining control of a synapse tends to be difficult for long periods of time. This makes plasticity experiments challenging, especially *in vivo*, as animals behave and learn new skills. Ultimately, understanding synaptic plasticity is inconsequential unless it can be connected to learning, development, and overall function. A more practical approach is to measure functional connectivity rather than synaptic strength.

Functional connections reflect the likelihood that activity from one recording site

influences activity at another recording site. Importantly, functional connectivity can be measured at different levels (e.g., connectivity between neurons or between brain regions). Focusing on the cellular level, the functional connectivity from 'Neuron A' to 'Neuron B' can be defined as the likelihood that 'Neuron B' spikes whenever 'Neuron A' has just fired. This is similar to the synaptic connectivity between 'Neuron A' and 'Neuron B', but there are a few important differences that I will highlight. First, instead of directly measuring the synaptic strength via intracellular postsynaptic currents, functional connectivity infers the connection strength by analyzing the statistics of observed neural activity (activity can be collected during spontaneous sessions, responses to stimuli, or after neural stimulation). Additionally, the functional connectivity between 'Neuron A' and 'Neuron B' represents the aggregate connectivity from many synapses and can include alternative paths from 'Neuron A' to 'Neuron B' (e.g., 'Neuron A' -> 'Neuron C' -> 'Neuron B').

Despite its lack of precision, functional connectivity has proven useful for studying neuroplasticity. At the level of synapses, we have a clear model for how changes in activity should lead to changes in synaptic strength (see section *1.2.2 Synaptic Plasticity Rules*). A few studies have combined synaptic plasticity models and functional connectivity measurements to test plasticity models *in vivo* and examine the effect of functional plasticity on behavior. In some seminal work, Jackson et al., used closed-loop electrical stimulation to induce connectivity between two unrelated recording channels on a multielectrode array (Jackson, Mavoori, & Fetz, 2006). In this study, stimulation at one site in M1 was triggered on the recorded spiking activity at another site in M1. After hours of repeated conditioning stimulation, they observed an increase in the functional connectivity between the two sites. Importantly, their increased functional connectivity affected the

tuning of motor output. Even though functional connectivity does not necessarily imply synaptic connectivity, this work used synaptic plasticity models to inform the timing of stimulation and the changes in functional connectivity that they observed were consistent with the expected changes in synaptic connectivity.

Other experiments have complemented this seminal study in an open loop paradigm, where functional connectivity between two sites was increased by repeatedly stimulating at the two cortical sites with a consistent time lag (J M Rebesco & Miller, 2011; J M Rebesco, Stevenson, Kording, Solla, & Miller, 2010). Functional plasticity matched the expected (unmeasured) synaptic plasticity, and resulted in a decrease in sensory detection thresholds, further linking functional connectivity to behavior. Both of these experiments, and subsequent ones using similar functional plasticity paradigms (Timothy H Lucas & Fetz, 2013; Nishimura, Perlmutter, Eaton, & Fetz, 2013; Seeman, Mogen, Fetz, & Perlmutter, 2017), relied on STDP models to inform the timing of conditioning stimulation and then used functional connectivity and behavior to measure neuroplasticity.

1.2 Sensorimotor Skill Learning

Neuroplasticity has been the focus of many experiments because it has the potential to mechanistically explain how animals and humans come to learn and adapt their behavior appropriately over time. There are many types of learning that humans and animals are capable of undertaking. Broadly, these types of learning can be organized into two categories: declarative learning and procedural learning (Dayan & Cohen, 2011; L. R. Squire & Zola, 1996; Larry R. Squire, 2004; Turner, Crossley, & Ashby, 2017). Declarative learning includes explicit spatial and episodic memory tasks (likely hippocampus-

dependent), whereas procedural learning covers tasks that require implicit acquisition (e.g., sequence learning or small visuomotor rotations; likely hippocampus-independent). While neuroplasticity can be studied across a variety of learning tasks, I focus on plasticity during sensorimotor skill learning, which has been studied extensively across a wide range of species. Acquiring sensorimotor skills likely relies on both declarative and procedural learning systems but may be biased toward one or the other at different learning stages.

One of the principal goals of the nervous system is to aid in the interaction of an organism, or more broadly a species, with its environment. Organisms need to sense changes in their environment and react to those changes accurately and swiftly in order to survive. Many species across the animal kingdom (and outside of it) have accordingly evolved specialized sensorimotor systems that provide organisms with the information and capability to effectively interact with their environments (e.g., Sachs 1988; Lovejoy et al. 2010; Jones & Teeling 2006; Jon H. Kaas 2004). In humans, one of the most investigated sensorimotor actions is the process of using one's hand to retrieve a nearby object (visually guided reaching and grasping). This seemingly simple motor action relies on multiple sensory systems for coordinated, dexterous control including visual feedback about the location and shape of the targeted object, proprioceptive information on the state of the human body (in particular the arm and hand), and somatosensory feedback when the hand touches and manipulates the object (for a review, see Stone & Gonzalez 2015). The integration of different sources of feedback with movement planning is crucial for behavior and is an area of intense study, but our understanding of the neural basis for such sensorimotor integration is still in its early stages.

1.2.1 Behavioral Markers of Sensorimotor Skills

Learning is often task- or goal-specific, but can be characterized across particular behaviors as general improvements in accuracy, speed, and consistency (Dayan & Cohen, 2011; Papale & Hooks, 2018; Ramanathan, Gulati, & Ganguly, 2015). The overall time-course of skill learning is relatively slow with task performance continuing to improve across multiple training sessions; however, initial skill acquisition tends to be fast, with rapid performance gains immediately after introduction to the task. Later skill acquisition, on the other hand, is slower, with more modest performance gains. Contributions to skill learning can be split into online contributions—within-session, or offline contributions—across or between training sessions. Importantly, the two stages of skill learning, and the different online and offline processes likely have different neural bases (Papale & Hooks, 2018).

1.2.2 Sensorimotor Anatomy

Even relatively simple actions are performed in the context of a complex, dynamic environment. To account for this variability, many organisms have developed intricate, complicated sensorimotor systems, which are distributed across a large, interconnected network of brain regions. In humans, a large portion of cortex is involved in sensorimotor behaviors including primary sensory cortex (S1), higher somatosensory areas (S2), primary motor cortex (M1), premotor cortex, higher premotor areas (e.g., the supplementary motor area and cingulate motor cortex), parietal cortex such as the parietal reach region (PRR), and others, but this network also includes important subcortical regions such as the thalamus, striatum, and cerebellum (Jon H. Kaas, 2004; Jon H Kaas,

2004; Stepniewska, Preuss, & Kaas, 1993).

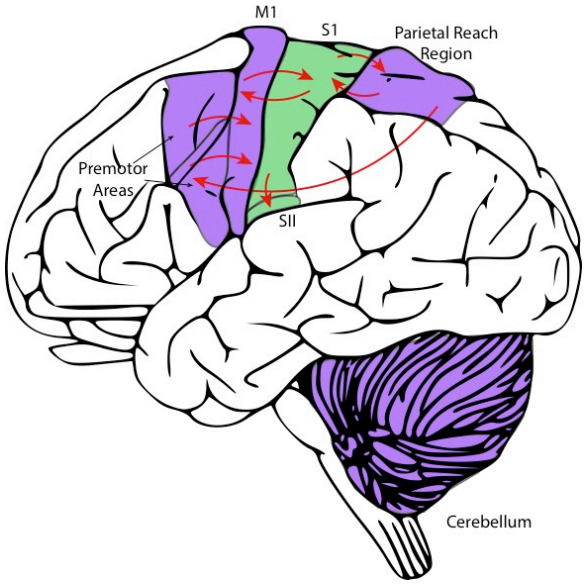

Figure 1.3 - Human Brain with some Sensorimotor Regions Highlighted. Green areas are nominally somatosensory and purple areas are nominally motor. Red arrows reflect a high degree of connectivity.

Given this complexity, the full sensorimotor system is incredibly difficult to study; in this work, I therefore focus on a subset of the sensorimotor network—S1 and M1. These areas both act as gatekeepers; somatosensory input primarily enters cortex via the thalamus and then S1 (E. G. Jones & Friedman, 1982; E. G. Jones & Powell, 1969, 1970; Jon H Kaas, 2004), whereas the motor output predominantly exits cortex through M1 (Fetz & Cheney, 1978; Graziano, Aflalo, & Cooke, 2005; Jon H Kaas, 2004). M1 and S1, like the rest of the sensorimotor system, are highly interconnected (Jon H Kaas, 2004; Mao et al., 2011; Papale & Hooks, 2018). Moreover, both areas show marked neuroplastic changes during motor skill learning (Greenough, Larson, & Withers, 1985; Kleim, Barbay, & Nudo, 1998; Nudo, Milliken, Jenkins, & Merzenich, 1996; Papale & Hooks, 2018). Finally, both of these areas are arranged side-by-side (Figure 1.4) and are accessible across different animal models (Figure 1.4). All of these properties make it feasible and worthwhile to study S1 and M1. Although I will focus on these areas, I believe that some of the principles of

sensorimotor connectivity and plasticity, which I explore in this work, provide useful insight into the broader sensorimotor system.

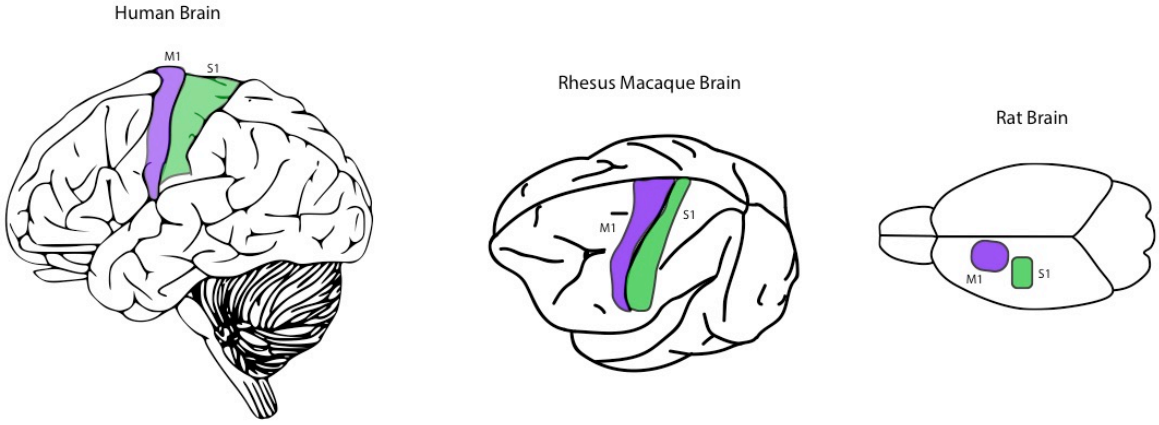

Figure 1.4 – Brains of Different Mammalian Species. Primary motor cortex (purple) and primary somatosensory cortex (green) highlighted.

1.2.3 Representational Plasticity in Sensorimotor Cortex

Much of the early work investigating changes in sensorimotor cortex relied on the stable functional relationships between individual neurons and aspects of behavior. Merzenich, and others, pioneered this work using somatosensory receptive fields in S1 (Buonomano & Merzenich, 1998; Clark, Allard, Jenkins, & Merzenich, 1988; M M Merzenich et al., 1983; Michael M. Merzenich et al., 1984; Michael M. Merzenich, Kaas, Sur, & Lin, 1978; Michael M. Merzenich & Jenkins, 1993; Recanzone, Merzenich, & Dinse, 1992a). Since it was, and still is, difficult to track neurons over long periods of time, initial studies focused on gross changes in receptive field maps across the cortex. Specifically, they focused on the cortical surface area in S1 that responded to tactile stimulation at the periphery. For example, they found the areas of S1 that responded to each of the digits on the hand. Then they intervened by amputating a digit, cutting afferent peripheral pathways, or sewing two digits together. They found that when a digit was amputated or afferent pathways were

transected, the neurons with interfered receptive fields were initially unresponsive but over time responded to the surrounding digits. When digits were sewed together they found that neurons that were initially responsive to one digit became responsive to both digits. This work was replicated across other of sensory systems (J H Kaas et al., 1990; Robertson & Irvine, 1989) and provided the basis for our current understanding that sensorimotor representations are not stable and that the neural modifications mirror changes at the periphery or even changes in behavior. More recent work has followed up on the Merzenich studies using rodent models and whisker plucking in barrel cortex (rodent S1). Intriguingly, these experiments demonstrated similar changes in cortical representation (Bender et al., 2006a; Feldman, 2009; Li et al., 2009), but combined the traditionaly methods with modern pharmacological interventions, which allowed them to discover that these changes in cortical representation are STDP-dependent.

Sensorimotor plasticity has also been characterized by analyzing changes in structure and spiking as animals learn sensorimotor skills. This work has revealed a large variety of changes in M1. For example, during skill learning, studies have used imaging techniques to find increases in synaptic growth that stabilize (Xu et al., 2009; Yang, Pan, & Gan, 2009), dynamic changes in single neuron variability (Kao, Doupe, & Brainard, 2005; Mandelblat-Cerf, Paz, & Vaadia, 2009; Zacksenhouse et al., 2007), and the emergence and stabilization of shared variability across neurons during slow skill learning (Athalye, Ganguly, Costa, & Carmena, 2017). While learning, population measures show that there is an increase in the proportion of shared variability across neurons. Intriguingly, there are also meso-scale markers of skill learning. In particular, low frequency oscillations in LFP have been shown to increase in power and stereotypy across trials as animals learn motor

tasks (Ramanathan et al., 2018). Intriguingly, neural stimulation has been shown to boost low frequency oscillations and improve task performance in animals that are recovering from stroke (Ramanathan et al., 2018).

1.3 Defining the Meso-scale

1.3.1 Recording Methods

Before understanding meso-scale plasticity, it is necessary to define the meso-scale in more detail. Neural activity can be recorded at different spatial scales, and the spatial resolution of the recording technology determines the number of neurons that contribute to each measurement. Technologies such as electroencephalography (EEG) record activity across ~ 6 cm^2 of neural activity; in humans this corresponds to roughly 100 million gyral neurons (Schwartz, Cui, Weber, & Moran, 2006). At the other end of the spectrum, patch-clamp, whole-cell recordings measure electrical fluctuations within a single neuron. The range of common electrophysiology recording methods is summarized in Figure 1.5, which has been adapted from (Schwartz et al., 2006).

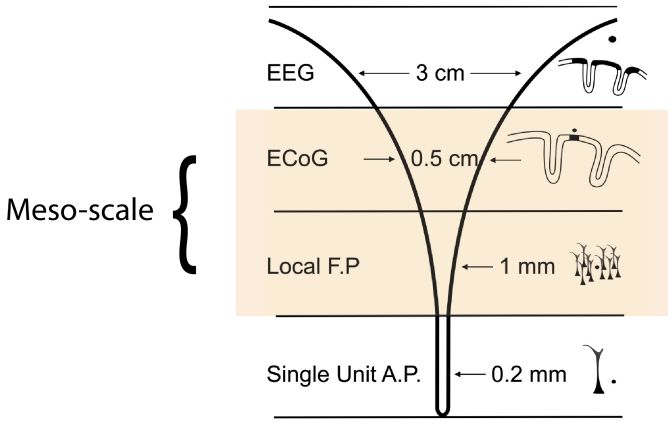

Figure 1.5 – Different Spatial Scales of Common Electrophysiology Recording Technologies. Adapted from (Schwartz et al., 2006) to specify technologies that record at the meso-scale (highlighted section).

The meso-scale refers to a midrange of neural populations of activity on the scale of organized neural populations. Meso-scale technologies maintain a higher spatial resolution than that of EEG, and a lower resolution than that of extracellular single or multiunit recordings, but the exact cut-offs have been loosely defined. For our purposes, I will define a range of recording technologies, from extracellular LFP to surface-level ECoG, which captures meso-scale activity. Importantly, within these technologies there can be a large degree of variability, particularly in ECoG recordings. ECoG grids that are used clinically have 10 mm spaces between electrodes and integrate activity across more neurons than ECoG grids that are used for research use electrode pitches of about 4 mm (Muller, Hamilton, Edwards, Bouchard, & Chang, 2016; Schwartz et al., 2006). These different spatial scales are all within the range of meso-scale plasticity but as we learn more about neural activity at the meso-scale we may need to make distinctions between these scales or even the different types of technologies that are used.

It is important to note that the technologies mentioned here do not form a comprehensive list of meso-scale neural recording technologies. Imaging modalities (such as function magnetic resonance imaging—*fMRI* or fiber photometry) and magnetic recording modalities (such as magnetoencephalography–*MEG*) measure population activity at comparable spatial scales and can be used to study meso-scale plasticity (though they have other limitations). In this work I focus on using electrophysiology methods due to their high temporal resolution, and their similarity across different special scales. These features enable us to relate meso-scale findings to the well-studied field of synaptic plasticity.

1.3.2 Oscillatory Neural Activity

In our discussion of synaptic plasticity we have treated spikes as significant, discrete events that influence plasticity. Moving beyond individual neurons, the aggregated activity across neural populations looks very different. Rather than discrete events, we record continuous signals; because these signals are continuous, it is difficult to extract meaning from them, but a few standardized techniques have been developed to identify important signals from background noise. First, meso-scale activity is broken down according to its frequency content using standard signal processing approaches. Biological signals (as do many other natural phenomena) have a stereotypical noise profile (He, Zempel, Snyder, & Raichle, 2010; Miller, Sorensen, Ojemann, & den Nijs, 2009; Voytek et al., 2015), where the signal power drops with increasing frequency; this noise is often referred to as *one-over-f* noise, since the spectral power can be modeled according to the following equation:

$$power = \frac{1}{freq^{(\alpha)}}$$

Deviations in spectral content from this model are typically used to identify frequencies containing biological signal. Importantly, frequency content in a signal varies dynamically across time. Sustained epochs of increased power in a specific frequency band are referred to as neural oscillations and are useful for neural markers that often relate to other known quantities such as spiking activity and behavior.

Comparing oscillations to spiking activity, brain states, and behavior has yielded a few significant insights into the role of meso-scale activity. First, oscillations in different frequency bands operate at different spatial gradations; lower frequency oscillations are far-reaching whereas higher frequency oscillations tend to be more local (Łęski, Lindén,

Tetzlaff, Pettersen, & Einevoll, 2013; Muller et al., 2016). High gamma power (~60 – 200 Hz) has even been shown to correlate strongly with local spiking activity (Liu & Newsome, 2006). Despite the differing spatial scope across the frequency spectrum, there is organization within oscillatory activity; lower frequency oscillations tend to be coupled with higher frequency oscillations (Cox, van Driel, de Boer, & Talamini, 2014; He et al., 2010). The power of different oscillatory components, and the coupling between different frequency bands, is not stable; rather it changes with external stimuli (Fries, Schröder, Roelfsema, Singer, & Engel, 2002; Liu & Newsome, 2006), brain state (He et al., 2010; Voytek et al., 2015), and long-term factors such as age (Voytek et al., 2015).

All of these observations have spurred a contentious debate about whether meso-scale oscillations are epiphenomenal (purely a result of spiking activity) or whether the oscillations themselves might influence spiking activity (Buzsáki, Anastassiou, & Koch, 2012). Without weighing in on this debate, it is important to acknowledge that boosting, inducing, and perturbing ongoing oscillations can have dramatic effects on spiking activity, behavioral performance, and offline memory consolidation (Binder, Rawohl, Born, & Marshall, 2014; Gulati, Guo, Ramanathan, Bodepudi, & Ganguly, 2017; Latchoumane, Ngo, Born, & Shin, 2017; Ramanathan et al., 2018). Regardless of the direction of causality, recording and generating oscillations is useful for understanding coordinated brain activity and complements observations of individual neurons and single synapses.

1.3.3 Meso-scale Connectivity

As discussed in *Section 1.1.3 Functional Connectivity & Plasticity*, functional connectivity lends itself well to studying connections at the meso-scale. Rather than

reporting large-scale anatomical or structural connections, which can be investigated though methods such as diffusion tensor imaging–*DTI*, functional connectivity focuses on co-activations of population activity. When recording activity at the meso-scale, functional connectivity maintains its definition and purpose—functional connectivity draws inferences from neural activity to measure the degree of influence that recording sites have over each other. When recording meso-scale activity, functional connectivity is inherently a measurement of the influence that neural populations have over each other.

Meso-scale connectivity is typically measured by borrowing a few more tools from the standard signal processing toolbox. Fundamentally, these methods fall into two categories: measuring activity correlations and measuring responses to interventions such as stimulation. Activity correlations can be measured in a variety of ways. The most common techniques decompose time-varying signals at two recording sites into their spectral content (amplitude and phase across time). Then measurements of correlated power across time (Fukushima & Sporns, 2018), phase locking, or cross-frequency phase-amplitude coupling can be used to infer functional connections (Bastos & Schoffelen, 2015). More complicated methods can be used to try and predict causality or account for other observed or unobserved activity (Bastos & Schoffelen, 2015), but these methods still rely on models that are built using observed activity correlations. To actually impute causation between recording sites, stimulation or other causal intervention is necessary. Neurostimulation has been used in conjunction with extracellular recordings to measure functional connectivity between a stimulation site and a recording site (Atsushi Iriki, Pavlides, & Keller, 2017; Keller et al., 2014; Matsumoto, Kunieda, & Nair, 2017; Seeman et al., 2017).

1.4 Thesis Goals and Chapter Previews

The immense complexity of synaptic plasticity and its relation to sensorimotor learning is overwhelming, and it is difficult to imagine how low-level changes percolate up to large-scale network reorganization to support sensorimotor learning. This is why it is critical to study plasticity at multiple levels. The work in this dissertation is motivated by the following open and unanswered questions—*Is there a set of learning rules for meso-scale plasticity? How do changes in meso-scale plasticity support sensorimotor learning?* These questions are broad and general, but they are also significant and far-reaching inquiries that have inspire this work. While trying to uncover answers, there are few important principles that previous neuroplasticity experiments have revealed, which I will use as a framework to think about and guide the analysis of plasticity at the meso-scale: (1) timing matters – the precise timing of neural activity leads to dramatically different outcomes in synaptic plasticity; (2) the rest of the network matters – cooperativity among other neurons changes the operating regime of STDP; and (3) brain state matters – neuromodulators and different states such as sleep vs. awake have a huge impact on learning windows. Therefore, in studying plasticity at any scale it is crucial to keep track of the precise, relative timing of neural activity across the network under study. In drawing conclusions, it is essential to think through the full impact of measured connectivity changes in the context of a network with many moving parts including a network that operates differently across a diverse range of brain states.

With these guiding principles in mind, I will use two different, but complementary paradigms to investigate meso-scale plasticity rules. In *Chapter 2*, I will explore functional connectivity and plasticity using artificial neurostimulation; this paradigm is a natural

extension of the functional plasticity literature outlined in section *1.2.3. Functional Connectivity & Plasticity*. In this work, I will measure functional connectivity at the meso-scale, apply conditioning stimulation, and analyze how functional connectivity changes in response to the conditioning stimulation. I will then use simple, linear models to predict plasticity from baseline connectivity and stimulation-induced activity correlations across the sensorimotor network. In *Chapter 3*, I will switch gears and study plasticity in the context of natural motor learning. Specifically, I will simultaneously measure LFP oscillations and spiking activity, and connect precise spike timing to ongoing oscillations during sleep. Using what we know about synaptic plasticity rules, I will use what we know about synaptic plasticity rules to interpret changes in spike correlations, and then I will relate these changes to the meso-scale oscillations that are thought to support learning.

Chapter 2
Using Low Frequency Stimulation to Probe Meso-scale Connectivity & Plasticity

2.1 Foreword

In this chapter, I extend synaptic plasticity paradigms to the meso-scale. Using a combination of optogenetic stimulation and micro-electrocorticography, we are able to simultaneously stimulate the brain while recording population activity across the sensorimotor network. First, we develop two methods to measure functional connectivity across S1 and M1 and see how these measures compare and contrast. Next we track the evolution of both measures of meso-scale functional connectivity during different neurostimulation patterns. The stimulation-induced changes are then interpreted in the context of synaptic plasticity models. Moreover, the stimulation-induced changes act differently across our two connectivity measures allowing us to speculate about the role of plasticity in stable and unstable networks. Though the stimulation and plasticity in this experiment are artificial, this work provides critical insight into how plasticity functions at the meso-scale.

2.2 Introduction

Many neurological and psychiatric disorders arise from dysfunctional neural dynamics at the network level, which in turn stem from aberrant neural connectivity (DeSalvo, Douw, Tanaka, Reinsberger, & Stufflebeam, 2014; Edwardson, Lucas, Carey, & Fetz, 2013; Skudlarski et al., 2010; Stam, 2014; Wu et al., 2016; Yahata et al., 2016). The brain shows marked plasticity across a variety of learning and memory tasks (T. V Bliss, Collingridge, & Morris, 2014; Takeuchi, Duszkiewicz, & Morris, 2014) and during recovery after brain injury or stroke (Edwardson et al., 2013; Hara, 2015; Murphy & Corbett, 2009), and many have proposed to take advantage of this innate plasticity to treat neural

disorders (Edwardson et al., 2013; Jackson et al., 2006; Timothy H Lucas & Fetz, 2013; James M. Rebesco, Stevenson, Koerding, Solla, & Miller, 2010). In principle, brain stimulation protocols can be designed to leverage this plasticity in order to rewire aberrant neural connectivity, potentially curing these disorders. Implementing such treatments requires a better understanding of how stimulation-induced plasticity drives changes in network connectivity and network dynamics.

The simple Hebbian model of plasticity (DO, 1949) and spike-timing dependent versions of it (G. Bi & Poo, 2001) explain a large body of data on activity-dependent plasticity, including both in vitro (Abrahamsson, Lalanne, Watt, & Sjostrom, 2016; Massobrio, Tessadori, Chiappalone, & Ghirardi, 2015), and in vivo (Andersen, Krauth, & Nabavi, 2017; Feldman, 2012; Shulz & Jacob, 2010) studies. Despite the extensive work studying Hebbian plasticity at the cellular level, it remains unclear how synaptic plasticity leads to large-scale functional reorganization. Recently, several studies have shown large-scale plasticity following brain stimulation that is consistent with Hebbian mechanisms (Jackson et al., 2006; Lajoie, Krouchev, Kalaska, Fairhall, & Fetz, 2017; Timothy H Lucas & Fetz, 2013; James M. Rebesco et al., 2010; W. Song, Kerr, Lytton, & Francis, 2013). In particular, Fetz and colleagues implemented an activity-dependent stimulation protocol, effectively introducing an artificial connection between two sites in the motor cortex (Jackson et al., 2006; Timothy H Lucas & Fetz, 2013). Continuous reinforcement of this artificial connection led to a stable functional change in stimulation-evoked movements, indicating that stimulation induces large-scale plasticity. Recently these results were reproduced in a modeling work at the network level (Lajoie et al., 2017). Similar results have been observed using open-loop stimulation protocols to induce targeted plasticity

between two cortical sites (James M Rebesco & Miller, 2011; Seeman et al., 2017). Notably, these papers have reported off-target effects, but these were either interpreted as global changes in excitability or were not explained. Although the results from closed-loop and open-loop experiments suggest large-scale plasticity, the underlying neural network changes remain unexplored.

Here, we measure connectivity across sensorimotor cortex and track changes in network connectivity in response to open-loop stimulation. This work takes advantage of a large-scale optogenetic interface (Yazdan-Shahmorad et al., 2016) that enables us to simultaneously stimulate populations of excitatory neurons while recording large-scale micro-electrocorticography (μECoG) activity across two brain areas, S1 and M1. We first establish and compare two measures of functional connectivity that provide complementary views of the mechanisms of plasticity. Next, we investigate how stimulation impacts connectivity between and within cortical areas. We then test whether stimulation-evoked activity drives large-scale network plasticity in a Hebbian manner. The goal of this work is to investigate large-scale functional reorganization following stimulation, which will inform future neurorehabilitation strategies.

2.3 Methods

2.3.1 IACUC

Two adult male rhesus monkeys (Monkey G: 8 years old, 17.5 kg; Monkey J: 7 years old, 16.5 kg) were used in this study. We used the same animals and interface published in (Yazdan-Shahmorad et al., 2016). All procedures were performed under the approval of the University of California, San Francisco Institutional Animal Care and Use Committee and

were compliant with the Guide for the Care and Use of Laboratory Animals.

2.3.2 Optogenetic Interface

Here, we give a brief summary of our large-scale optogenetic interface (details in Yazdan-Shahmorad et al. 2016). We combined three existing techniques to implement a practical, large-scale interface for both manipulation and recording from the surface of the brain in two rhesus macaque monkeys. We used an efficient technique for infusion of the optogenetics viral vector (AAV5.CamKIIa.C1V1(E122T/E162T).TS.eYFP.WPRE.hGH, 2.5×10^{12} virus molecules/ml; Penn Vector Core, University of Pennsylvania, PA, USA, Addgene number: 35499) into primary somatosensory (S1) and motor (M1) cortices based on convection-enhanced delivery (CED). We infused 200 µl of virus in four sites (two in M1 and two in S1) in Monkey G and 250 µl in five sites (two in M1 and three in S1) of Monkey J. Infusion rate started at 1 µl/min and was increased to 5 µl/min. Following infusion we used an artificial dura to protect the brain while maintaining optical access. We verified expression with epifluorescent imaging. To record the evoked responses in both M1 and S1, we used either one or two 96 channel micro-electrocorticography (µECoG) arrays that were designed to allow minimally attenuated (Ledochowitsch et al., 2015) optical access. In both animals, we observed reliable light evoked neural responses from the large channel-expressing areas.

2.3.3 Data Acquisition

Optical stimulation was applied using a fiber optic (core/cladding diameter: 62.5/125 um, Fiber Systems, TX, USA) that was connected to a 488 nm laser (PhoxX 488-

60, Omicron-Laserage, Germany) and was positioned above the array (Figure 2.1). We used a Tucker-Davis Technologies system (FL, USA) for μECoG recording and to control the laser stimulus. We chose the locations of stimulation based on the results of our epifluorescent imaging (see Yazdan-Shahmorad et al. 2016) and whether we got a physiological response to light stimulation. Based on detailed histological and electrophysiological analyses, (see Yazdan-Shahmorad et al. 2016) we estimated that there was a uniform distribution of expression around our infusion sites and at the locations that we chose to stimulate. The monkeys were awake sitting in primate chairs for the duration of experiments. To ensure that the monkey remained awake, we provided random reward as well as audio and visual stimuli (displaying a cartoon). In addition, we visually monitored the animal throughout the experiment.

2.3.4 Stimulation & Recording Protocol

Our stimulation and recording paradigms consisted of baseline recording, testing and conditioning blocks (Figure 2.5A and 2.12A). During baseline recording blocks we collected 5 minutes or 30 seconds of baseline activity. In test blocks we delivered 100 light pulses (5 ms duration at 5 Hz) to each laser; pulses were alternated between laser sites in blocks of 10 pulses each in order to reduce the effects of habituation. For the conditioning blocks we stimulated through one or two lasers at a frequency of 5 or 7 Hz for 10 minutes. We repeated this protocol 5 times for each experiment. Our experiments consisted of single- and two-site stimulation with different laser configurations: (1) one or two lasers in M1, (2) one or two lasers in S1 and (3) one laser in M1 and one laser in S1. For two-site stimulation, the delay between the two lasers was 10, 30, 70 or 100 ms. We also included

control sessions in which we kept the structure of passive recording and active testing blocks, but did not stimulate during the conditioning blocks.

	Monkey G	Monkey J
Number of sessions	37	33
Number of sessions analyzed	29	15
Number of control sessions	3	2
Number of single-site and long latency sessions	6	9
Number of short-latency sessions	20	4

Table 2.1 - Summary of Analyzed Sessions Broken Down by Monkey. The data were collected in two to three week periods for each animal. Depending on the health of the animal and the quality of the neural recordings, one to four experiments were performed per day.

2.3.5 Data Analysis

5.2.3.1 Preprocessing

All processing and statistical analyses were performed using custom MATLAB (MA, USA) code. After signal acquisition, broadband surface potentials (sampled at 24kHz) were visually inspected, and faulty recording sites were removed from further analysis. Next, stimulation-triggered responses were visually inspected for photoelectric artifacts, which were characterized by their timing and amplitude; (see Yazdan-Shahmorad et al. 2016; Ledochowitsch et al. 2015) and subsequently removed from further analysis. Next, two measures of connectivity were defined to quantify the strength of functional connectivity between M1 and S1.

5.2.3.2 Stimulus evoked response ratio

During test blocks, we delivered 100 stimulus pulses at 5 Hz through each laser. To capture high-fidelity timing of the signal and avoid significant phase distortions we

acausally band-pass filtered [0.1 - 500 Hz] the broadband signal. We first calculated the delay between the onset of the stimulus pulse and the response measured at each electrode. We calculated the delay from an average waveform across the distribution of evoked responses (bootstrapped 1000 times). This bootstrapped averaging was less susceptible to artifacts and noise that could dominate the signal on individual trials. The average evoked response lasted no longer than 30ms with the trough of the response occurring in the first 5-15 ms. We defined the delay of the evoked response as the time of the trough within 20 ms of stimulation onset (Figure 2A&2E). The distribution of delays across recording sites (Figure 2B&2F) was bimodal with a delay greater than 1.5 ms between the modes. We manually set a threshold based on this delay distribution to distinguish between "primary" and "secondary" sites, i.e. those in the same area as the stimulation site and those in the other area. Sessions without a secondary response were excluded from further analysis (see Table 2.1).

Next, we calculated the amplitude of each evoked response. We downsampled the raw surface field potentials to 1kHz (after applying a lowpass Chebychev filter for anti-aliasing) and applied an acausal band-pass filter to capture high gamma activity (60 - 200 Hz), which is known to be representative of neural activity of local cortical columns (Suzuki & Larkum, 2017; Yazdan-Shahmorad, Kipke, & Lehmkuhle, 2013). The peak (maximum) and trough (minimum) of the evoked waveform were identified within a 20ms window after each laser pulse (Figure 2.3AE, dashed, inset boxes) and the peak-to-trough difference was computed. The SERR was then defined as the average (across 100 repeated laser pulses) of the ratio of two peak-to-trough amplitudes, the one for the secondary site over the one for the stimulation site. We considered other measures for calculating SERR,

including commonly used measures such as the amplitude or slope of the broadband stimulus-evoked response (data not shown), but the high gamma peak-to-trough amplitude yielded the most robust results.

Changes in SERR were calculated by taking the difference in SERR between the initial and final testing blocks for each secondary site. A paired, two-sided t-test was used to determine whether SERR changes were significant across secondary recording sites. Similarly, a paired, two-sided t-test was applied to determine significance across sessions separately for controls (in which no stimulation was applied during conditioning), stimulation in M1, and stimulation in S1. Unpaired, two-sided t-tests were used to directly compare the change in SERR between all stimulation sessions and all control sessions. P-values less than .05 were considered significant.

5.2.3.3 Coherence

First, broadband surface potentials were downsampled to 1 kHz (after applying a lowpass Chebychev filter for anti-aliasing). Then pairwise coherences were calculated in 10s Hamming windows in 4 Hz frequency bands. The coherence between channels x and y is defined as:

$$C_{xy}(f) = \frac{|G_{xy}(f)|^2}{G_{xx}(f)G_{yy}(f)}$$

where G_{xx} and G_{yy} refer to power spectral density of channels x and y respectively, and G_{xy} refers to their cross-spectral density. For simplicity, in future equations we refer to the coherence between channels x and y as C. Coherence in the theta band (4 - 8 Hz) was used in all analyses, unless otherwise specified.

Changes in inter-area coherence were calculated by taking the difference in

coherence between the initial and final recording blocks. Significant changes after individual sessions were detected with a paired, two-sided t-test across secondary sites. Across sessions, changes in inter-area coherence in each frequency band were compared between stimulation sessions and control sessions with an unpaired, two-sided t-test, applying the Bonferroni correction for multiple comparisons.

5.2.3.4 Connectivity dynamics

Both the SERR and inter-area theta coherence were measured for each baseline recording and each test block during the experiment. To measure connectivity during the conditioning block, the block was sectioned into 100-pulse segments, and the SERR was calculated for each of these segments. To understand the connectivity trends, we pooled experiments together and calculated the average connectivity measures. To account for initial differences in connectivity across sessions and monkeys, we calculated the average changes in each connectivity measure with respect to its initial value and computed their Pearson correlation across sessions and blocks.

5.2.3.5 Network analysis

Since stimulation evokes a network-wide pattern of activity, the coherence during conditioning blocks (C_c) is different than the coherence calculated during baseline recording blocks (C_r). This difference, or "stimulus-evoked coherence" captures the correlations introduced through stimulation. We propose that these correlations drive plasticity in a Hebbian manner, so that the change in baseline coherence between pre- and post-conditioning reflects the stimulus-evoked coherence. Therefore, we assessed how well the change in recording coherence is predicted, across blocks and electrode pairs, by the

stimulus-evoked coherence using linear regression,

$$(C_c - C_{r,pre}) = \hat{B} * (C_{r,post} - C_{r,pre}),$$

where $\hat{B}$ are the fit regression parameters. However, this simple regression analysis is biased, since the same C_r is used on the left and right side of this equation. To avoid the spurious correlations that would therefore arise, we split the recording blocks into 5s non-overlapping windows and calculated the coherence separately for the odd (α) and even (β) windows. We then averaged the coherence for each set of windows ($C_{r,\alpha}$ and $C_{r,\beta}$) and used these as independent measures of baseline coherence for the regression,

$$(C_c - C_{r,pre,\alpha}) = \hat{B} * (C_{r,post,\alpha} - C_{r,pre,\beta}).$$

5.2.3.6 Histology analysis

Monkeys were deeply sedated (per above Surgical procedures) and perfused transcardially with heparinized phosphate buffered saline (PBS) followed by cold 4% paraformaldehyde in phosphate buffer. The brain was extracted and post-fixed in the same fixative for 24 hours at 4°C and then dissected into twelve 6 mm-thick coronal blocks using a custom matrix. After 7-10 days incubation in 30% sucrose, blocks were frozen and cut on a cryostat (Microm, Germany) into 50 μm thick sections. Representative sections were selected from each block and processed for EYFP immunocytochemistry using a free-floating technique. Sections were initially washed in PBS, incubated in 3% hydrogen peroxide in PBS for 10 min to quench endogenous peroxidase activity, and then rinsed in two changes of 50% ethanol followed by three changes of PBS for 5 min each. Next we incubated the sections in 5% normal donkey serum in PBS for 1 hour to block non-specific binding. Primary rabbit polyclonal anti-GFP antibody (Abcam, RRID: AB_303395) was

diluted 1:15,000 in PBS containing 0.01% Triton X-100 and was applied to the tissue for 48 hours at 4oC. Sections were then rinsed in PBS, incubated in biotinylated donkey anti-rabbit antibody, (1:2,000, Jackson Immunoresearch) for 12 hours at 4oC, rinsed and incubated in ExtrAvidin (Sigma-Aldrich, 1:5,000) for 5-6 hours in room temperature. Peroxidase was detected using a diaminobenzidine (DAB) chromogenic reagent (Sigma-Aldrich). Sections were rinsed in PBS, mounted on gelatin-coated slides, air-dried, dehydrated in graded alcohols, cleared in xylene and coverslipped with D.P.X. mounting media (Sigma-Aldrich). Additional adjacent sections were stained with cresyl violet (Nissl) using standard techniques, to reveal cortical cytoarchitecture.

Double immunofluorescence was performed using the similar to above approach although using combinations of primary antibodies from different host species for GFP and for the interneuron markers. Primary antibodies were: rabbit polyclonal anti-GFP antibody (1:10,000, Abcam, GFP ab290); goat polyclonal anti-GFP antibody (1:1,000, Abcam, ab5450). These were paired with one of the following primary antibodies: mouse monoclonal antibody for parvalbumin, (1;1,000, Sigma-Aldrich P3088); mouse monoclonal antibody for calbindin (1:800 Sigma-Aldrich CB-955); rat monoclonal antibody for somatostatin (1;200, Millipore MAP354); rabbit polyclonal anti-GAD65/67 (1:500, Millipore AB1511); a cocktail of mouse monoclonal antibody for GAD67 (1:500, Millipore MAB5406) and mouse monoclonal antibody for GAD65 (1:1,000, Sigma-Aldrich G1166). Sections were blocked using normal donkey serum and processed for 4-6 hours in a mix of matching secondary antibodies; all of which were raised in donkey (Thermo Fisher Scientific, 1:300): anti-rabbit Alexa Fluor 488, anti-goat Alexa Fluor 488 (both for GFP). Interneuron markers were visualized using anti-mouse Alexa Fluor 594, anti-rat Alexa

Fluor 594; anti-rabbit Alexa Fluor 594. Sections were mounted using Vectashield (Vector Labs) and imaged using Zeiss LSM510 Meta confocal microscope (Zeiss, Germany) and 90i imaging system equipped with a CCD camera (Nikon, Japan).

2.4 Results

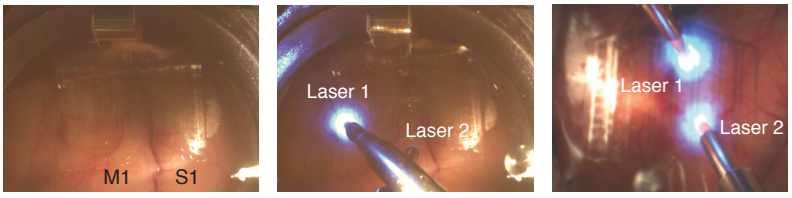

Figure 2.1 - Stimulation and Recording Setup. Photo of the μECoG array placement over M1 and S1 in Monkey G (left panel) and the placement of lasers on top of the array in two different configurations (middle panel: Monkey G and right panel: Monkey J).

We performed optogenetic stimulation via laser illumination of the cortical surface while simultaneously recording surface potentials (μECoG) from about 1.5 cm^2 of primary somatosensory (S1) and motor (M1) cortices (Figure 2.1). The viral vector used to obtain opsin expression targeted excitatory neurons (Figure 2.2). These neurons have strong projections within and between the two brain areas (Kinnischtzke, Simons, & Fanselow, 2014; Murray & Keller, 2011; Weiler, Wood, Yu, Solla, & Shepherd, 2008). Activating these excitatory cells should increase the excitability of the underlying network, improving the likelihood of neuroplastic change (A Iriki, Pavlides, Keller, & Asanuma, 1989). In two macaque monkeys we explored stimulation-induced changes in network connectivity.

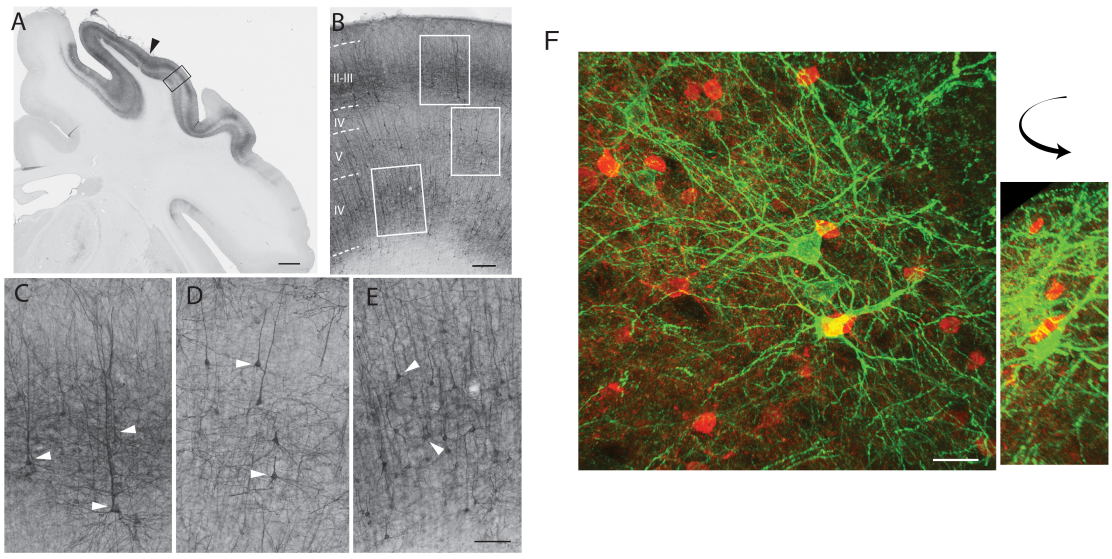

Figure 2.2 - Opsin Expression was Observed Only in Pyramidal Neurons. Using immunohistochemistry, we confirmed that EYFP reporter expression was present in both S1 and M1. YFP-positive cells in these areas were morphologically identified as pyramidal neurons, and were in layers II-III, V-VI. To confirm this we further analyzed the tissue using markers for inhibitory neurons such as parvalbumin, somatostatin and calbindin as well as GAD67/65 and examined cell bodies and axonal terminals at the areas around the site of infusion. We found no evidence for co-localization between EGFP and any of the inter-neuronal markers therefore excluding the possibility for significant opsin presence in cells other than pyramidal neurons, and confirmed that optical stimulation was selectively activating excitatory pyramidal neurons in M1 and S1. **(A)** Low magnification image of the coronal section processed anti-GFP antibody showing the medio-lateral aspect of YFP expression in the somatosensory cortex of Monkey J (areas 1, 2, 3). The black arrowhead indicates the location of the injector needle track; the adjacent tissue (black frame) is microscopically enlarged in **(B)** to show laminar distribution of the YFP-positive cells. In Monkey G (Yazdan-Shahmorad et al., 2016), we saw some cortical thinning post-mortem. In contrast, the histology for Monkey J, presented here does not show change in the cortical thickness. B: Densely YFP-positive cells are located predominantly in layers II-III and V-VI, and also show typical pyramidal morphology (cells in white frames are further enlarged in panels C-E). White arrowheads on bottom panels C-E point to typical pyramidal cells in layers II-III **(C)**; layer V **(D)**; and layer VI **(E)**. Scale bars: A, 2 mm; B, 200 μm; C-E, 100 μm. **(F)** Stained parvalbumin neurons (red) and YFP-expressing cells (green), no overlap indicated high specificity. Scale bar: 30 μm.

2.4.1 Two Measures of M1-S1 Functional Connectivity are Correlated

We quantified inter-area functional connectivity between S1 and M1 using two different measures—one based on the network response to optogenetic stimulation and the other based on spontaneous neural activity.

The first measure focused on how the response to optical stimulation propagates through the network. Optogenetic stimulation in S1 and M1 evoked responses across both

cortical areas, and we characterized these responses according to their amplitudes and delays (Figure 2.3A & 2.3E). The delays exhibited a bimodal distribution with a 3-6 ms separation between the early and late responses (Figure 2.3B & 2.3F). Classifying the responses across electrodes by delay recovers the spatial separation between the two cortical areas—following a boundary along the central sulcus—with shorter delays in the area being stimulated (S1 or M1; primary responses) and longer delays in the other area (M1 or S1; secondary responses) (Figure 2.3C & 2.3G). The short delay of the secondary responses suggests close functional connectivity between these areas. We quantified this connectivity with the stimulus-evoked response ratio (SERR). SERR was defined as the peak-to-trough amplitude of the secondary responses (A2) normalized by the peak-to-trough amplitude of the primary response (A1), with both responses measured in the high gamma (60 - 200 Hz) filtered signal (see dashed-line inset boxes in Figure 2.3A & 2.3E). The SERR is a measure of the connectivity between the site of stimulation and sites in the other cortical area (see Figure 2.3D & 2.3H).

The second measure evaluates functional connectivity during spontaneous activity. We focused on the coherence in field potential recordings between electrodes, a widely used measure of connectivity that captures the degree of phase locking between two signals (Bastos & Schoffelen, 2015; Lang, Tome, Keck, Gorriz-Saez, & Puntonet, 2012). We calculated the coherence between the site of stimulation and the secondary sites across different frequency bands.

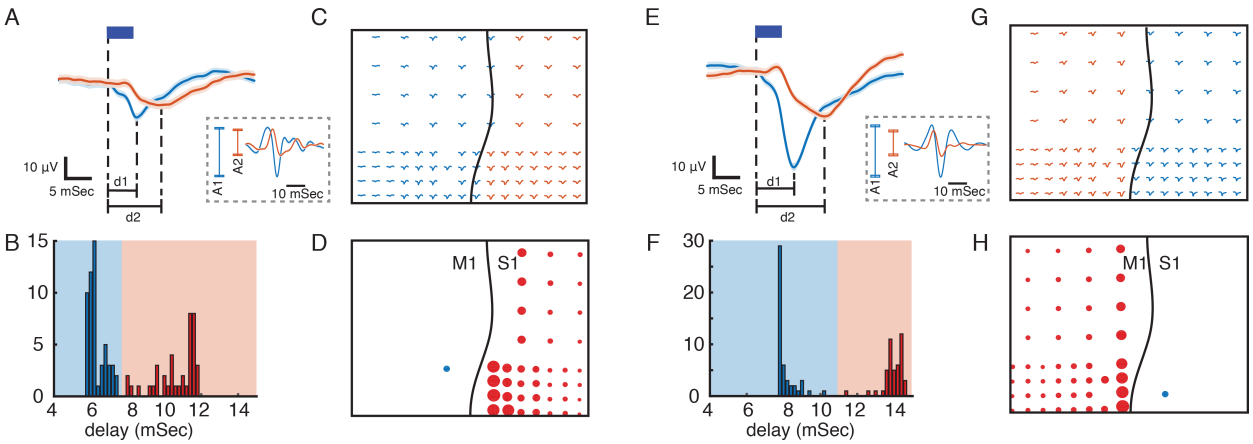

Figure 2.3 - Using Evoked Responses to Measure Connectivity across M1 and S1. (A) Primary (blue) and secondary (orange) evoked responses to light stimulation. The dark blue rectangle represents the duration of light stimulation. Shaded areas show standard error. The delays of evoked responses are calculated as the time difference between the onset of stimulation and the time of the response trough (vertical lines at d1 and d2). (B) Distribution of the delays color-coded based on the primary and secondary light-evoked responses. (C) Evoked responses across the array, color-coded based on the delays. As shown here there is a spatial separation of primary and secondary responses that corresponds to the locations of M1 and S1. This suggests that due to functional connectivity between M1 and S1, we see a delayed (secondary) response in S1 to light stimulation in M1. (D) S1 connectivity with the site of stimulation. The blue circle shows the location of stimulation in M1. The black line shows the location of central sulcus with respect to the recording array. The size of the red circles represents the strength of connectivity between each site and the stimulation location for the recording sites with secondary responses across S1. SERR Connectivity is defined as the peak-to-trough of filtered high gamma responses (60-200 Hz: trace plots shown on the dashed rectangle on (A) at each site (orange) normalized to the peak-to-trough of high gamma at the site of stimulation (blue). (E-H) Same as (A-D) with the same array placement but S1 stimulation.

Because SERR and coherence are derived independently from recordings with and without simultaneous stimulation, they might reflect fundamentally different aspects of network connectivity. Therefore, in each experiment we compared these two measures prior to conditioning stimulation. Figure 2.4A shows two examples comparing SERR to coherence in the theta band (4 - 8 Hz). To quantify the relationship between these measures, we performed linear regression between them across channels (Figure 2.4B-C). An example showing a strong relationship between SERR and theta band coherence is shown in Figure 2.4B. This analysis was repeated for all coherence frequency bands and for all experiments; the mean and standard error of the regression parameters are shown in

Figure 2.4C. The two measures of functional connectivity were highly correlated: across sessions, the distribution of regression slopes was significantly different from 0 (t-test: p=3.73e-06), and 69% of individual experiments showed significant linear regressions ($p < 0.05$). This finding indicates that at baseline, before any conditioning, the stimulation-evoked response reflects network dynamics across the frequency spectrum.

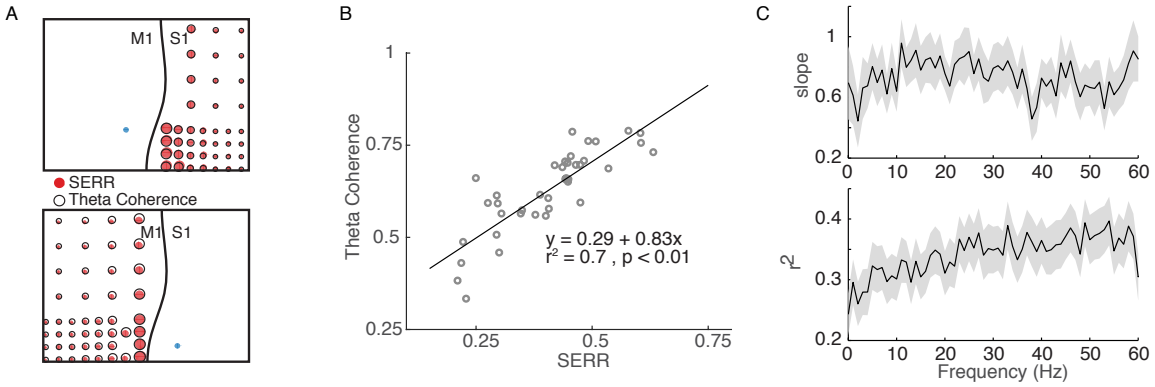

Figure 2.4 - Coherence Measure of Inter-Area Connectivity Correlates with SERR. (A) S1 (top panel) and M1 (bottom panel) theta coherence and SERR connectivity with the site of stimulation in Monkey G. The blue circle shows the location of stimulation. The black line shows the location of central sulcus with respect to the recording array. The size of the red and white circles represents the strength of connectivity between each secondary site and the stimulation location. **(B)** An example session showing relationship between SERR and theta coherence. The black line shows the linear regression fit. **(C)** Linear relationship between evoked response and coherence across different frequencies. Summary data showing the mean and standard error (shaded region) of regression parameters (shown in B) across frequencies.

2.4.2 Stimulation Strengthens Inter-area Connectivity

We used a simple stimulation protocol—we delivered 5 ms laser light pulses at a frequency of 5 or 7 Hz at either one or two cortical sites (Figure 2.5A). For our initial analyses, and unless otherwise noted, when stimulating at two cortical sites, we alternated stimulation between the two light sources to avoid any interference between the evoked responses from each light source. In each experiment, conditioning stimulation was applied for 50 minutes, and we evaluated functional connectivity every 10 minutes during blocks of passive baseline recording and active testing (100 light pulses delivered through each

laser). We also conducted control sessions with the same passive recording and testing blocks, but with no stimulation during the conditioning blocks.

A representative example experiment, shown in Figure 2.5B, shows that stimulation in M1 leads to significant increases (paired t-test, p=3.88e-11) in mean SERR. Similar increases in SERR were observed for a majority of experiments (14 out of 18 experiments) across both monkeys (Figure 2.5C). The increases in connectivity were symmetric between the two cortical areas (Figure 2.5D). No significant change in mean SERR was observed in control sessions without conditioning stimulation (Figure 2.5D). Furthermore, the increase in SERR was significantly larger for stimulation sessions than control sessions (unpaired t-test: p=0.0036)

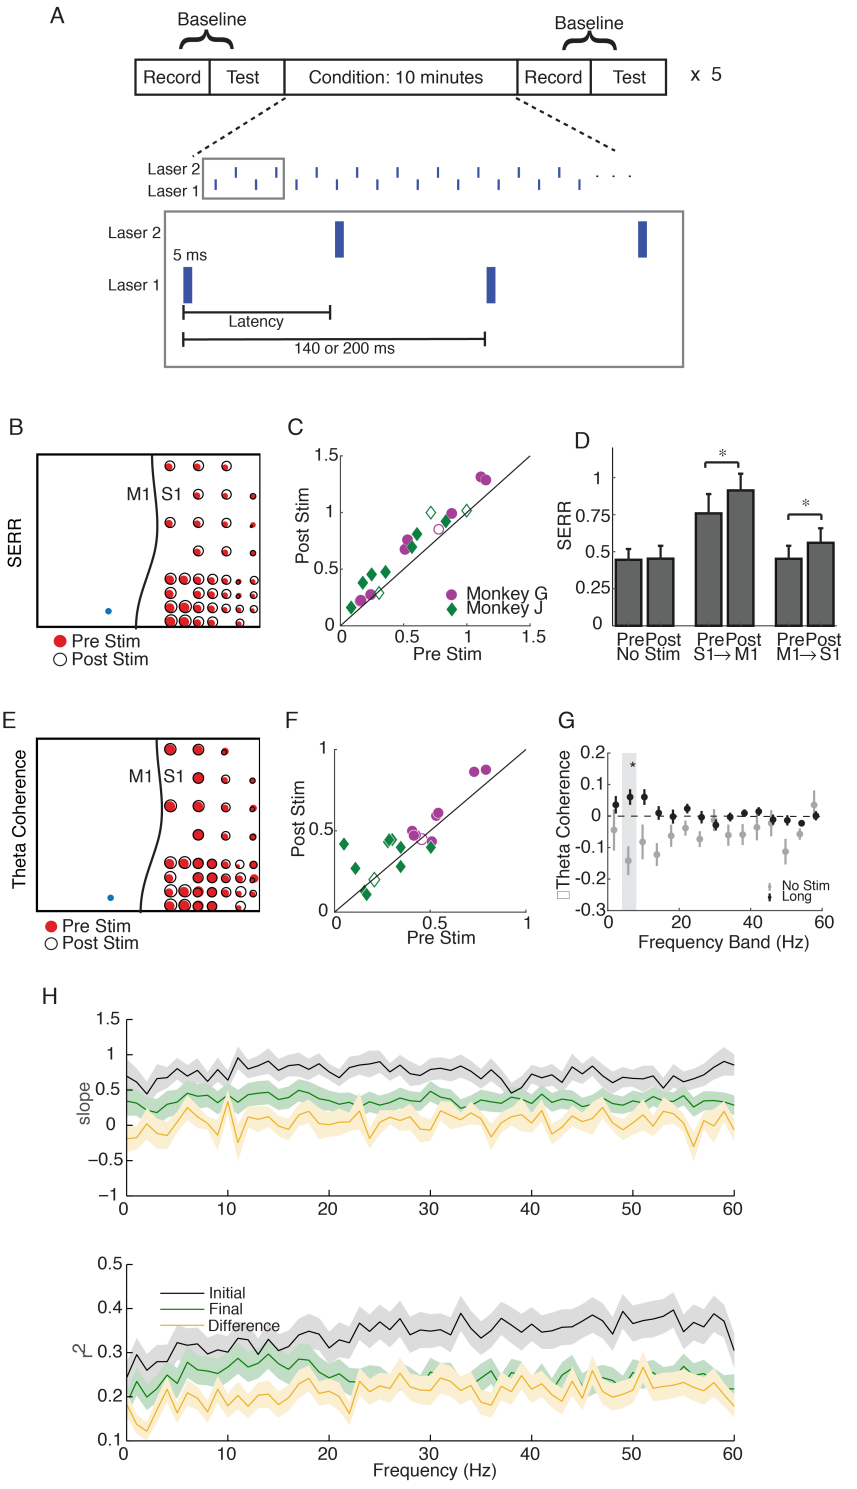

Figure 2.5 - Single Site and Long-Latency Stimulation Increase the Functional Connectivity between M1 and S1. (A) Experimental protocol. Conditioning stimulation was interrupted by periodic connectivity measurements including passive recording and active testing. Either one or two non-interfering lasers were used. **(B)** Examples of changes in SERR across the recording array in Monkey G. Red circles show initial connectivity and white circles show connectivity after 50 min. of conditioning. **(C)** Summary of SERR changes

across all experiments. Each symbol represents the connectivity averaged over all secondary channels for one experiment. Filled markers show significant changes (paired t-test; $p<0.05$). **(D)** Changes in SERR when stimulating in either S1 or M1 in comparison to control. Error bars represent standard error and asterisks show significant changes (paired t-test; Control: $p=0.8$, laser in S1: $p=0.01$, laser in M1: $p=1.6e-04$). **(E)** Examples of changes in theta (4-8 Hz) coherence across the recording array. **(F)** Summary of theta coherence changes across all experiments. Each marker represents the connectivity averaged over all channels for one experiment. Filled markers show significant changes (paired t-test; $p<0.05$). **(G)** Change in coherence across different frequency bands in comparison to controls. Asterisk show significant difference between the two groups (unpaired t-test, Bonferroni corrected; $p<0.05$). **(H)** Linear relationship between SERR and coherence across different frequencies for pre-stim, post-stim and the change in both measures. Summary data showing the mean and standard error (shaded region) of regression slope and r^2 as a function of coherence frequency (see the example regression for the theta-band in Fig. 2.4B).

We next asked whether conditioning increased coherence-based measures of inter-area connectivity. We did observe increases in mean theta band (4 - 8Hz) coherence following conditioning, as shown in Figure 2.5E for the same dataset used in Figure 2.5B (paired t-test, $p=0.03$). The increase in theta coherence was significant for the majority of experiments (13 out of 20; Figure 2.5E), and the effect was localized to the theta band (Figure 2.5G; paired t-test, $p=0.017$; Bonferroni correction for multiple comparisons). We also looked for change in the theta-band power as a result of stimulation and did not see any significant changes, supporting the conclusion that changes in coherence reflect changes in functional connectivity (Figure 2.6).

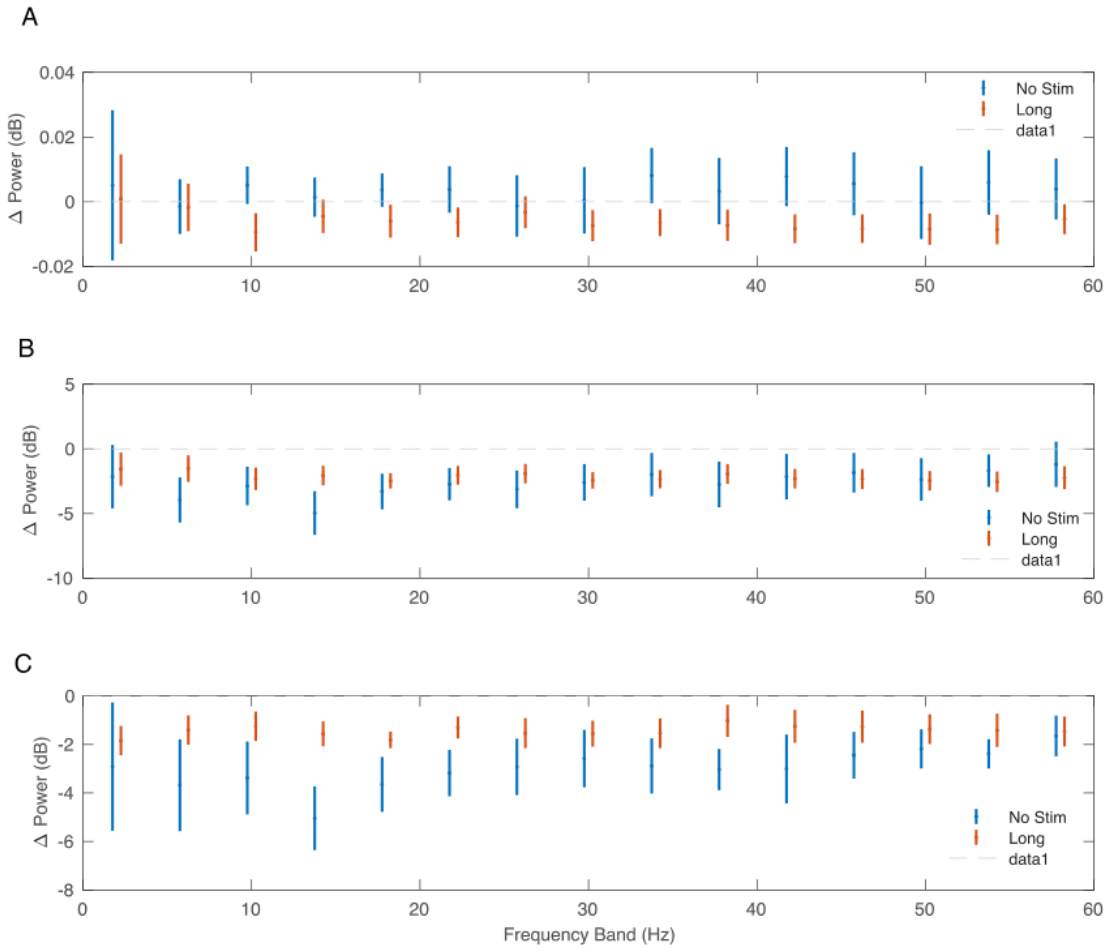

Figure 2.6 - Change in Power following Stimulation. Average change in power from pre- to post-conditioning across monkeys and experiments. **(A)** Average change in power across secondary recording sites normalized to the change in power at the stimulation site. **(B)** Average change in power across secondary recording sites. **(C)** Average change in power at the stimulation site.

2.4.3 Stimulation Weakens Correlation between Connectivity Measures

We showed in Figure 2.4C above that SERR and coherence are correlated measures of functional connectivity in the baseline condition. Here we ask how this relationship is changed by stimulation. Figure 2.5H replicates the baseline data from Figure 2.4C (black lines), along with new data showing the correlation between these measures at the end of each experiment (green lines) and between their conditioning-induced changes (yellow lines). After conditioning, the cross-channel correlation between these measures

significantly decreased (paired t-test, p=6.3e-10), and the changes in connectivity were uncorrelated (t-test on distribution of regression slopes, p=0.85). This suggests that the stimulation affects the two measures differently.

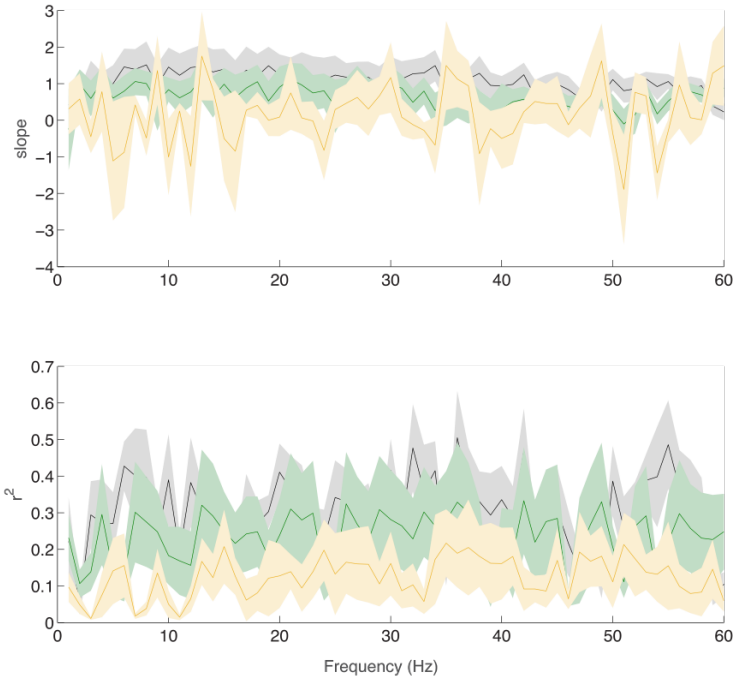

Figure 2.7 - Linear Relationship between SERR and Coherence across Different Frequencies for Control Sessions. Replication of Figure 2.5H for control (no stimulation) sessions. Linear relationship between SERR and coherence across different frequencies for pre-stim, post-stim and the change in both measures. Summary data showing the mean and standard error (shaded region) of regression parameters (slope—upper, r^2—lower) across frequencies.

To further address this, we replicated this analysis for our control (no stimulation) sessions. In this case, we did not see a significant decrease in the correlation between measures after control sessions (paired t-test, p=0.36; see Figure 2.7). Additionally, a direct comparison between the median decrease in the correlation for stimulation versus control sessions showed a significant difference when all frequencies were considered (p=0.0025), though the difference was not significant when only compared in the theta band. Together, these results show that the stimulation-induced changes in these two measures of

connectivity do not correlate with each other at a fine-scale. This may reflect the fact that the two measures are not functionally equivalent despite their correlation; they capture related, but not identical aspects of connectivity.

2.4.4 Stimulation Drives Increases in Inter-area Connectivity within Minutes

As shown in Figure 2.5A, our stimulation protocol included 5 repetitions of baseline recording, testing, and conditioning. This design allowed us to track connectivity changes across time, after each 10 minute increment of conditioning. Additionally, since neural activity was recorded throughout the experiment, we were able to measure changes in SERR during conditioning blocks. Figure 2.8A shows the evolution of mean inter-area connectivity for both measures in an example session. To account for differences in network connectivity across monkeys and sessions, we also analyzed the changes in each measure with respect to the baseline pre-conditioning measurements (Figure 2.8B). For this example session (Figure 2.8A), and on average across all sessions (Figure 2.8B), there is a trend of increasing connectivity across the session that starts within the first 10 minute conditioning block. These increases in inter-area connectivity measured with SERR and coherence (Figure 2.8B) were correlated across time (Pearson correlation coefficient=0.26). Additionally, the rate of strengthening was consistent across conditioning blocks; we detected a significant increase in SERR ($p<0.05$) at ~4.5 minutes of stimulation in 4 out of 5 conditioning blocks (see Figure 2.9 for comparison with control sessions).

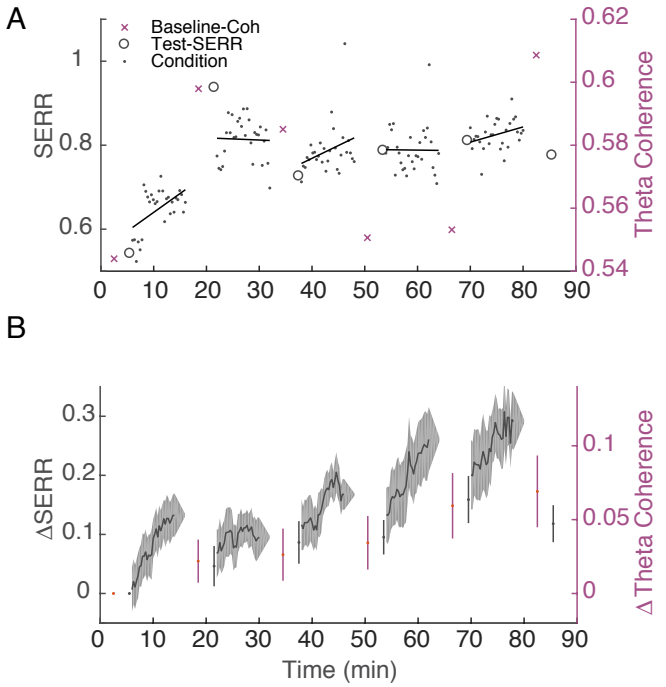

Figure 2.8 - Stimulation Induces an Increase in Inter-Area Connectivity across Time. (A) An example of the dynamics of mean inter-area SERR and coherence in an experiment. Black lines show linear regression to each conditioning block. **(B)** Dynamics of change in SERR and theta coherence with respect to baseline connectivity across all experiments (shaded area show standard error).

Note that the stimulation-induced plasticity was not reinforced during baseline recording blocks. Given the length of the conditioning (10min) and baseline recording (5min in most experiments) blocks, one might expect that stimulation-induced changes would revert back to baseline. We did observe a significant decrease (paired t-test, p=2.9e-05) in the SERR between the last 100 pulses of the conditioning block and the following test block, indicating some, but not complete, unlearning of the stimulation-induced connectivity changes.

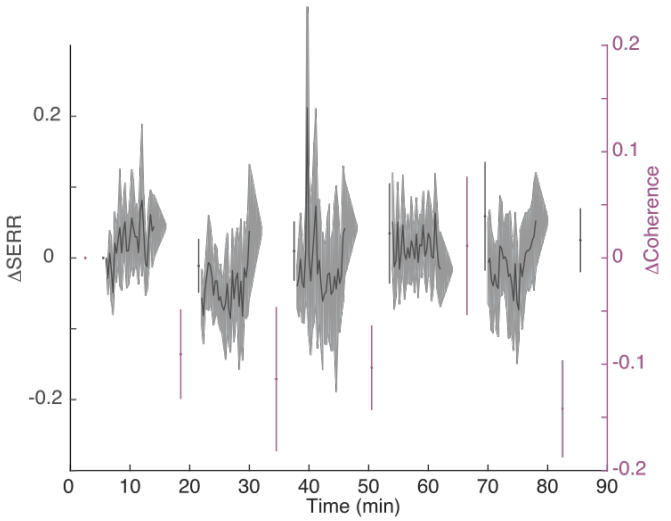

Figure 2.9 - Dynamics of Changes in Connectivity for Control Sessions. Replication of Figure 5B for control (no stimulation) sessions. Dynamics of changes in SERR and theta coherence with respect to baseline connectivity across all experiments (shaded area show standard error).

2.4.5 Stimulation Drives Fine-scale Changes across the Network that are Consistent with Hebbian Learning

We next evaluated the fine-scale effects of conditioning across the entire network. Since SERR is restricted to connectivity from the stimulation site, this analysis was conducted using only changes in pairwise coherence. Figure 2.10A shows data from an example session, with connectivity changes represented as a set of heatmaps: the heatmap at each recording site represents the change in pairwise coherence between that electrode and all of the other electrodes. A magnified example heatmap is shown in Figure 2.10A, right; top panel. This example shows heterogeneous changes in connectivity across the network.

We next asked whether such heterogeneous changes are to be expected. A Hebbian plasticity model suggests that the fine-scale changes in connectivity should reflect the statistics of plasticity-inducing activity. In this case, we predicted that

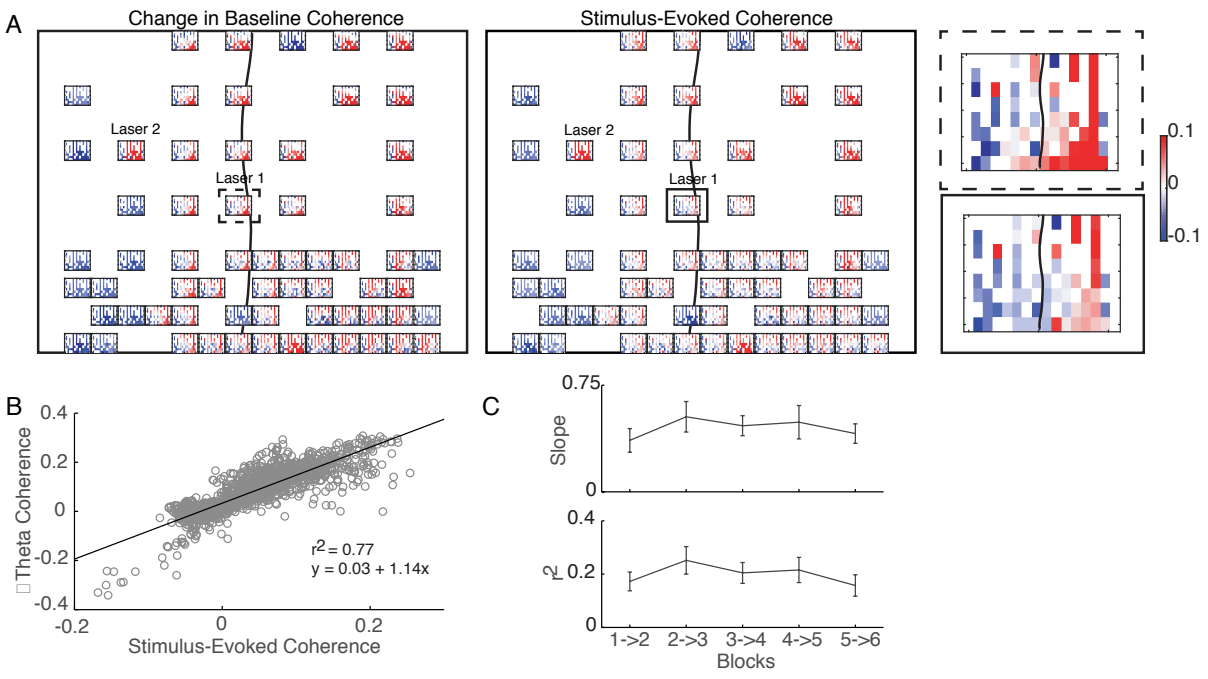

Figure 2.10 - Hebbian Plasticity Models Explain Stimulation-Induced Fine-scale Network Connectivity Changes. (A) An example session highlighting the similarity between the stimulus-evoked coherence (middle panel) and the change in baseline coherence (left panel) across the array in Monkey J. At each recording site across the array, a heatmap represents the respective coherence between that location with all other recording sites. Enlarged examples for a single location are compared in the right panel. The black line shows the location of central sulcus on the array and rectangles show the location of the magnified examples. The average value across the array was subtracted for visualization. **(B)** Linear regression between stimulus-evoked coherence and the change in baseline coherence for the example session shown in A. **(C)** Summary of regression parameters across single-site and non-interference experiments. Errorbars show standard error.

correlated stimulation-evoked activity between two cortical sites would strengthen the functional connectivity between them, while uncorrelated or anti-correlated stimulation-evoked activity would weaken their functional connectivity. We quantified stimulation-induced correlations by measuring the pairwise coherence during the conditioning block and subtracting out the coherence calculated during baseline recording (thought to be representative of initial connectivity). We refer to this difference as the stimulus-evoked coherence. Indeed, in the example session of Figure 2.10A, one can see a similarity between the pattern of changes in baseline coherence across the array (left panel) with the stimulus-

evoked coherence (middle panel).

To quantify this similarity, we used linear regression to predict changes in baseline coherence across sites based on the stimulus-evoked coherence. Importantly, both of these quantities use the initial baseline coherence as a reference. So to avoid spurious correlations introduced by subtracting the same data from the two regression variables, we split the baseline recording blocks into two intervals and used the data from only one interval for each regressor (see *Section 2.3 Methods* for more details). Data from the example session in Figure 2.10A is shown as a scatter plot in Figure 2.10B, along with the calculated regression parameters. The plot shows a strong correlation between the input coherence and changes in baseline coherence. This relationship was consistent across our data (Figure 2.10C): the distribution of linear regression slopes across experimental sessions and monkeys was significant (paired t-test: $p=2.41e-19$; 101 out of 105 experimental blocks had significant ($p<0.05$) linear regression models). These results support a Hebbian model for network-wide changes in connectivity.

To further test this model, we conducted some experimental sessions using a more complex spatio-temporal pattern of optical stimulation. Specifically, by reducing the latency between the two light sources to either 10 or 30 ms (Figure 2.12A), we introduce interference between evoked responses from the two light sources (Figure 2.12B). We repeated the analysis of mean inter-area coherence changes (c.f. Figure 2.5F) for these sessions, and we did not see a consistent increase in mean inter-area connectivity (Figure 2.12C). Nevertheless, we found that stimulus-evoked coherence continued to predict changes in baseline coherence (paired t-test on the distribution of regression slopes: $p=1.37e-28$; linear regression slopes with $p<0.05$: 160 out 170 blocks; Figure 2.12D, see

Figure 2.11 for an example session). Furthermore, the regression parameters were similar to those obtained with the simple stimulation patterns (Figure 2.10C), and both simple and complex stimulation experiments yielded regression slopes that were significantly larger than those obtained from control sessions (long-latency vs. control sessions, ranksum test: p=0.048; short-latency vs. control sessions, ranksum test: p=0.037; see Figure 2.13).

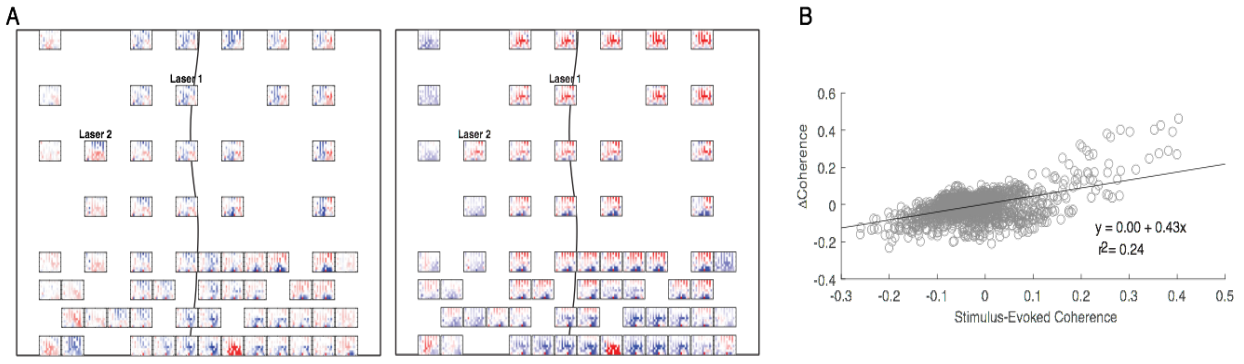

Figure 2.11 - Short-Latency Example Showing the Relationship between Input Coherence and Change in Baseline Coherence. (A) An example of similarity between the input coherence (right panel) and the change in baseline coherence (left panel) in Monkey J. **(B)** Linear regression between input coherence and the change in baseline coherence for the example session shown in A.

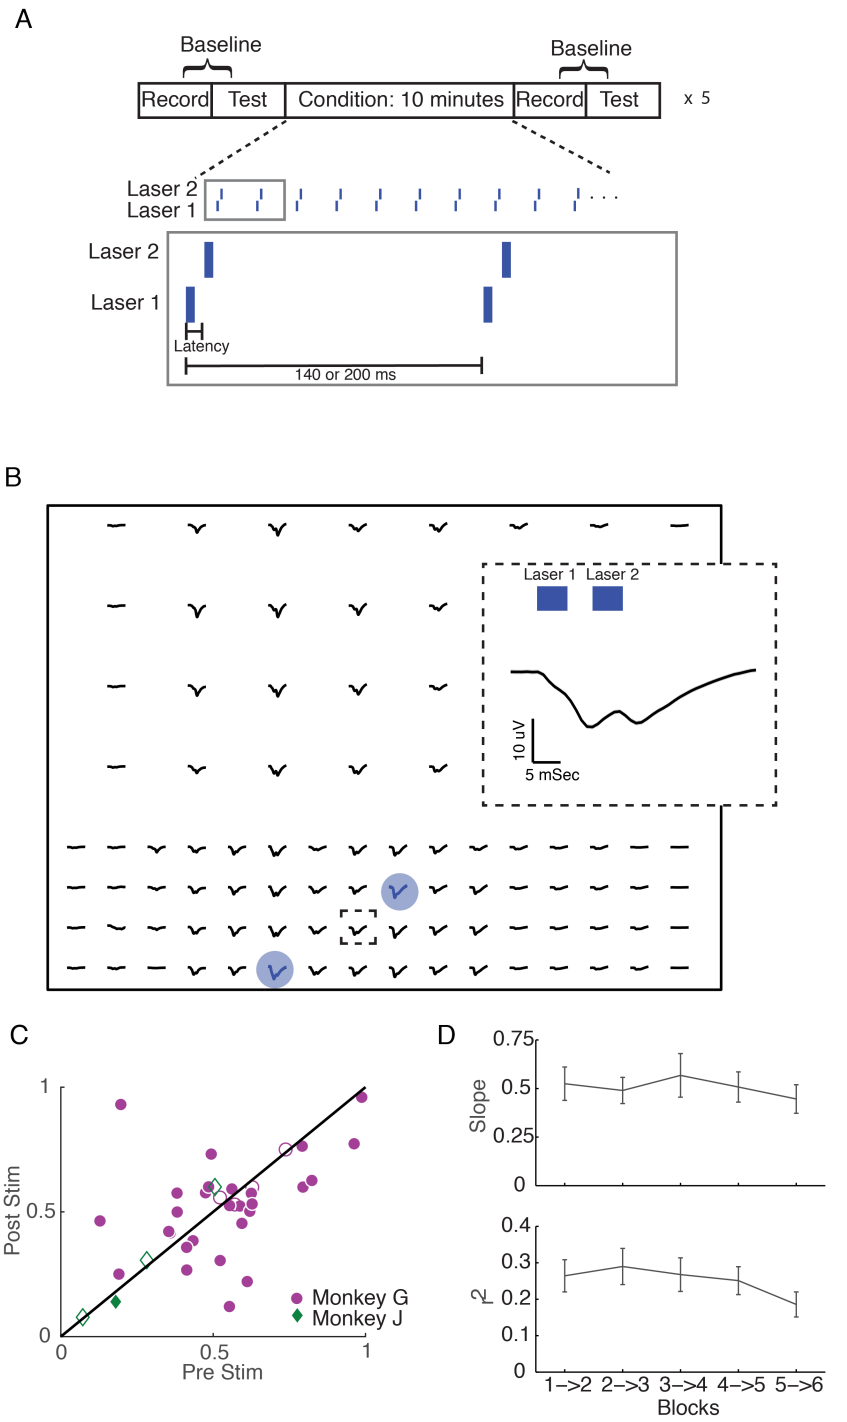

Figure 2.12 - Hebbian Plasticity Models Explain Fine-scale Network Connectivity Changes Driven by Complex Spatio-Temporal Stimulation Patterns. (A) Same stimulation protocol as described in Figure 2.5A. However here we reduced the latency between to the two lasers to 10 ms or 30 ms, creating a more complicated pattern of stimulation. **(B)** An example of the stimulus-evoked activity across the array from Monkey G. Blue circles show the locations of stimulation. The inset shows the enlarged pattern of evoked response at the framed electrode, which is located close to both lasers. **(C)** Summary of inter-area theta coherence changes for all interference experiments across both monkeys. **(D)** Summary of regression parameters across all interference experiments. Errorbars show standard error.

These results demonstrate a robust relationship between stimulus-evoked correlations and changes in baseline connectivity across the network. These findings are well explained by a Hebbian model of large-scale stimulation-induced plasticity.

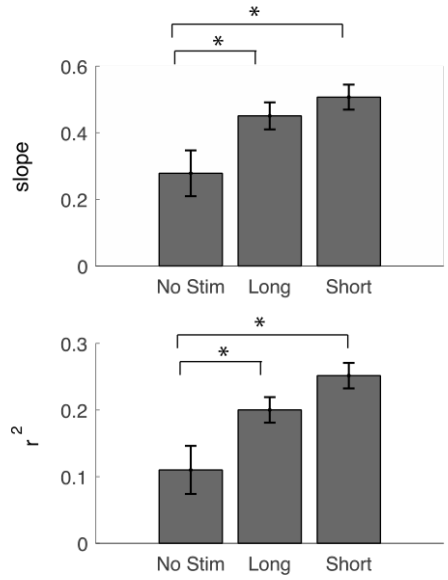

Figure 2.13 - Comparing the Effects of Stimulation on Network Connectivity between Stimulation and Controls. Summary of regression parameters across all experiments. Errorbars show standard error.

2.5 Discussion

In this study we investigated how the functional connectivity between and across sensory and motor areas changes in response to stimulation. With optogenetics, we selectively manipulated local populations of excitatory neurons within this sensorimotor network. We compared two different methods to measure functional connectivity at both gross- and fine-scales, and demonstrated how these measures change in response to conditioning stimulation. We showed that these changes are consistent with Hebbian synaptic plasticity rules, extending Hebbian models of stimulation-driven plasticity to large-scale networks. This work demonstrates the feasibility of driving targeted plasticity with optogenetic stimulation. This framework is a starting point for designing principled

approaches to large-scale neuroplasticity and stimulation-based therapies for neurological and neuropsychiatric disorders.

2.5.1 Measures of Functional Connectivity between M1 & S1

Functional connectivity between M1 and S1 has been demonstrated in fMRI (Matsui et al., 2011; McGregor & Gribble, 2015), electrophysiology (A Iriki et al., 1989; Prsa, Galinanes, & Huber, 2017; Yazdan-Shahmorad et al., 2016), and anatomical studies (Kosar, Waters, Tsukahara, & Asanuma, 1985; Petrof, Viaene, & Sherman, 2015). To evaluate stimulation-induced changes across these areas we investigated two measures of functional connectivity, one based on stimulation and one based on natural neural processing. Variants of SERR (Feldmeyer & Sakmann, 2000; Klinshov, Teramae, Nekorkin, & Fukai, 2014; Seeman et al., 2017) and coherence (Bastos & Schoffelen, 2015; Lang et al., 2012) have been used previously to evaluate functional connectivity. However, to our knowledge they have never been used in combination or compared.

In principle, these two measures reflect different aspects of functional connectivity. SERR is a more direct measure of the projections from the stimulation site. Importantly, stimulation likely evokes activity from both terminals and cell bodies located at the stimulation site, though C1V1 does not express well down at the axon (Rajasethupathy et al., 2015). Secondary responses arise from a combination of synaptic, antridromic, and indirect (network) effects. That said, based on the timing of the evoked responses and the limited expression of C1V1 in axon terminals, we expect the amplitude of secondary responses normalized to primary response amplitudes to reflect synaptic connectivity. Coherence, on the other hand, measures the broader effects of network-wide dynamics.

Despite these differences, we saw that the measures are robustly correlated across channels at the beginning of our experiments. This would be expected in a network where stability was maintained with a Hebbian mechanism. In other words, direct connectivity, as reflected in SERR, drives correlated activity, which is reflected in coherence; conversely, correlated activity drives changes in connectivity (see *Network effects* and *Network stability*). These coupled mechanisms lead the network to a stable, steady state (Fox & Stryker, 2017; Toyoizumi, Kaneko, Stryker, & Miller, 2014; Zenke & Gerstner, 2017).

Notably, however, our conditioning protocol caused the two measures of connectivity to become less correlated across electrodes. This might seem surprising given that conditioning led to a significant mean increase in connectivity between S1 and M1 for both measures and that the changes in coherence were robustly predicted by the stimulus-evoked coherence. However, this decrease in the correlation between the measures after conditioning seems to be limited to high frequencies and may reflect a transient instability of the network due to the stimulation-induced perturbations in connectivity.

2.5.2 Stimulation Changes the Functional Connectivity between M1 & S1

Previous work has demonstrated that either activity dependent (Jackson et al., 2006; Timothy H Lucas & Fetz, 2013; Nishimura et al., 2013; James M. Rebesco et al., 2010; W. Song et al., 2013) or paired electrical stimulation (James M Rebesco & Miller, 2011; Seeman et al., 2017) can lead to plastic changes in both primate and rodent cortex. Here we demonstrated that optogenetic stimulation at even a single location strengthens inter-area functional connectivity between brain areas. A similar observation was reported for one

session of single-site electrical stimulation (Seeman et al., 2017). This result is consistent with spike-timing dependent plasticity (STDP) rules (G. Bi & Poo, 2001). In particular, since the time difference between the activation of the primary (pre-synaptic) and secondary (post-synaptic) responses to optical stimulation is about 3 – 6 ms; repetitive stimulation should strengthen the connectivity from the site of stimulation to the other area. While stimulation-induced long-term potentiation (LTP) and STDP have been observed within local circuits (Feldman, 2012; A Iriki et al., 1989; Shulz & Jacob, 2010), and have been posited as the mechanism of stimulation-induced plasticity within M1 (Lajoie et al., 2017) this is the first study explaining large-scale changes across cortical networks. Notably, conditioning drove increases in coherence between S1 and M1 only in the theta band (4 – 8 Hz). This result is difficult to interpret because the conditioning stimulation frequency (5 or 7 Hz) itself lies within the range of the theta band. It is possible that Hebbian mechanisms are selectively enhancing connectivity at the conditioning frequency. It's also possible that theta band coherence best reflects the plastic connections between these areas. Since we did not vary stimulation frequency outside of this range, no strong conclusions can be drawn.

2.5.3 Dynamics of Connectivity Changes

The temporal dynamics of both measures reveal a trend of increasing inter-area connectivity throughout the experiment. This trend suggests that neuroplastic changes start almost immediately after the start of stimulation. Interestingly, we observed unlearning of the stimulation-induced changes during passive recording, though some of the stimulation-induced changes persisted. Previous studies have shown that longer

conditioning sessions (3 – 48 hours) result in changes that are stable over timescales comparable to the length of conditioning (Jackson et al., 2006; Timothy H Lucas & Fetz, 2013; Nishimura et al., 2013; James M Rebesco & Miller, 2011; Seeman et al., 2017). Further experiments are required to estimate the stability of plasticity following optogenetic stimulation over longer time scales.

2.5.4 Network Effects

Given that Hebbian rules have been used extensively to explain synaptic plasticity, we wanted to explore the idea that large-scale stimulation-induced changes are consistent with Hebbian plasticity. The strongest test of the Hebbian model is to compare the correlations directly induced by stimulation with the changes in network correlations observed after conditioning. Our ability to record artifact-free signals during stimulation allowed us to perform this test. We did find a predictive linear relationship between stimulus-evoked coherence and the changes in baseline coherence after conditioning. Furthermore, the model held for both simple and complex spatiotemporal patterns of stimulation-evoked activity.

2.5.5 Network Stability

There are a few important observations that we have touched upon but wanted to highlight here:

1) SERR is a more direct measure of connectivity from one area to the other.
2) Theta coherence measures how correlated activity is between M1 and S1.
3) At the beginning of experiments SERR and theta coherence are correlated

across electrodes (Figure 2.4).

4) This correlation decreases during the experiment. The change in SERR is not correlated to the change in theta coherence across electrodes (Figure 2.5).
5) The functional connectivity between M1 and S1 increases almost immediately with stimulation. Importantly, these increases are not permanent. During "Recording Blocks", in which passive recordings are taking place and stimulation is temporarily halted, the inter-area connectivity decreases (Figure 2.8).

These observations provide key insight into neural networks and suggest that the sensorimotor cortex is a stable network under normal conditions. A few computational papers have examined neural networks with ongoing Hebbian plasticity (Fox & Stryker, 2017; Toyoizumi et al., 2014; Zenke & Gerstner, 2017). In these networks connectivity shapes activity, and activity, in turn, drives changes in connectivity. More explicitly, strong connections between specific neurons lead to correlated activity between these neurons whereas weak connections do not generally lead to correlated activity. Hebbian plasticity suggests that correlated activity leads to strengthening connections and uncorrelated activity leads to weakening connections (DO, 1949). In simulation, these networks move towards an equilibrium point and become stable; at this stable point, correlations are reflective of the underlying connectivity in the network. In the context of this body of work, the observation that SERR and Coherence are correlated at the beginning of experiments is unsurprising. If Hebbian plasticity is ongoing, the sensorimotor cortex comprises a stable network where correlations (coherence) are reflective of underlying connectivity (SERR). During stimulation, activity patterns change, and different correlations are introduced to

the network. These correlations then drive connectivity changes (to a first order approximation, this is a Hebbian process; see *Network analysis*). However, the changes that are made cause the SERR and Coherence to be less correlated; furthermore, these changes do not last indefinitely. This suggests that our stimulation protocol pushed the network from its stable equilibrium to an unstable state, and then the network was pulled back to the equilibrium point after stimulation ended.

Many experiments have used targeted stimulation protocols to induce small changes in a network (G. Bi & Poo, 2001; Jackson et al., 2006; T H Lucas & Fetz, 2013; Seeman et al., 2017). These changes last at least as long as the stimulation itself, and in the case of slice experiments examining the connectivity at a single synapse, often much longer. However, changing targeted connections is different than changing broad connectivity across a network. These broad changes may be unstable without sufficiently long stimulation that pushes the network to a new stable equilibrium point.

2.5.6 Comparison to Representational Plasticity Literature

There is a large body of research on sensory cortical plasticity (Recanzone et al., 1992a; Wu et al., 2016). The present study differs from this body of work in two key dimensions. First, previous work has largely focused on how changes in the statistics of sensory input and/or motor output affect the cortical representation in sensory cortex. In contrast, we have directly manipulated cortical circuits, without a detailed knowledge of the somatotopic map. Also, earlier work studied plasticity on the time scale of days to months, while we are studying changes across minutes. Despite these differences, both appear to support a key role for Hebbian learning in sensorimotor cortical plasticity. Still, a

deeper understanding of the relationship between these bodies of work would require a study of intermediate time-frames as well, along with detailed cortical mapping. These are important questions for future work.

2.5.7 Limitations of the Experiment

The large-scale interface used in this study enabled robust estimation of functional connectivity and stimulation-evoked correlations. However, there are several limitations in our experimental setup. First, surface stimulation combined with high blue light absorption and scattering in tissue limits the depth of light penetration (Yizhar, Fenno, Davidson, Mogri, & Deisseroth, 2011). The ability to target specific cortical layers would help us understand the anatomical basis of the plasticity we observed. Second, surface μECoG recordings reflect a summation of nearby neural activity, making comparisons to explicit synaptic learning rules difficult. Depth recordings would reveal more detailed information about spike timing, the role of different cortical layers, and the relationship between synaptic and large-scale plasticity. Lastly, the μECoG arrays were placed acutely for each experiment. To investigate the long-term effects of stimulation we need stable, chronic recordings.

2.5.8 Clinical Applications

Our results offer a proof-of-concept that optogenetic stimulation can drive predictable changes in network-scale connectivity. The success of optogenetics in NHP represents an important step forward for translational use (Doroudchi et al., 2011; Han, 2012). The continued rapid progress in the field has the potential to bring cell-type specific

neuromodulation therapies to the clinic. However, even where therapeutic approaches ultimately wind up using electrical stimulation, the advantages of optogenetics make it a powerful tool for better understanding the underlying mechanisms of plasticity and, thus, for the development of therapeutic applications of neuromodulation. This framework also provides scientific insight into the mechanisms of neural plasticity. Future efforts should focus on linking the plasticity results presented here with improvements in motor function or sensory perception, which will have important implications for stimulation-based therapies.

Chapter 3
Meso-scale Oscillations Coordinate Spiking Activity to Support Sensorimotor Learning

3.1 Foreword

In Chapter 2, I used an artificial neurostimulation paradigm to model meso-scale connectivity changes as a function of baseline connectivity and stimulation-induced correlations. This work stressed the importance of induced correlations, reinforcing principles attained from synaptic plasticity experiments. In this chapter, we try to connect synaptic plasticity to meso-scale oscillations during natural sensorimotor learning. We take advantage of meso-scale oscillations that are thought to support sensorimotor learning (sleep spindles and slow oscillations) and measure the timing and structure of spiking activity dynamically during these oscillations. We the connect the changes in spiking structure during oscillations to expected changes in synaptic plasticity, changes with sensorimotor learning, and changes during sleep. Since the spiking structure is specifically measured in relation to meso-scale activity, these observations provide insight into the role of meso-scale oscillations in plasticity.

3.2 Introduction

Sleep-dependent offline processing is required for the consolidation of new memories and skills (Gulati et al., 2017; Latchoumane et al., 2017; Miyamoto et al., 2016; Rasch & Born, 2013). In particular, thalamocortical spindles, a hallmark of light sleep (Rasch & Born, 2013; M. Steriade, 2000; M Steriade, McCormick, & Sejnowski, 1993), have been linked to offline processing across a variety of paradigms ranging from declarative memory tasks (Z. Clemens, Fabó, & Halász, 2005; Zsófia Clemens, Fabó, & Halász, 2006; Eschenko, Molle, Born, & Sara, 2006; Gais, Mölle, Helms, & Born, 2002; Logothetis et al., 2012; Mölle, Eschenko, Gais, Sara, & Born, 2009) to motor learning paradigms (Barakat et

al., 2011; Fogel & Smith, 2006; Johnson et al., 2012; Nishida & Walker, 2007; Ramanathan et al., 2015). Moreover, slow oscillations (SO) during NREM sleep have been linked to offline processing of awake experiences during sleep (Gulati et al., 2017; Gulati, Ramanathan, Wong, & Ganguly, 2014; Marshall, Helgadóttir, Mölle, & Born, 2006; Miyamoto et al., 2016; Molle, Marshall, Gais, & Born, 2004; Ramanathan et al., 2015; Mircea Steriade & Timofeev, 2003).

There is growing evidence that "nesting" of SO with spindles may be particularly important for offline processing. For putative hippocampus dependent tasks, the coordination of thalamocortical spindles (10 – 16 Hz), cortical SO (< 1 Hz), and hippocampal sharp-wave ripples (150 – 250 Hz) appear to promote consolidation through hierarchical nesting, in which higher frequency oscillations are embedded ("*nested*") within lower frequency oscillations (Cox et al., 2014; Mölle et al., 2009; Mölle, Marshall, Gais, & Born, 2002; Niknazar, Krishnan, Bazhenov, & Mednick, 2015; M. Steriade, Contreras, Curró Dossi, & Nuñez, 1993). Interestingly, a recent study selectively silenced or induced sleep spindles during a contextual fear conditioning task and found concordant changes in behavior (Latchoumane et al., 2017), confirming that nesting plays a causal role in learning. In motor learning paradigms, both spindles and SO are correlated with the reactivation of awake experiences and offline performance gains (Ramanathan et al., 2015); it remains unclear, however, if such cortical nesting is associated with hippocampal sharp-wave ripples (Genzel & Robertson, 2015).

Despite the extensive work linking sleep spindles, SO, and the phenomenon of nesting to memory consolidation, little is understood about the relationship between these oscillations, spiking activity, and neuroplasticity. In general, the temporal precision of

neural spiking is important for regulating changes in synaptic efficacy (G. Q. Bi & Poo, 1998; Hebb, 1949), and network correlation structure is predictive of large-scale functional reorganization (Michael M. Merzenich et al., 1984; Recanzone, Merzenich, & Dinse, 1992b; Yazdan-Shahmorad, Silversmith, Kharazia, & Sabes, 2018). A few studies have found that neurons in prefrontal cortex are phase-locked to sleep spindles (Gardner, Hughes, & Jones, 2013; Peyrache, Battaglia, & Destexhe, 2011; Sela, Vyazovskiy, Cirelli, Tononi, & Nir, 2016). However, despite the importance of neural correlations for driving synaptic plasticity, the effect of spindles on the correlation structure of neural firing remains largely unexplored.

To that end, we simultaneously recorded local field potentials (LFPs) and spiking activity from electrode arrays in the primary motor cortex (M1) of sleeping rats. This allowed us to examine the precise relationship between spike timing relative to ongoing spindle oscillations and SO. By parsing spindles into their component cycles we were able to analyze the dynamics of spiking during the evolution of spindles. This analysis further revealed a waxing and waning of fine-scale structure, which featured increased spiking, increased phase locking, and increased local synchrony that reached a maximum at the peak of spindles. This analysis also revealed that SO and spindles interact—the distribution of spindle-induced synchrony narrowed when spindles were in closer temporal proximity to a SO. By understanding the connection between spindles, SO, and spiking correlation structure, we hope to gain insight into the neurophysiological basis of offline processing driven by spindle-SO interactions.

3.3 Methods

3.3.1 Animals and Surgical Procedures

This study was performed in strict accordance with guidelines from the USDA Animal Welfare Act and United States Public Health Science Policy. The protocol was approved by the San Francisco VA Medical Center Institutional Animal Care and Use Committee (IACUC, Protocol Number 13–006). We used 4 adult Long–Evans male rats (approximately 8 wks old). Animals were kept under controlled temperature and a 12 hr light/dark cycle with lights on at 06:00 A.M. Probes were implanted during a recovery surgery performed under isofluorane (1%–3%) anesthesia. The post–operative recovery regimen included administration of buprenorphine at 0.02 mg/kg b.w. and meloxicam at 0.2 mg/kg b.w. Dexamethasone at 0.5 mg/kg b.w. and Trimethoprim sulfadiazine at 15 mg/kg b.w. were also administered post–operatively for 5 days. All animals were allowed to recover for 5 days prior to the start of experiments.

3.3.2 Electrophysiology

We recorded extracellular neural activity using tungsten microwire electrode arrays (MEAs, n = 2 rats, Tucker–Davis Technologies or TDT, FL), tetrodes (n = 1 rat, NeuroNexus, Michigan), and custom probes (n = 1 rat). We implanted arrays in the caudal forelimb area of the primary motor cortex (M1), centered at 3–4 mm lateral, 0.5 mm anterior to bregma. We recorded spike and LFP activity using a 128– and 256–channel TDT–RZ2 system (Tucker–Davies Technologies). Spike data was sampled at 24,414 Hz and LFP data at 1,018 Hz. ZIF–clip based analog headstages with a unity gain and high impedance (~1 GΩ) were

used. Only clearly identifiable single units with consistent waveforms and high signal-to-noise were analyzed. MEA recordings were sorted offline using PCA-based algorithms followed by manual cluster-cutting using TDT's OpenSorter software. Tetrodes were sorted using the "UltraMegaSort" toolbox (available at https://physics.ucsd.edu/neurophysics/software.php), a set of MATLAB based scripts for tetrode sorting described in detail previously (Gulati et al., 2014; Ramanathan et al., 2015). Sorting on the 256-channel custom probes was done using MountainSort (Chung et al., 2017). Briefly, MountainSort is spike-sorting software that uses an automatic algorithm, which compares clusters of data using one-dimensional projections. If data along these projections are bimodal then the clusters are considered distinct. We used a consolidation factor of 0.9 and a noise overlap threshold of 0.03 to identify clusters as single units. We performed a minimal amount of manual merging and rejection of clusters to correct for drift during long recordings.

3.3.3 Analysis

Data analysis was performed in MATLAB using a combination of custom scripts and available neural data analysis Toolboxes. For all analyses, disconnected/high impedance LFP channels were removed from the analysis.

3.3.3.1 Behavior

Prior to surgery, animals were handled and acclimated to behavioral boxes. Next, they were oriented to the pellet tray for 1 wk, and then we determined handedness for each rat by evaluating behavior on 10 trials of the Whishaw forelimb reach to grasp pellet task. Then we implanted multielectrode arrays into the forelimb region of the contralateral

motor cortex. After five days of recovery, animals were food-restricted for 2 d, after which animals attained a fixed amount of food during the course of training (2 average sized food pellets/day). We trained animals on the Whishaw reach to grasp task in a clear plexiglass chamber, with a 1.5 cm slit for animals to reach through to grasp a 45 g pellet on a shallow dish that was 1.5 cm away from the behavioral chamber. All reaches were videotaped for post- hoc analysis of accuracy, kinematics, and dynamics. All behavioral sessions began in the morning and consisted of a 2 h block of spontaneous recording (to record a "baseline" sleep period, Sleep1); a block motor skill learning (Reach1); and a second 2-h block of spontaneous recording (Sleep2). Differences between Sleep1 and Sleep2 can be used to assess learning-related changes.

3.3.3.2 Sleep Classification

During the spontaneous sleep blocks, each LFP channel was segmented into non-overlapping 6 sec windows. In each window the power spectral density was computed and averaged over the delta (0.1–4 Hz) and gamma (30– 60 Hz) frequency bands. Then a K-means classifier was used to cluster epochs into two clusters, NREM sleep and REM/awake (Figure 3.1). Only long (>30sec) epochs of sleep were analyzed. The identified NREM sleep epochs were verified by post-hoc visual inspection of the LFP activity.

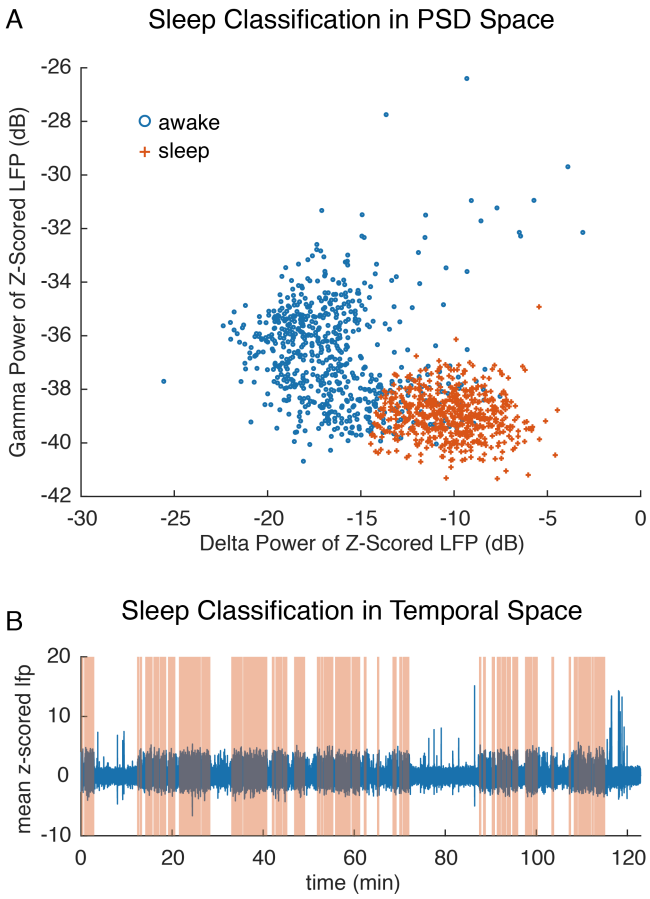

Figure 3.1 - Example Sleep Classification. (A) PSD scatter plot. Each dot represents the power spectral density in the delta (0.1–4 Hz) and gamma (30–60 Hz) frequency bands during 6 sec windows. A K-means classifier was used to cluster epochs into two clusters. Blue dots belong to the cluster representing NREM sleep and orange pluses are used for the remaining dots. **(B)** Sleep classification across time. Average LFP trace is plotted across a 2 hr sleep session (blue). Identified sleep epochs are highlighted with red boxes.

3.3.3.3 Spindle Detection

The spindle detection that applied here is similar to the algorithm used by (Sela et al., 2016). Good channels were first z-scored and averaged to form a virtual LFP channel. This signal was filtered in the spindle band (10-16Hz) using a zero-phase shifted, third order Butterworth filter. A smoothed envelope was calculated by computing the magnitude of the Hilbert transform of this signal then convolving it with a Gaussian window ($\alpha = 2.5$). Next, we determined two thresholds for spindle detection based on the mean (μ) and standard deviation (s.d.) of the spindle band envelope during NREM sleep (

$lower = \mu + 1.5\,s.d.$; $upper = \mu + 2.5\,s.d.$). Epochs in which the spindle power exceeded the upper threshold for at least one sample and the spindle power exceeded the lower threshold for at least 500 ms were considered spindles. Finally, spindles that were sufficiently close in time (<300 ms) were combined. For each spindle epoch, the peak of the spindle band LFP was identified. Spindles were aligned to this peak for generating average spindle waveforms, spectrograms, and spike rasters (Figures 3.2 and 3.3).

3.3.3.4 Spectrogram Generation

Spindles epochs are typically 0.5 to 1 sec, though they can last up to 3 sec. Therefore, we used a window from -4 to +4 sec around each spindle peak to generate spectrograms. Then we used the MATLAB spectrogram function to compute the spectrogram across the entire window (500 ms Hamming window, 350 ms overlap) during each spindle. Next, we subdivided this window into two 1-sec epochs: (1) Baseline: -2.5 to -1.5 sec (Figure 3.2C, dashed line); (2) Spindle: -0.5 to +0.5 sec (Figure 3.2C, solid line). During each spindle, we computed the power for each frequency across the baseline epoch and spindle epoch. Average power across the baseline and spindle epochs is plotted in Figure 3.2C, left plot. Then we calculated the % Power Change from the baseline epoch for each frequency band:

$$\%\ Power\ Change = 100 * (Pspindle - Pbase)/Pbase$$

The % Power Change is shown in Figure 3.2C, right plot.

3.3.3.5 Spike Phase Extraction

Similar methods for spike phase extraction were used to assess the spiking structure within a spindle cycle and across spindles. For both analyses, we first computed the same virtual signal used in spindle detection then filtered the data in the spindle band

(10-16Hz). Next we applied the Hilbert transform and took the angle at each sample to get a continuous representation of the relative spindle phase (within cycles). To assess the absolute phase across a spindle epoch, the relative phase was unwrapped and centered such that the phase was 0 at the peak of the spindle and 10π after the 5th cycle from the peak (Figures 3.3 and 3.4). For each spindle epoch and each neuron, the nearest phase was collected at each spike event.

3.3.3.6 Phase Locking Value and Preferred Spindle Phase

We calculated the phase locking value in order to assess the degree of phase consistency of spiking within spindle cycles. Briefly, for a given neuron, across all spindle epochs, the phase of the spindle band LFP signal was collected for each detected action potential within a given cycle (e.g., [0 , 2pi]) yielding a distribution of spike phases. Each phase value in this distribution was treated as a vector of magnitude one and angle equal to the phase.

$$Average\ Spike\ Phase\ Vector = \frac{1}{n}\sum_{i=0}^{n} e^{\emptyset_i}$$

The average phase vector was computed according to the above equation. From this vector, we attained the phase locking value (vector magnitude) and the preferred spindle phase (vector angle). We calculated these measures for each neuron and each spindle cycle.

3.3.3.7 Cross Correlation Histogram and Pairwise Synchrony

For each spindle epoch we segmented the LFP into it's individual cycles. Next we collected all spike times in each of these cycles. To compute the cross correlation histogram for a pair of neurons we adapted methods from (Engelhard, Ozeri, Israel, Bergman, & Vaadia, 2013). Briefly, both spike trains were binned into 1ms time bins, and one neuron

was initially treated as the trigger neuron. For each spike from the trigger neuron we generated a count vector of the other neuron's spiking (t) relative to the trigger spikes, where $t \epsilon \{-50\ ms, +50\ ms\}$. Each count vector was concatenated into a matrix. Next we swapped trigger neurons and built a similar count vector matrix, but flipped each count vector (matrix row) to reflect any asymmetries in firing between the neurons. The resulting matrix of 0's and 1's had 101 columns (1 ms time bins from -50 to +50 ms) and rows equal to the total number of spikes from both neurons. A Gaussian kernel (5 ms) was used to smooth each row of this matrix and then the average smoothed count vector was computed. This count vector represents the normalized frequency of co-firing between the two neurons. The *raw* pairwise synchrony was taken to be the peak of this function.

Despite normalizing the cross correlation histogram, an increase in pairwise synchrony can arise from multiple sources including true increases in pairwise correlations but also changes in excitability, firing rate, and others. Since these first-order effects might influence second-order measures, we sought to find a method to isolate second-order changes (Brody, 1999; Palm, Aertsen, & Gerstein, 1988). To accomplish this we first used a shuffling method to isolate the first-order effects on synchrony and then subtracted them off. As before, we collected all spike times in each individual spindle cycle and selected one neuron within a given pair to be the trigger neuron. However, we then shuffled the second unit's spindle epochs. This maintained a consistent timing and phase relationship of the second unit with the ongoing spindle, but broke the statistical relationship between the two neurons under examination. We repeated this shuffling 25 times and then computed the same normalized, smoothed count vector for the shuffled condition. This count vector represents the expected normalized frequency of the co-firing given both neurons' average

spike rates and phase locking with the specific spindle cycle. Finally, we subtract this expected count vector (*shuffled*) from the unshuffled (*raw*) count vector, resulting in a *corrected* cross correlation histogram (Figure 3.6B). It is important to note that the first-order measures change as the spindle progresses (see Section *3.4.3 Spindles Increase Single Unit Phase Locking* and Figure 3.4), but this shuffling is done within each cycle and therefore accounts for these changes across the spindle. The synchrony for each pair of neurons was then taken to be the peak of the corrected cross correlation histogram.

3.3.3.8 Slow Oscillation Detection

The slow oscillation detection used here is similar to algorithms that have been previously used (Sela et al., 2016). To detect the < 1 Hz SO, a virtual LFP channel was constructed by averaging the LFP across all recording channels (excluding bad/disconnected channels). Next this virtual signal was filtered in a low frequency band (2^{nd} order, zero phase shifted, high pass Butterworth filter with a cutoff at .1Hz followed by a 5^{th} order, zero phase shifted, low pass Butterworth filter with a cutoff at 4Hz). Next, all positive-to-negative zero crossings were identified during NREM sleep, along with the previous peaks, the following troughs, and the surrounding negative-to-positive zero crossings. Each identified epoch was considered a slow oscillation if the peak was in the top 85% of peaks, the trough was in the top 40% of troughs and the time between the negative-to-positive zero crossings was greater than 300 ms but did not exceed 1 sec.

3.3.3.9 Spindle Nesting Analysis

Each spindle was binned according to its timing relative to SO. For each spindle, the preceding slow oscillation was collected. We computed the time difference between the

down state (peak delta band LFP) and the spindle (peak spindle band LFP). For the nesting analysis, spindles were characterized as nested if they had a SO precede them within 1.5 secs and unnested otherwise. For the interaction analysis, the time difference was binned in 1.5 sec bins with 1 sec overlap (Figure 3.8A-B). Since this distribution of time delays is skewed, there are more spindles in the bins closer to SO than in subsequent bins. In order to accurately compare the distribution of pairwise correlations across these bins, we used a subsampling procedure in which we randomly sampled the nested, unnested, and binned spindles using the minimum number of spindles across all bins. We repeated this sampling 100 times in order to generate a bootstrapped distribution of pairwise correlations and then used the mean of this distribution for each pair of neurons.

Two methods were used to analyze differences in the distributions of pairwise correlations across bins. The first method was designed to analyze trends across bins. We computed the mean and standard deviation of the pairwise correlations for each spindle bin. Then we performed linear regression, where the mean or standard deviation was a function of the time bins. The second method was designed to detect if the distributions significantly changed. We focused on the differences between distributions in the first bin, [0,1.5 sec], and the last bin [4.5,6 sec]. To detect broad differences between the distributions we used a Kolmogorov–Smirnov test. If the Kolmogorov–Smirnov test was significant, we followed up with a Shift test (Rousselet, Pernet, & Wilcox, 2017) to assess how the distributions differed. Briefly, the sextiles of each distribution are computed using the Harrell-Davis quantile estimator. Then differences in sextiles between the two distributions are computed with a bootstrapped estimate of confidence intervals for each sextile difference. A multiple comparisons correction was then used to account for the five

different estimators.

3.3.3.10 Random Spindle Epochs

There is a certain degree of spiking, phase locking, and synchrony that would be expected even if neurons were not modulated by spindles. To account for extraneous measurements, we generated a random spindle distribution that had similar statistics to the true spindle epoch distribution. Briefly, for each spindle epoch, a random offset was computed 5-10 sec back from the true spindle peak, i.e., $T \sim U(-10, -5)$. The nearest peak in the spindle band LFP was taken as the peak of the random spindle epoch. Each analysis was jointly computed for the true spindle epochs (blue in Figures 3.3-3.7) and the random spindle epochs (black in Figures 3.3-3.7).

3.4 Results

3.4.1 Neural Oscillation Detection

We recorded extracellular LFP and spiking activity from M1 in four rats (see Section *3.3 Methods* for description of electrophysiology and recording probe details; Figure 3.2A). We intermittently recorded activity during "Sleep Blocks" in which animals were given the opportunity to sleep for ~2 hrs. During NREM sleep we identified ongoing spindles and SO using standard algorithms for automatic detection (Sela et al., 2016). Briefly, LFP channels were z-scored to standardize activity levels. Next, the average activity across electrodes was used as a virtual LFP signal (Figure 3.2B).

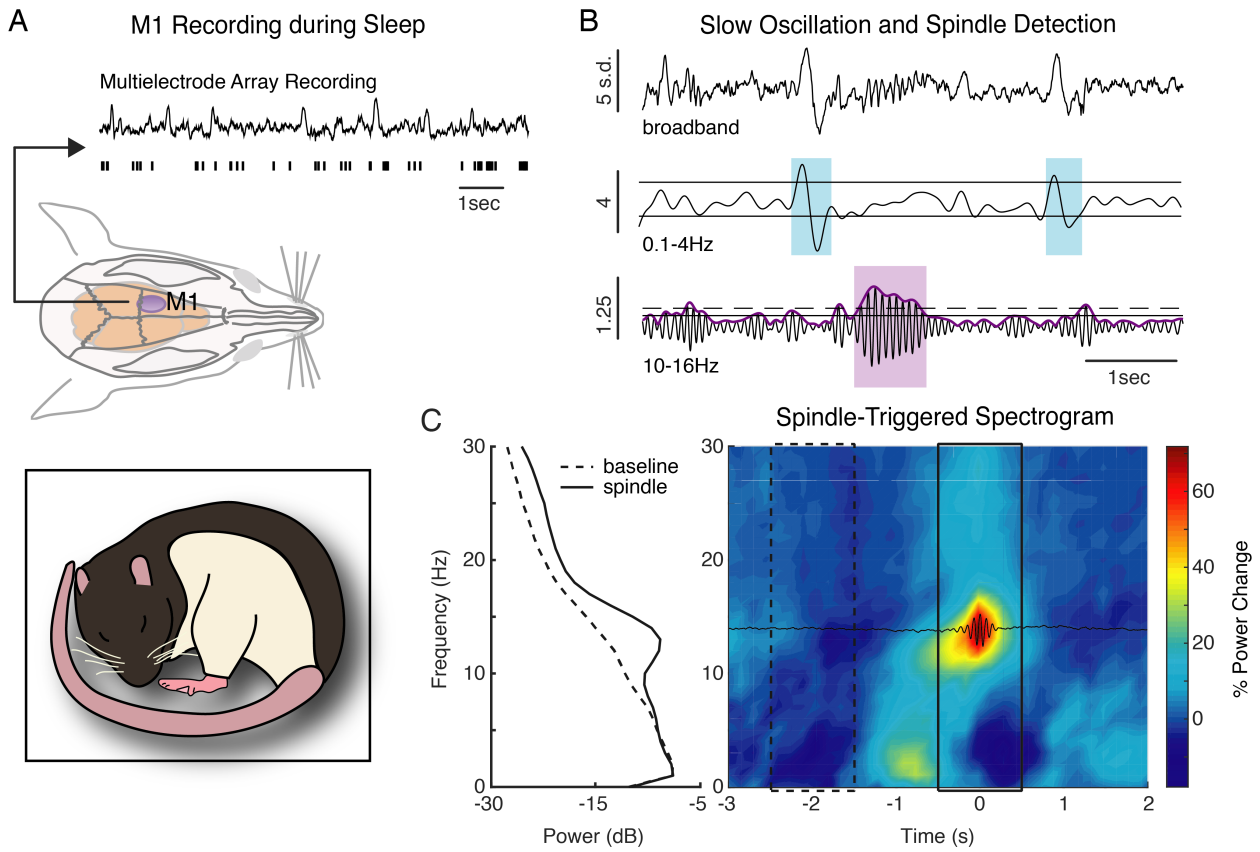

Figure 3.2 - Recording Setup and Oscillation Detection. (A) Sleeping rat along with anatomical location of multielectrode arrays and example LFP and spiking data from one recording channel. **(B)** Examples of detected spindles and slow oscillations, highlighting the automatic methods used for detection. The broadband LFP (top) is decomposed into a low-frequency band for SO detection (middle) and spindle band (bottom) components. SO must have had sufficient positive and negative amplitudes (black lines; middle) and sufficiently slow durations (highlighted blue). The spindle band envelope (purple line) must have exceeded an upper threshold (solid black line) for one sample and a lower threshold (dashed black line) for at least 500 ms. The detected spindle duration is highlighted in purple. **(C)** Average spindle-triggered waveform (black line; right) and spectrogram (heat map; right). Average spindle-triggered power spectrum (black line; left) and baseline power spectrum (gray line; left). Black and gray boxes on the heat map display the timing of the spindle and baseline periods used to calculate the power spectra.

This signal was filtered in the spindle band (10-16Hz) and a lower frequency band (0.1-4Hz). Periods in which spindle power exceeded an upper threshold for at least one sample and a lower threshold for at least 500 ms were identified as spindles (see *Section 3.3 Methods*; Figure 3.2B). Indeed, the average spectrogram of identified spindles (aligned to the spindle peak) revealed a time-frequency specific bump in spindle power compared to random epochs (Figure 3.2C). Notably, this spectrogram also shows a bump in low

frequency power, which precedes spindles, in line with previous research on the close timing of SO and spindles (Cox et al., 2014; Mölle et al., 2002; M. Steriade et al., 1993). To detect SO, we identified all positive-to-negative zero crossings in the delta band along with the previous peak and following trough. The positive-to-negative zero crossings in which the time from the peak to the trough was at least 300 ms were considered SO. To minimize false detections, we focused on high amplitude SO with large peaks and troughs (*Section 3.3 Methods*; Figure 3.2B).

3.4.2 Spindles Modulate Single Unit Spiking

To observe spindle-neuron interactions, we aligned spike rasters to the peak of identified spindles. We noted that the average oscillatory firing rate of the representative example unit in Figure 3.3A-B closely matched (with a phase shift) the average spindle waveform. The similarity of firing rate and LFP during spindles suggested that spike timing is modulated during ongoing spindles. To quantify spindle modulation of neural spiking, we extracted the spindle phase at each recorded action potential (Figure 3.3C). To get the spindle phase, we used the angle of the Hilbert spindle band LFP (Figure 3.3C, Phase Plot). Then we collected the phase triggered on each spike occurring within one cycle of the spindle peak (Figure 3.3C, highlighted blue portion). This yielded a spike phase distribution (Figure 3.3D, upper), which was used to calculate the degree of phase locking for each single unit. Briefly, each spike-triggered phase was converted to a vector of unit magnitude and in the direction of the triggered phase. Then the average vector was computed and the magnitude of this vector was taken as the phase locking value, whereas the direction of this vector was taken as the preferred spindle phase (Figure 3.3D, lower). Most units had a

preferred spindle phase between $\pi/2$ and π –i.e., during the second half of the downward component of the spindle cycle (Figure 3.3D, upper). We found that neurons were significantly more phase locked to spindles compared to random epochs (Kolmogorov–Smirnov test, p=1.02e-14; Figure 3.3D, lower).

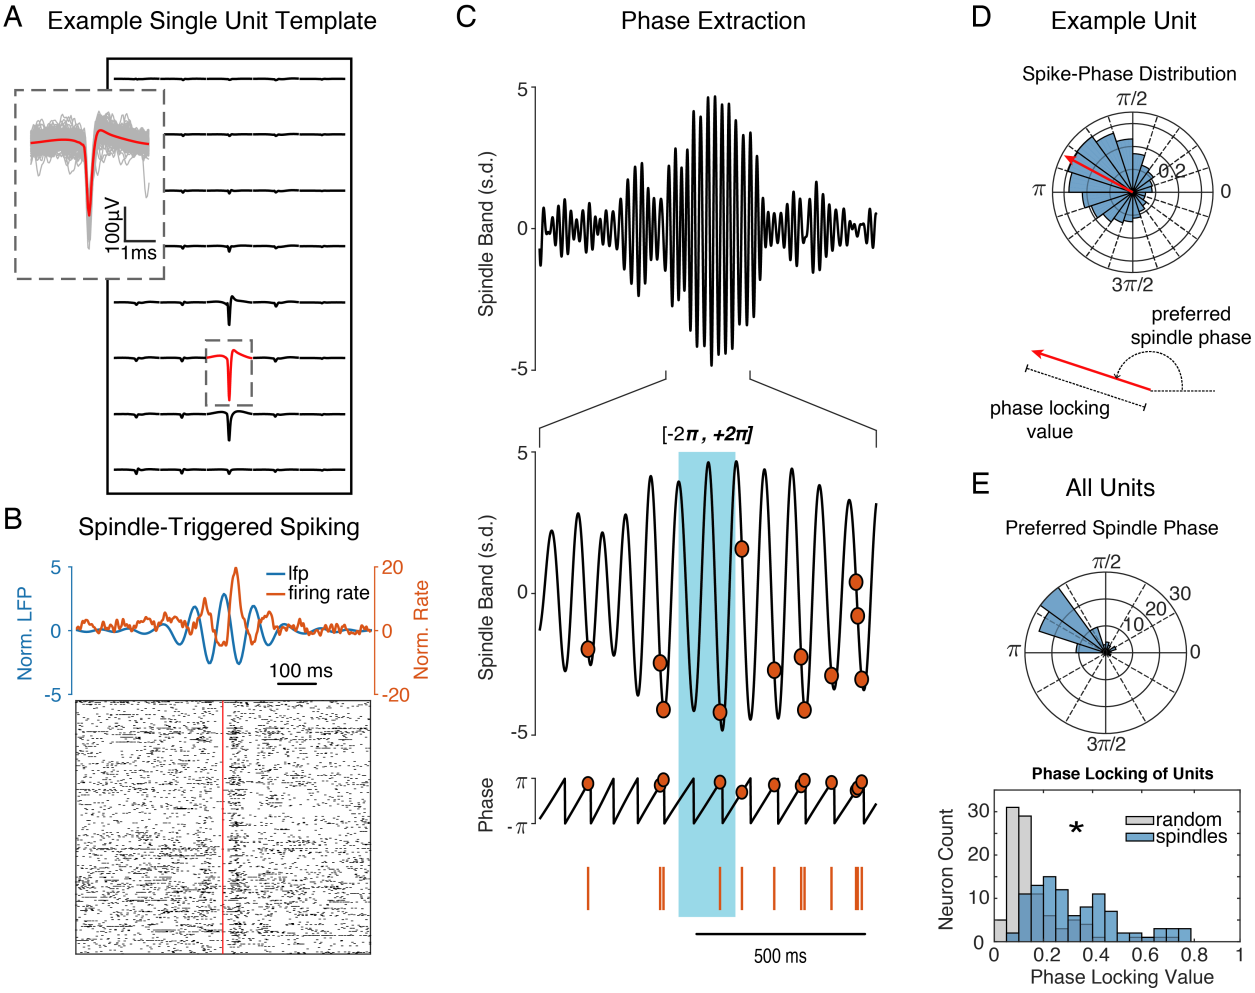

Figure 3.3 - Spindle Modulation of Spiking. (A) An example of the average waveform for a single unit recorded across many channels using a custom polytrode probe. Inset shows the spike waveforms (gray) on one channel of many spiking events along with the average waveform (red). This example unit is used for all panels in this Figure. **(B)** Spindle-triggered spiking for the example unit. The average spindle waveform (blue) is plotted with the average normalized firing rate (orange). A raster of spike times is displayed below. **(C)** Phase extraction methods. Spindle band (10-16 Hz) LFP is plotted during a detected spindle (top) and is re-plotted immediately below (2nd from the top) with a finer time resolution. The Hilbert phase of the spindle activity is plotted below (2nd from the bottom) and fluctuates between $[-\pi,\pi]$. The spiking activity is displayed on the bottom (bars) and as dots in the spindle band and phase subplots. The highlighted portion of these plots extends from $[-2\pi,2\pi]$ and shows spikes that are used to compute the phase locking value (C-D). **(D)** The spike phase distribution for the example unit, plotted as a circular histogram (blue). The average

phase vector is overlaid in red and copied below for clarity. The magnitude and direction of this vector are defined as the phase locking value and preferred spindle phase, which are collected for all units. **(E)** Summary of the phase locking value (bottom) and preferred spindle direction (top) for all units. Gray bars (bottom) represent the phase locking to the spindle band during random epochs.

3.4.3 Spindles Increase Single Unit Phase Locking

We next quantified changes in phase locking dynamically across spindles. We segmented spindles (Figure 3.4A, right) and random epochs (Figure 3.4A, left) into their component cycles and separately generated a spike phase distribution for each cycle. This process is presented for an example unit in Figure 3.4A, in which spike phase distributions were generated and displayed for two different spindle cycles, one at the peak of spindles ($[-2\pi,0]$; Figure 3.4B, left) and one at the tail of spindles ($[8\pi,10\pi]$; Figure 3.4B, right). The average preferred phase across neurons did not change during spindles, but the average spike count and phase locking increased near the spindle peaks (Figure 3.4C). We quantified these dynamics by grouping cycles into three categories: (1) RAND—the two cycles at the center of the random epochs; (2) TAIL—the two cycles farthest from the spindle peaks; and (3) PEAK—the two cycles nearest the spindle peaks (Figure 3.4D). A repeated measures ANOVA confirmed that these categories were significantly different for spike counts (p=6.283e-07) and phase locking (p=1.298e-42), but not for preferred phase (p=0.868). *Post hoc,* paired t-tests confirmed that spike counts were significantly larger near the peak of spindles (PEAK vs. TAIL t-test, p=2.136e-03); even at the spindle tails counts were larger than during random epochs (TAIL vs. RAND t-test, p=6.611e-05). Likewise, phase locking increased near spindle peaks (PEAK vs. TAIL t-test, p=4.299e-25) and was significantly larger than random epochs even 5 cycles away from the spindle peaks (TAIL vs. RAND t-test, p=6.574e-04).

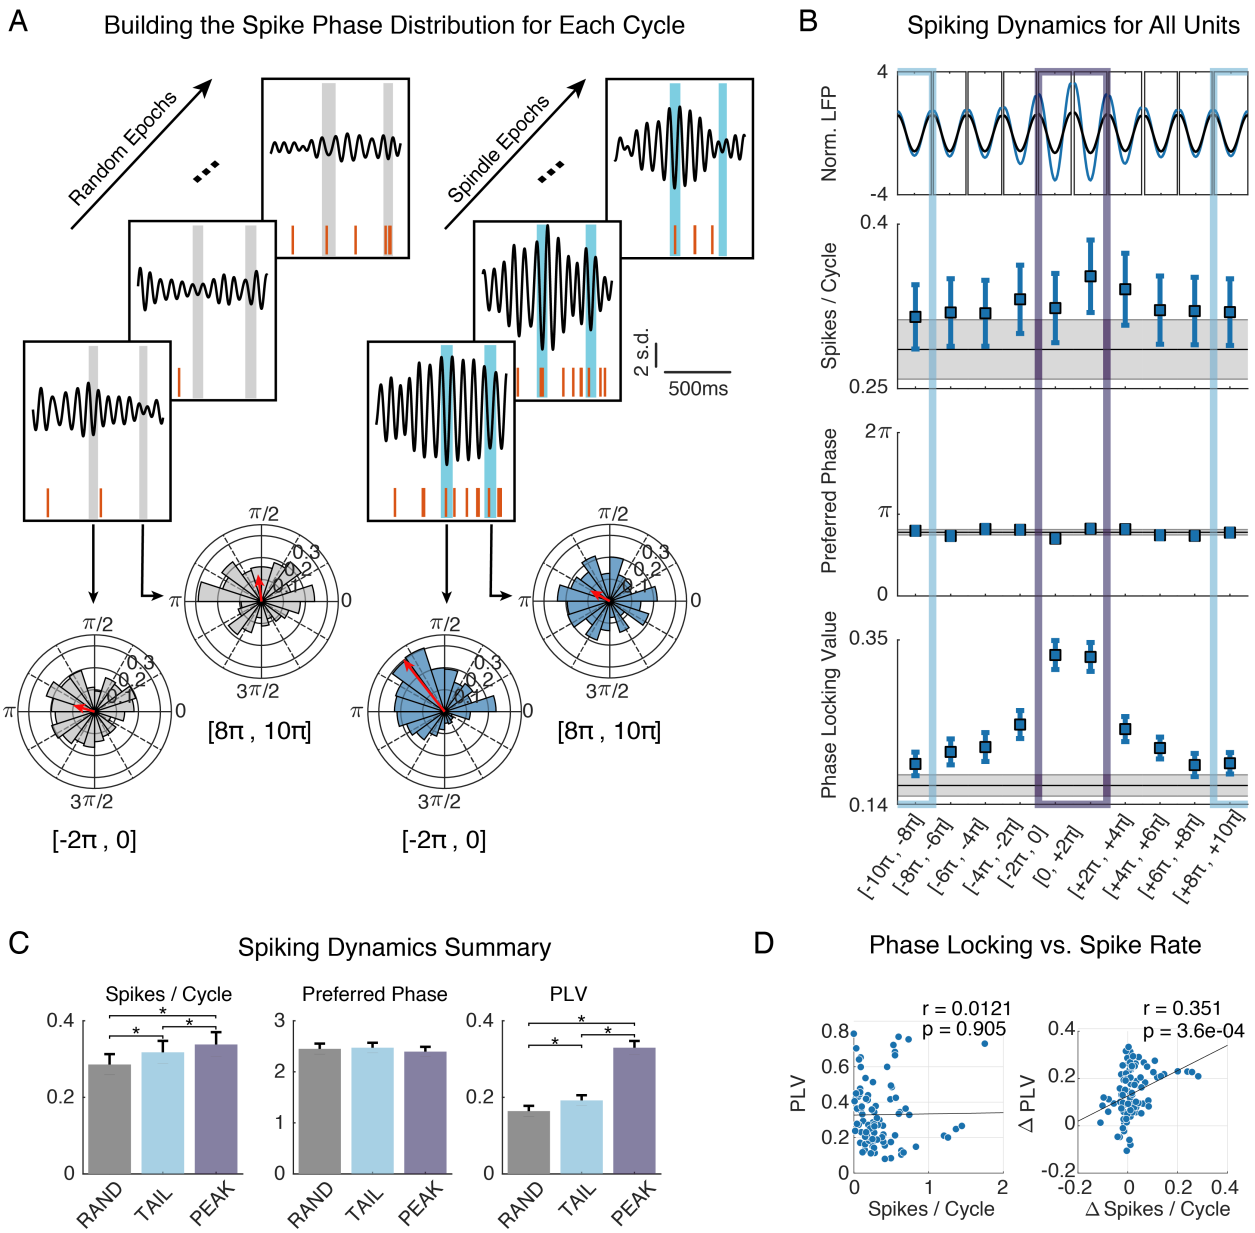

Figure 3.4 - Spindle Cycle Analysis of Phase Locking. (A) Generation of the spike phase distribution across spindle cycles. Spike-triggered phases are extracted from single unit spiking during specific spindle cycles. Phases are aggregated across actual spindle epochs (right) or random epochs (left). **(B)** Summary of all neurons' spike phase distribution statistics across spindle cycles. The average spindle band waveform for each spindle cycle (top) is plotted along with the average spike rate (2nd from top), preferred spindle direction (2nd from bottom), and phase locking value (bottom). Blue lines reflect averages during actual spindle epochs and black/gray lines reflect averages during random epochs. Error bars represent the standard error of the mean. **(C)** Summary of spiking dynamics. Spike rates, preferred spindle directions, and phase locking values are combined into three categories: (1) RAND—the two cycles at the center of the random epochs; (2) TAIL—the two cycles farthest from the spindle peaks; and (3) PEAK—the two cycles nearest the spindle peaks. The newly categorized data is reproduced as bar plots and significant differences between the categories are highlighted with asterisks. **(D)** Relationship between phase locking and spike count across neurons. Each neuron's phase locking value and spiking rate in the PEAK is shown as a scatter

plot (left) along with each neuron's change in phase locking value and spiking rate (PEAK – TAIL; scatter plot, right). Regression lines are superimposed on the scatter plots along with Pearson correlation coefficients and associated p-values.

We were curious whether the neurons that increased their firing during spindles also increased their phase locking or if these were two separate neural populations. To address this question, we ran a correlation analysis and found that there was a significant correlation between the change (PEAK – TAIL) in spike count and the change in phase locking value (Pearson's r=0.351, p=3.632e-04; Figure 3.4D, right). This implies that there is significant overlap in the populations of neurons that modulate their firing rates and spike timing during spindles. Importantly, there is no correlation between the raw spike count and the raw phase locking value of neurons (Pearson's r=0.0121, p=0.905; Figure 3.4D, left), assuaging any concerns that this finding is driven by higher estimation of phase locking value for higher firing neurons, which is known to happen with small sample sizes (Vinck, van Wingerden, Womelsdorf, Fries, & Pennartz, 2010).

3.4.4 Spindles Increase Pairwise Synchrony

To quantify pairwise synchrony changes as spindles evolve, we followed a similar procedure of segmenting spindles into their component cycles, then independently generated a cross correlation histogram (CCH) for every pair of neurons in each spindle cycle (*Section 3.3 Methods*; Figure 3.5A and 3.6A, left). Each bin in the CCH represents the probability of the two neurons under examination co-firing with a specific time difference; the peak of the CCH (Figure 3.5B and 3.6B, left) is a normalized measure of the degree of co-firing for a pair of neurons. Given the increase in phase locking near the peak of spindles, one would expect a corresponding increase in the raw pairwise co-firing

probability, which is what we observed (Figure 3.5).

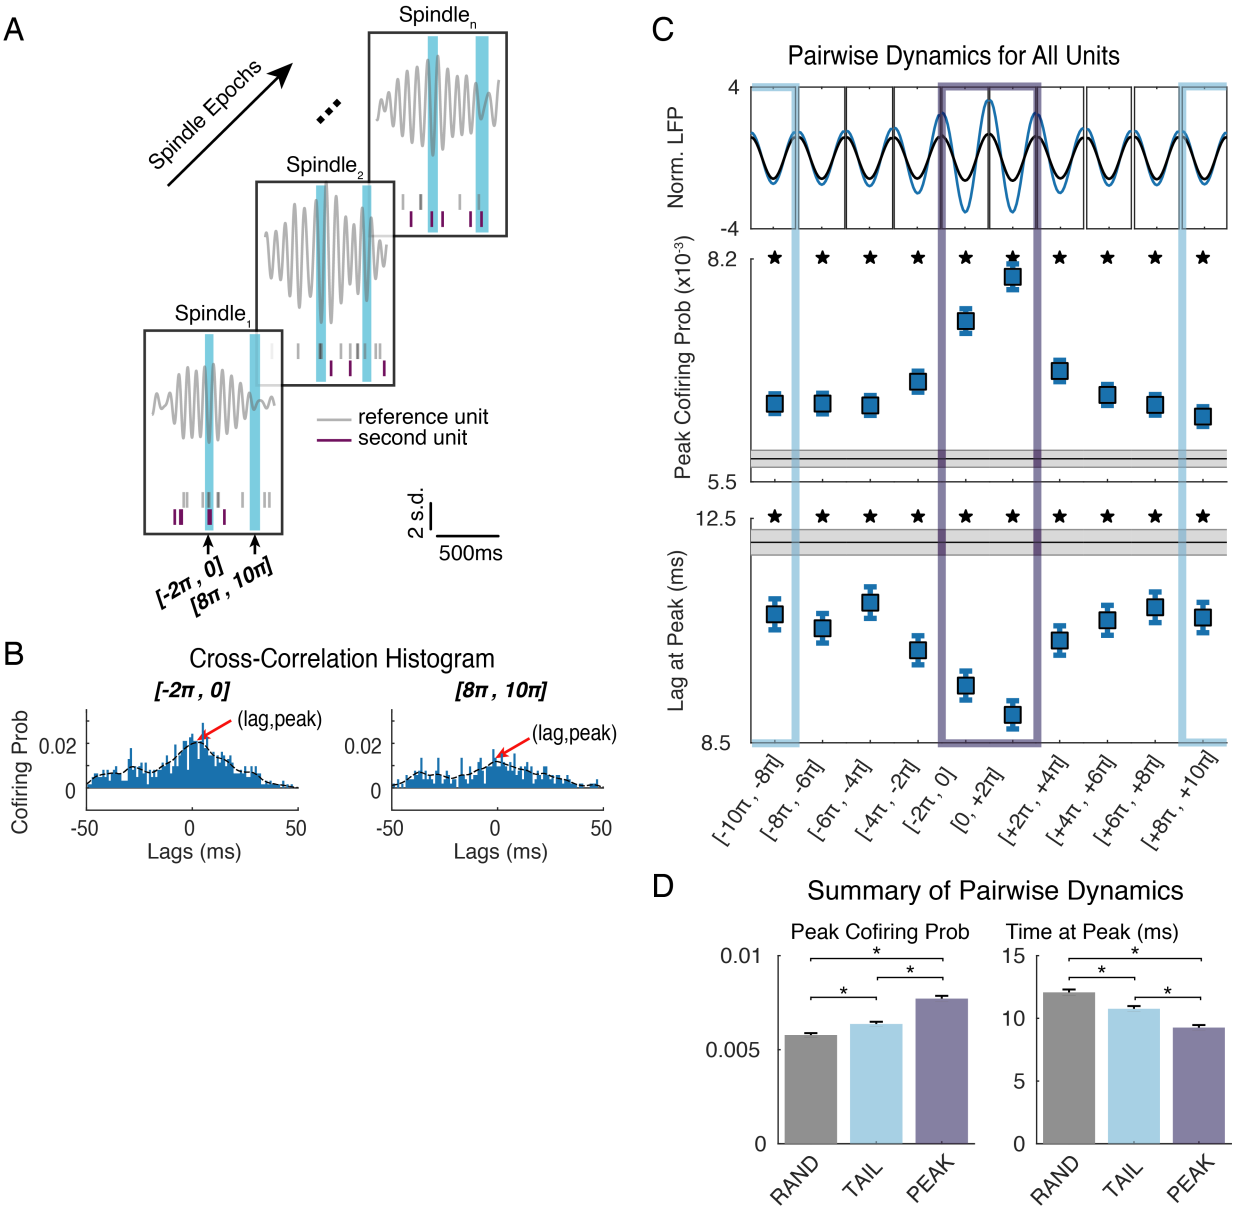

Figure 3.5 - Spindle Cycle Analysis of Synchrony Using the Raw Cross Correlation Histogram. (A) Generation of the cross correlation histogram. For a pair of units, relative spike times are extracted during specific spindle cycles. Relative spike times are aggregated across spindle epochs. **(B)** Examples of the cross correlation histogram for two spindle cycles ($[-2\pi, 0]$, left; $[+8\pi, +10\pi]$, right). The peak, and time of peak (red arrows) are collected for each pair of neurons in each spindle cycle. **(C)** Summary of all neurons' CCH statistics across spindle cycles. The average spindle band waveform for each spindle cycle (top) is plotted along with the average CCH peak (middle) and time of peak (bottom). Blue lines reflect averages during actual spindle epochs and black/gray lines reflect averages during random epochs. Error bars represent the standard error of the mean. **(D)** The peak and time of peak are reproduced for the same categories as in Figure 3D and Figure 4D. Significant differences between the categories are highlighted with asterisks.

However, changes in firing rate or phase locking can influence this measure, and we

wanted to know whether spindles modulate any additional correlation structure across neurons. We isolated changes in pairwise correlations during spindles by using a shuffling procedure, which maintained firing rates and phase locking within spindle cycles but broke correlation structure across neurons (*Section 3.3 Methods*; Figure 3.6A, right).

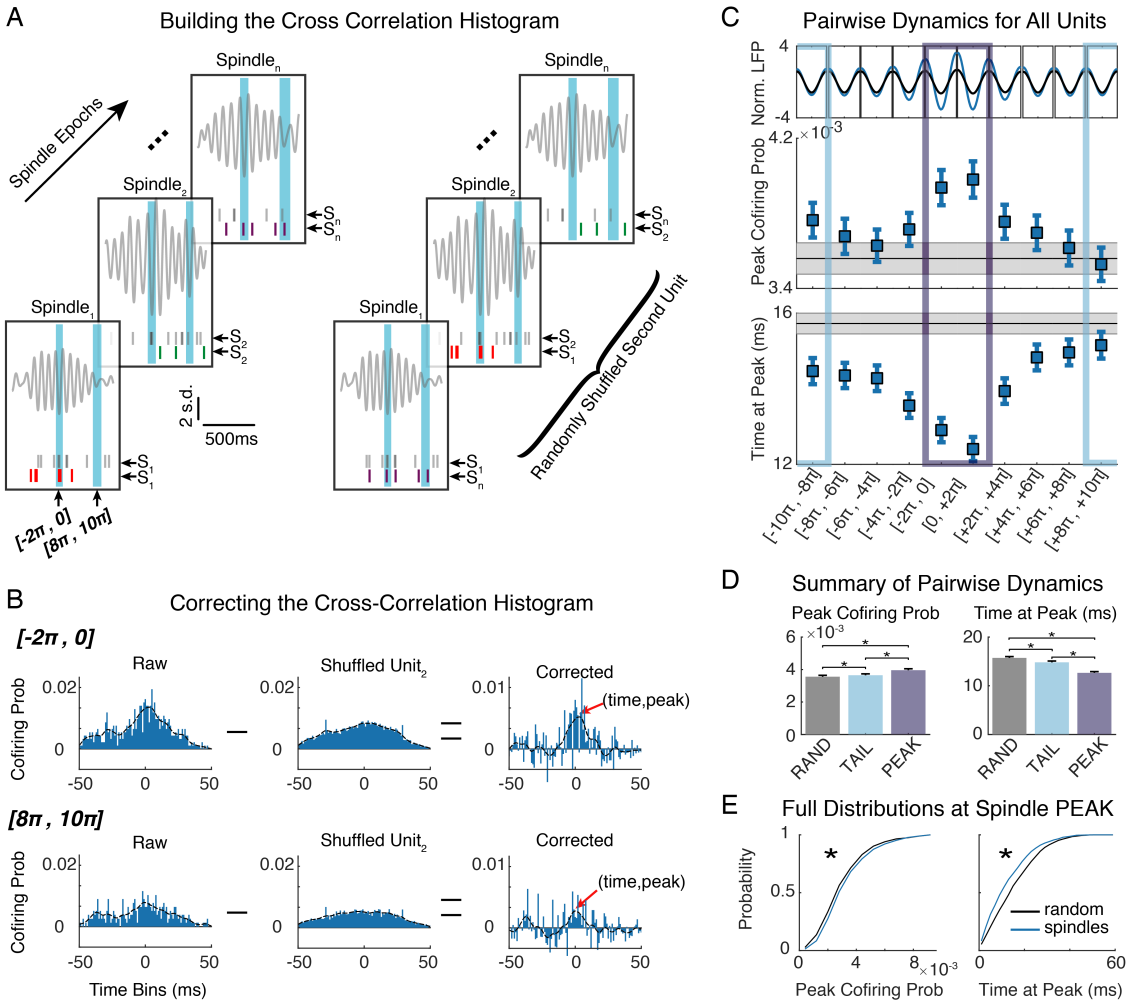

Figure 3.6 - Spindle Cycle Analysis of Synchrony. (A) Generation of the cross correlation histogram. For a pair of units, relative spike times are extracted during specific spindle cycles. Relative spike times are aggregated across spindle epochs (Raw CCH; left). Additionally, a shuffled cross correlation histogram is constructed using a similar procedure. Relative spike times are extracted during specific spindle cycles, but one neuron's relative spike times are shuffled across spindle epochs (Shuffled CCH; right). **(B)** Cross correlation histogram correction. Within each spindle cycle, the Shuffled CCH is subtracted from the Raw CCH to generate a Corrected CCH. The peak, and time of peak (red arrows) are collected for each pair of neurons in each spindle cycle. **(C)** Summary of all neurons' Corrected CCH statistics across spindle cycles. The average spindle band waveform for each spindle cycle (top) is plotted along with the average Corrected CCH peak

(middle) and time of peak (bottom). Blue lines reflect averages during actual spindle epochs and black/gray lines reflect averages during random epochs. Error bars represent the standard error of the mean. **(D)** The peak and time of peak are reproduced for the same categories as in Figure 3.4D. Significant differences between the categories are highlighted with asterisks (KS tests, $p<0.05$). **(E)** The full distributions of Peak Cofiring Probability (left) and Time at Peak Cofiring Probability (right) plotted as CDFs.

This procedure yielded a shuffled CCH (Figure 3.6B, middle), which reflected the expected pairwise correlations in a spindle cycle given the firing rates and spindle phase relationships of the two neurons being examined. We then subtracted the shuffled CCH from the raw CCH, to construct a corrected CCH. The peak of this corrected CCH measured the degree of co-firing that exclusively resulted from pairwise correlations (Figure 3.6B, right).

Surprisingly, the corrected pairwise correlations revealed similar dynamics during spindles. As before, we compared three groups: (1) RAND—the two cycles at the center of the random epochs; (2) TAIL—the two cycles farthest from the spindle peaks; and (3) PEAK—the two cycles nearest the spindle peaks (Figure 4D). We used a repeated measures ANOVA to determine that there were differences in the peak of the CCH ($p=1.020e-21$) and the time of this peak ($p=5.617e-23$) across the three categories. From the TAIL to the PEAK of spindles the co-firing probability increased (paired t-test, $p=6.302e-14$) and the absolute time of peak co-firing decreased (paired t-test, $p=2.130e-12$). Likewise, both measures were significantly different from random epochs (RAND) even at the spindle TAILs, (paired t-tests, peak co-firing: $p=0.0320$; time of peak co-firing: $p=0.0055$). Taken together, these dynamics reflect an increase in pairwise synchrony during spindles that surpasses that expected from independent neuron changes. Notably, the increase in neuron synchrony was not limited to only highly correlated neuron pairs; rather, all pairwise correlations increase during spindles (Figure 3.7).

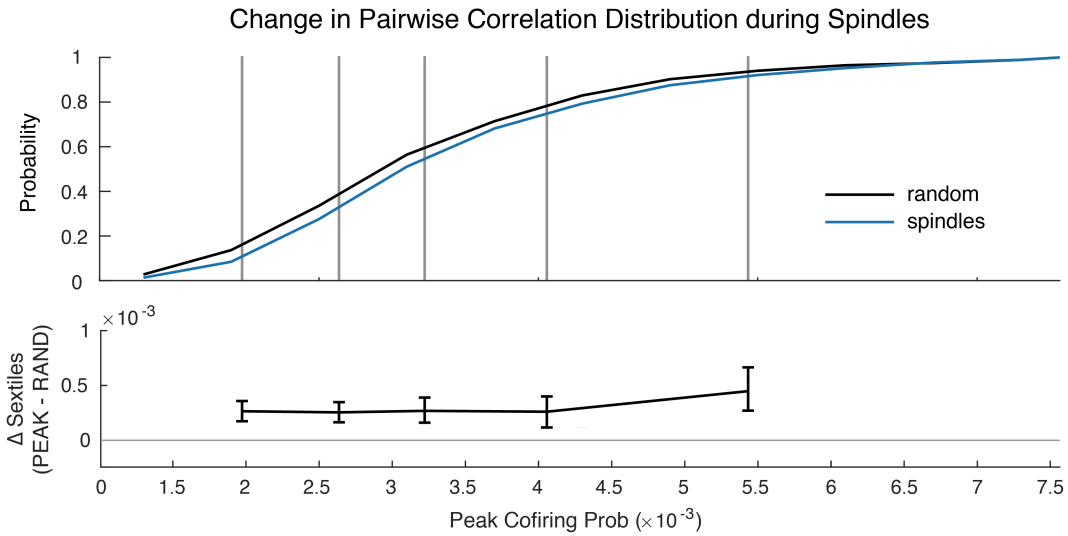

Figure 3.7 - Spindles Increase Pairwise Correlations Across the Distribution of Correlations. The distributions of induced pairwise correlations are reproduced during random epochs (black; upper) and spindle peaks (blue; upper). Gray lines divide the distributions into sextiles (upper). The difference in sextile dividers is plotted with 95% confidence intervals as error bars (lower).

3.4.5 Slow Oscillations Narrow the Distribution of Spindle-Induced Correlations

A large body of work has focused on the relationship between spindles and SO, suggesting that *nested* spindles, which occur during the up state of SO, are particularly important for learning and neuroplasticity (Latchoumane et al., 2017; Mölle et al., 2009; Niknazar et al., 2015). This work led us to hypothesize that the spiking correlation structure might be different for nested and unnested spindles. To investigate the role of SO/spindle interactions, we followed two different procedures. In the first approach we specifically asked if nested spindles organize spikes differently than unnested spindles. We first investigated the coupling of SO and spindles and found that, triggered on SO down

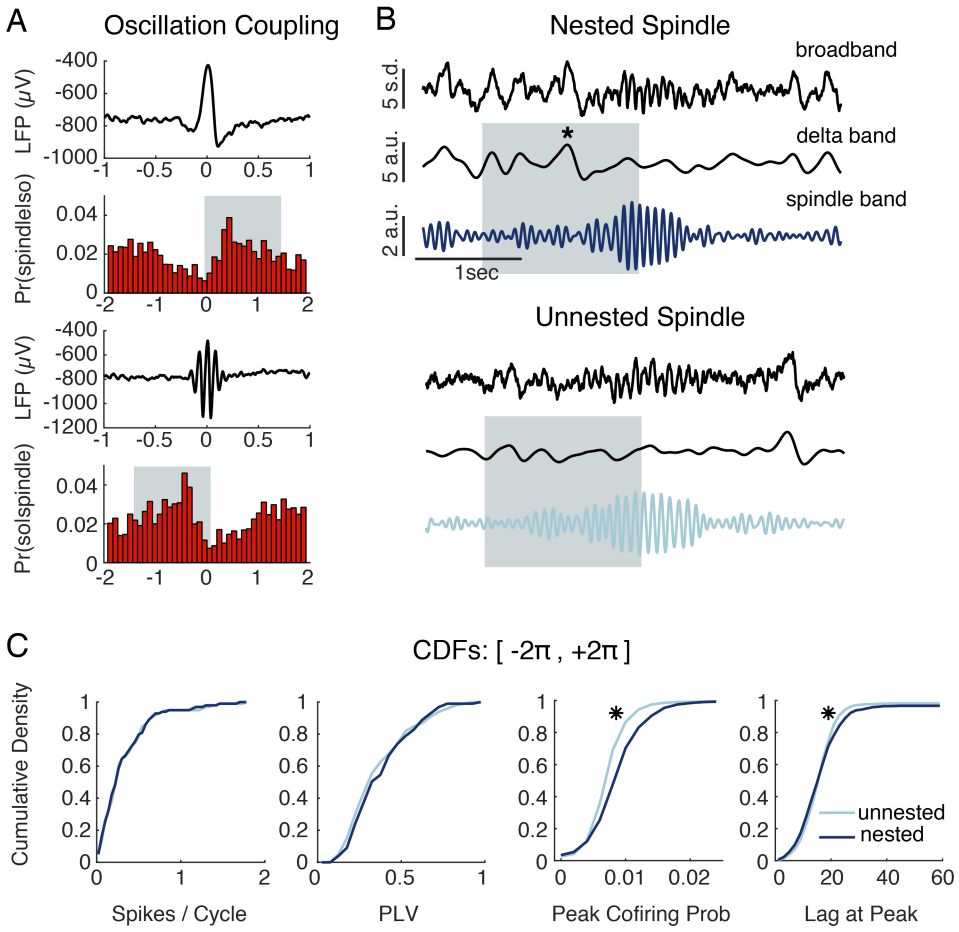

Figure 3.8 – Nested Spindles Appear to Induce Higher Pairwise Correlations. (A) SO and spindles are coupled. Average SO waveforem is plotted with SO-triggered spindle frequencies (upper). Average spindle waveforem is plotted with spindle-triggered SO frequencies (lower). Shaded region indicates an uptick in triggered frequecies. **(B)** Categorization of spindles as nested or unnested. Example traces (broadband, delta band, and spindle band) are plotted. A 1.5 sec window (shaded region) is generated for each spindle. In the upper plot a SO was present in the window so the spindle was categorized as nested (dark blue). In the lower plot a SO was absent in from the window so the spindle was categorized as unnested (light blue). **(C)** CDFs of Spike Counts, PLV, Peak Cofiring Probability, and Time at Peak Cofiring are generated during the spindle peaks. CDFs are shown separately for nested (dark blue) and unnested (light blue) spindles. Asterisks indicate significant differences in the distributions (KS tests, $p < .05$).

states, there is an increase in the likelihood of spindles for about 1.5 secs (Figure 3.8A). We therefore split all spindles into two groups: *nested*—spindles within 1.5 secs after a SO, and *unnested*—spindles outside of this window (Figure 3.8B). Focusing on the spindle peaks (where spiking activity is the most structured) we analyzed whether there were

differences in the distribution of spiking relative to the spindle oscillation and relative to other neurons. Intriguingly, there was no difference in the first-order spiking statistics. Spike counts and phase locking was not significantly different between nested and unnested spindles (Figure 3.8C, Kolmogorov–Smirnov tests: spike count, p=0.9994; PLV, p=0.5493). However, there was a significant difference in the correlation structure. Nested spindles had increased peak cofiring (Kolmogorov–Smirnov test, p=4.542e-32) and different inter-neuron spike intervals (Kolmogorov–Smirnov test, p=0.0059). This result is intriguing and in line with previous results finding differences between nested and unnested spindles. However, this approach assumes that there is a difference between nested and unnested spindles. We wanted to use a less biased method to ask if there are truly any differences between nested and unnested spindles.

In the next approach, we binned spindles (bin width: 1.5 sec; step size: 500 ms) based on their timing relative to SO (Figure 3.9A-B). The first bin (0-1.5 sec) includes *nested* spindles and subsequent bins do not. If there were a difference between nested and unnested spindles then we would expect to see differences between the first bin and the remaining bins. Accordingly, we examined the distribution of pairwise correlations (peak cofiring probabilities of the corrected CCH) during the spindle peak ($[-2\pi,+2\pi]$) in each time bin; these distributions are plotted as CDFs in Figure 3.9C.

Qualitatively, the distributions appear to narrow for spindles closer to SO (Figure 3.9C). To quantify trends in the distributions, we performed linear regression on the mean and standard deviation of the pairwise correlations across time bins: the average pairwise correlation did not change linearly across time bins (Figure 3.9C, left inset; linear regression model, p=0.070), but the standard deviation of the pairwise correlations

decreased across time bins (Figure 3.9C, right inset; linear regression model, $s.d. = 0.00041 * timeFromSO + 0.0065$, p=7.52E-08).

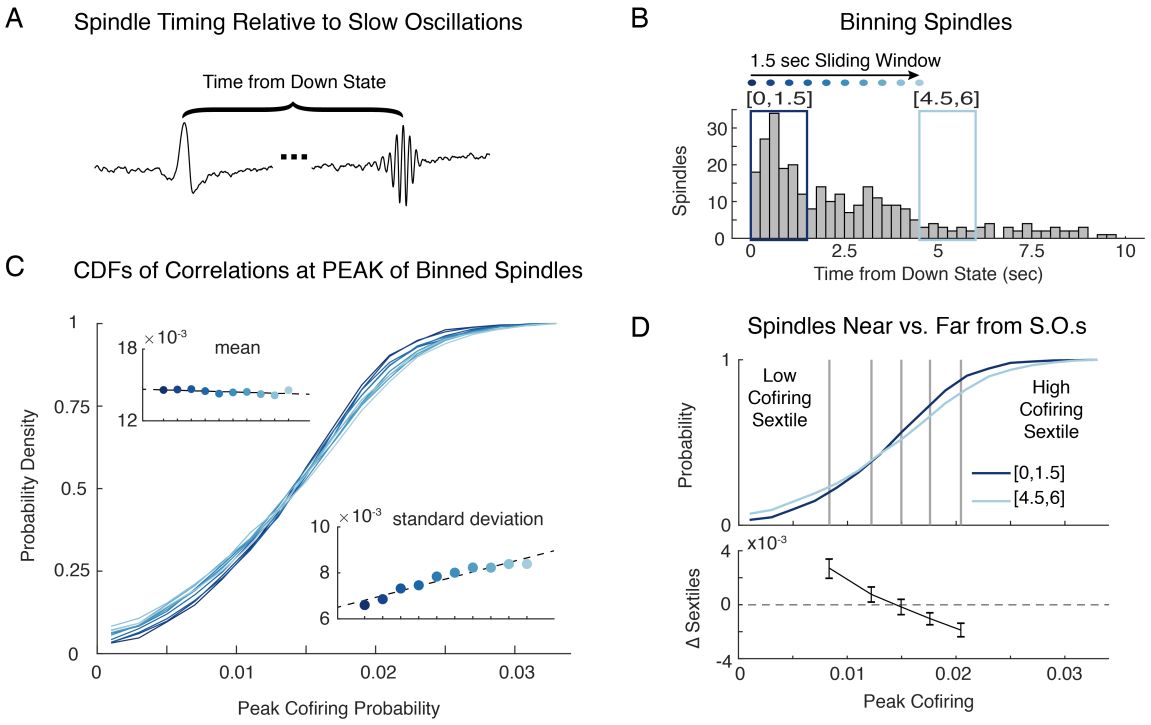

Figure 3.9 – Interaction of SO and Spindles. (A) Timing of spindles relative to SO. **(B)** Example of binning spindles based on timing relative to SO in one animal. The distribution of time delays is plotted as a bar graph. Windows representing the spindle bin closest to (dark blue rectangle), and farthest from (light blue rectangle) SO are overlaid. Intermediate window starts are represented by blue dots. The color gets lighter as the window start gets farther from the slow oscillation. **(C)** The distribution of spindle induced pairwise correlations for each bin. Colors match the window colors in B. The average pairwise correlation across bins is plotted with a linear fit (left inset), and the standard deviation of pairwise correlations across bins in plotted with a linear fit (right inset). **(D)** The distributions of spindle induced pairwise correlations are reproduced for the bins closest to SO, [0,1.5], and farthest from SO, [4.5,6], (upper). Gray lines divide the distributions into sextiles (upper). The difference in sextile dividers is plotted with 95% confidence intervals as error bars (lower).

Additionally, we used nonparametric statistics to detect differences in the distributions of pairwise correlations between the bin closest to SO, (dark blue, [0,1.5 sec]), and the bin farthest from SO, (light blue, [4.5,6 sec]). The CDFs for these distributions are reproduced in Figure 3.9D for clarity. A Kolmogorov–Smirnov test revealed a significant difference between these distributions (p=3.98E-05). We then split the distributions of

pairwise correlations into sextiles (Figure 3.9D, upper, gray lines) and computed the difference in the sextile dividers (Figure 3.9D, lower). This revealed a pattern in which the lowest correlation pairs significantly increased during spindles near SO, but the highest correlation pairs significantly decreased during spindles near SO.

3.4.6 Spindle-Induced Correlations Decay During Sleep

Many papers have noted changes in neurophysiology that occur during sleep. Many of these studies have focused on changes in delta power, SO rates, and other oscillatory phenomena (Tononi & Cirelli, 2003), but some work has also focused on changes in firing rates (Levenstein, Watson, & Rinzel, 2017; Watson, Levenstein, Greene, Gelinas, & Buzsáki, 2016). We were interested to see whether higher-order spiking structure changes dynamically across sleep. Accordingly, we split the sleep sessions into thirds, and asked whether there were significant differences in spindle-induced spiking structure between the first third of sleep and the last third of sleep (Figure 3.10A). We found that spindle-induced spike counts, phase locking, and inter-neuron spike intervals did not change across sleep (Kolmogorov–Smirnov Tests: spike count, $p=0.9592$; phase locking value, $p=0.9574$; time at peak, $p=0.3606$), but the correlation decreased (Kolmogorov–Smirnov Test: peak cofiring probability, $p=0.0307$).

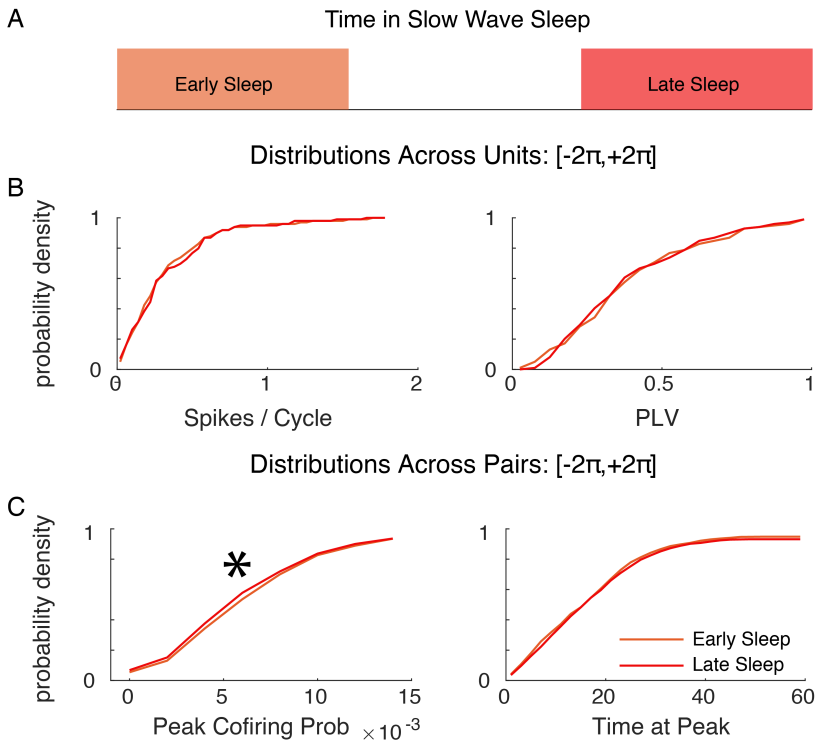

Figure 3.10 – Spindle Induced Structure Across Sleep. (A) Slow wave sleep is split into early sleep (orange) and late sleep (red). **(B)** Distributions of Spike Counts (left) and PLV (right) are plotted as CDFs for early and late sleep. **(C)** Distributions of Peak Cofiring Probabiltiy (left) and Time at Peak Cofiring (right) are plotted as CDFs for early and late sleep. For B and C, asterisks indicate significant differences between the distributions (KS test, $p<0.05$).

3.4.7 Spindle-Induced Correlations Decrease After Learning

We next asked whether spindles or spindle induced spiking structure changed with learning. To assess the role of learning, we compared baseline sleep blocks to sleep blocks immediately after exposure to a pellet retrieval motor task (Figure 3.11A). Previously, we have found that rats increase their accuracy and speed in this task (Ramanathan et al., 2015). Notably, we demonstrated that sleep supports offline improvements in speed in this task. Across ¾ of the rats in this study, spindle rates increased after learning (Figure 3.11B), and the rate of spindles nested within SOs also increased after learning (Figure 3.11C). We next focused our analysis on changes in spindle induced spiking structure. We found that there was no change in spindle-induced spike counts (Figure 3.11D, left,

Kolmogorov–Smirnov Test, p=0.9912), phase locking (Figure 3.11D, right, Kolmogorov–Smirnov Test, p=0.1457), or inter-neuron spike intervals (Figure 3.11E, right, Kolmogorov–Smirnov Test, p=0.3606), but we found a significant decrease in spindle-induced correlations after learning (Figure 3.11E left, Kolmogorov–Smirnov Test, p=0.0307).

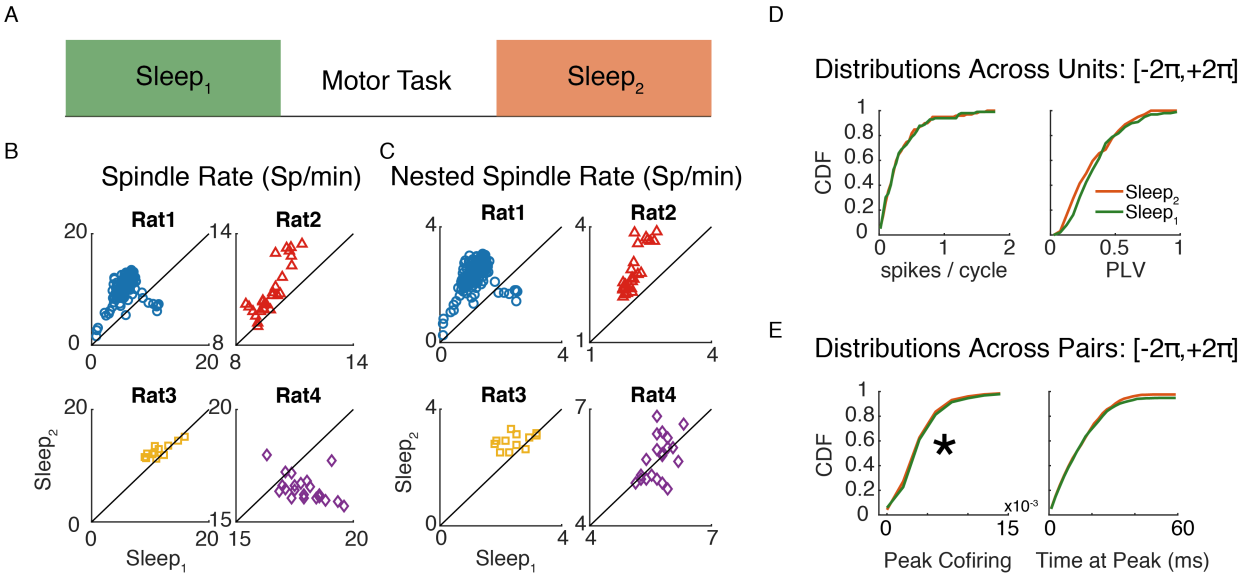

Figure 3.11 – Spindle Induced Structure Changes with Learning. (A) Baseline sleep (green) is compared to sleep after learning a motor task (orange) to assess learning-related changes. **(B)** Spindle rates detected on individual recording channels in each rat are shown as scatter plots, where x-values represent rates during baseline sleep and y-values represent rates during post-learning sleep. Black lines indicate x=y. **(C)** Same as B, but only detected the rate of nested spindles. **(D)** Distributions of Spike Counts and PLV are plotted as a CDFs for baseline and post-learning sleep. **(E)** Distributions of Peak Cofiring Probabiltiy and Time at Peak Cofiring are plotted as a CDFs for baseline and post-learning sleep. For D and E, asterisks indicate significant differences between the distributions (KS test, p<0.05).

3.5 Discussion

In this study we investigated the relationship between M1 neural firing and LFP oscillations during sleep. We focused on the structure of neural spiking during thalamocortical spindles, which are thought to be important for consolidation and for promoting neural plasticity after learning a new skill (Barakat et al., 2011; Fogel & Smith, 2006; Johnson et al., 2012; Nishida & Walker, 2007; Ramanathan et al., 2015). We found that individual neurons in M1 fired at a preferred phase of spindle cycles; by segmenting

spindles into their component cycles, we determined that such phase locking was more pronounced at spindle peaks (i.e., around the highest spindle power). Moreover, neural synchrony, a measure of interneuron relationships, shared similar dynamics during spindles; the distribution of pairwise correlations also reached its maximum at the peak of a spindle event. Interestingly, this correlation structure was modified by a spindle's proximity to a SO. Spindles that were closer to SOs exhibited a narrower distribution of correlations. In contrast, the distribution of pairwise correlations was broader during spindles that were temporally far from SOs.

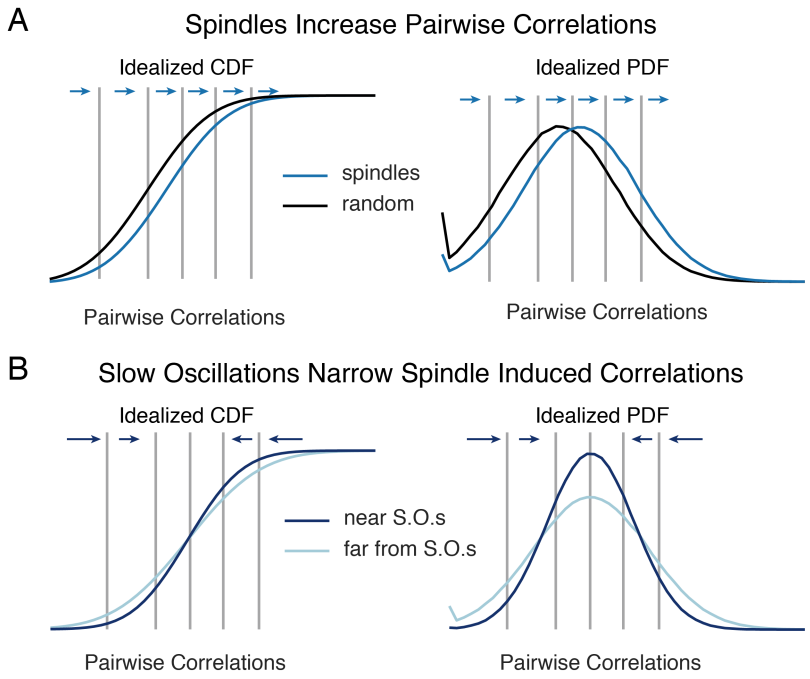

Figure 3.12 - Summary of Spindle-Induced Correlations and SO Interactions. (A) Idealized distributions (left, CDF; right, PDF) of pairwise correlations during random epochs (black) and spindle epochs (blue). Sextile dividers are plotted as gray lines and blue arrows indicate a rightward shift in the distribution of correlations across all sextiles during spindles. **(B)** Idealized distributions (left, CDF; right, PDF) of spindle induced pairwise correlations near SO (dark blue) and far from SO (light blue). Sextile dividers are plotted as gray lines and dark blue arrows indicate a tightening of the distribution of correlations for spindles nearer to SO.

3.5.1 Spindles & Changes in Neural Synchrony

One of our key results is that there are changes in both the phase locking and the

correlation structure of spiking during spindle cycle dynamics. We found that spiking activity in M1 became significantly more structured with each spindle cycle, leading to maximum changes at spindle peaks. Spike correlations have long been thought to drive neuroplasticity (Hebb, 1949), and more recently, spike timing dependent plasticity (STDP) models (G. Bi & Poo, 2001; G. Q. Bi & Poo, 1998; Feldman, 2012; Markram et al., 1997; Shulz & Jacob, 2010) have been developed, which emphasize the role of precise spike timing in neuroplasticity; importantly, pairs of spikes occurring within a STDP learning window (~ 50 ms) lead to direct, predictable changes in synaptic efficacy, but additional factors such as firing rates and network activity also modify this learning window. In this study we demonstrated that spindles increase firing rates (Figure 3.4B-C) and modulate the precise (~ 5 ms) timing of spiking activity relative to ongoing spindles (Figure 3.4B-C) and relative to other M1 neurons (Figure 3.6C-D, 3.12A). This spindle modulation likely engages known synaptic plasticity mechanisms to support learning.

The reported changes in neural synchrony during spindles are also consistent with previous findings in which spindles are linked to reactivations of task-related neural activity patterns (Ramanathan et al., 2015). However, we see an increase across the distribution of pairwise correlations (Figure 3.7), which suggests that in addition to reactivating specific activity patterns, spindles trigger a general increase in local functional connectivity. In support of this local functional connectivity function, the rate of spindles increases after learning. However, perhaps counter-intuitively, the average spindle-induced increase in spike correlations is larger before learning than after learning (Figure 3.11). Taken together, the increase in local spindles certainly suggests that after learning there would be an overall increase in local functional connectivity, but the role of each

individual spindle might be less impactful.

3.5.2 Sources of Spindle-Induced Synchrony

Correlation measurements are notoriously sensitive to a variety of factors including firing rates, excitability, sample sizes, and other first-order effects such as spike timing relative to external stimuli or internal LFP oscillations (Barreiro & Ly, 2017; Brody, 1999; Engelhard et al., 2013; Tchumatchenko, Geisel, Volgushev, & Wolf, 2011). It is difficult to disentangle second-order correlations from the impact of first-order changes, but we took several steps to do just that. First, we compared correlations for the same neuron pairs across conditions (i.e., during random epochs and during spindle cycles). Second, we used a shuffling procedure that preserved first-order statistics, including spike counts and phase-locking values of individual neurons, while eliminating second-order correlations across neuron pairs. This shuffling procedure allowed us to analyze (1) raw correlations, which are influenced by both first- and second-order statistics; (2) shuffled correlations, which reflect the expected correlations given the first-order statistics; and (3) corrected correlations, which reflect the pure second-order statistics.

These three correlation measures give us insight into the sources of spindle-induced correlations. The raw correlations increase during spindles; importantly, these raw values reflect the true activity correlations that neurons experience and likely influence neuroplasticity. Predictably, the shuffled correlations also increase during spindles. This suggests that first-order changes in spiking, such as the increase in firing rates and phase locking (Figure 3.6B-C), drive some of the increases in correlations during spindles. Surprisingly, the corrected correlations also increase during spindles. This suggests that in

addition to the impact from first-order changes, there are also second-order changes that increase pairwise correlations. These likely reflect increases in shared input to the M1 neurons during spindles (Destexhe, Contreras, Destexhe, Sejnowski, & Steriade, 1997; M. Steriade, 2000; Tchumatchenko et al., 2011), but could also reflect changes in brain state or synaptic efficacy (Tchumatchenko et al., 2011).

3.5.3 Interaction between Slow Oscillations & Spindles

SO are one of the most prominent signals during sleep, and several theories have been put forth about their functional roles (Genzel, Kroes, Dresler, & Battaglia, 2014; Rasch & Born, 2013). One framework that connects learning, SO, spindles, and their interaction is the *active system consolidation hypothesis*. This theory posits that a functional role for sleep oscillations is to coordinate and organize spiking activity across different brain regions (see Rasch and Born, 2013 for a comprehensive review). This has mostly been studied in the declarative memory system, where cortical SO nest thalamocortical spindles, which nest hippocampal sharp-wave ripples. Such hierarchical nesting is proposed to coordinate activity across brain regions, thereby enabling the transfer of stereotyped spiking patterns from short-term memory in the hippocampus to long-term memory storage in the prefrontal cortex (Latchoumane et al., 2017; Rasch & Born, 2013). A possible parallel to this memory transfer system has been found in the motor system, where SO and spindles have been shown to be time-locked to reactivated activity patterns from recently learned motor tasks (Ramanathan et al., 2015); it is unclear whether this phenomenon is related to hippocampal sharp-wave ripples.

In our study, we found that SO and spindles are coupled (Figures 3.2C, 3.8B, and

3.9B), and that the temporal proximity of these oscillations modifies the distribution of neural correlations (Figure 3.9C-E). When spindles were temporally distant from a SO, the distribution of pairwise correlations contained both larger and smaller values. In contrast, when spindles were temporally close to a SO, the distribution of pairwise correlations transiently narrowed. Interestingly, there appeared to be a linear relationship between the exact temporal proximity and the extent of distribution tightening. To understand the functional role of this narrowing, it is helpful to focus on the tails of the correlation distribution. Neurons that are rarely active together (low correlation pairs) begin to fire together more often, whereas more stereotyped neural activity (high correlation pairs) is reduced. The effect of narrowing the correlation distribution is consistent with exploration of novel neural activity patterns. This novel exploration function is in line with research that demonstrates a link between sleep and improvements in generalization and insight (Djonlagic et al., 2009; Ellenbogen, Hu, Payne, Titone, & Walker, 2007; Fenn, Nusbaum, & Margoliash, 2003; Wagner, Gais, Haider, Verleger, & Born, 2004). Mechanistically, SO are known to reflect a depolarizing current that causes a brief increase in neural firing (Luczak, Bartho, Marguet, Buzsaki, & Harris, 2007). We have confirmed a transient increase in spiking after SO in our data. Interestingly, this increase is significantly more substantial for nested SO (i.e., SOs that contain a spindle peak within 1.5s after the down state, Figure 3.13). It is possible that this transient depolarization modifies the correlation structure by adding energy to the network, causing an increase in new cofiring patterns. In contrast, the spiking correlation structure during independent spindles likely reflects the 'baseline' cortical connectivity.

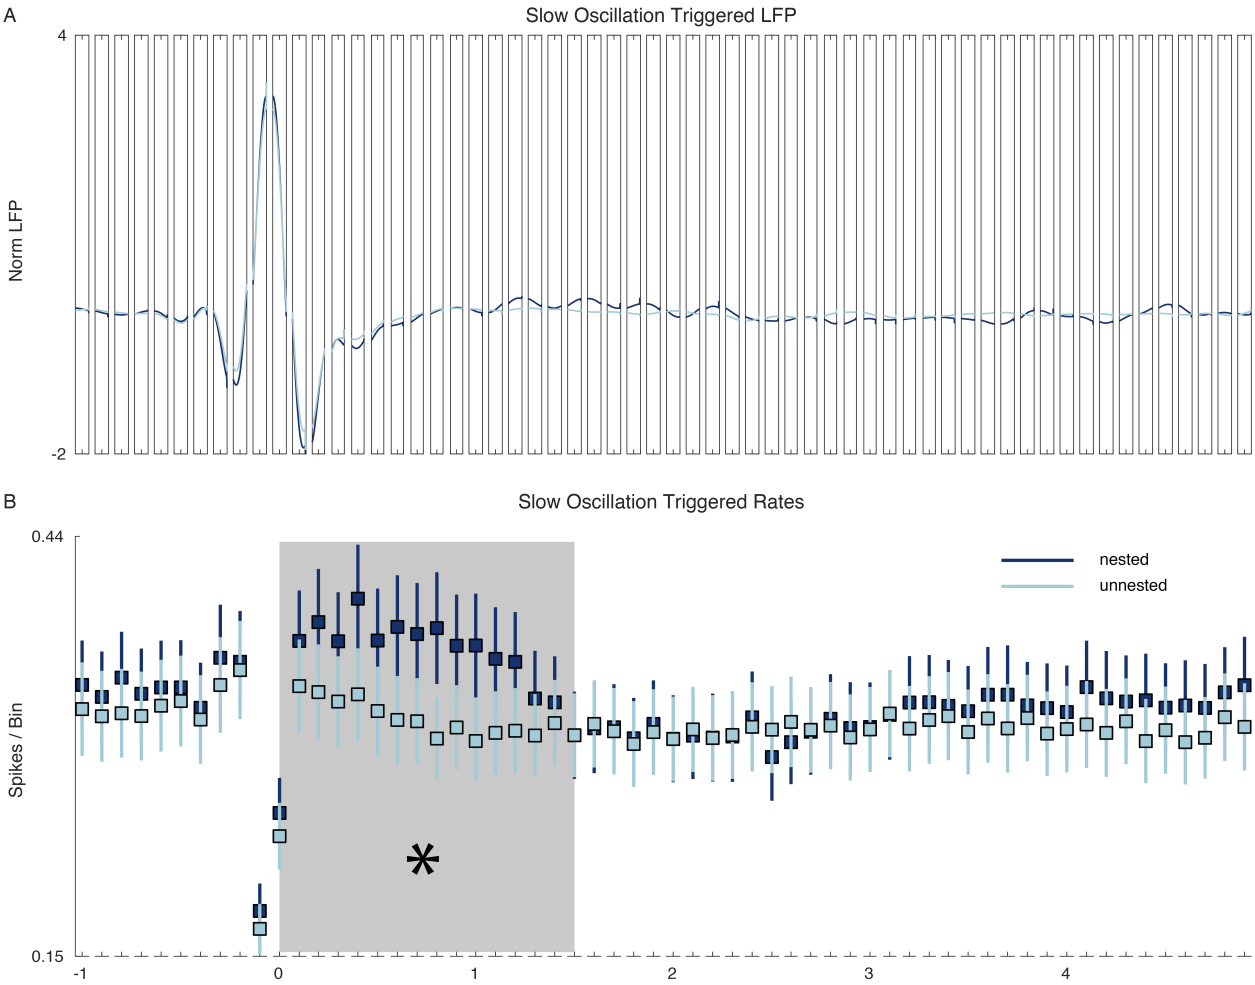

Figure 3.13 – Slow Oscillation Triggered Spiking. (A) Average SO waveforms broken down into 100 ms bins. Waveforms are plotted separately for nested (dark blue) and unnested SOs. **(B)** Average spike counts in each 100 ms bin, also plotted separately for nested and unnested SOs. Error bars reflect standard errors across all recorded neurons. Shaded region is a 1.5 sec window used to organize SOs as nested and unnested. This region also reflects nominal SO up states. Asterisk signifies a significant difference in firing rates between nested and unnested SOs in the 1.5 sec window (paired t-test, p=5.8612e-06).

Additionally, the active system consolidation framework proposes that spindles nested within SO are a distinct phenomenon from non-overlapping SO or spindles. Several studies have accordingly categorized SO and spindles as nested or unnested and observed significant differences between these groups in prefrontal cortex (Latchoumane et al., 2017; Niethard, Ngo, Ehrlich, & Born, 2018). Rather than just using a binary categorization, we decided to compare spindles across a range of time delays after SO. Using this organization we did not find a stark difference between nested and barely unnested

spindles; rather, we observed a relatively smooth progression, where the distribution of pairwise correlations appears to become more narrow as spindles occur closer to SO.

3.5.4 Implications for Homeostasis

In addition to active consolidation, sleep is thought to subserve a homeostatic role. This is, in part, based on observed changes over the course of sleep in LFP power and synaptic strength (González-Rueda et al., 2018; Tononi & Cirelli, 2003). This work has emphasized the role of slow waves and SO up-states in synaptic downscaling. Recent work has demonstrated that the distribution of firing rates homogenizes during sleep (Watson et al., 2016). Watson et al. showed that low-firing neurons increase their activity and high-firing neurons decrease their activity after NREM sleep episodes, and this homogenization of neural activity is apparent during SO up-states.

Since SO play a role in homeostasis, we wanted to reinterpret our results within this framework. We focused on the role of SO in modifying spindle-induced correlations and asked whether the narrowing of correlations that we observed could play a role in homeostasis during sleep. Though this is a complicated network, we know that spike correlations drive neuroplasticity (G. Bi & Poo, 2001; Hebb, 1949), and that there is strong evidence that spike correlations are positively correlated with neural firing rates (Bair, Zohary, & Newsome, 2001; Barreiro & Ly, 2017; Lin, Okun, Carandini, & Harris, 2015; Schulz, Sahani, & Carandini, 2015). Given these observations, the narrowing of correlations that we observe is entirely consistent with the homogenization of firing rates during NREM sleep. Moreover, even in our relatively short (~2 hour sleep blocks) the measured spindle-induced correlations decrease across sleep. Interestingly, we do not see changes in spindle-

induced firing rates during this time. More work is necessary to illuminate whether the narrowing of correlations could cause changes to LFP power or firing rates outside the influence of spindles.

Chapter 4
Conclusions

4.1 Summary of Contributions

In this thesis, I had the sweeping objective of exploring the neural basis of how we are able to learn new sensorimotor skills. I began with a brief review of two ways that neuroscientists have traditionally studied sensorimotor learning: external influences on synaptic plasticity and behavioral modification during learning. Connecting these two fields of research is critical to make progress towards understanding how neural activity supports sensorimotor learning. At its essence, this thesis is an attempt to try to bridge the two levels by exploring connectivity and plasticity at the meso-scale in which neuron populations coordinate their activity and connectivity in order to support functional reorganization and sensorimotor learning. In this thesis, I use two vastly different approaches and two different animal models to study what functional reorganization during learning looks like at the meso-scale.

4.1.1 Using Low Frequency Stimulation to Probe Meso-scale Connectivity & Plasticity

In my first set of experiments, I extended classic synaptic plasticity experimental methods to study changes in functional connectivity in the sensorimotor system. We used a meso-scale optogenetic interface to simultaneously stimulate and record from population activity. Little is understood about connectivity at this scale. Therefore, our first endeavor was to develop two methods to measure functional connectivity across a network. These methods relied on both analyzing statistical dependencies in spontaneous activity and analyzing how activity propagated though the sensorimotor network in response to focal stimulation. We tracked these measures of network connectivity while employing simple

and eventually more complicated stimulation protocols. Within minutes of stimulating either area, the inter-area functional connectivity strengthened, but upon closer examination of the connectivity patterns, we discovered that stimulation led to variable connectivity changes across the network. We then tried to figure out why stimulation had such heterogeneous effects. From the vast synaptic plasticity literature, we reasoned that the neuroplasticity that we observed might depend on activity-dependent correlations. In fact, the models that successfully predicted functional plasticity relied on both activity-dependent correlations and knowledge of the baseline functional connectivity across the network. This work provides strong evidence that Hebbian plasticity models can be applied to meso-scale circuits; we were able to show that the specific correlations introduced by stimulation-evoked activity led to predictable changes in functional connectivity, consistent with models of Hebbian plasticity.

4.1.2 Meso-scale Oscillations Coordinate Spiking Activity to Support Sensorimotor Learning

In my second set of experiments, I sought to study natural sensorimotor learning and connect changes at the cellular level to changes in meso-scale plasticity. To uncover this link I took advantage of two neural oscillations that are known to support learning and consolidation in the sensorimotor system; during sleep, spindles and slow oscillations (SOs) are both thought to provide an important function in memory consolidation. The temporal overlap between the two oscillatory events has been theorized to be particularly important for learning and consolidation. Spiking activity during these meso-scale oscillations has been characterized at the level of individual neurons; however, the effect of

spindles on the correlation structure of neural firing was largely unstudied. Given the known relevance of neural correlations in synaptic plasticity, we reasoned that these meso-scale oscillations might support learning by changing the correlation structure of spiking activity. To investigate the relationship between the meso-scale and neuron correlations, we simultaneously recorded local field potentials (LFPs) and spiking activity from electrode arrays in M1 of four sleeping rats. This enabled us to examine the precise relationship between spike timing relative to ongoing spindle oscillations and SO. By deconstructing spindles into separate cycles we were able to track spiking structure dynamically during the evolution of spindles. This analysis revealed changes in spiking structure that followed the waxing and waning of the spindle itself. During spindles spike rates increased, timing was more phase-locked to the spindle frequency band, and local synchrony resonated to a maximum at the peak of spindles. This analysis also revealed that SO and spindles interact—the distribution of spindle-induced correlations narrowed when spindles were modified by SO. This modification of the correlation structure is consistent with exploration of novel neural states and may be a key mechanism through which the interaction of meso-scale oscillations can support sensorimotor consolidation.

4.2 Future Directions

To gain a deeper understanding of how we learn to better interact with the world, we need to relate neural activity at different levels to improvements in behavior. In this thesis I attempted to both connect meso-scale activity to the activity of individual neurons and to directly model plasticity at the meso-scale. This work is important, but remains far from finished. There are two critical next steps that should be taken: (1) correlation

analysis and causal perturbation analysis of meso-scale activity during behavior and throughout learning (2) measure and track meso-scale connectivity across different brain regions. One of the dominant theories for why we have population-level oscillations is to coordinate activity across different brain areas. Although I did relate functional plasticity within and across M1 and S1, the field of neuroscience needs to expand meso-scale measurements of coordinated activity between more cortical areas and between cortical and subcortical areas. To develop a more comprehensive view of the neural underpinnings of sensorimotor control and learning, understanding the relationship between all nodes in the sensorimotor network is imperative.

Additionally, this work has significant implications for clinical research and clinical treatment. In this thesis, we have only indirectly touched upon the diverse field of neurostimulation for clinical disorders. I'll briefly mention that neurostimulation has been used to treat a diverse palette of clinical conditions including cochlear implants (Peterson, Pisoni, & Miyamoto, 2010), stroke rehabilitation (Kubis, 2016; Ramanathan et al., 2018), motor disorders (Collomb-Clerc & Welter, 2015; Herrington, Cheng, & Eskandar, 2016), psychological disorders (Herrington et al., 2016; Kellner et al., 2012; Oluigbo, Salma, & Rezai, 2012), and a variety of other conditions. Notably, in many of these applications, neurostimulation has a small initial impact; rather, its effects evolve over the course of weeks or months, implying that stimulation is likely engaging neuroplasticity mechanisms. Yet, we have almost no understanding of how neural stimulation effects connectivity locally or impacts connectivity across brain networks. Understanding the mechanisms of neuroplasticity in response to neural stimulation and during natural oscillations will enable researchers and clinicians to develop more effective, targeted stimulation technologies and

protocols, which could have an immediate and tangible impact on society.

Bibliography

Abrahamsson, T., Lalanne, T., Watt, A. J., & Sjostrom, P. J. (2016). In Vitro Investigation of Synaptic Plasticity. *Cold Spring Harb Protoc*, *2016*(6), pdb top087262. http://doi.org/10.1101/pdb.top087262

Andersen, N., Krauth, N., & Nabavi, S. (2017). Hebbian plasticity in vivo: relevance and induction. *Curr Opin Neurobiol*, *45*, 188–192. http://doi.org/10.1016/j.conb.2017.06.001

Athalye, V. R., Ganguly, K., Costa, R. M., & Carmena, J. M. (2017). Emergence of Coordinated Neural Dynamics Underlies Neuroprosthetic Learning and Skillful Control. *Neuron*, *93*(4), 955–970.e5. http://doi.org/10.1016/j.neuron.2017.01.016

Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: time scales and relationship to behavior. *Journal of Neuroscience*, *21*(5), 1676–97. http://doi.org/21/5/1676 [pii]

Barakat, M., Doyon, J., Debas, K., Vandewalle, G., Morin, A., Poirier, G., … Carrier, J. (2011). Fast and slow spindle involvement in the consolidation of a new motor sequence. *Behavioural Brain Research*, *217*(1), 117–121. http://doi.org/10.1016/j.bbr.2010.10.019

Baranyi, A., & Fehér, O. (1981). Synaptic facilitation requires paired activation of convergent pathways in the neocortex. *Nature*, *290*(5805), 413–5. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6261140

Barreiro, A. K., & Ly, C. (2017). When do correlations increase with firing rates in recurrent networks? *PLoS Computational Biology*, *13*(4), 1–30. http://doi.org/10.1371/journal.pcbi.1005506

Bastos, A. M., & Schoffelen, J. M. (2015). A Tutorial Review of Functional Connectivity Analysis Methods and Their Interpretational Pitfalls. *Front Syst Neurosci*, *9*, 175. http://doi.org/10.3389/fnsys.2015.00175

Bi, G., & Poo, M. (2001). Synaptic modification by correlated activity: Hebb's postulate revisited. *Annu Rev Neurosci*, *24*, 139–166. http://doi.org/10.1146/annurev.neuro.24.1.139

Bi, G. Q., & Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *18*(24), 10464–10472. http://doi.org/10.1038/25665

Binder, S., Rawohl, J., Born, J., & Marshall, L. (2014). Transcranial slow oscillation stimulation during NREM sleep enhances acquisition of the radial maze task and modulates cortical network activity in rats. *Frontiers in Behavioral Neuroscience*, *7*, 220. http://doi.org/10.3389/fnbeh.2013.00220

Bliss, T. V. P., & Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path.

The Journal of Physiology, *232*(2), 331–356. http://doi.org/10.1113/jphysiol.1973.sp010273

Bliss, T. V, Collingridge, G. L., & Morris, R. G. (2014). Synaptic plasticity in health and disease: introduction and overview. *Philos Trans R Soc Lond B Biol Sci*, *369*(1633), 20130129. http://doi.org/10.1098/rstb.2013.0129

Brody, C. D. (1999). Correlations Without Synchrony. *Neural Computation*, *11*(7), 1537–1551. http://doi.org/10.1162/089976699300016133

Buonomano, D. V., & Merzenich, M. M. (1998). CORTICAL PLASTICITY: From Synapses to Maps. *Annual Review of Neuroscience*, *21*(1), 149–186. http://doi.org/10.1146/annurev.neuro.21.1.149

Buzsáki, G., Anastassiou, C. A., & Koch, C. (2012). The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nature Reviews. Neuroscience*, *13*(6), 407–20. http://doi.org/10.1038/nrn3241

Cassenaer, S., & Laurent, G. (2012). Conditional modulation of spike-timing-dependent plasticity for olfactory learning. *Nature*, *482*(7383), 47–52. http://doi.org/10.1038/nature10776

Chung, J. E., Magland, J. F., Barnett, A. H., Tolosa, V. M., Tooker, A. C., Lee, K. Y., … Greengard, L. F. (2017). A Fully Automated Approach to Spike Sorting. *Neuron*, *95*(6), 1381–1394.e6. http://doi.org/10.1016/j.neuron.2017.08.030

Clark, S. A., Allard, T., Jenkins, W. M., & Merzenich, M. M. (1988). Receptive fields in the body-surface map in adult cortex defined by temporally correlated inputs. *Nature*, *332*(6163), 444–445. http://doi.org/10.1038/332444a0

Clemens, Z., Fabó, D., & Halász, P. (2005). Overnight verbal memory retention correlates with the number of sleep spindles. *Neuroscience*, *132*(2), 529–535. http://doi.org/10.1016/j.neuroscience.2005.01.011

Clemens, Z., Fabó, D., & Halász, P. (2006). Twenty-four hours retention of visuospatial memory correlates with the number of parietal sleep spindles. *Neuroscience Letters*, *403*(1–2), 52–56. http://doi.org/10.1016/j.neulet.2006.04.035

Collomb-Clerc, A., & Welter, M.-L. (2015). Effects of deep brain stimulation on balance and gait in patients with Parkinson's disease: A systematic neurophysiological review. *Neurophysiologie Clinique/Clinical Neurophysiology*, *45*(4–5), 371–388. http://doi.org/10.1016/j.neucli.2015.07.001

Cox, R., van Driel, J., de Boer, M., & Talamini, L. M. (2014). Slow Oscillations during Sleep Coordinate Interregional Communication in Cortical Networks. *Journal of Neuroscience*, *34*(50), 16890–16901. http://doi.org/10.1523/JNEUROSCI.1953-14.2014

Dayan, E., & Cohen, L. G. (2011). Neuroplasticity Subserving Motor Skill Learning. *Neuron*, *72*(3), 443–454. http://doi.org/10.1016/J.NEURON.2011.10.008

Debanne, D., Gähwiler, B. H., & Thompson, S. M. (1994). Asynchronous pre- and postsynaptic activity induces associative long-term depression in area CA1 of the rat

hippocampus in vitro. *Proceedings of the National Academy of Sciences of the United States of America*, *91*(3), 1148–52. http://doi.org/10.1073/PNAS.91.3.1148

Debanne, D., Gähwiler, B. H., & Thompson, S. M. (1997). Bidirectional Associative Plasticity of Unitary CA3-CA1 EPSPs in the Rat Hippocampus In Vitro. *Journal of Neurophysiology*, *77*(5), 2851–2855. http://doi.org/10.1152/jn.1997.77.5.2851

DeSalvo, M. N., Douw, L., Tanaka, N., Reinsberger, C., & Stufflebeam, S. M. (2014). Altered structural connectome in temporal lobe epilepsy. *Radiology*, *270*(3), 842–848. http://doi.org/10.1148/radiol.13131044

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

Djonlagic, I., Rosenfeld, A., Shohamy, D., Myers, C., Gluck, M., & Stickgold, R. (2009). Sleep enhances category learning. *Learning and Memory*, *16*(12), 751–755. http://doi.org/10.1101/lm.1634509

DO, H. (1949). *The organization of behavior. A neuropsychological theory.* Wiley, New York.

Doroudchi, M. M., Greenberg, K. P., Liu, J., Silka, K. A., Boyden, E. S., Lockridge, J. A., … Horsager, A. (2011). Virally delivered channelrhodopsin-2 safely and effectively restores visual function in multiple mouse models of blindness. *Mol Ther*, *19*(7), 1220–1229. http://doi.org/10.1038/mt.2011.69

Edwardson, M. A., Lucas, T. H., Carey, J. R., & Fetz, E. E. (2013). New modalities of brain stimulation for stroke rehabilitation. *Exp Brain Res*, *224*(3), 335–358. http://doi.org/10.1007/s00221-012-3315-1

Ellenbogen, J. M., Hu, P. T., Payne, J. D., Titone, D., & Walker, M. P. (2007). Human relational memory requires time and sleep. *Proceedings of the National Academy of Sciences*, *104*(18), 7723–7728. http://doi.org/10.1073/pnas.0700094104

Engelhard, B., Ozeri, N., Israel, Z., Bergman, H., & Vaadia, E. (2013). Inducing Gamma Oscillations and Precise Spike Synchrony by Operant Conditioning via Brain-Machine Interface. *Neuron*, *77*(2), 361–375. http://doi.org/10.1016/j.neuron.2012.11.015

Eschenko, O., Molle, M., Born, J., & Sara, S. J. (2006). Elevated Sleep Spindle Density after Learning or after Retrieval in Rats. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *26*(50), 12914–12920. http://doi.org/10.1523/JNEUROSCI.3175-06.2006

Feldman, D. E. (2012). The Spike-Timing Dependence of Plasticity. *Neuron*, *75*(4), 556–571. http://doi.org/10.1016/j.neuron.2012.08.001

Feldmeyer, D., & Sakmann, B. (2000). Synaptic efficacy and reliability of excitatory connections between the principal neurones of the input (layer 4) and output layer (layer 5) of the neocortex. *J Physiol*, *525 Pt 1*, 31–39. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/10811722

Fenn, K. M., Nusbaum, H. C., & Margoliash, D. (2003). Consolidation during sleep of perceptual learning of spoken language. *Nature*, *425*(6958), 614–616.

http://doi.org/10.1038/nature01951

Fetz, E. E., & Cheney, P. D. (1978). Muscle fields of primate corticomotoneuronal cells. *Journal de Physiologie*, *74*(3), 239–45. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/102773

Fogel, S. M., & Smith, C. T. (2006). Learning-dependent changes in sleep spindles and Stage 2 sleep. *Journal of Sleep Research*, *15*(3), 250–255. http://doi.org/10.1111/j.1365-2869.2006.00522.x

Fox, K., & Stryker, M. (2017). Integrating Hebbian and homeostatic plasticity: introduction. *Philos Trans R Soc Lond B Biol Sci*, *372*(1715). http://doi.org/10.1098/rstb.2016.0413

Fries, P., Schröder, J.-H., Roelfsema, P. R., Singer, W., & Engel, A. K. (2002). Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *22*(9), 3739–54. http://doi.org/20026318

Fukushima, M., & Sporns, O. (2018). Comparison of fluctuations in global network topology of modeled and empirical brain functional connectivity. *PLoS Computational Biology*, *14*(9), e1006497. http://doi.org/10.1371/journal.pcbi.1006497

Gais, S., Mölle, M., Helms, K., & Born, J. (2002). Learning-dependent increases in sleep spindle density. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *22*(15), 6830–4. http://doi.org/20026697

Gardner, R. J., Hughes, S. W., & Jones, M. W. (2013). Differential Spike Timing and Phase Dynamics of Reticular Thalamic and Prefrontal Cortical Neuronal Populations during Sleep Spindles. *Journal of Neuroscience*, *33*(47), 18469–18480. http://doi.org/10.1523/JNEUROSCI.2197-13.2013

Genzel, L., Kroes, M. C. W., Dresler, M., & Battaglia, F. P. (2014). Light sleep versus slow wave sleep in memory consolidation: a question of global versus local processes? *Trends in Neurosciences*, *37*(1), 10–19. http://doi.org/10.1016/J.TINS.2013.10.002

Genzel, L., & Robertson, E. M. (2015). To Replay, Perchance to Consolidate. *PLoS Biology*, *13*(10), 1–9. http://doi.org/10.1371/journal.pbio.1002285

González-Rueda, A., Pedrosa, V., Feord, R. C., Clopath, C., & Paulsen, O. (2018). Activity-Dependent Downscaling of Subthreshold Synaptic Inputs during Slow-Wave-Sleep-like Activity In Vivo. *Neuron*, *97*(6), 1244–1252.e5. http://doi.org/10.1016/j.neuron.2018.01.047

Graziano, M. S. A., Aflalo, T. N. S., & Cooke, D. F. (2005). Arm Movements Evoked by Electrical Stimulation in the Motor Cortex of Monkeys. *Journal of Neurophysiology*, *94*(6), 4209–4223. http://doi.org/10.1152/jn.01303.2004

Greenough, W. T., Larson, J. R., & Withers, G. S. (1985). Effects of unilateral and bilateral training in a reaching task on dendritic branching of neurons in the rat motor-sensory forelimb cortex. *Behavioral and Neural Biology*, *44*(2), 301–14. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2415103

Gulati, T., Guo, L., Ramanathan, D. S., Bodepudi, A., & Ganguly, K. (2017). Neural

reactivations during sleep determine network credit assignment. *Nature Neuroscience*, *20*(9), 1277–1284. http://doi.org/10.1038/nn.4601

Gulati, T., Ramanathan, D. S., Wong, C. C., & Ganguly, K. (2014). Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning. *Nature Neuroscience*, *17*(8), 1107–1113. http://doi.org/10.1038/nn.3759

Gustafsson, B., Wigström, H., Abraham, W. C., Huang, Y.-Y. Y., & Wigstrom, H. (1987). Long-term potentiation in the hippocampus using depolarizing current pulses as the conditioning stimulus. *J. Neurosci.*, *7*(3), 774–780. http://doi.org/10.1038/nature13294

Han, X. (2012). Optogenetics in the nonhuman primate. *Prog Brain Res*, *196*, 215–233. http://doi.org/10.1016/B978-0-444-59426-6.00011-2

Hara, Y. (2015). Brain plasticity and rehabilitation in stroke patients. *J Nippon Med Sch*, *82*(1), 4–13. http://doi.org/10.1272/jnms.82.4

He, B. J., Zempel, J. M., Snyder, A. Z., & Raichle, M. E. (2010). The Temporal Structures and Functional Significance of Scale-free Brain Activity. *Neuron*, *66*(3), 353–369. http://doi.org/10.1016/J.NEURON.2010.04.020

Hebb, D. (1949). *The Organization of Behavior*. New York: Wiley.

Herrington, T. M., Cheng, J. J., & Eskandar, E. N. (2016). Mechanisms of deep brain stimulation. *Journal of Neurophysiology*, *115*(1), 19–38. http://doi.org/10.1152/jn.00281.2015

Iriki, A., Pavlides, C., & Keller, A. (2017). Long-Term Potentiation in the Motor Cortex Author (s): Atsushi Iriki , Constantine Pavlides , Asaf Keller and Hiroshi Asanuma Published by : American Association for the Advancement of Science Stable URL : http://www.jstor.org/stable/1704276 REFERENCES , *245*(4924), 1385–1387.

Iriki, A., Pavlides, C., Keller, A., & Asanuma, H. (1989). Long-term potentiation in the motor cortex. *Science*, *245*(4924), 1385–1387. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2551038

Jackson, A., Mavoori, J., & Fetz, E. E. (2006). Long-term motor cortex plasticity induced by an electronic neural implant. *Nature*, *444*(7115), 56–60. http://doi.org/10.1038/nature05226

Johnson, L. A., Blakely, T., Hermes, D., Hakimian, S., Ramsey, N. F., & Ojemann, J. G. (2012). Sleep spindles are locally modulated by training on a brain-computer interface. *Proceedings of the National Academy of Sciences*, *109*(45), 18583–18588. http://doi.org/10.1073/pnas.1207532109

Jones, E. G., & Friedman, D. P. (1982). Projection pattern of functional components of thalamic ventrobasal complex on monkey somatosensory cortex. *Journal of Neurophysiology*, *48*(2), 521–44. http://doi.org/10.1152/jn.1982.48.2.521

Jones, E. G., & Powell, T. P. S. (1969). CONNEXIONS OF THE SOMATIC SENSORY CORTEX OF THE RHESUS MONKEY. *Brain*, *92*(3), 477–502. http://doi.org/10.1093/brain/92.3.477

Jones, E. G., & Powell, T. P. S. (1970). AN ANATOMICAL STUDY OF CONVERGING SENSORY PATHWAYS. *Brain*, *93*, 793–820.

Jones, G., & Teeling, E. C. (2006). The evolution of echolocation in bats. *Trends in Ecology & Evolution*, *21*(3), 149–156. http://doi.org/10.1016/J.TREE.2006.01.001

Kaas, J. H. (2004). Evolution of somatosensory and motor cortex in primates. *The Anatomical Record*, *281A*(1), 1148–1156. http://doi.org/10.1002/ar.a.20120

Kaas, J. H. (2004). *Evolution of the Large, Complex Sensorimotor Systems of Anthropoid Primates*. *International Journal of Comparative Psychology* (Vol. 17). Retrieved from https://escholarship.org/uc/item/2gd6p7fj

Kaas, J. H., Krubitzer, L. A., Chino, Y. M., Langston, A. L., Polley, E. H., & Blair, N. (1990). Reorganization of retinotopic cortical maps in adult mammals after lesions of the retina. *Science (New York, N.Y.)*, *248*(4952), 229–31. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2326637

Kao, M. H., Doupe, A. J., & Brainard, M. S. (2005). Contributions of an avian basal ganglia–forebrain circuit to real-time modulation of song. *Nature*, *433*(7026), 638–643. http://doi.org/10.1038/nature03127

Keller, C. J., Honey, C. J., Mégevand, P., Entz, L., Ulbert, I., & Mehta, A. D. (2014). Mapping human brain networks with cortico-cortical evoked potentials. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *369*(1653). http://doi.org/10.1098/rstb.2013.0528

Kellner, C. H., Greenberg, R. M., Murrough, J. W., Bryson, E. O., Briggs, M. C., & Pasculli, R. M. (2012). ECT in Treatment-Resistant Depression. *American Journal of Psychiatry*, *169*(12), 1238–1244. http://doi.org/10.1176/appi.ajp.2012.12050648

Kinnischtzke, A. K., Simons, D. J., & Fanselow, E. E. (2014). Motor cortex broadly engages excitatory and inhibitory neurons in somatosensory barrel cortex. *Cereb Cortex*, *24*(8), 2237–2248. http://doi.org/10.1093/cercor/bht085

Kleim, J. A., Barbay, S., & Nudo, R. J. (1998). Functional Reorganization of the Rat Motor Cortex Following Motor Skill Learning. *Journal of Neurophysiology*, *80*(6), 3321–3325. http://doi.org/10.1152/jn.1998.80.6.3321

Klinshov, V. V, Teramae, J. N., Nekorkin, V. I., & Fukai, T. (2014). Dense neuron clustering explains connectivity statistics in cortical microcircuits. *PLoS One*, *9*(4), e94292. http://doi.org/10.1371/journal.pone.0094292

Kosar, E., Waters, R. S., Tsukahara, N., & Asanuma, H. (1985). Anatomical and physiological properties of the projection from the sensory cortex to the motor cortex in normal cats: the difference between corticocortical and thalamocortical projections. *Brain Res*, *345*(1), 68–78. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/2998549

Kubis, N. (2016). Non-Invasive Brain Stimulation to Enhance Post-Stroke Recovery. *Frontiers in Neural Circuits*, *10*, 56. http://doi.org/10.3389/fncir.2016.00056

Lajoie, G., Krouchev, N. I., Kalaska, J. F., Fairhall, A. L., & Fetz, E. E. (2017). Correlation-based model of artificially induced plasticity in motor cortex by a bidirectional brain-

computer interface. *PLoS Comput Biol*, *13*(2), e1005343. http://doi.org/10.1371/journal.pcbi.1005343

Lang, E. W., Tome, A. M., Keck, I. R., Gorriz-Saez, J. M., & Puntonet, C. G. (2012). Brain connectivity analysis: a short survey. *Comput Intell Neurosci*, *2012*, 412512. http://doi.org/10.1155/2012/412512

Latchoumane, C.-F. V., Ngo, H.-V. V., Born, J., & Shin, H.-S. (2017). Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms. *Neuron*, 1–12. http://doi.org/10.1016/j.neuron.2017.06.025

Ledochowitsch, P., Yazdan-Shahmorad, A., Bouchard, K. E., Diaz-Botia, C., Hanson, T. L., He, J. W., … Maharbiz, M. M. (2015). Strategies for optical control and simultaneous electrical readout of extended cortical circuits. *J Neurosci Methods*, *256*, 220–231. http://doi.org/10.1016/j.jneumeth.2015.07.028

Łęski, S., Lindén, H., Tetzlaff, T., Pettersen, K. H., & Einevoll, G. T. (2013). Frequency Dependence of Signal Power and Spatial Reach of the Local Field Potential. *PLoS Computational Biology*, *9*(7), e1003137. http://doi.org/10.1371/journal.pcbi.1003137

Levenstein, D., Watson, B. O., & Rinzel, J. (2017). ScienceDirect Sleep regulation of the distribution of cortical firing rates. http://doi.org/10.1016/j.conb.2017.02.013

Levy, W. B., & Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, *8*(4), 791–797. http://doi.org/10.1016/0306-4522(83)90010-6

Lin, I. C., Okun, M., Carandini, M., & Harris, K. D. (2015). The Nature of Shared Cortical Variability. *Neuron*, *87*(3), 645–657. http://doi.org/10.1016/j.neuron.2015.06.035

Liu, J., & Newsome, W. T. (2006). Local Field Potential in Cortical Area MT: Stimulus Tuning and Behavioral Correlations. *Journal of Neuroscience*, *26*(30), 7779–7790. http://doi.org/10.1523/JNEUROSCI.5052-05.2006

Logothetis, N. K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H. C., … Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. *Nature*, *491*(7425), 547–553. http://doi.org/10.1038/nature11618

Lovejoy, N. R., Lester, K., Crampton, W. G. R., Marques, F. P. L., & Albert, J. S. (2010). Phylogeny, biogeography, and electric signal evolution of Neotropical knifefishes of the genus Gymnotus (Osteichthyes: Gymnotidae). *Molecular Phylogenetics and Evolution*, *54*(1), 278–290. http://doi.org/10.1016/J.YMPEV.2009.09.017

Lucas, T. H., & Fetz, E. E. (2013). Myo-cortical crossed feedback reorganizes primate motor cortex output. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(12), 5261–74. http://doi.org/10.1523/JNEUROSCI.4683-12.2013

Lucas, T. H., & Fetz, E. E. (2013). Myo-cortical crossed feedback reorganizes primate motor cortex output. *J Neurosci*, *33*(12), 5261–5274. http://doi.org/10.1523/JNEUROSCI.4683-12.2013

Luczak, A., Bartho, P., Marguet, S. L., Buzsaki, G., & Harris, K. D. (2007). Sequential structure

of neocortical spontaneous activity in vivo. *Proceedings of the National Academy of Sciences*, *104*(1), 347–352. http://doi.org/10.1073/pnas.0605643104

Mandelblat-Cerf, Y., Paz, R., & Vaadia, E. (2009). Trial-to-Trial Variability of Single Cells in Motor Cortices Is Dynamically Modified during Visuomotor Adaptation. *Journal of Neuroscience*, *29*(48), 15053–15062. http://doi.org/10.1523/JNEUROSCI.3011-09.2009

Mao, T., Kusefoglu, D., Hooks, B. M., Huber, D., Petreanu, L., & Svoboda, K. (2011). Long-Range Neuronal Circuits Underlying the Interaction between Sensory and Motor Cortex. *Neuron*, *72*(1), 111–123. http://doi.org/10.1016/J.NEURON.2011.07.029

Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science (New York, N.Y.)*, *275*(5297), 213–5. http://doi.org/10.1126/SCIENCE.275.5297.213

Marshall, L., Helgadóttir, H., Mölle, M., & Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, *444*(7119), 610–613. http://doi.org/10.1038/nature05278

Massobrio, P., Tessadori, J., Chiappalone, M., & Ghirardi, M. (2015). In vitro studies of neuronal networks and synaptic plasticity in invertebrates and in mammals using multielectrode arrays. *Neural Plast*, *2015*, 196195. http://doi.org/10.1155/2015/196195

Matsui, T., Tamura, K., Koyano, K. W., Takeuchi, D., Adachi, Y., Osada, T., & Miyashita, Y. (2011). Direct comparison of spontaneous functional connectivity and effective connectivity measured by intracortical microstimulation: an fMRI study in macaque monkeys. *Cereb Cortex*, *21*(10), 2348–2356. http://doi.org/10.1093/cercor/bhr019

Matsumoto, R., Kunieda, T., & Nair, D. (2017). Single pulse electrical stimulation to probe functional and pathological connectivity in epilepsy. *Seizure*, *44*, 27–36. http://doi.org/10.1016/j.seizure.2016.11.003

McGregor, H. R., & Gribble, P. L. (2015). Changes in Visual and Sensory-Motor Resting-State Functional Connectivity Support Motor Learning by Observing. *J Neurophysiol*, jn 00286 2015. http://doi.org/10.1152/jn.00286.2015

McNaughton, B. L., Douglas, R. M., & Goddard, G. V. (1978). Synaptic enhancement in fascia dentata: Cooperativity among coactive afferents. *Brain Research*, *157*(2), 277–293. http://doi.org/10.1016/0006-8993(78)90030-6

Merzenich, M. M., & Jenkins, W. M. (1993). Reorganization of Cortical Representations of the Hand Following Alterations of Skin Inputs Induced by Nerve Injury, Skin Island Transfers, and Experience. *Journal of Hand Therapy*, *6*(2), 89–104. http://doi.org/10.1016/S0894-1130(12)80290-0

Merzenich, M. M., Kaas, J. H., Sur, M., & Lin, C.-S. (1978). Double representation of the body surface within cytoarchitectonic area 3b and 1 in "SI" in the owl monkey (aotus trivirgatus). *The Journal of Comparative Neurology*, *181*(1), 41–73. http://doi.org/10.1002/cne.901810104

Merzenich, M. M., Kaas, J. H., Wall, J., Nelson, R. J., Sur, M., & Felleman, D. (1983). Topographic reorganization of somatosensory cortical areas 3b and 1 in adult monkeys following restricted deafferentation. *Neuroscience*, *8*(1), 33–55. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/6835522

Merzenich, M. M., Nelson, R. J., Stryker, M. P., Cynader, M. S., Schoppmann, A., & Zook, J. M. (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *Journal of Comparative Neurology*, *224*(4), 591–605. http://doi.org/10.1002/cne.902240408

Miller, K. J., Sorensen, L. B., Ojemann, J. G., & den Nijs, M. (2009). Power-Law Scaling in the Brain Surface Electric Potential. *PLoS Computational Biology*, *5*(12), e1000609. http://doi.org/10.1371/journal.pcbi.1000609

Miyamoto, D., Hirai, D., Fung, C. C. A., Inutsuka, A., Odagawa, M., Suzuki, T., … Murayama, M. (2016). Top-down cortical input during NREM sleep consolidates perceptual memory. *Science*, *352*(6291), 1315–1318. http://doi.org/10.1126/science.aaf0902

Mölle, M., Eschenko, O., Gais, S., Sara, S. J., & Born, J. (2009). The influence of learning on sleep slow oscillations and associated spindles and ripples in humans and rats. *European Journal of Neuroscience*, *29*(5), 1071–1081. http://doi.org/10.1111/j.1460-9568.2009.06654.x

Molle, M., Marshall, L., Gais, S., & Born, J. (2004). Learning increases human electroencephalographic coherence during subsequent slow sleep oscillations. *Proceedings of the National Academy of Sciences*, *101*(38), 13963–13968. http://doi.org/10.1073/pnas.0402820101

Mölle, M., Marshall, L., Gais, S., & Born, J. (2002). Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. *The Journal of Neuroscience*, *22*(24), 10941–10947. http://doi.org/22/24/10941 [pii]

Morrison, A., Diesmann, M., & Gerstner, W. (2008). Phenomenological models of synaptic plasticity based on spike timing. *Biological Cybernetics*, *98*(6), 459–478. http://doi.org/10.1007/s00422-008-0233-1

Muller, L., Hamilton, L. S., Edwards, E., Bouchard, K. E., & Chang, E. F. (2016). Spatial resolution dependence on spectral frequency in human speech cortex electrocorticography. *Journal of Neural Engineering*, *13*(5), 056013. http://doi.org/10.1088/1741-2560/13/5/056013

Murphy, T. H., & Corbett, D. (2009). Plasticity during stroke recovery: from synapse to behaviour. *Nat Rev Neurosci*, *10*(12), 861–872. http://doi.org/10.1038/nrn2735

Murray, P. D., & Keller, A. (2011). Somatosensory response properties of excitatory and inhibitory neurons in rat motor cortex. *J Neurophysiol*, *106*(3), 1355–1362. http://doi.org/10.1152/jn.01089.2010

Niethard, N., Ngo, H.-V. V., Ehrlich, I., & Born, J. (2018). Cortical circuit activity underlying sleep slow oscillations and spindles. *Proceedings of the National Academy of Sciences*, 201805517. http://doi.org/10.1073/pnas.1805517115

Niknazar, M., Krishnan, G. P., Bazhenov, M., & Mednick, S. C. (2015). Coupling of thalamocortical sleep oscillations are important for memory consolidation in humans. *PLoS ONE*, *10*(12), 1–14. http://doi.org/10.1371/journal.pone.0144720

Nishida, M., & Walker, M. P. (2007). Daytime naps, motor memory consolidation and regionally specific sleep spindles. *PLoS ONE*, *2*(4). http://doi.org/10.1371/journal.pone.0000341

Nishimura, Y., Perlmutter, S. I., Eaton, R. W., & Fetz, E. E. (2013). Spike-timing-dependent plasticity in primate corticospinal connections induced during free behavior. *Neuron*, *80*(5), 1301–1309. http://doi.org/10.1016/j.neuron.2013.08.028

Nudo, R. J., Milliken, G. W., Jenkins, W. M., & Merzenich, M. M. (1996). Use-dependent alterations of movement representations in primary motor cortex of adult squirrel monkeys. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *16*(2), 785–807. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8551360

Oluigbo, C. O., Salma, A., & Rezai, A. R. (2012). Deep brain stimulation for neurological disorders. *IEEE Rev Biomed Eng*, *5*, 88–99. http://doi.org/10.1109/RBME.2012.2197745

Palm, G., Aertsen, A. M., & Gerstein, G. L. (1988). On the significance of correlations among neuronal spike trains. *Biological Cybernetics*, *59*(1), 1–11. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/3401513

Papale, A. E., & Hooks, B. M. (2018). Circuit Changes in Motor Cortex During Motor Skill Learning. *Neuroscience*, *368*, 283–297. http://doi.org/10.1016/J.NEUROSCIENCE.2017.09.010

Pawlak, V., & Kerr, J. N. D. (2008). Dopamine receptor activation is required for corticostriatal spike-timing-dependent plasticity. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *28*(10), 2435–46. http://doi.org/10.1523/JNEUROSCI.4402-07.2008

Peterson, N. R., Pisoni, D. B., & Miyamoto, R. T. (2010). Cochlear implants and spoken language processing abilities: Review and assessment of the literature. *Restorative Neurology and Neuroscience*, *28*(2), 237–250. http://doi.org/10.3233/RNN-2010-0535

Petrof, I., Viaene, A. N., & Sherman, S. M. (2015). Properties of the primary somatosensory cortex projection to the primary motor cortex in the mouse. *J Neurophysiol*, *113*(7), 2400–2407. http://doi.org/10.1152/jn.00949.2014

Peyrache, A., Battaglia, F. P., & Destexhe, A. (2011). Inhibition recruitment in prefrontal cortex during sleep spindles and gating of hippocampal inputs. *Proceedings of the National Academy of Sciences*, *108*(41), 17207–17212. http://doi.org/10.1073/pnas.1103612108

Prsa, M., Galinanes, G. L., & Huber, D. (2017). Rapid Integration of Artificial Sensory Feedback during Operant Conditioning of Motor Cortex Neurons. *Neuron*, *93*(4), 929–939 e6. http://doi.org/10.1016/j.neuron.2017.01.023

Rajasethupathy, P., Sankaran, S., Marshel, J. H., Kim, C. K., Ferenczi, E., Lee, S. Y., … Deisseroth, K. (2015). Projections from neocortex mediate top-down control of memory retrieval. *Nature*, *526*(7575), 653–659. http://doi.org/10.1038/nature15389

Ramanathan, D. S., Gulati, T., & Ganguly, K. (2015). Sleep-Dependent Reactivation of Ensembles in Motor Cortex Promotes Skill Consolidation. *PLoS Biology*, *13*(9), 1–25. http://doi.org/10.1371/journal.pbio.1002263

Ramanathan, D. S., Guo, L., Gulati, T., Davidson, G., Hishinuma, A. K., Won, S.-J., … Ganguly, K. (2018). Low-frequency cortical activity is a neuromodulatory target that tracks recovery after stroke. *Nature Medicine*, *24*(8), 1257–1267. http://doi.org/10.1038/s41591-018-0058-y

Rasch, B., & Born, J. (2013). About Sleep's Role in Memory. *Physiol Rev*, *93*(2), 681–766. http://doi.org/10.1152/Physrev.00032.2012

Rebesco, J. M., & Miller, L. E. (2011). Enhanced detection threshold for in vivo cortical stimulation produced by Hebbian conditioning. *Journal of Neural Engineering*, *8*(1), 016011. http://doi.org/10.1088/1741-2560/8/1/016011

Rebesco, J. M., & Miller, L. E. (2011). Enhanced detection threshold for in vivo cortical stimulation produced by Hebbian conditioning. *J Neural Eng*, *8*(1), 16011. http://doi.org/10.1088/1741-2560/8/1/016011

Rebesco, J. M., Stevenson, I. H., Koerding, K., Solla, S. A., & Miller, L. E. (2010). Rewiring neural interactions by micro-stimulation. *Frontiers in Systems Neuroscience*, *4*(August), 39. http://doi.org/10.3389/fnsys.2010.00039

Rebesco, J. M., Stevenson, I. H., Kording, K. P., Solla, S. A., & Miller, L. E. (2010). Rewiring neural interactions by micro-stimulation. *Front Syst Neurosci*, *4*. http://doi.org/10.3389/fnsys.2010.00039

Recanzone, G. H., Merzenich, M. M., & Dinse, H. R. (1992a). Expansion of the cortical representation of a specific skin field in primary somatosensory cortex by intracortical microstimulation. *Cerebral Cortex (New York, N.Y. : 1991)*, *2*(3), 181–96. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1511220

Recanzone, G. H., Merzenich, M. M., & Dinse, H. R. (1992b). Expansion of the cortical representation of a specific skin field in primary somatosensory cortex by intracortical microstimulation. *Cerebral Cortex (New York, N.Y. : 1991)*, *2*(3), 181–96. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/1511220

Robertson, D., & Irvine, D. R. F. (1989). Plasticity of frequency organization in auditory cortex of guinea pigs with partial unilateral deafness. *The Journal of Comparative Neurology*, *282*(3), 456–471. http://doi.org/10.1002/cne.902820311

Rousselet, G. A., Pernet, C. R., & Wilcox, R. R. (2017). Beyond differences in means: robust graphical methods to compare two groups in neuroscience. *European Journal of Neuroscience*, *46*(2), 1738–1748. http://doi.org/10.1111/ejn.13610

Sachs, F. (1988). Mechanical transduction in biological systems. *Critical Reviews in Biomedical Engineering*, *16*(2), 141–69. Retrieved from

http://www.ncbi.nlm.nih.gov/pubmed/2460290

Schulz, D. P. A., Sahani, M., & Carandini, M. (2015). Five key factors determining pairwise correlations in visual cortex. *Journal of Neurophysiology*, *114*(2), 1022–1033. http://doi.org/10.1152/jn.00094.2015

Schwartz, A. B., Cui, X. T., Weber, D. J., & Moran, D. W. (2006). Brain-Controlled Interfaces: Movement Restoration with Neural Prosthetics. *Neuron*, *52*(1), 205–220. http://doi.org/10.1016/J.NEURON.2006.09.019

Seeman, S. C., Mogen, B. J., Fetz, E. E., & Perlmutter, S. I. (2017). Paired Stimulation for Spike-Timing-Dependent Plasticity in Primate Sensorimotor Cortex. *J Neurosci*, *37*(7), 1935–1949. http://doi.org/10.1523/JNEUROSCI.2046-16.2017

Sela, Y., Vyazovskiy, V. V., Cirelli, C., Tononi, G., & Nir, Y. (2016). Responses in Rat Core Auditory Cortex are Preserved during Sleep Spindle Oscillations. *Sleep*, *39*(5), 1069–1082. http://doi.org/10.5665/sleep.5758

Seol, G. H., Ziburkus, J., Huang, S., Song, L., Kim, I. T., Takamiya, K., … Kirkwood, A. (2007). Neuromodulators Control the Polarity of Spike-Timing-Dependent Synaptic Plasticity. *Neuron*, *55*(6), 919–929. http://doi.org/10.1016/J.NEURON.2007.08.013

Shen, W., Flajolet, M., Greengard, P., & Surmeier, D. J. (2008). Dichotomous dopaminergic control of striatal synaptic plasticity. *Science (New York, N.Y.)*, *321*(5890), 848–51. http://doi.org/10.1126/science.1160575

Shulz, D. E., & Jacob, V. (2010). Spike-timing-dependent plasticity in the intact brain: counteracting spurious spike coincidences. *Front Synaptic Neurosci*, *2*, 137. http://doi.org/10.3389/fnsyn.2010.00137

Sjöström, P. J., Turrigiano, G. G., & Nelson, S. B. (2001). Rate, Timing, and Cooperativity Jointly Determine Cortical Synaptic Plasticity. *Neuron*, *32*(6), 1149–1164. http://doi.org/10.1016/S0896-6273(01)00542-6

Skudlarski, P., Jagannathan, K., Anderson, K., Stevens, M. C., Calhoun, V. D., Skudlarska, B. A., & Pearlson, G. (2010). Brain connectivity is not only lower but different in schizophrenia: a combined anatomical and functional approach. *Biol Psychiatry*, *68*(1), 61–69. http://doi.org/10.1016/j.biopsych.2010.03.035

Song, S., Miller, K. D., & Abbott, L. F. (2000). Competitive Hebbian learning through spike-timing-dependent synapticplasticity. *Nature Neuroscience*, *3*(9), 919–926. http://doi.org/10.1038/78829

Song, W., Kerr, C. C., Lytton, W. W., & Francis, J. T. (2013). Cortical plasticity induced by spike-triggered microstimulation in primate somatosensory cortex. *PLoS One*, *8*(3), e57453. http://doi.org/10.1371/journal.pone.0057453

Squire, L. R. (2004). Memory systems of the brain: A brief history and current perspective. *Neurobiology of Learning and Memory*, *82*(3), 171–177. http://doi.org/10.1016/J.NLM.2004.06.005

Squire, L. R., & Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences*, *93*(24), 13515–

13522. http://doi.org/10.1073/pnas.93.24.13515

Stam, C. J. (2014). Modern network science of neurological disorders. *Nat Rev Neurosci*, *15*(10), 683–695. http://doi.org/10.1038/nrn3801

Stepniewska, I., Preuss, T. M., & Kaas, J. H. (1993). Architectionis, somatotopic organization, and ipsilateral cortical connections of the primary motor area (M1) of owl monkeys. *The Journal of Comparative Neurology*, *330*(2), 238–271. http://doi.org/10.1002/cne.903300207

Steriade, M. (2000). Corticothalamic resonance, states of vigilance and mentation. *Neuroscience*, *101*(2), 243–276. http://doi.org/10.1016/S0306-4522(00)00353-5

Steriade, M., Contreras, D., Curró Dossi, R., & Nuñez, A. (1993). The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *13*(August), 3284–99. http://doi.org/11124996

Steriade, M., McCormick, D., & Sejnowski, T. (1993). Thalamocortical oscillations in the sleeping and aroused brain. *Science*, *262*(5134), 679–685. http://doi.org/10.1126/science.8235588

Steriade, M., & Timofeev, I. (2003). Neuronal plasticity in thalamocortical networks during sleep and waking oscillations. *Neuron*, *37*(4), 563–76. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/12597855

Stone, K. D., & Gonzalez, C. L. R. (2015). The contributions of vision and haptics to reaching and grasping. *Frontiers in Psychology*, *6*, 1403. http://doi.org/10.3389/fpsyg.2015.01403

Suzuki, M., & Larkum, M. E. (2017). Dendritic calcium spikes are clearly detectable at the cortical surface /631/378 /631/443 article. *Nature Communications*, *8*(1), 1–10. http://doi.org/10.1038/s41467-017-00282-4

Takeuchi, T., Duszkiewicz, A. J., & Morris, R. G. (2014). The synaptic plasticity and memory hypothesis: encoding, storage and persistence. *Philos Trans R Soc Lond B Biol Sci*, *369*(1633), 20130288. http://doi.org/10.1098/rstb.2013.0288

Tchumatchenko, T., Geisel, T., Volgushev, M., & Wolf, F. (2011). Spike correlations - What can they tell about synchrony? *Frontiers in Neuroscience*, *5*(MAY), 1–9. http://doi.org/10.3389/fnins.2011.00068

Tononi, G., & Cirelli, C. (2003). Sleep and synaptic homeostasis: A hypothesis. *Brain Research Bulletin*, *62*(2), 143–150. http://doi.org/10.1016/j.brainresbull.2003.09.004

Toyoizumi, T., Kaneko, M., Stryker, M. P., & Miller, K. D. (2014). Modeling the dynamic interaction of Hebbian and homeostatic plasticity. *Neuron*, *84*(2), 497–510. http://doi.org/10.1016/j.neuron.2014.09.036

Turner, B. O., Crossley, M. J., & Ashby, F. G. (2017). Hierarchical control of procedural and declarative category-learning systems. *NeuroImage*, *150*, 150–161. http://doi.org/10.1016/J.NEUROIMAGE.2017.02.039

Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., & Pennartz, C. M. A. (2010). The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. *NeuroImage*, *51*(1), 112–122. http://doi.org/10.1016/j.neuroimage.2010.01.073

Voytek, B., Kramer, M. A., Case, J., Lepage, K. Q., Tempesta, Z. R., Knight, R. T., & Gazzaley, A. (2015). Age-Related Changes in 1/f Neural Electrophysiological Noise. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *35*(38), 13257–65. http://doi.org/10.1523/JNEUROSCI.2332-14.2015

Wagner, U., Gais, S., Haider, H., Verleger, R., & Born, J. (2004). Wagner 2004 Sleep inspires insight, *427*(January), 21–24. http://doi.org/10.1038/nature02223

Watson, B. O., Levenstein, D., Greene, J. P., Gelinas, J. N., & Buzsáki, G. (2016). Network Homeostasis and State Dynamics of Neocortical Sleep. *Neuron*, *90*(4), 839–852. http://doi.org/10.1016/j.neuron.2016.03.036

Weiler, N., Wood, L., Yu, J., Solla, S. A., & Shepherd, G. M. (2008). Top-down laminar organization of the excitatory network in motor cortex. *Nat Neurosci*, *11*(3), 360–366. http://doi.org/10.1038/nn2049

Wu, P., Zhou, Y. M., Zeng, F., Li, Z. J., Luo, L., Li, Y. X., … Liang, F. R. (2016). Regional brain structural abnormality in ischemic stroke patients: a voxel-based morphometry study. *Neural Regen Res*, *11*(9), 1424–1430. http://doi.org/10.4103/1673-5374.191215

Xu, T., Yu, X., Perlik, A. J., Tobin, W. F., Zweig, J. A., Tennant, K., … Zuo, Y. (2009). Rapid formation and selective stabilization of synapses for enduring motor memories. *Nature*, *462*(7275), 915–919. http://doi.org/10.1038/nature08389

Yahata, N., Morimoto, J., Hashimoto, R., Lisi, G., Shibata, K., Kawakubo, Y., … Kawato, M. (2016). A small number of abnormal brain connections predicts adult autism spectrum disorder. *Nat Commun*, *7*, 11254. http://doi.org/10.1038/ncomms11254

Yang, G., Pan, F., & Gan, W.-B. (2009). Stably maintained dendritic spines are associated with lifelong memories. *Nature*, *462*(7275), 920–924. http://doi.org/10.1038/nature08577

Yazdan-Shahmorad, A., Diaz-Botia, C., Hanson, T. L., Kharazia, V., Ledochowitsch, P., Maharbiz, M. M., & Sabes, P. N. (2016). A Large-Scale Interface for Optogenetic Stimulation and Recording in Nonhuman Primates. *Neuron*, *89*(5), 927–939. http://doi.org/10.1016/j.neuron.2016.01.013

Yazdan-Shahmorad, A., Kipke, D. R., & Lehmkuhle, M. J. (2013). High gamma power in ECoG reflects cortical electrical stimulation effects on unit activity in layers V/VI. *J Neural Eng*, *10*(6), 66002. http://doi.org/10.1088/1741-2560/10/6/066002

Yazdan-Shahmorad, A., Silversmith, D. B., Kharazia, V., & Sabes, P. N. (2018). Targeted cortical reorganization using optogenetics in non-human primates. *Elife*, *7*. http://doi.org/10.7554/eLife.31034

Yizhar, O., Fenno, L. E., Davidson, T. J., Mogri, M., & Deisseroth, K. (2011). Optogenetics in neural systems. *Neuron*, *71*(1), 9–34. http://doi.org/10.1016/j.neuron.2011.06.004

Zacksenhouse, M., Lebedev, M. A., Carmena, J. M., O'Doherty, J. E., Henriquez, C., & Nicolelis, M. A. L. (2007). Cortical Modulations Increase in Early Sessions with Brain-Machine Interface. *PLoS ONE*, *2*(7), e619. http://doi.org/10.1371/journal.pone.0000619

Zenke, F., & Gerstner, W. (2017). Hebbian plasticity requires compensatory processes on multiple timescales. *Philos Trans R Soc Lond B Biol Sci*, *372*(1715). http://doi.org/10.1098/rstb.2016.0259

Publishing Agreement

It is the policy of the University to encourage the distribution of all theses, dissertations, and manuscripts. Copies of all UCSF theses, dissertations, and manuscripts will be routed to the library via the Graduate Division. The library will make all theses, dissertations, and manuscripts accessible to the public and will preserve these to the best of their abilities, in perpetuity.

Please sign the following statement:

I hereby grant permission to the Graduate Division of the University of California, San Francisco to release copies of my thesis, dissertation, or manuscript to the Campus Library to provide access and preservation, in whole or in part, in perpetuity.

Daniel Schmidt

Author Signature

12/11/2018

Date