Causal evidence for the behavioral impact of oscillations in neocortex and hippocampus

by

Joshua H. Siegle
Sc.B., Brown University (2007)

Submitted to the Department of Brain and Cognitive Sciences
in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2014

Signature redacted

Signature of Author .....
Department of Brain and Cognitive Sciences
July 31, 2014

Signature redacted

Certified by .....
Matthew A. Wilson
Sherman Fairchild Professor of Neuroscience
Thesis Supervisor

Signature redacted

Accepted by .....
Matthew A. Wilson
Sherman Fairchild Professor of Neuroscience
Director of Graduate Education for Brain and Cognitive Sciences

Causal evidence for the behavioral impact of oscillations in neocortex and hippocampus

by

Joshua H. Siegle

Submitted to the Department of Brain and Cognitive Sciences on May 7, 2014
in Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy in Neuroscience

Abstract

Neuroscientists hold widely divergent opinions on the behavioral relevance of oscillatory brain states. Some consider them to be a side effect of anatomical connectivity, with little or no role in guiding action. Others view them as a fundamental feature of the network states that underlie perception and cognition. In this thesis, I take a systematic approach to studying two of the most prominent types of oscillations, gamma rhythms in the neocortex (30-80 Hz) and theta rhythms in the hippocampus (4-12 Hz). In both cases, I use light-gated ion channels to manipulate spike activity on a cycle-by-cycle basis in awake, behaving mice. By rhythmically stimulating fast-spiking interneurons in somatosensory cortex, I can emulate the activity patterns that define gamma oscillations under natural conditions. Emulating gamma enhances the detection of threshold-level vibrissae deflections, analogous to the behavioral effects of shifting attention. By triggering stimulation of fast-spiking interneurons in the hippocampus on peaks and troughs of endogenous rhythms, I can reduce spike activity at specific phases of theta. In the context of a spatial navigation task, I find that the ability of inhibition to enhance decision-making accuracy depends on both the theta phase and the task segment in which it occurs. Both of these experiments provide novel causal evidence for the behavioral impact of oscillations, which offers a much more compelling argument for their utility than traditional correlative measures. Finally, I present a new platform for extracellular electrophysiology. This platform, called Open Ephys, makes the closed-loop experiments that are ideal for studying oscillations accessible to a wider audience.

Thesis supervisor: Matthew A. Wilson, Ph.D.
Title: Sherman Fairchild Professor of Neuroscience

Table of contents

Acknowledgments

First and foremost, I'd like to extend deep gratitude to my advisors, Matthew Wilson and Christopher Moore. Completing scientifically diverse projects across two different labs was a risky endeavor, but they gave me the freedom and unflagging support required to accomplish my goals. If even a tiny fraction of their insight, thoroughness, and charisma has rubbed off on me, I am much better for it.

Mark Andermann, my outside committee member, was always generous with his time and was an adept and forthright critic. I am continually impressed by the creativity and rigor of his science, which will serve as a model throughout my career.

My time in graduate school would not have been nearly as fruitful without the help of my three "co-first" authors. Dominique Pritchett was an invaluable partner in carrying out the experiments in Chapter 2; we each brought complementary skills to the table, which made an otherwise daunting project tractable. Open Ephys would not have happened without Jakob Voigts; our shared passions for science, engineering, and design made our collaboration effortless and exceedingly productive. Joining forces with Mike Halassa to selectively activate the thalamic reticular nucleus (work not included in this thesis) resulted in a pioneering publication that opened the door to many fascinating follow-up studies.

The experiments contained in Chapters 2 and 3 depended on dedicated and capable research assistants: Kristine Kim, Richard LeCoultre, Rachel Clary, Jenelle Feather, Hannah Farrow, Nancy Padilla, Rebecca Lichtin, Sophie Bechek, and Justin Klee.

Early in my graduate career, I had the opportunity to collaborate with four incredibly talented postdocs, who now run labs of their own. Jason Ritt, Jess Cardin, Marie Carlén, and Dinos Meletis taught me a great deal about how to be a better scientist.

I'm extremely grateful to Greg Hale, Stuart Layton, and Fabian Kloosterman for introducing me to the wild world of extracellular electrophysiology, and to all the members of the Wilson and Moore labs for their mentoring and friendship.

I'd like to thank Nikhil Bhatla, Rosa Cao, Ulf Knoblich, Dan Bendor, Sam Cooke, and Andy Bolton for our enlightening discussions over the years.

I'm grateful to everyone who contributed to the Open Ephys initiative, and for those promoting the development and adoption of open-source tools in neuroscience.

Thanks to everyone from the BCS incoming class of 2008 for being so awesome. I can't imagine a better cohort with which to experience the trials and travails of graduate school.

Finally, I'd like to thank my parents for their love and support, and Lizzie for raising me well.

Chapter 1: Introduction

> "In the occurrence of these relatively stable, *continuous*, *spontaneous*, *rhythmic* patterns of activity we are presented with an intriguing physiological problem."
>
> – T. H. Bullock, 1945

1.1 Overview of neural oscillations

1.1.1 The view from 1945

The words above were written by T.H. Bullock, a leading neuroethologist, as he took stock of the oscillations routinely observed in nervous tissue (Bullock, 1945). Scientists in Bullock's time recognized that oscillations recorded extracellularly or extracranially were ubiquitous across regions and across species, and that they reliably correlated with behavior. These rhythms, Bullock noted, could provide "clear indications of sleep, of light ether anesthesia, and, under the proper circumstances, of thinking as opposed to the resting state."

But what concerned him most were issues of *timing*. The most widely discussed oscillations of his era were in the range of 1 to 50 Hz—much faster than behavioral changes on the organismal level, but much slower than the action potentials that were the archetype of brain activity. What was the mechanism underlying the "coordinated beat of cell masses," Bullock wondered, and what relationship—if any—did they have to the function of the nervous system as a whole?

Nearly 70 years later, the same questions remain topical (Buzsáki et al., 2013). Although we now know a great deal more about their mechanistic details, the functional relevance of oscillations is an issue of considerable debate. Some view rhythms as a meaningless byproduct of neural activity; they believe that studying them distracts us from what the brain *actually* cares about, which is the integration of spikes over time (Shadlen and Movshon, 1999). Others see oscillations as inherently meaningful, and seek a direct mapping between rhythmic states and cognitive processes (Singer, 1999; Jensen et al., 2007).

The controversy stems from several sources. Practically speaking, a substantial portion of the debate originates in the widely varying recording techniques and analysis methods used by

different laboratories. But equally important is the fact that oscillations lie, as Bullock observed, in the nebulous middle ground between spikes and behavior. To bridge the gap between the two, neuroscientists have leaned heavily on correlative measures. There are numerous examples of papers that attempt to deduce the function of oscillations by observing the organism-level activities with which they co-occur—whether they be moving (Vanderwolf, 1969), perceiving (Fries et al., 1997), attending (Gruber et al., 1999), learning (Miltner et al., 1999), or speaking (Arnal et al., 2011). This approach is ultimately unsatisfying, however, since it cannot establish a causal link between rhythmic activity and behavior (Singer, 1999). To better understand the meaning of oscillations, we need to bring them under causal control. The ability to induce and perturb oscillations in a biologically plausible way—and on a physiologically relevant timescale —has the potential to resolve debates that began in the first half of the 20th century and continue to this day.

1.1.2 The scope of the debate today

Research into the nature of oscillatory brain activity dates back at least to the work of Richard Caton, who was the first to record "great fluctuations" in the electroencephalogram of mammals (Caton, 1877; Ahmed and Cash, 2013). In the century that followed, researchers characterized specific oscillatory phenomena in greater detail: alpha oscillations (8-12 Hz) recorded over visual cortex of quietly resting human subjects (Berger, 1933), spindle oscillations (10-16 Hz) that appeared in the EEG in the early stages of sleep (Loomis et al., 1935), and theta waves (4-12 Hz) that dominated the hippocampal EEG of freely moving rats (Vanderwolf, 1969). But oscillations did not become a mainstream controversy until the characterization of gamma waves (30-80 Hz) in the visual cortex by Wolf Singer, Charles Gray, and colleagues in the late 1980s (Gray et al., 1989; Gray and Singer, 1989).

Those in the "pro-gamma" camp found that neurons in V1 had a striking tendency to fire in synchrony with 40 Hz oscillations in the local field potential (Gray and Singer, 1989). Further work found gamma-range synchronization between cortical columns (Gray et al., 1989), between areas (Engel et al., 1991b), and across hemispheres (Engel et al., 1991a). But the proponents of gamma were too liberal in their interpretation of the function of this rhythm. They wished to

promote gamma as the mechanism of a high-level perceptual phenomenon: the "binding" of disparate stimuli into unified objects (Singer and Gray, 1995). While the story behind gamma's role in binding was seductive, the actual evidence for it was weak or nonexistent, as critics were keen to point out (Shadlen and Movshon, 1999). They took issue with the idea that synchrony in the gamma range could itself be a *signal* for object identity, especially when standard hierarchical processing models already seemed to solve the binding problem on their own.

As with most scientific debates, both sides of the argument were partially correct. Indeed, gamma's role in binding was not supported by the available evidence, and there have yet to be any conclusive experiments showing that gamma and binding are related (Ray and Maunsell, 2010). But there is plenty of evidence that the brain is highly sensitive to temporal correlations on the timescale of gamma (Mainen and Sejnowski, 1995; Markram et al., 1997b; Stevens and Zador, 1998), and that such correlations could benefit signal transmission (Dan et al., 1998; Jia et al., 2013). However, we still do not know how these correlations affect downstream regions in the context of behavior.

The early debates surrounding gamma exemplify many of the features of the general dialog around oscillations. Proponents tend to place too much emphasis on finding a direct link between oscillations and higher perceptual and cognitive processes, trying to imbue them with causal powers that are divorced from reality. Detractors—in their quest to treat oscillations as epiphenomena—tend to ignore basic biophysical principles that explain how oscillations impact circuit function. In order to make the debate more fruitful, we need to bring our focus back to the underlying mechanisms of rhythmic activity, and to place more emphasis on interventions that bring oscillations under experimental control.

1.1.3 Oscillations are repeated patterns of temporally organized spikes

When oscillations are observed in the local field potential or electroencephalogram, it is convenient to classify them in terms of the frequency bands of large-amplitude voltage fluctuations, rather than the patterns of cellular currents that generate them (Buzsáki et al., 2012). This approach has led to substantial confusion. For one, it makes it easier to write off oscillations as meaningless when there's no mechanism attached. In addition, it is inevitable that unrelated

phenomena will be grouped together under the same label, simply because they occur at a similar frequency. Movement artifacts (Yuval-Greenberg et al., 2008), spectral leakage (Ray and Maunsell, 2011; Schomburg et al., 2012; Scheffer-Teixeira et al., 2013), and volume conduction (Srinivasan et al., 2007; Xing et al., 2009; Kajikawa and Schroeder, 2011) exacerbate the problem by creating apparent power or synchrony in a frequency band when none actually exists.

To avoid these pitfalls, we should classify oscillations in terms of the patterns of spike activity that generate them, rather than their frequency. Just as we label extracellularly recorded units as "putative interneurons" if they have narrow waveforms, it would behoove us to label LFP and EEG-based power in the 30-80 Hz range as "putative gamma" until we known the mechanism behind it. For most of the widely studied frequency bands, a cellular mechanism is known or hypothesized. The classic example is gamma oscillations, which arise from synchronous barrages of inhibition creating windows of opportunity for nearby principal cells to spike (Buzsáki and Wang, 2012). On each cycle, fast-spiking interneurons fire a burst of action potentials, which results in ~25 ms of local inhibition. When this inhibition wears off, excitatory cells become active in unison, synchronously relaying signals to downstream areas and depolarizing the interneurons that drive the next cycle (Fries et al., 2007). This mechanism has been supported by *in vitro* (Whittington et al., 1995), *in vivo* (Hasenstaub et al., 2005), computational (Wang and Buzsáki, 1996), and, recently, causal (Cardin et al., 2009) evidence. However, the inhibitory cells driving gamma are not recorded in most preparations; experimenters presume that this mechanism is at work, simply because they detect 30-80 Hz oscillations in the LFP. Drawing attention to the fact that oscillations arise from repeated patterns of temporally organized spikes will make their efficacy harder to dismiss.

If oscillations are nothing more than temporally organized spike activity, what is the point of rhythmicity? Couldn't the same goals be accomplished through aperiodic mechanisms? Such questions expose a perspective that is all too common in neuroscience—trying to abstract away from the dynamics of nervous tissue to expose "computational principles" of the brain. It may be true that what *actually matters* is the relative timing of inputs and outputs to a given region, but it's also true that the brain must establish these temporal relationships in the context of physical structures with *momentum*. As an analogy, consider the heartbeat, another rhythmic phenomenon

whose function is not tied to its periodicity. If the time between beats of an artificial heart were completely irregular, blood would still flow. But implementing this randomness in biological tissue is not plausible. Reliably coordinating large cell assemblies (in this case, heart myocytes) is best accomplished through a mechanism involving an oscillation.

Rhythms are the signature of the brain at work. Just like the heartbeat, periodicity may not be essential to their function, but it is essential to their *implementation*. We cannot understand how cell assemblies become organized without understanding the physical dynamics that allow them to coordinate their activity. Therefore, when we study rhythms experimentally, our toolbox must include cell type–specific manipulations on the timescale of single cycles.

1.1.4 Causal approaches to understanding oscillations

In order to study oscillations in a manner that respects their intrinsic mechanisms while shedding light on their function, we need to employ some type of causal intervention. Many previous attempts to experimentally manipulate oscillations have failed to control them directly. For example, experiments that use lesions (Winson, 1978; Lee et al., 1994), pharmacology (Middleton et al., 2008; Hakami et al., 2009), or genetics (Buhl et al., 2003; Carlén et al., 2011) to disrupt or enhance oscillations change some underlying variable that also happens to affect rhythmic activity. If behavioral changes also occur, it does not imply that the oscillations themselves are responsible.

Other studies have used entrainment to better understand oscillations (Lakatos et al., 2008; Wimber et al., 2012; Sieben et al., 2013; Spaak et al., 2014). In this case, some external drive (either sensory stimuli or electrical potentials) is used to synchronize endogenous rhythmic activity. Such manipulations are still agnostic to the actual implementation, however, and do little to help us interpret the cellular mechanisms by which oscillations may benefit circuit function.

Only recently has the advent of optogenetics made it possible to interact with the brain in a cell type–specific manner on the timescale of milliseconds (Zhang et al., 2006; Yizhar et al., 2011). This advance makes it possible to study oscillations more directly. If we consider oscillations as repeated patterns of temporally organized spikes, we can use optogenetics to perturb those patterns (or induce them artificially) on a cycle-by-cycle basis. Some restraint is

still in order; stimulating at a given frequency does not necessarily constitute a physiologically realistic rhythm (Kim et al., 2012). But if our manipulation is informed by the cell types and activity patterns involved in endogenous rhythms, optogenetics gives us the opportunity to explore the functional relevance of oscillations with unprecedented precision (Moore et al., 2010; Isaacson and Scanziani, 2011).

Optogenetics becomes an even more powerful tool for dissecting rhythmic activity when combined with closed-loop feedback. Closed-loop experiments are the *de facto* way for engineers to analyze unknown systems, but they are decidedly under-appreciated in neuroscience. The major exception is whole-cell recordings, which use closed-loop feedback in the context of voltage clamp, current clamp, and dynamic clamp experiments to gain a much more detailed picture of the inputs and outputs of individual cells (Blanton et al., 1989). In the study of distributed rhythms, however, experiments involving stimulation triggered on neural states have been relative rare. Electrical stimulation triggered on individual cycles of hippocampal theta has been used to study changes in network plasticity as a function of phase (Hyman et al., 2003; Kwag and Paulsen, 2009). The behavioral role of ripples has been studied by disrupting the hippocampus whenever these events are detected (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; Jadhav et al., 2012). Closed-loop control is also a major focus of treatments to prevent seizures, which are oscillations gone out of control (Berényi et al., 2012; Paz et al., 2013). In humans, online analysis of neural state has been used to bias the detection of sensory stimuli (Andermann et al., 2012). But the vast majority of experiments in systems neuroscience are still "open-loop."

1.1.5 Summary of the thesis

In this thesis, I describe a set of experiments that employ optogenetic interventions to enhance our understanding of the role of rhythmic activity in perception and cognition. In Chapter 2, I manipulate the relative timing of excitatory and inhibitory spikes in somatosensory cortex to show that artificially induced gamma oscillations (40 Hz) can benefit processes related to the detection of tactile stimuli. In Chapter 3, I use closed-loop optogenetics to inhibit the

hippocampus at specific times during the hippocampal theta rhythm (4-12 Hz) to show that activity related to encoding and retrieval processes occurs preferentially at different phases.

In the process of carrying out these experiments, I found that the available tools for setting up experiments involving closed-loop control were inadequate in many ways. The commercially available hardware was closed-source and based on exclusive standards, making custom protocols cumbersome to share. Existing open-source systems were limited by their reliance on specific programming languages and commercial digitization hardware (Rolston et al., 2010; Newman et al., 2012). Rather than cobble together a custom rig that would be used once and discarded, I wanted to create a platform that others could easily adopt and extend. In Chapter 4, I describe the tools I developed as part of the Open Ephys initiative, which were instrumental in helping me carry out the experiments in Chapter 3.

1.2 Gamma oscillations in sensory neocortex

1.2.1 Background

Gamma oscillations are typically defined as 30-80 Hz oscillations in the local field potential, but a more appropriate definition takes into account the mechanism by which they are generated: synchronous volleys of excitation and inhibition between reciprocally connected cells (Bartos et al., 2002; Fries et al., 2007; Cardin et al., 2009; Buzsáki and Wang, 2012). The decay time of each bout of inhibition (16 to 33 ms) determines the cycle length (Wang and Buzsáki, 1996; Traub et al., 1996), and hence the range of frequencies that constitute the gamma band.

Gamma rhythms have been characterized in a variety of species and a variety of regions, but they have been studied in most detail in mammalian neocortex. Famous examples come from the visual cortex of primates, where sustained visual stimulation evokes an unmistakable LFP rhythm, as well as spikes phase-locked to the trough of each cycle (Gray and Singer, 1989; Fries et al., 2001; Roberts et al., 2013). Gamma has also been characterized in visual cortex of humans (Hall et al., 2005; Hoogenboom et al., 2010; van Pelt et al., 2012) and rodents (Nase et al., 2003; Niell and Stryker, 2010), auditory cortex of humans (Pantev et al., 1991; Edwards et al., 2005), primates (Brosch et al., 2002; Lakatos et al., 2005; Fukushima et al., 2012), and rodents (Sukov and Barth, 1998; Macdonald et al., 1998; Headley and Weinberger, 2011), and somatosensory

cortex of humans (Bauer et al., 2006; Gross et al., 2007; Zhang et al., 2012) and rodents (Jones and Barth, 1997; Hamada et al., 1999; Sirota et al., 2008), among other regions. Reciprocally connected populations of excitatory and inhibitory cells are a general feature of cortical circuits, and we should expect to find gamma anywhere this architecture is present (Bartos et al., 2007).

There are two aspects of gamma that are most likely to entail its function: (1) its effect on local synchronization and (2) its association with inter-areal coherence. The former occurs by definition, if we choose to classify gamma as synchronous bursts of excitation and inhibition. Local synchronization provides many potential benefits for cortical processing, which can occur in the absence of any cross-regional coupling. The experiments presented in Chapter 2 aim to explore this aspect of gamma, with a focus on its role as a mechanism of attention.

Inter-areal coherence in the gamma range can occur as a consequence of multiple regions simultaneously expressing gamma rhythms; their degree of coordination (or lack thereof) could mediate differential interactions between brain regions. The idea of "dynamic routing" via gamma synchronization is a popular one (Fries, 2009), but only recently has correlative evidence started to emerge (Buschman and Miller, 2007; Gregoriou et al., 2009; Bosman et al., 2012; Roberts et al., 2013). Because of the lack of convincing evidence, the importance of gamma for the routing of information and, in turn, the binding of associated cell assemblies, has been the main target of the debate surrounding gamma.

1.2.2 Controversy over gamma

It is perhaps unfortunate that the early proponents of gamma were so focused on its role in the binding problem (Singer, 1999; Engel et al., 1999; Engel and Singer, 2001), since it stirred up a backlash against research on oscillations in general (Shadlen and Movshon, 1999). But marketing gamma as a potential solution to a fundamental issue in the generation of conscious percepts may have helped it garner recognition as something worth studying in the first place. In his book, *Rhythms of the Brain*, György Buzsáki called Wolf Singer's 1993 symposium on "binding by synchrony" at the Society for Neuroscience meeting the end of a "long vacuum in systems [neuroscience] research" (Buzsáki, 2006). Not only did Singer's theory make gamma

fashionable, it abruptly and irreversibly transitioned the field away from single cells and toward the role of distributed network activity in cognition.

In recent years, the focus has shifted from gamma's role in binding to its role in attention, spearheaded by Pascal Fries, a disciple of Wolf Singer, and Robert Desimone, a researcher at the NIH and MIT. Their seminal study in 2001 showed that the degree of local gamma synchrony was correlated with attentional shifts (Fries et al., 2001). Later work showed that gamma-range synchrony was associated with improved reaction times (Womelsdorf et al., 2006), suggesting that the behavioral benefits of attention could be mediated by changes in gamma. The basic hypothesis is that areas that become more synchronized in the gamma range create efficient channels of communication, at the same time blocking input from areas that are asynchronous or out-of-phase (Womelsdorf and Fries, 2007). This idea is especially enticing due to its potential to explain the bottleneck that turns the incoming flood of sensory signals into the "stream" of consciousness.

Skeptics have repeatedly attempted to discount the "binding by synchrony" and "communication through coherence" theories by showing that the properties of gamma are not regular enough to mediate inter-areal coordination. For example, they found that stimulus-evoked gamma oscillations in V1 are not autocoherent, meaning that the phase and frequency are not consistent from cycle to cycle (Burns et al., 2011; Xing et al., 2012). Another study looked at the mean frequency of evoked gamma in response to different stimuli, and found that it was highly dependent on stimulus contrast (Ray and Maunsell, 2010). If the frequency in a given area is inconsistent across different conditions, gamma oscillations cannot be used to set up efficient channels of communication. However, gamma's potential role in communication may not necessitate that it maintain the same frequency over time if it is simultaneously coherent across areas. A recent study showed that, indeed, the frequency of evoked gamma changes with stimulus contrast, but that frequency fluctuations were highly correlated between V1 and V2 (Roberts et al., 2013).

In general, gamma's critics are right to point out that many theories about gamma's function are too high-level. For example, gamma's role in attention could be mediated by its ability to alter gain locally (Tiesinga et al., 2004; Börgers et al., 2005; Jia et al., 2013), without resorting to

the need for inter-areal synchronization. It may be necessary to assess simpler hypotheses about gamma—such as its role in normalization (Ray et al., 2013)—before we make claims about its capacity to coordinate networks distributed across the entire brain.

1.2.3 Proposed function of gamma

Debates around the function of gamma often overlap with debates about the degree to which synchronization impacts neural circuits. Until recently, it was thought that synchrony was a property of "offline" brain states, since large, synchronous population waves are a signature of inactivity (Berger, 1933; Harris and Thiele, 2011). The "desynchronized state" was taken to be indicative of wakefulness, to differentiate it from the synchronized resting state (Poulet and Petersen, 2008; Ahmed and Cash, 2013). The desynchronizing effects of attention have been the focus of recent work (Cohen and Maunsell, 2009). However, research on gamma has demonstrated that considerable synchrony exists in the awake state (Fries et al., 2001; Bosman et al., 2012; Brunet et al., 2014). And studies that record in multiple areas at once have shown that synchronization can have a profound impact on signal transmission (Alonso et al., 1996; Bruno and Sakmann, 2006; Jia et al., 2013). Because gamma involves the compression of spikes into smaller temporal windows, its main purpose could be the creation of synchrony. Such synchrony could have a variety of beneficial effects on circuit function.

Computational models have played an important role in our understanding of gamma. Network models have demonstrated that gamma depends on inhibitory synchrony (Wang and Buzsáki, 1996; Traub et al., 1997; Börgers and Kopell, 2003), although the degree to which excitatory cells are necessary for gamma is still debated (Tiesinga and Sejnowski, 2009). Further work with simulations has shown that gamma rhythmicity is sufficient to increase the gain of neural circuits, as neurons released synchronously from inhibition are more excitable than those that are tonically inhibited (Tiesinga et al., 2004). Networks undergoing gamma oscillations can mediate stimulus competition, amplifying strong inputs while suppressing weaker ones (Tiesinga, 2005; Börgers et al., 2005; Börgers and Kopell, 2008). In sum, simulations have demonstrated that gamma can (1) increase sensitivity to weak excitatory stimuli, (2) allow more strongly driven stimuli to suppress less efficient ones, and (3) create synchronous spiking which

is easier to process downstream (Börgers and Kopell, 2008; Azouz and Gray, 2000). All of these proposed functions of gamma were based on computational models, due to the difficulty of recording from large populations of cells *in vivo*. But by bringing inhibitory synchrony under causal control, we can come closer to understanding the actual impact of gamma on circuit function. Optogenetics opens up a new class of experiments that have the potential to add causal evidence to the debate.

1.2.4 Causal approaches to understanding gamma

Prior to the advent of optogenetics, the preferred methods for experimentally inducing gamma were pharmacology and electrical stimulation. The introduction of kainate (a glutamate receptor agonist) in slices (Hormuzdi et al., 2001; Cunningham et al., 2003) or ketamine (an NMDA antagonist) *in vivo* (Hakami et al., 2009) resulted in elevated gamma power across the network. Electrical stimulation of the nucleus basalis activates acetylcholine release throughout the neocortex, which also increases gamma power (Metherate et al., 1992). However, these methods enhance gamma indirectly via diverse mechanisms.

The application of light-gated ion channels in neuroscience represents a qualitative advance in our ability to interact with circuits on the timescale of gamma. We can now can directly control the precise timing of the spikes that underlie gamma. Of course, with great power comes great responsibility: not every manipulation at 40 Hz will result in a biologically relevant gamma rhythm (Cardin et al., 2009). And even the most precise optogenetic manipulations remain crude in comparison to the subtlety of intact circuits; given our current technologies, every manipulation is going to be a-physiological in some important way. But the key point is that we are now able to causally test hypotheses related to the impact of spike timing—which we should take to be the defining feature of gamma.

The first study to causally manipulate gamma with optogenetics was Cardin et al. (2009). In this work, the authors expressed the gene for channelrhodopsin-2 (ChR2) (Nagel et al., 2003) in parvalbumin-positive interneurons of somatosensory cortex. Stimulating these cells—which are uniformly fast-spiking interneurons in the upper layers of neocortex—at a range of frequencies resulted in preferential resonance in the gamma range, as measured by LFP power ratio relative

to baseline. Stimulation of nearby pyramidal cells (via CaMKII-dependent expression of ChR2) produced power at low frequencies (<20 Hz), but minimal changes in the gamma range. Optogenetically induced fast-spiking gamma continued after the end of stimulation and was able to phase-reset ongoing rhythms brought about by nucleus basalis stimulation. As a final proof of concept, the authors showed that changing the offset between whisker-driven excitation and optogenetic gamma could change the synchrony of the evoked response without altering the evoked firing rate. Therefore, this manipulation gives us the opportunity to bring gamma-range synchrony under experimental control.

This last result was modeled by Knoblich and colleagues, who demonstrated that the optogenetically induced gamma made downstream communication more effective and more efficient (Knoblich et al., 2010). Enhanced "effectiveness" was a result of the gamma-range synchrony; a more concentrated packet of spikes could drive more spikes downstream. Enhanced "efficiency," on the other hand, resulted from the cyclical nature of gamma; spikes that occurred too late to have a substantial impact on downstream regions (due to the limited integration window) were removed by the subsequent burst of inhibition. Thus, gamma's utility depends on its ability to synchronize local spikes as well as its rhythmic properties.

In the same issue of *Nature* as Cardin et al. (2009), Sohal and colleagues demonstrated that inhibiting parvalbumin-positive interneurons reduced gamma power in cortical slices (Sohal et al., 2009). Since then, other optogenetic interventions have been used to study the role of fast inhibition in network function (Atallah et al., 2012; Lee et al., 2012; Royer et al., 2012; Stark et al., 2013; Buetfering et al., 2014). Others have explored the importance of synchrony in perception by directly activating pyramidal cells (Histed and Maunsell, 2014). However, none of these studies specifically focused on the relevance of gamma to perceptual processes in awake animals.

1.2.5 Causal interventions in the context of behavior

In Chapter 2, I describe a study in which optogenetic manipulations were used to generate gamma rhythms in awake, behaving mice. We chose to carry out this study in mice due to the availability of genetically modified strains for targeting new genes to specific cell types. We also

needed a sensory system with fast, precise signal transmission from the periphery to the cortex, for which mouse barrel cortex was a perfect candidate. By inserting the gene for ChR2 into parvalbumin-positive interneurons, we were able to bring inhibition under causal control. In conjunction, we used fast whisker deflections to control excitation in the same network. Changing the relative timing between the two gave us unprecedented control over gamma-range inhibitory and excitatory synchrony.

We trained mice to lick in response to deflections of the whisker pad. Thus, we could use the mice's licking accuracy to measure the detectability of tactile stimuli under different conditions. We found that weak stimuli could be enhanced by the causal induction of gamma rhythms, consistent with the role of gamma in attention (Reynolds et al., 2000). However, this enhancement was sensitive to the relative timing of excitation and inhibition. Only conditions that reinforced the dynamics of endogenous gamma resulted in improved performance. Shifting the temporal offset between sensory excitation and optogenetic inhibition by only 5 ms resulted in significantly impaired performance. This study offers the first causal evidence that gamma can impact perception, an important existence proof that gamma can benefit circuit function in a mammalian model system.

1.3 Theta oscillations in hippocampus

1.3.1 Background

Unlike gamma, which can occur relatively locally, theta oscillations are coherent throughout a distributed network that includes the hippocampus, cortex, thalamic nuclei, and subcortical regions (Vertes et al., 2004). Each cycle of theta involves the intricate coordination of spike times across multiple structures. A mechanistic definition of theta cannot be as succinct as that for gamma, but it is tenable to devise one. Similarly to gamma, theta is dependent upon inhibition. Bidirectional GABAergic projections link the hippocampus and the medial septum and serve as the theta "pacemaker" *in vivo* (Köhler et al., 1984; Tóth et al., 1993; Jinno et al., 2007; Hangya et al., 2009). Within the hippocampus, different types of interneurons are active at different phases of theta (Klausberger et al., 2003; Klausberger et al., 2005; Somogyi and Klausberger, 2005). These inhibitory cells differentially modulate inputs to the dendrites and soma (Klausberger and

Somogyi, 2008), and determine when principal cells will become active on each cycle (Mehta et al., 2002).

Theta was first described by Case Vanderwolf in electrical recordings from the hippocampus of freely moving rats (Vanderwolf, 1969). Theta rhythms have since been documented in rabbits (Kramis et al., 1975), mice (Buzsáki et al., 2003), bats (Ulanovsky and Moss, 2007), monkeys (Stewart and Fox, 1991), and humans (Kahana et al., 1999), and therefore appear to be a general feature of mammalian hippocampal circuits. Theta's association with the hippocampus—and states of active exploration and REM sleep—has inspired arguments that its function is primarily mnemonic (Vertes et al., 2001; Colgin, 2013; Buzsáki and Moser, 2013). But this does not constitute a satisfying explanation on its own. More work is needed to understand precisely how theta coordinates memory formation on a cellular level.

1.3.2 Proposed function of theta

A few years after theta was first characterized, John O'Keefe discovered that cells in the hippocampus are spatially selective (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976). When rats navigate through an environment, cells will consistently respond to the same location, known as a "place field." Two decades later, following technical advances that allowed dozens of cells to be recorded simultaneously, researchers accurately decoded a rat's location based on information contained in hippocampal firing patterns (Wilson and McNaughton, 1993). Whereas previous lesion studies had shown that the hippocampus was *necessary* for navigation (Morris et al., 1982), this finding implied that it contained information that was *sufficient* for guiding spatial behaviors.

How is the spatial selectivity of the hippocampus tied to the theta rhythm? Studies of phase precession (O'Keefe and Recce, 1993; Skaggs et al., 1996), showed that the information content of hippocampal activity changes as a function of theta phase. Earlier in the cycle (relative to the peak of multiunit firing), active cells represent recently departed locations. Later in the cycle (after the peak of multiunit firing), active cells represent locations the animal is about to enter. Furthermore, the nature of the representation changes throughout the cycle. At earlier phases, the correlation between phase and location is less pronounced, and thus firing rate is the best

predictor of location. At later phases, phase and location are highly correlated, making relative spike times more informative (Mehta et al., 2002).

These observations have two main implications. First, within the hippocampus, processes attuned to relative timing must occur later in the theta cycle, when phase and location and correlated. Writing new memories depends on long-term potentiation (LTP), which is sensitive to millisecond-timescale shifts in spike timing (Markram et al., 1997a; Bi and Poo, 1998). Therefore, encoding is likely to occur at later phases, when relative spike times are informative. Second, regions that receive projections from hippocampus can use differential phase-locking to access different content. Information related to past experience is available early in the theta cycle, whereas information about upcoming trajectories is available later in the cycle. In other words, the spatial layout of the environment is "temporally compressed" into theta sequences (Skaggs et al., 1996). Taken together, these observations suggest that theta may optimize the hippocampus for learning by keeping information that is being actively encoded from interfering with previously stored information that may be necessary to guide behavior. Encoding and retrieval are continuous processes, which occur more or less simultaneously on behavioral timescales. Theta could be the way the hippocampus deals avoids content interference, a fundamental problem for any system involved in memory formation.

This idea has been formalized by Michael Hasselmo, who developed computational models to explore the segregation of encoding and retrieval processes as a function of theta phase (Hasselmo et al., 2002; Hasselmo and Eichenbaum, 2005; Kunec et al., 2005). In these models, theta emerges from the coordination of inputs to the CA1 region of hippocampus from CA3 and entorhinal cortex (EC). EC acts as the relay station for information about the current state of the world, which must be encoded at phases optimized for LTP (Hyman et al., 2003; Kwag and Paulsen, 2009). CA3 stores information in its recurrent connections, which can be broadcast to the cortex via CA1's divergent projections (Cenquizca and Swanson, 2007). Hasselmo's model predicts that CA1 activity at earlier phases of theta should be CA3-dominant and associated with retrieval, while CA1 activity at later phases of theta should be EC-dominant and associated with encoding.

Correlative evidence supports this hypothesis. In a study involving simultaneous recordings from CA1, CA3, and EC, the authors found that synchronization between pairs of structures occurred at different phases of theta (Colgin et al., 2009). CA1–EC synchrony occurred later in the cycle, and was associated with enhanced coherence in the high gamma range (60-80 Hz). CA1–CA3 synchrony occurred earlier in the cycle, and was associated with enhanced coherence in the low gamma range (25-50 Hz). This is consistent with Hasselmo's model of encoding and retrieval.

A separate study found that placing rats in a novel environment shifted multiunit firing later in the theta cycle—toward the phase involved in encoding (Lever et al., 2010). The opposite shift (away from the encoding phase) occurred in the presence of an acetylcholine antagonist (Douchamps et al., 2013). Acetylcholine release is associated with novelty, and could be the mechanism that biases the hippocampus toward an encoding state (Hasselmo, 2006).

These correlative studies are noteworthy for their congruence with prior modeling work. But they cannot tell us whether the timing of spikes relative to the theta cycle is truly important for behavior. It is possible that structures involved in guiding action integrate hippocampal spikes over windows much longer than a single theta cycle. In that case, overall firing rate would matter more than relationship between spikes and theta phase. To answer this question more definitively, we need to bring hippocampal activity under causal control.

1.3.3 Causal approaches to studying theta

The traditional approaches to experimentally manipulating theta employ lesions, pharmacology, or genetics. Lesions of the medial septum abolish theta rhythmicity in the hippocampus, and are associated with deficits in navigation performance (Hagan et al., 1988; Berger-Sweeney et al., 1994; Lee et al., 1994; Yoder and Pang, 2005). Classic studies involved the application of atropine, an acetylcholine antagonist, to identify theta rhythms that depend on cholinergic inputs to hippocampus (Kramis et al., 1975; Sainsbury and Montoya, 1984). Abnormal theta rhythms have been characterized in various genetically modified mouse lines, in order to explore the effects of specific genes on theta and associated deficits in spatial navigation (Wulff et al., 2009; Korotkova et al., 2010; Sigurdsson et al., 2010). Of course, as in the case of

gamma rhythms, such correlational studies cannot show that this oscillation has a direct impact on behavior. More precise manipulations are required to make such claims.

An alternate approach is to stimulate the hippocampus as function of theta phase. Online detection of peaks and troughs can be used to target stimulation to specific times in the theta cycle. This has been done in slices (Huerta and Lisman, 1995) and *in vivo* (Pavlides et al., 1988; Hölscher et al., 1997; Hyman et al., 2003), with the main outcome being that opposite phases promote synpatic potentiation and depotentiation. Early results to this effect inspired models of segregated encoding and retrieval, which predict that encoding will preferentially occur at phases associated with enhanced long-term potentiation (Hasselmo et al., 2002). Retrieval, on the other hand, must avoid "over-writing" previously stored information, and should therefore occur when LTP is weakest.

In recent years, causal manipulation of the hippocampus has been carried out with optogenetic techniques (Zhang et al., 2006; Yizhar et al., 2011). Tonic inhibition of specific cell types has been used to explore the role of CA1 in context retrieval (Goshen et al., 2011) and the impact of inhibitory cell types on phase precession (Royer et al., 2012). Rhythmic activation of parvalbumin-positive interneurons (not phase-locked to theta) revealed a post-inhibitory rebound in pyramidal cells that explained their propensity to resonate in the theta-frequency range (Stark et al., 2013).

The combination of optogenetics and closed-loop feedback has the potential to offer more insight than either method in isolation. Just as with gamma, the defining feature of theta is not its frequency, but its association with the intricate temporal coordination of different cell types. Activating or inhibiting specific cell types on theta timescales allows us to test whether patterns of activity within each cycle are important for guiding behavior.

1.3.4 Closed-loop interventions in the context of behavior

In Chapter 3, I describe a study in which closed-loop optogenetic manipulations were used to disrupt the hippocampus in a theta phase–specific manner. This manipulation was carried out in the context of a spatial navigation task. Each trial included an "encoding segment," in which information about upcoming reward location was available, and a "retrieval segment," in which

behavior must be guided by internal states. This was the first study to look at the effects of closed-loop optogenetic feedback on behavior. We found that our precise intervention affected task performance in a way that depended on both theta phase and task segment. Stimulation in the encoding segment enhanced performance when inhibition was triggered by the peak of theta. Conversely, stimulation in the retrieval segment enhanced performance when inhibition was triggered by the trough of theta. These results suggest that processes related to encoding and retrieval are most active at different phases of theta, in agreement with prior computational models and correlative studies.

The modular tools developed to carry out these experiments are described in Chapter 4. These tools, which include both software and hardware, are easy to repurpose for other experiments involving real-time feedback. Closed-loop manipulations of the brain are a powerful but under-utilized way of probing the function of neural circuits. We can lower the barrier to entry for such experiments by making the required tools more accessible.

1.4 References

Ahmed OJ, Cash SS (2013) Finding synchrony in the desynchronized EEG: the history and interpretation of gamma rhythms. *Front Integr Neurosci* **7**:58.

Alonso JM, Usrey WM, Reid RC (1996) Precisely correlated firing in cells of the lateral geniculate nucleus. *Nature* **383**: 815-819.

Andermann M, Kauramäki J, Palomäki T, Moore C, Hari R, Jääskeläinen I, Sams M (2012) Brain state-triggered stimulus delivery: An efficient tool for probing ongoing brain activity. *Open J Neurosci* **2**.

Arnal LH, Wyart V, Giraud AL (2011) Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nat Neurosci* **14**: 797-801.

Atallah BV, Bruns W, Carandini M, Scanziani M (2012) Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron* **73**: 159-170.

Azouz R, Gray CM (2000) Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc Natl Acad Sci U S A* **97**: 8110-8115.

Bartos M, Vida I, Frotscher M, Meyer A, Monyer H, Geiger JR, Jonas P (2002) Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *Proc Natl Acad Sci U S A* **99**: 13222-13227.

Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nat Rev Neurosci* **8**: 45-56.

Bauer M, Oostenveld R, Peeters M, Fries P (2006) Tactile spatial attention enhances gamma-band activity in somatosensory cortex and reduces low-frequency activity in parieto-occipital areas. *J Neurosci* **26**: 490-501.

Berényi A, Belluscio M, Mao D, Buzsáki G (2012) Closed-loop control of epilepsy by transcranial electrical stimulation. *Science* **337**: 735-737.

Berger H (1933) Über das Elektroencephalogram des Menschen. *Arch Psychiatr Nervenkr* **99**: 555-574.

Berger-Sweeney J, Heckers S, Mesulam M-M, Wiley RG, Lappi DA, Sharma M (1994) Differential effects on spatial navigation of immunotoxin-induced cholinergic lesions of the medial septal area and nucleus basalis magnocellularis. *The Journal of neuroscience* **14**: 4507-4519.

Bi G, Poo M (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci* **18**: 10464-10472.

Blanton MG, Lo Turco JJ, Kriegstein AR (1989) Whole cell recording from neurons in slices of reptilian and mammalian cerebral cortex. *J Neurosci Methods* **30**: 203-210.

Bosman CA, Schoffelen JM, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* **75**: 875-888.

Börgers C, Epstein S, Kopell NJ (2005) Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc Natl Acad Sci U S A* **102**: 7002-7007.

Börgers C, Kopell N (2003) Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Comput* **15**: 509-538.

Börgers C, Kopell NJ (2008) Gamma oscillations and stimulus selection. *Neural Comput* **20**: 383-414.

Brosch M, Budinger E, Scheich H (2002) Stimulus-related gamma oscillations in primate auditory cortex. *J Neurophysiol* **87**: 2715-2725.

Brunet NM, Bosman CA, Vinck M, Roberts M, Oostenveld R, Desimone R, De Weerd P, Fries P (2014) Stimulus repetition modulates gamma-band synchronization in primate visual cortex. *Proc Natl Acad Sci U S A* **111**: 3626-3631.

Bruno RM, Sakmann B (2006) Cortex is driven by weak but synchronously active thalamocortical synapses. *Science* **312**: 1622-1627.

Buetfering C, Allen K, Monyer H (2014) Parvalbumin interneurons provide grid cell-driven recurrent inhibition in the medial entorhinal cortex. *Nat Neurosci.* **17**: 710-718.

Buhl DL, Harris KD, Hormuzdi SG, Monyer H, Buzsáki G (2003) Selective impairment of hippocampal gamma oscillations in connexin-36 knock-out mouse in vivo. *J Neurosci* **23**: 1013-1018.

Bullock TH (1945) Problems in the comparative study of brain waves. *Yale J Biol Med* **17**: 657-679.

Burns SP, Xing D, Shapley RM (2011) Is gamma-band activity in the local field potential of V1 cortex a "clock" or filtered noise? *J Neurosci* **31**: 9658-9664.

Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* **315**: 1860-1862.

Buzsáki G (2006) Rhythms of the Brain. Oxford: Oxford Univ Press.

Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents - EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* **13**: 407-420.

Buzsáki G, Buhl DL, Harris KD, Csicsvari J, Czéh B, Morozov A (2003) Hippocampal network patterns of activity in the mouse. *Neuroscience* **116**: 201-211.

Buzsáki G, Logothetis N, Singer W (2013) Scaling brain size, keeping timing: evolutionary preservation of brain rhythms. *Neuron* **80**: 751-764.

Buzsáki G, Moser EI (2013) Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nat Neurosci* **16**: 130-138.

Buzsáki G, Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci.* **35**: 203-25.

Cardin JA, Carlén M, Meletis K, Knoblich U, Zhang F, Deisseroth K, Tsai LH, Moore CI (2009) Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* **459**: 663-667.

Carlén M, Meletis K, Siegle JH, Cardin JA, Futai K, Vierling-Claassen D, Rühlmann C, Jones SR, Deisseroth K, Sheng M, Moore CI, Tsai LH (2011) A critical role for NMDA receptors in parvalbumin interneurons for gamma rhythm induction and behavior. *Mol Psychiatry.* **17**.

Caton R (1877) Interim report on investigation of the electric currents of the brain. *Br Med J (Suppl I)*: 62-65.

Cenquizca LA, Swanson LW (2007) Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Res Rev* **56**: 1-26.

Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* **12**: 1594-1600.

Colgin LL (2013) Mechanisms and functions of theta rhythms. *Annu Rev Neurosci* **36**: 295-312.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**: 353-357.

Cunningham MO, Davies CH, Buhl EH, Kopell N, Whittington MA (2003) Gamma oscillations induced by kainate receptor activation in the entorhinal cortex in vitro. *J Neurosci* **23**: 9761-9769.

Dan Y, Alonso JM, Usrey WM, Reid RC (1998) Coding of visual information by precisely correlated spikes in the lateral geniculate nucleus. *Nat Neurosci* **1**: 501-507.

Douchamps V, Jeewajee A, Blundell P, Burgess N, Lever C (2013) Evidence for encoding versus retrieval scheduling in the hippocampus by theta phase and acetylcholine. *J Neurosci* **33**: 8689-8704.

Edwards E, Soltani M, Deouell LY, Berger MS, Knight RT (2005) High gamma activity in response to deviant auditory stimuli recorded directly from human cortex. *J Neurophysiol* **94**: 4269-4280.

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

Engel AK, Fries P, König P, Brecht M, Singer W (1999) Temporal binding, binocular rivalry, and consciousness. *Conscious Cogn* **8**: 128-151.

Engel AK, König P, Kreiter AK, Singer W (1991a) Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science* **252**: 1177-1179.

Engel AK, Kreiter AK, König P, Singer W (1991b) Synchronization of oscillatory neuronal responses between striate and extrastriate visual cortical areas of the cat. *Proc Natl Acad Sci U S A* **88**: 6048-6052.

Engel AK, Singer W (2001) Temporal binding and the neural correlates of sensory awareness. *Trends Cogn Sci* **5**: 16-25.

Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu Rev Neurosci* **32**: 209-224.

Fries P, Nikolić D, Singer W (2007) The gamma cycle. *Trends Neurosci* **30**: 309-316.

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**: 1560-1563.

Fries P, Roelfsema PR, Engel AK, König P, Singer W (1997) Synchronization of oscillatory responses in visual cortex correlates with perception in interocular rivalry. *Proc Natl Acad Sci U S A* **94**: 12699-12704.

Fukushima M, Saunders RC, Leopold DA, Mishkin M, Averbeck BB (2012) Spontaneous high-gamma band activity reflects functional organization of auditory cortex in the awake macaque. *Neuron* **74**: 899-910.

Girardeau G, Benchenane K, Wiener SI, Buzsáki G, Zugaro MB (2009) Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* **12**: 1222-1223.

Goshen I, Brodsky M, Prakash R, Wallace J, Gradinaru V, Ramakrishnan C, Deisseroth K (2011) Dynamics of retrieval strategies for remote memories. *Cell* **147**: 678-689.

Gray CM, König P, Engel AK, Singer W (1989) Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**: 334-337.

Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proc Natl Acad Sci U S A* **86**: 1698-1702.

Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* **324**: 1207-1210.

Gross J, Schnitzler A, Timmermann L, Ploner M (2007) Gamma oscillations in human primary somatosensory cortex reflect pain perception. *PLoS Biol* **5**:e133.

Gruber T, Müller MM, Keil A, Elbert T (1999) Selective visual-spatial attention alters induced gamma band responses in the human EEG. *Clin Neurophysiol* **110**: 2074-2085.

Hagan JJ, Salamone JD, Simpson J, Iversen SD, Morris RG (1988) Place navigation in rats is impaired by lesions of medial septum and diagonal band but not nucleus basalis magnocellularis. *Behav Brain Res* **27**: 9-20.

Hakami T, Jones NC, Tolmacheva EA, Gaudias J, Chaumont J, Salzberg M, O'Brien TJ, Pinault D (2009) NMDA receptor hypofunction leads to generalized and persistent aberrant gamma oscillations independent of hyperlocomotion and the state of consciousness. *PLoS ONE* **4**:e6755.

Hall SD, Holliday IE, Hillebrand A, Singh KD, Furlong PL, Hadjipapas A, Barnes GR (2005) The missing link: analogous human and primate cortical gamma oscillations. *Neuroimage* **26**: 13-17.

Hamada Y, Miyashita E, Tanaka H (1999) Gamma-band oscillations in the "barrel cortex" precede rat's exploratory whisking. *Neuroscience* **88**: 667-671.

Hangya B, Borhegyi Z, Szilágyi N, Freund TF, Varga V (2009) GABAergic neurons of the medial septum lead the hippocampal network during theta activity. *J Neurosci* **29**: 8094-8102.

Harris KD, Thiele A (2011) Cortical state and attention. *Nat Rev Neurosci* **12**: 509-523.

Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* **47**: 423-435.

Hasselmo ME (2006) The role of acetylcholine in learning and memory. *Curr Opin Neurobiol* **16**: 710-715.

Hasselmo ME, Bodelón C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput* **14**: 793-817.

Hasselmo ME, Eichenbaum H (2005) Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw* **18**: 1172-1190.

Headley DB, Weinberger NM (2011) Gamma-band activation predicts both associative memory and cortical plasticity. *J Neurosci* **31**: 12748-12758.

Histed MH, Maunsell JH (2014) Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proc Natl Acad Sci U S A* **111**: E178-E187.

Hoogenboom N, Schoffelen JM, Oostenveld R, Fries P (2010) Visually induced gamma-band activity predicts speed of change detection in humans. *Neuroimage* **51**: 1162-1167.

Hormuzdi SG, Pais I, LeBeau FE, Towers SK, Rozov A, Buhl EH, Whittington MA, Monyer H (2001) Impaired electrical signaling disrupts gamma frequency oscillations in connexin 36-deficient mice. *Neuron* **31**: 487-495.

Hölscher C, Anwyl R, Rowan MJ (1997) Stimulation on the positive phase of hippocampal theta rhythm induces long-term potentiation that can be depotentiated by stimulation on the negative phase in area CA1 in vivo. *The Journal of Neuroscience* **17**: 6470-6477.

Huerta PT, Lisman JE (1995) Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. *Neuron* **15**: 1053-1063.

Hyman JM, Wyble BP, Goyal V, Rossi CA, Hasselmo ME (2003) Stimulation in hippocampal region CA1 in behaving rats yields long-term potentiation when delivered to the peak of theta and long-term depression when delivered to the trough. *J Neurosci* **23**: 11725-11731.

Isaacson JS, Scanziani M (2011) How inhibition shapes cortical activity. *Neuron* **72**: 231-243.

Jadhav SP, Kemere C, German PW, Frank LM (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**: 1454-1458.

Jensen O, Kaiser J, Lachaux JP (2007) Human gamma-frequency oscillations associated with attention and memory. *Trends Neurosci* **30**: 317-324.

Jia X, Tanabe S, Kohn A (2013) Gamma and the coordination of spiking activity in early visual cortex. *Neuron* **77**: 762-774.

Jinno S, Klausberger T, Marton LF, Dalezios Y, Roberts JD, Fuentealba P, Bushong EA, Henze D, Buzsáki G, Somogyi P (2007) Neuronal diversity in GABAergic long-range projections from the hippocampus. *J Neurosci* **27**: 8790-8804.

Jones MS, Barth DS (1997) Sensory-evoked high-frequency (gamma-band) oscillating potentials in somatosensory cortex of the unanesthetized rat. *Brain Res* **768**: 167-176.

Kahana MJ, Sekuler R, Caplan JB, Kirschen M, Madsen JR (1999) Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature* **399**: 781-784.

Kajikawa Y, Schroeder CE (2011) How local is the local field potential? *Neuron* **72**: 847-858.

Kim A, Latchoumane C, Lee S, Kim GB, Cheong E, Augustine GJ, Shin HS (2012) Optogenetically induced sleep spindle rhythms alter sleep architectures in mice. *Proc Natl Acad Sci U S A* **109**: 20673-20678.

Klausberger T, Magill PJ, Márton LF, Roberts JD, Cobden PM, Buzsáki G, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* **421**: 844-848.

Klausberger T, Marton LF, O'Neill J, Huck JH, Dalezios Y, Fuentealba P, Suen WY, Papp E, Kaneko T, Watanabe M, Csicsvari J, Somogyi P (2005) Complementary roles of cholecystokinin- and parvalbumin-expressing GABAergic neurons in hippocampal network oscillations. *J Neurosci* **25**: 9782-9793.

Klausberger T, Somogyi P (2008) Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science* **321**: 53-57.

Knoblich U, Siegle JH, Pritchett DL, Moore CI (2010) What do we gain from gamma? Local dynamic gain modulation drives enhanced efficacy and efficiency of signal transmission. *Front Hum Neurosci* **4**: 1-12.

Korotkova T, Fuchs EC, Ponomarenko A, von Engelhardt J, Monyer H (2010) NMDA receptor ablation on parvalbumin-positive interneurons impairs hippocampal synchrony, spatial representations, and working memory. *Neuron* **68**: 557-569.

Köhler C, Chan-Palay V, Wu JY (1984) Septal neurons containing glutamic acid decarboxylase immunoreactivity project to the hippocampal region in the rat brain. *Anat Embryol (Berl)* **169**: 41-44.

Kramis R_, Vanderwolf CH, Bland BH (1975) Two types of hippocampal rhythmical slow activity in both the rabbit and the rat: relations to behavior and effects of atropine, diethyl ether, urethane, and pentobarbital. *Exp Neurol* **49**: 58-85.

Kunec S, Hasselmo ME, Kopell N (2005) Encoding and retrieval in the CA3 region of the hippocampus: a model of theta-phase separation. *J Neurophysiol* **94**: 70-82.

Kwag J, Paulsen O (2009) The timing of external input controls the sign of plasticity at local synapses. *Nat Neurosci* **12**: 1219-1221.

Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* **320**: 110-113.

Lakatos P, Shah AS, Knuth KH, Ulbert I, Karmos G, Schroeder CE (2005) An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *J Neurophysiol* **94**: 1904-1911.

Lee MG, Chrobak JJ, Sik A, Wiley RG, Buzsáki G (1994) Hippocampal theta activity following selective lesion of the septal cholinergic system. *Neuroscience* **62**: 1033-1047.

Lee SH, Kwan AC, Zhang S, Phoumthipphavong V, Flannery JG, Masmanidis SC, Taniguchi H, Huang ZJ, Zhang F, Boyden ES, Deisseroth K, Dan Y (2012) Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature* **488**: 379-383.

Lever C, Burton S, Jeewajee A, Wills TJ, Cacucci F, Burgess N, O'Keefe J (2010) Environmental novelty elicits a later theta phase of firing in CA1 but not subiculum. *Hippocampus* **20**: 229-234.

Loomis AL, Harvey EN, Hobart G (1935) Potential rhythms of the cerebral cortex during sleep. *Science* **81**: 597-598.

Macdonald KD, Fifkova E, Jones MS, Barth DS (1998) Focal stimulation of the thalamic reticular nucleus induces focal gamma waves in cortex. *J Neurophysiol* **79**: 474-477.

Mainen ZF, Sejnowski TJ (1995) Reliability of spike timing in neocortical neurons. *Science* **268**: 1503-1506.

Markram H, Lubke J, Frotscher M, Sakmann B (1997a) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**:213.

Markram H, Lübke J, Frotscher M, Sakmann B (1997b) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science* **275**: 213-215.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* **417**: 741-746.

Metherate R, Cox CL, Ashe JH (1992) Cellular bases of neocortical activation: modulation of neural oscillations by the nucleus basalis and endogenous acetylcholine. *J Neurosci* **12**: 4701-4711.

Middleton S, Jalics J, Kispersky T, Lebeau FE, Roopun AK, Kopell NJ, Whittington MA, Cunningham MO (2008) NMDA receptor-dependent switching between different gamma rhythm-generating microcircuits in entorhinal cortex. *Proc Natl Acad Sci U S A* **105**: 18572-18577.

Miltner WH, Braun C, Arnold M, Witte H, Taub E (1999) Coherence of gamma-band EEG activity as a basis for associative learning. *Nature* **397**: 434-436.

Moore CI, Carlen M, Knoblich U, Cardin JA (2010) Neocortical interneurons: from diversity, strength. *Cell* **142**: 189-193.

Morris RGM, Garrud P, Rawlins J, O'Keefe JO (1982) Place navigation impaired in rats with hippocampal lesions. *Nature* **297**: 681-683.

Nagel G, Szellas T, Huhn W, Kateriya S, Adeishvili N, Berthold P, Ollig D, Hegemann P, Bamberg E (2003) Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proc Natl Acad Sci U S A* **100**: 13940-13945.

Nase G, Singer W, Monyer H, Engel AK (2003) Features of neuronal synchrony in mouse visual cortex. *J Neurophysiol* **90**: 1115-1123.

Newman JP, Zeller-Townson R, Fong MF, Arcot Desai S, Gross RE, Potter SM (2012) Closed-loop, multichannel experimentation using the open-source NeuroRighter electrophysiology platform. *Front Neural Circuits* **6**:98.

Niell CM, Stryker MP (2010) Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* **65**: 472-479.

O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurology* **51**: 78-109.

O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* **34**: 171-175.

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

Pantev C, Makeig S, Hoke M, Galambos R, Hampson S, Gallen C (1991) Human auditory evoked gamma-band magnetic fields. *Proc Natl Acad Sci U S A* **88**: 8996-9000.

Pavlides C, Greenstein YJ, Grudman M, Winson J (1988) Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of θ-rhythm. *Brain Res* **439**: 383-387.

Paz JT, Davidson TJ, Frechette ES, Delord B, Parada I, Peng K, Deisseroth K, Huguenard JR (2013) Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat Neurosci* **16**: 64-70.

Poulet JF, Petersen CC (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* **454**: 881-885.

Ray S, Maunsell JH (2010) Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* **67**: 885-896.

Ray S, Maunsell JH (2011) Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol* **9**:e1000610.

Ray S, Ni AM, Maunsell JH (2013) Strength of gamma rhythm depends on normalization. *PLoS Biol* **11**:e1001477.

Reynolds JH, Pasternak T, Desimone R (2000) Attention increases sensitivity of V4 neurons. *Neuron* **26**: 703-714.

Roberts MJ, Lowet E, Brunet NM, Ter Wal M, Tiesinga P, Fries P, De Weerd P (2013) Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron* **78**: 523-536.

Rolston JD, Gross RE, Potter SM (2010) Closed-loop, open-source electrophysiology. *Front Neurosci* **4**: 31.

Royer S, Zemelman BV, Losonczy A, Kim J, Chance F, Magee JC, Buzsáki G (2012) Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat Neurosci* **15**: 769-775.

Sainsbury RS, Montoya CP (1984) The relationship between type 2 theta and behavior. *Physiol Behav* **33**: 621-626.

Scheffer-Teixeira R, Belchior H, Leão RN, Ribeiro S, Tort AB (2013) On high-frequency field oscillations (>100 Hz) and the spectral leakage of spiking activity. *J Neurosci* **33**: 1535-1539.

Schomburg EW, Anastassiou CA, Buzsáki G, Koch C (2012) The spiking component of oscillatory extracellular potentials in the rat hippocampus. *J Neurosci* **32**: 11798-11811.

Shadlen MN, Movshon JA (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**: 67-77, 111-25.

Sieben K, Röder B, Hanganu-Opatz IL (2013) Oscillatory entrainment of primary somatosensory cortex encodes visual control of tactile processing. *J Neurosci* **33**: 5736-5749.

Sigurdsson T, Stark KL, Karayiorgou M, Gogos JA, Gordon JA (2010) Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. *Nature* **464**: 763-767.

Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? *Neuron* **24**: 49-65, 111-25.

Singer W, Gray CM (1995) Visual feature integration and the temporal correlation hypothesis. *Annu Rev Neurosci* **18**: 555-586.

Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* **60**: 683-697.

Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus* **6**: 149-172.

Sohal VS, Zhang F, Yizhar O, Deisseroth K (2009) Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature* **459**: 698-702.

Somogyi P, Klausberger T (2005) Defined types of cortical interneurone structure space and spike timing in the hippocampus. *J Physiol* **562**: 9-26.

Spaak E, de Lange FP, Jensen O (2014) Local entrainment of alpha oscillations by visual stimuli causes cyclic modulation of perception. *J Neurosci* **34**: 3536-3544.

Srinivasan R, Winter WR, Ding J, Nunez PL (2007) EEG and MEG coherence: measures of functional connectivity at distinct spatial scales of neocortical dynamics. *J Neurosci Methods* **166**: 41-52.

Stark E, Eichler R, Roux L, Fujisawa S, Rotstein H, Buzsáki G (2013) Inhibition-induced theta resonance in cortical circuits. *Neuron* **80**: 1263-1276.

Stevens CF, Zador AM (1998) Input synchrony and the irregular firing of cortical neurons. *Nat Neurosci* **1**: 210-217.

Stewart M, Fox SE (1991) Hippocampal theta activity in monkeys. *Brain Res* **538**: 59-63.

Sukov W, Barth DS (1998) Three-dimensional analysis of spontaneous and thalamically evoked gamma oscillations in auditory cortex. *J Neurophysiol* **79**: 2875-2884.

Tiesinga P, Sejnowski TJ (2009) Cortical enlightenment: are attentional gamma oscillations driven by ING or PING? *Neuron* **63**: 727-732.

Tiesinga PH (2005) Stimulus competition by inhibitory interference. *Neural Comput* **17**: 2421-2453.

Tiesinga PH, Fellous JM, Salinas E, José JV, Sejnowski TJ (2004) Inhibitory synchrony as a mechanism for attentional gain modulation. *J Physiol Paris* **98**: 296-314.

Tóth K, Borhegyi Z, Freund TF (1993) Postsynaptic targets of GABAergic hippocampal neurons in the medial septum-diagonal band of broca complex. *J Neurosci* **13**: 3712-3724.

Traub RD, Jefferys JG, Whittington MA (1997) Simulation of gamma rhythms in networks of interneurons and pyramidal cells. *J Comput Neurosci* **4**: 141-150.

Traub RD, Whittington MA, Colling SB, Buzsáki G, Jefferys JG (1996) Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *J Physiol* **493 (Pt 2)**: 471-484.

Ulanovsky N, Moss CF (2007) Hippocampal cellular and network activity in freely moving echolocating bats. *Nat Neurosci* **10**: 224-233.

Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol* **26**: 407-418.

van Pelt S, Boomsma DI, Fries P (2012) Magnetoencephalography in twins reveals a strong genetic determination of the peak frequency of visually induced gamma-band synchronization. *J Neurosci* **32**: 3388-3392.

Vertes RP, Albo Z, Viana Di Prisco G (2001) Theta-rhythmically firing neurons in the anterior thalamus: implications for mnemonic functions of Papez's circuit. *Neuroscience* **104**: 619-625.

Vertes RP, Hoover WB, Viana Di Prisco G (2004) Theta rhythm of the hippocampus: subcortical control and functional significance. *Behav Cogn Neurosci Rev* **3**: 173-200.

Wang XJ, Buzsáki G (1996) Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J Neurosci* **16**: 6402-6413.

Whittington MA, Traub RD, Jefferys JG (1995) Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature* **373**: 612-615.

Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* **261**: 1055-1058.

Wimber M, Maaß A, Staudigl T, Richardson-Klavehn A, Hanslmayr S (2012) Rapid memory reactivation revealed by oscillatory entrainment. *Curr Biol* **22**: 1482-1486.

Winson J (1978) Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science* **201**: 160-163.

Womelsdorf T, Fries P (2007) The role of neuronal synchronization in selective attention. *Curr Opin Neurobiol* **17**: 154-160.

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**: 733-736.

Wulff P, Ponomarenko AA, Bartos M, Korotkova TM, Fuchs EC, Bähner F, Both M, Tort AB, Kopell NJ, Wisden W, Monyer H (2009) Hippocampal theta rhythm and its coupling with gamma oscillations require fast inhibition onto parvalbumin-positive interneurons. *Proc Natl Acad Sci U S A* **106**: 3561-3566.

Xing D, Shen Y, Burns S, Yeh CI, Shapley R, Li W (2012) Stochastic generation of gamma-band activity in primary visual cortex of awake and anesthetized monkeys. *J Neurosci* **32**: 13873-13880.

Xing D, Yeh CI, Shapley RM (2009) Spatial spread of the local field potential and its laminar variation in visual cortex. *J Neurosci* **29**: 11540-11549.

Yizhar O, Fenno LE, Davidson TJ, Mogri M, Deisseroth K (2011) Optogenetics in neural systems. *Neuron* **71**: 9-34.

Yoder RM, Pang KC (2005) Involvement of GABAergic and cholinergic medial septal neurons in hippocampal theta rhythm. *Hippocampus* **15**: 381-392.

Yuval-Greenberg S, Tomer O, Keren AS, Nelken I, Deouell LY (2008) Transient induced gamma-band response in EEG as a manifestation of miniature saccades. *Neuron* **58**: 429-441.

Zhang F, Wang LP, Boyden ES, Deisseroth K (2006) Channelrhodopsin-2 and optical control of excitable cells. *Nat Methods* **3**: 785-792.

Zhang ZG, Hu L, Hung YS, Mouraux A, Iannetti GD (2012) Gamma-band oscillations in the primary somatosensory cortex—a direct and obligatory correlate of subjective pain intensity. *J Neurosci* **32**: 7429-7438.

Chapter 2: Gamma-range synchronization of fast-spiking interneurons can enhance detection of tactile stimuli[1]

2.1 Abstract

We tested the sensory impact of repeated synchronization of fast-spiking interneurons (FS), an activity pattern thought to underlie neocortical gamma oscillations. We optogenetically drove "FS-gamma" while mice detected naturalistic vibrissal stimuli and found enhanced detection of less salient stimuli and impaired detection of more salient ones. Prior studies have predicted that the benefit of FS-gamma is generated when sensory neocortical excitation arrives in a specific temporal window 20-25 ms after FS synchronization. To systematically test this prediction, we aligned periodic tactile and optogenetic stimulation. We found that the detection of less salient stimuli was improved only when peripheral drive led to the arrival of excitation 20-25 ms after synchronization and that other temporal alignments either had no effects or impaired detection. These results provide causal evidence that FS-gamma can enhance processing of less salient stimuli, those that benefit from the allocation of attention.

[1]The findings in this chapter were in press at the time of thesis submission (Siegle, Pritchett, and Moore, *Nature Neuroscience*, 2014).

2.2 Introduction

Neocortical oscillations in the gamma range (~30–80 Hz) are hypothesized to benefit sensory processing, a prediction that has generated significant debate. In support of this view, elevated gamma activity is correlated with the allocation of attention (Fries et al., 2001; Fries et al., 2008; Bosman et al., 2012; Ray et al., 2013), and predicts enhanced performance on sensory detection tasks (Meador et al., 2002; Womelsdorf et al., 2006; Hoogenboom et al., 2010). Further support comes from a variety of computational studies that have found that realistic modeling of gamma —emerging from synchronous activity in fast-spiking interneurons—can enhance the relay of signals in simulated neural networks (Wang and Buzsáki, 1996; Börgers and Kopell, 2003; Tiesinga et al., 2004; Knoblich et al., 2010).

These different lines of evidence suggest that gamma expression could explain some of the functional consequences of shifting attention. Specifically, the allocation of attention augments neural and perceptual responses to subthreshold or liminal inputs, the stimuli that require an additional boost for effective processing. However, as emphasized in a number of studies, the relationship between gamma and enhanced processing associated with attention may be purely correlative, an epiphenomenal consequence of increased excitability in the neocortex (Shadlen and Movshon, 1999; Ray and Maunsell, 2010; Ray et al., 2013). This debate surrounding gamma has, to date, only been informed by correlative data. There has not yet been a direct causal examination of the impact—either positive or negative—of realistic gamma oscillations on perception.

A mechanistic understanding of the circuit-level activity patterns driving neocortical gamma is essential to considering its potential role in sensory processing, as well as to guiding causal interventions to test this dynamic. Correlative, causal, and computational evidence indicate that gamma in the 30–80 Hz range depends on synchronous activity of fast-spiking interneurons (FS) (Penttonen et al., 1998; Whittington et al., 2000; Tiesinga et al., 2004; Hasenstaub et al., 2005; Börgers et al., 2008; Cardin et al., 2009; Sohal et al., 2009; Buzsáki and Wang, 2012). Coincident FS activity drives a period of concentrated inhibition, the relaxation from which creates a "window of opportunity" for adjacent pyramidal neurons to relay signals to downstream targets

(Moore et al., 2010). These excitatory spikes recruit nearby FS, thereby initiating the next gamma cycle (Buzsáki and Wang, 2012). The time constant of this inhibition (on the order of ~15 ms) explains why gamma occurs in the 30–80 Hz range (Traub et al., 1996). *In vivo* data supporting this mechanism have shown phase-locking of FS and pyramidal neurons during periods of high-amplitude gamma activity (Hasenstaub et al., 2005). Further, selective optogenetic drive of FS *in vivo* entrains a rhythm that emulates many properties of endogenous gamma (Cardin et al., 2009; Sohal et al., 2009; Cardin et al., 2010; Histed and Maunsell, 2014). Because the mechanism generating gamma is inherent to understanding this phenomenon (Tiesinga et al., 2004; Buzsáki and Wang, 2012), throughout we refer to "FS-gamma." The present findings pertain to gamma emerging through FS synchronization, and not simply to electrophysiological activity in this frequency range. The LFP measurements that are typically used to characterize gamma are considered to be a consequence of this underlying mechanism.

In the present study, we examined the behavioral impact of directly entraining FS-gamma in primary sensory neocortex. To test this question in a well-controlled model, we selectively activated FS at 40 Hz in the barrel cortex of mice engaged in detecting vibrissal deflections. Barrel cortex is required for detection of such passive vibrissal stimuli (Cohen and Castro-Alamancos, 2010; Miyashita and Feldman, 2013; Sachidhanandam et al., 2013), making this target appropriate for measuring the impact of realistic FS-gamma on perception. We found that endogenous gamma at stimulus onset predicted detection of naturalistic tactile stimuli, specifically for the detection of less salient input. When we used light pulses to entrain peri-stimulus FS, we similarly found improved detection of less salient naturalistic stimuli. To test whether this improvement was a product of more general transformations resulting from increased inhibitory tone, or whether it required access to the cyclical window of optimal processing generated by FS-gamma, we tested the detection of 40 Hz vibrissae deflections presented at temporal lags relative to the induced FS synchrony. In this artificial condition, gamma also selectively enhanced detection of less salient stimuli. Further, improved detection was only observed when the relative timing of feed-forward, sensory-driven excitation arrived in the window of opportunity created by prior FS synchronization. A limited sample of single unit recordings from layer II/III showed that induction of FS-gamma led to temporal sharpening that

predicted behavioral benefit, but that firing rate averaged across this population did not increase, indicating that generic increases in rate are not a sufficient explanation for our findings. Neither rate nor temporal changes predicted a behavioral deficit that was observed at one temporal lag in the artificial stimulus condition, indicating that population-level transformations in layer II/III cannot be simplistically interpreted as the sole driver of the perceptual transformations observed. In sum, these data provide correlative and causal evidence that local FS-gamma can enhance neocortical processing for harder-to-perceive stimuli, consistent with a role for this dynamic in behavioral states such as increased attention.

2.3 Results

2.3.1 Mice can detect vibrissal deflections

We trained head-fixed mice to respond to 400 ms sequences of vibrissae deflections. These sequences were either stochastic ("naturalistic" stimulation) or 40 Hz pulses of uniform amplitude. If mice licked a reward spout within 500 ms of the first deflection, a drop of water was delivered (Fig. 1a–b). Mice sustained performance of the detection task consistently over weeks to months, with an average of 495 ± 96 trials per session ($N = 8$ mice; 27,534 ± 11,483 total trials per mouse). In interleaved trials, optogenetic drive was applied to FS, as described in greater detail below. Mice did not respond directly to this stimulation, showing nearly identical false alarm rates on catch trials with light present or absent (Fig. 1c; $N = 8$ mice; 0.22 ± 0.04 versus 0.21 ± 0.06; $P = 0.84$, Wilcoxon signed-rank test). Further, reaction times on "hit" trials were also indistinguishable during light-present and light-absent trials (211 ± 32 ms versus 210 ± 32 ms, respectively).

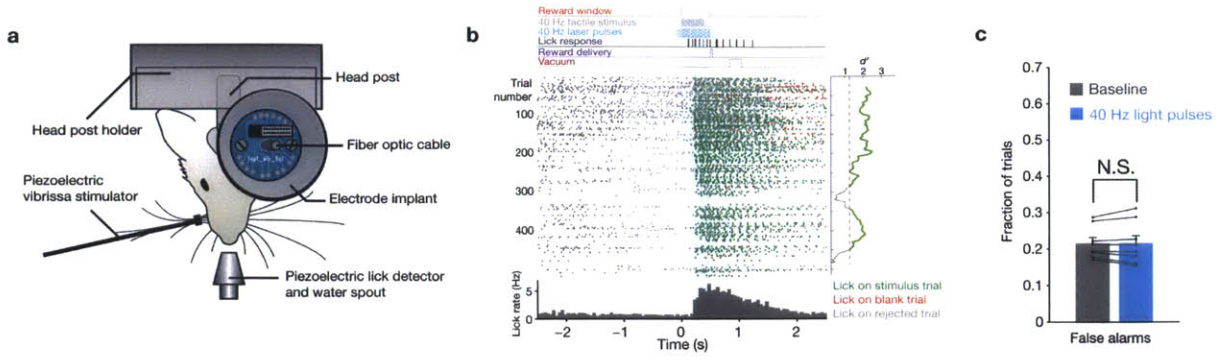

Figure 1 | Task structure. a, Overview of the behavioral setup. Head-fixed mice licked a reward spout to report detection of vibrissal motion. An array of chronically implanted multi-wire electrodes and a fiber optic cable were used to simultaneously monitor local neural activity and optogenetically drive fast-spiking (FS) interneurons. **b**, Top: Structure of an example trial, indicating the timing of the vibrissal stimulus relative to laser pulses, licking behavior, and the reward window. Middle: raster plot of lick times for an example session, aligned to the start of the reward window (gray shading). Each dot indicates a single lick, with color representing the trial type (hit, false alarm, or rejected). Licks before time zero are either false alarms or licks from the end of the previous trial. The mean d' calculated over 50-trial periods is shown to the right of the raster plot, with the threshold of 1.25 indicated. Bottom: Histogram of lick times for this session. **c**, The false alarm rate on trials with and without light presentation ($N = 8$ mice, 141,343 total trials, error bars indicate mean ± s.e.m.) indicate that mice did not respond to optical stimulation. Values for individual mice are superimposed in black.

2.3.2 Gamma is present in spontaneous SI activity

Gamma oscillations have previously been observed in rodent primary somatosensory neocortex (SI), both during baseline conditions (Hamada et al., 1999; Sirota et al., 2008) and in response to sensory drive (Jones and Barth, 1997; Shaw and Chew, 2003). These oscillations also arise in SI *in vitro* when stimulating pyramidal cells of the upper layers with ramping optogenetic stimuli (Adesnik and Scanziani, 2010; Shao et al., 2013). Using chronically implanted stereotrodes in layer II/III, we recorded local field potential (LFP) activity while mice performed the detection task. The average pre-stimulus power spectra for 2 s intervals of inter-trial data averaged over an entire session showed a typical 1/f falloff, and did not contain peaks in the gamma range (Fig. 2a). This finding replicates recent reports in mouse barrel neocortex using this analysis method (Baranauskas et al., 2012). When looking at finer temporal scales, however, distinct bouts of spontaneously occurring increases in 30–80 Hz power were evident during the

pre-stimulus period. Examples of these distinct epochs are illustrated in Fig. 2b and Supplementary Fig. 1. These bouts of gamma were identifiable in the raw LFP traces and in time-frequency spectrograms computed by either wavelet and Fourier transforms. They were detected at a rate of 0.89 Hz during inter-trial intervals without licking.

Multiunit and single-unit firing patterns were strongly modulated as a function of the phase of spontaneous gamma LFP power (Fig. 2c). Multiunit spike activity was concentrated at the trough of the LFP (mean phase across electrodes = 182.2 ± 26.9°, 180° indicates trough). This alignment and increase in phase-locking replicates many prior studies (Fries et al., 2001; Hasenstaub et al., 2005; Womelsdorf et al., 2006; Sirota et al., 2008; Ray et al., 2013; Jia et al., 2013), indicating that the endogenous gamma that occurs in mouse barrel neocortex is associated with the same transformations observed in monkey higher visual neocortex (Fries et al., 2001) and ferret frontal neocortex (Hasenstaub et al., 2005). This pattern of phase-locking of FS is also consistent with computational studies indicating the importance of synchrony in this cell class for the generation of naturally-occurring gamma (Wang and Buzsáki, 1996; Tiesinga et al., 2004; Börgers et al., 2005; Buzsáki and Wang, 2012), and with recent optogenetic data from mouse SI showing that selective drive of FS generates naturalistic gamma *in vivo* (Cardin et al., 2009; Carlén et al., 2011; Sohal et al., 2009).

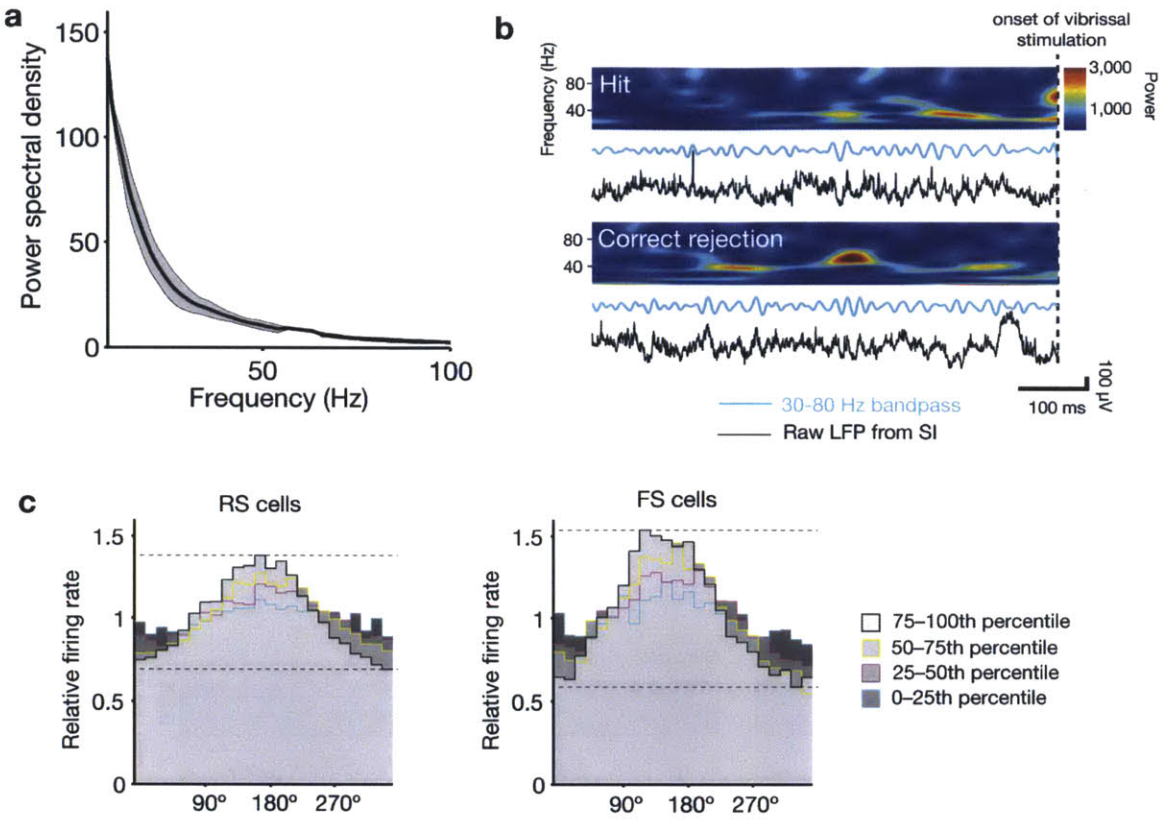

Figure 2 | Gamma is expressed in distinct bouts of activity in the local field potential (LFP), associated with enhanced phase locking of spiking activity. a, Average LFP power spectral density for 3 electrodes from 3 mice (multi-taper spectra with time-bandwidth product of 3 and taper count of 5, computed over 2 s pre-stimulus intervals for an entire session, error bars indicate mean ± s.e.m.). A peak in the gamma band is not apparent when averaging over these longer intervals. **b**, In contrast, when more refined temporal analyses are applied (see text), spontaneously arising bouts of ~40 Hz oscillations in the local field potential are common during the pre-stimulus period. Such events are apparent in the time-frequency spectrograms (top), band-passed LFP (blue lines), and raw LFP (black lines), and occurred approximately once per second. See also Supplementary Fig. 1 for examples from all mice. **c**, Single-unit cycle histograms normalized by baseline firing rate and plotted as a function of the endogenous level of 30–80 Hz oscillations in the pre-stimulus period. Each detected interval was sorted into quartiles based on the amplitude of the LFP signal filtered in the gamma band. Periods with the highest gamma power (75^{th}–100^{th} percentile) show the highest modulation of spike activity. Units were classified as regular-spiking (RS) or fast-spiking (RS) according to their peak-to-trough ratio and spread (Supplementary Fig. 4).

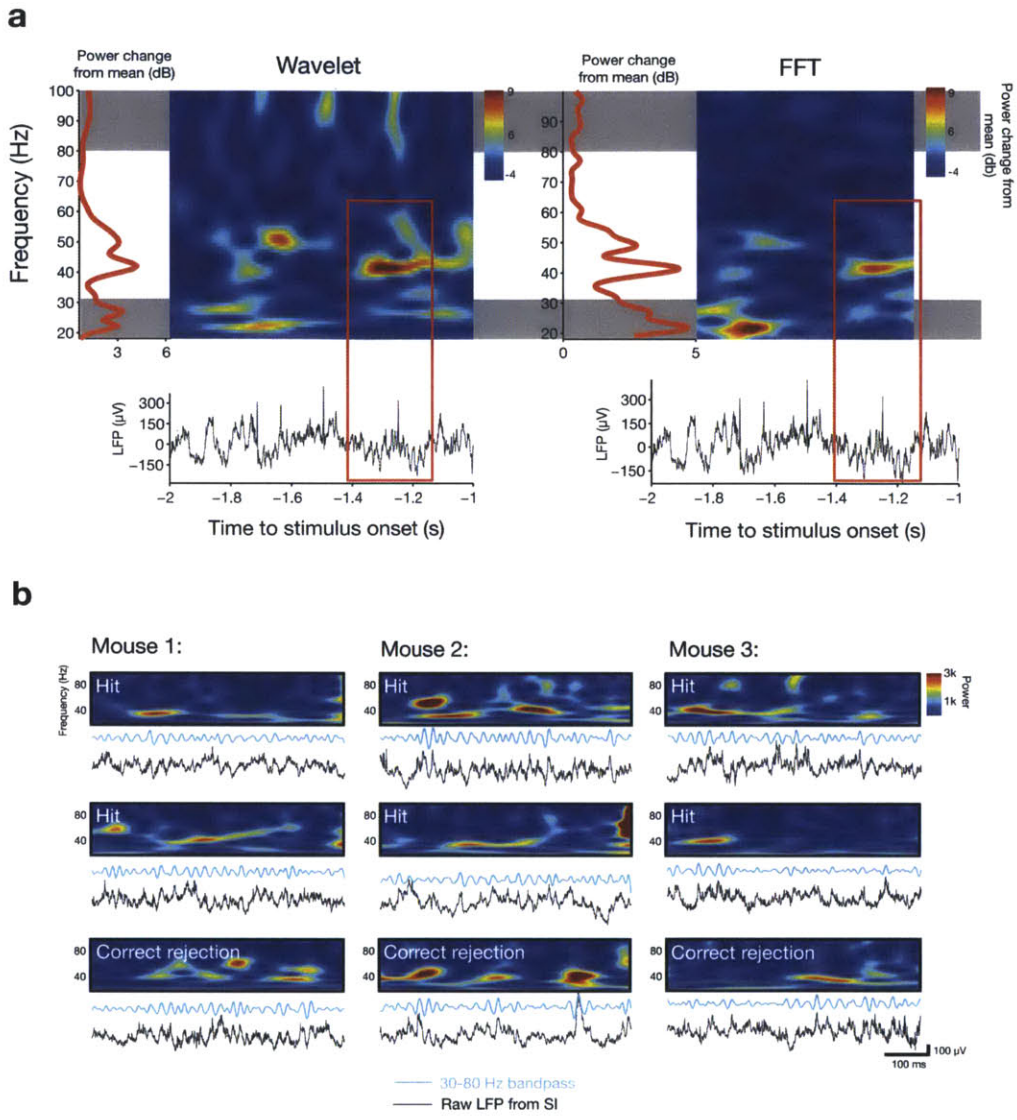

Supplementary Figure 1 | Quantifying gamma events. a, An example of a gamma event in the local field potential (bottom), wavelet power spectrum and spectrogram (top left), and Fourier transform power spectrum and spectrogram (top right). Note that the highlighted peak in the spectrogram is associated with a clearly visible 40 Hz oscillation in the raw voltage trace. The 30–80 Hz range for event centroids is shown in white. Similar events were observed in all animals, with an average rate of occurrence of 0.89 Hz during periods in which the animals were not licking. **b**, Examples of spontaneously arising gamma oscillations during the pre-stimulus period (similar to Figure 2b in the main text). Mouse 1 is the same subject used in Figure 2b. Gamma events also occurred on "miss" trials, but at a lower frequency. Mouse 1: 1.430 events/trial on hits vs. 1.407 events/trial on misses; Mouse 2: 1.309 events/trial on hits vs. 1.272 events/trial on misses; Mouse 3: 1.245 events/trial on hits vs. 1.205 events/trial on misses.

2.3.3 Optogenetically entrained gamma in awake mice

We brought gamma oscillations under experimental control by expressing the gene for channelrhodopsin-2, a light-gated cation channel (H134R variant), in parvalbumin-positive (PV) cells in the upper layers of the neocortex (Supplementary Fig. 2). In this layer, PV neurons are uniformly FS (Cardin et al., 2009). We emulated naturally occurring bouts of gamma by presenting 1 ms light pulses at 40 Hz for 600 ms. The light stimulus began 100 ms prior to the onset of vibrissal deflections, and ended 100 ms after (Fig. 1b). Light power at the surface of the cortex was 30 mW/mm^2, a lower level than that used in a previous study from our lab (Cardin et al., 2009).

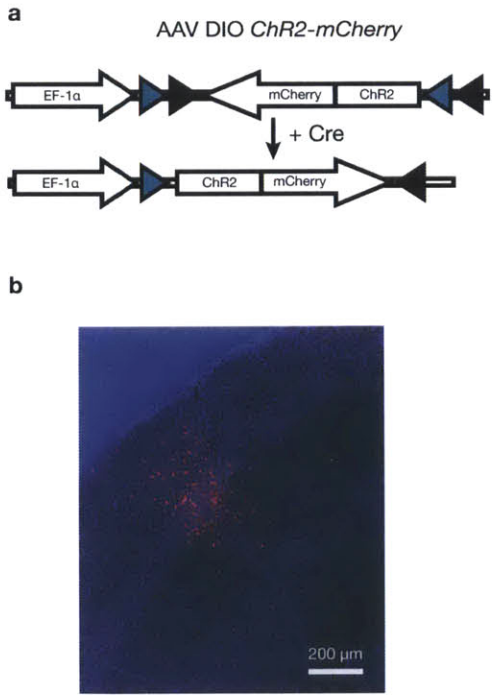

Supplementary Figure 2 | AAV and histology. a, Schematic of the viral vector, which interacts with Cre in parvalbumin-positive cells to induce expression of ChR2-mCherry. **b**, Example histological section, showing viral expression primarily in the upper layers of barrel cortex. Expression was typically distributed over 1–3 barrel columns (300–500 μm medial–lateral spread, 500–1000 μm anterior–posterior spread). Blue = DAPI stain, red = mCherry fluorescence. Viral expression was confirmed in all experimental animals.

This optical input drove robust gamma entrainment in the LFP of awake mice (Fig. 3a) and short-latency action potentials in the multiunit activity (Fig. 3b). Optogenetic stimulation induced a 40 Hz peak in the LFP power spectral density (Fig. 3c), consistent with prior findings in anesthetized preparations (Cardin et al., 2009; Sohal et al., 2009; Carlén et al., 2011). Light pulses activated well-isolated FS, with evoked spikes concentrated between 2–3 ms after the onset of the laser (Fig. 3d), similar to previous findings (Cardin et al., 2009; Carlén et al., 2011). We refer to these effects as "entrained" FS-gamma because the network impact of these oscillations on spiking mimics physiological gamma *in vivo* (Hasenstaub et al., 2005; Cardin et al., 2009)(Fig. 2c) and optogenetic manipulation of this type can be used to phase-reset ongoing endogenous gamma oscillations (Cardin et al., 2009). The observed effects are consistent with FS activation, rather than a light-induced artifact at the electrode interface, as such artifacts are positive-going in input-inverted recordings (same direction as spikes), begin precisely at the time of light onset, and do not suppress spike activity (Cardin et al., 2010).

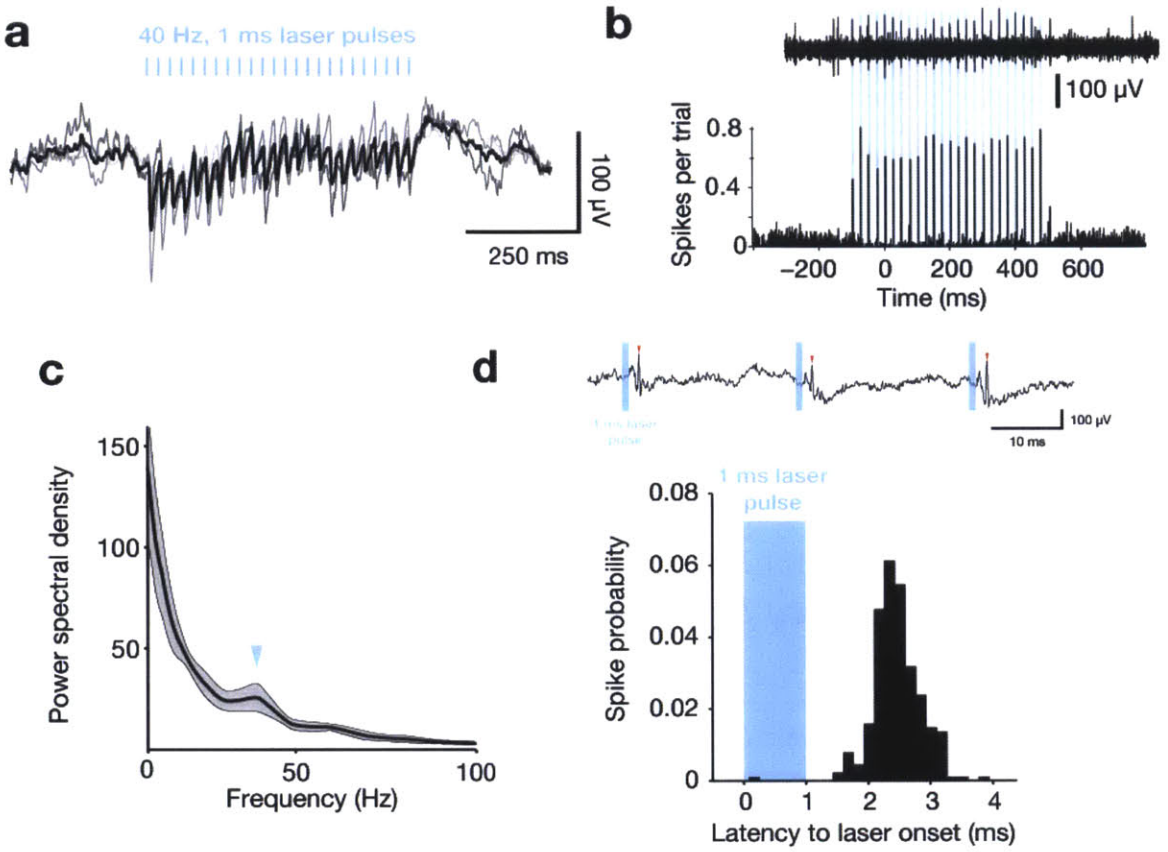

Figure 3 | Optogenetic stimulation to emulate FS-gamma. a, The mean (thick line) and individual (thin lines) light-evoked responses in the input-inverted (spikes positive) local field potential to 40-Hz optogenetic drive of FS. Thin black lines represent mean response for each of 3 mice, thick black line represents the average across mice; vertical blue lines indicate timing of 1 ms laser pulses. **b**, An example of light-evoked spikes from a single trial (top) and the average of 96 trials from a single session (bottom). Vertical blue lines indicate timing of 1 ms laser pulses. **c**, Average LFP power spectral density for 3 electrodes from 3 mice, computed using the same parameters as Fig. 2a, except during epochs of 40 Hz laser pulses only. **d**, Example of a well-isolated, light-driven FS cell, showing an unfiltered trace with identified spikes (top) and histogram of spikes relative to laser onset (bottom). The vast majority of spikes occur between 2 and 3 ms after the start of the light pulse, replicating prior findings in anesthetized mice (Cardin et al., 2009).

2.3.4 Acquisition and characterization of vibrissal stimuli

To test the association of endogenous and entrained gamma with the detection of vibrissae deflections possessing natural stimulus statistics, we recorded the motions of an *ex vivo* vibrissa contacting a textured, rotating surface, and presented these patterns to behaving mice (Fig. 4a; Supplementary Fig. 3a). Using this approach, we generated 17 distinct "movies" of vibrissal motion to replay for psychophysical testing. High-velocity micro-motions are a common feature of vibrissal motion during surface contact, including during free behavior (Ritt et al., 2008; Wolfe et al., 2008), and are widely regarded as a key driver of salience in the vibrissa sensory system. Neurons in SI are more effectively driven by higher velocity vibrissal motions (Pinto et al., 2000), higher velocity of small amplitude motions predicts increased behavioral detection probability (Gerdjikov et al., 2010), and individual high velocity micro-motion events drive spiking in SI awake, freely behaving rodents (Jadhav et al., 2009). The frequency of vibrissal motion for our stimuli was 180–200 Hz (Fig. 4b). The 17 natural stimuli movies were typified by the presence of high-frequency micro-motions, with a mean peak velocity of $1.75 \pm 0.41 \times 10^3$ degrees/s and maximal peak velocity of $5.53 \pm 1.73 \times 10^3$ degrees/s, (Fig. 4c). The occurrences of these micro-motions were evenly distributed relative to the timing of entrained FS-gamma, with no bias in alignment to a phase or temporal delay of optogenetic drive (Fig. 4d).

Four mice were presented with these 17 naturalistic velocity profiles interleaved with the periodic stimuli described below. To ensure that all trials were drawn from epochs in which mice were actively performing the task, we only analyzed blocks of trials in which behavioral performance was high (d' was above 1.25; Fig. 1b, Supplementary Fig. 4). We also eliminated trials in which mice licked during a 1.5-second interval before trial onset. These filtering steps left $27.9 \pm 9.0\%$ of the original trials in our analysis, with a minimum of 469 trials per mouse in the naturalistic stimulation condition.

The 17 natural stimuli were not equivalent in terms of their detectability, with mean d' across animals ranging from 1.57 ± 0.47 to 2.66 ± 0.55. Consistent with prior psychophysical studies[41], stimulus detectability was correlated with higher velocity of vibrissal motion, as the integrated velocity envelope during the first 100 ms of stimulation significantly predicted d' (Fig. 4e–f, $N = 17$ stimuli; $R^2 = 0.4468$; $P = 0.0034$; mean reaction time of 206 ± 20 ms).

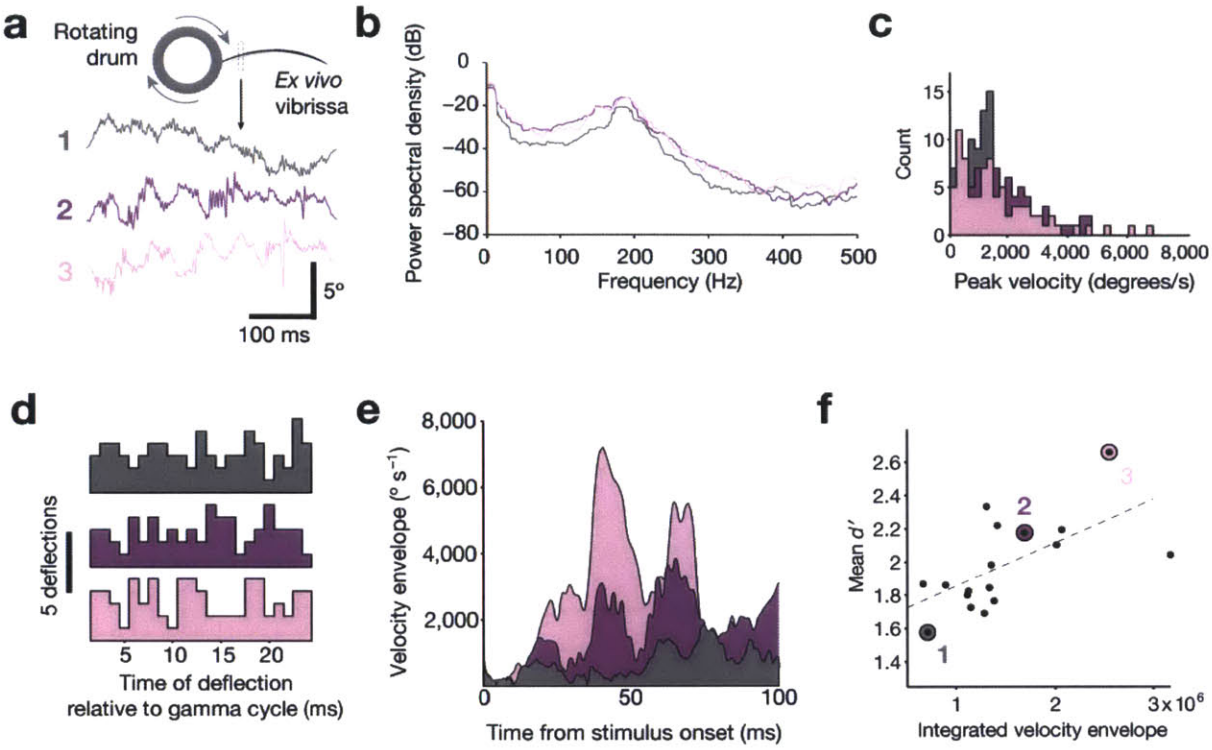

Figure 4 | Generation and characterization of naturalistic vibrissal stimuli a, Movements of an *ex vivo* B2 vibrissa during contact with a sandpaper-covered rotating drum were recorded to obtain naturalistic stimulation patterns (3 examples shown here, see text for details). **b**, Power spectral densities computed for the three naturalistic stimuli shown in panel a, revealing a peak at ~200 Hz. **c**, Histograms of the peak micro-motion velocities for all deflections in the three example waveforms. The range is comparable to velocities measured from the vibrissae of freely behaving rats[38]. **d,** Distribution of deflections within cycles of optogenetically entrained gamma. For each stimulus, deflections with a peak velocity over 100°/s were aligned according to their latency from pulses of 40 Hz light. These high-velocity micro-motion events were evenly distributed throughout the gamma cycle, preventing any preferential phase alignment. **e**, Velocity envelope of the 3 naturalistic stimuli shown in panel a. **f**, Relationship between the integrated velocity envelope for the first 100 ms of each naturalistic stimulus and the mean detectability of that stimulus across $N = 4$ mice (1981 total trials), detection rates for the 3 traces shown in panel a are indicated.

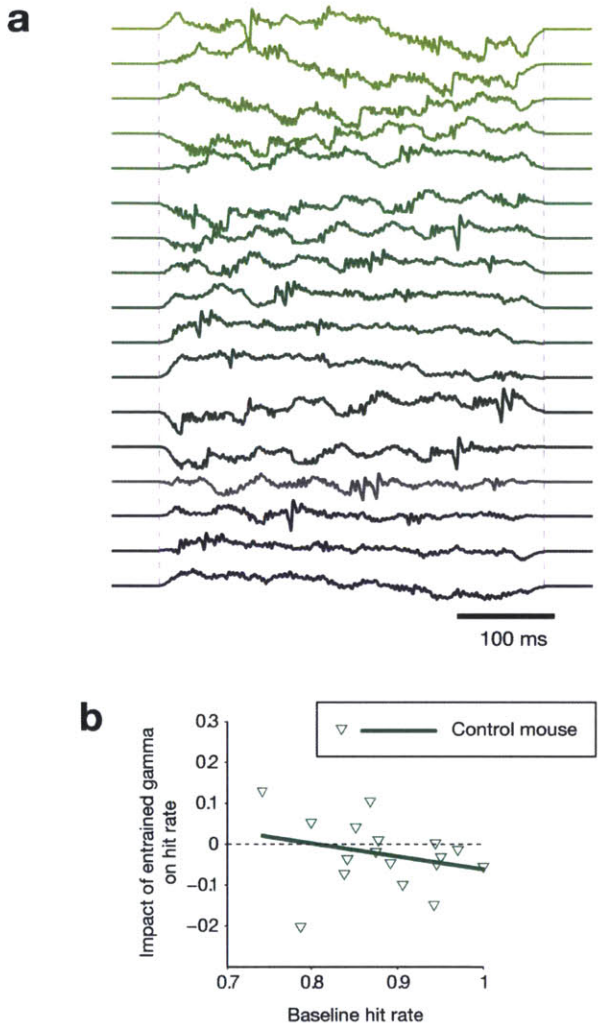

Supplementary Figure 3 | Naturalistic stimuli. a, The 17 motion sequences of an *ex vivo* B2 vibrissa contacting a rotating drum covered in sandpaper that were used for natural stimulus presentation in the detection task. **b**, Impact of optical stimulation in a control (Cre-negative littermate) animal on the natural scenes task. This mouse was injected with ChR2 and subjected to the same training protocol as all other mice. In contrast with the four Cre-positive animals, no significant modulation by the laser was observed ($R^2 = 0.073$, $P = 0.30$).

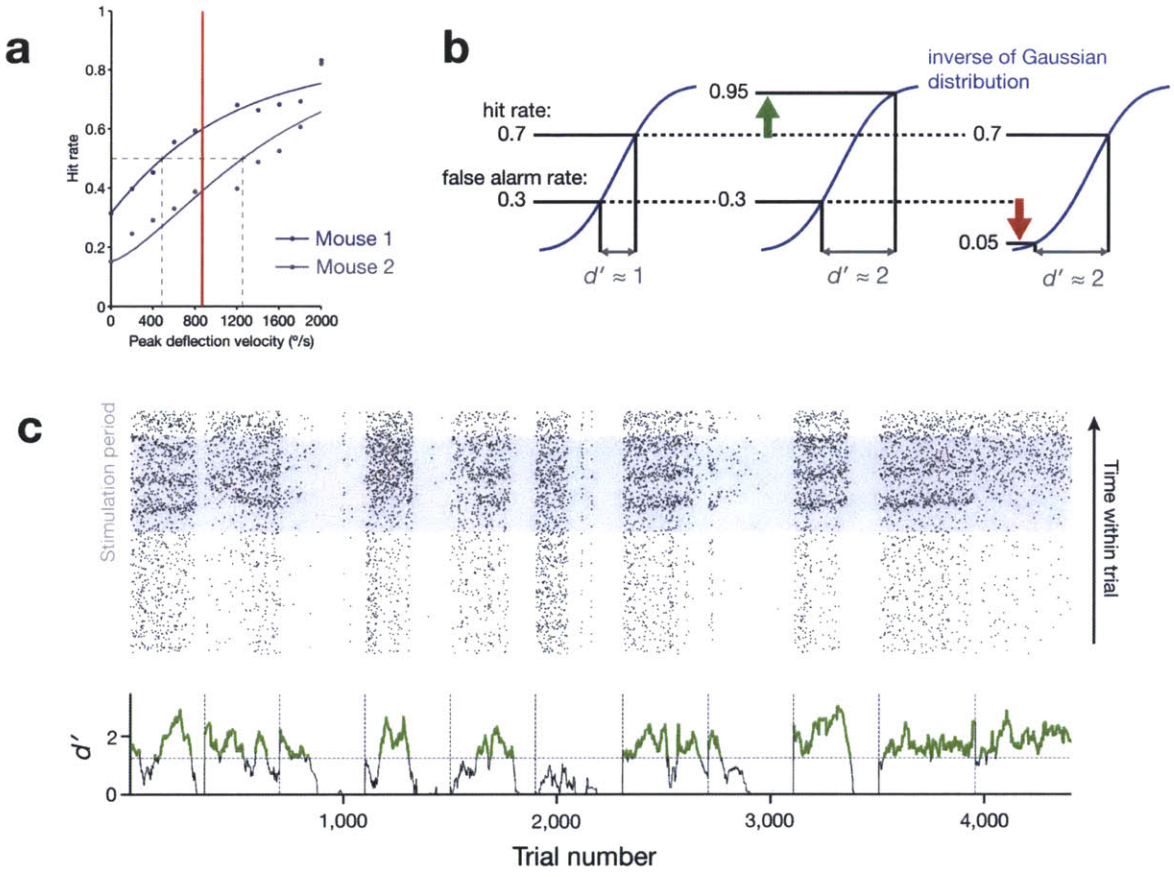

Supplementary Figure 4 | *d′* as an indicator of detection performance. a, Illustration of how *d′* converts hit rate and false alarm rate into a single measure of perceptual acuity. b, Performance of one animal over the course of 4407 trials from 11 sessions. Top plot shows lick times on individual trials (black dots), bottom plot shows change in *d′* for blocks of 50 trials. The animal transitioned between periods of constant licking, accurately licking to the stimulus, and no licking. Only blocks of trials in which *d′* was above a threshold of 1.25 (green lines) were included in further analysis.

2.3.5 Endogenous gamma predicts detection of less salient stimuli

Increased probability of sensory detection, and increased salience of sensory stimuli, have both been found to correlate with increased peri-stimulus gamma (Womelsdorf et al., 2006; Gross et al., 2007). We similarly found that peri-stimulus gamma in SI predicted detection of vibrissal motion for less salient sensory stimuli. Our natural stimuli showed a non-significant broadband increase in LFP power during the peri-stimulus onset period (–200 to 50 ms, relative to vibrissal onset at 0) versus a baseline period (–450 to –200 ms) (Fig. 5a). When we compared the power spectra for hit and miss trials, gamma (30–50 Hz) was higher on hit trials in the peri-stimulus onset period for the least detectable sensory stimuli (Fig. 5b, bottom 50% of stimuli ranked according to detectability; $N = 2$ mice, 7488 trials; 2 mice rejected due to motion artifact in LFP; Bonferroni-corrected t-tests revealed significant difference in the 30–50 Hz band between hits and misses [$P = 4.3 \times 10^{-5}$ for mouse 1, $P = 6.8 \times 10^{-8}$ for mouse 2]).

2.3.6 Entrained FS-gamma and naturalistic stimulation

This result hints at the possibility of an effect of endogenous gamma on detection performance, but is far from conclusive. More extensive sampling of barrel cortex (ideally across identified layers) in more animals will be necessary to support this correlative claim. We also observed differences in the beta band between hits and misses, suggesting that this frequency range could also be involved in detection. These observations mainly serve to highlight the limitations of a correlative approach to studying oscillations, which is why our main goal was to directly test the impact of FS-gamma by optogenetically driving FS in conjunction with naturalistic vibrissal stimulation. On trials where FS were driven at 40 Hz, all mice showed a selective enhancement in detection of less salient stimuli (Fig. 5c). This observation is reflected in a negative correlation between the d' for the 17 natural stimuli (i.e., whether they were more salient for a given mouse) and the impact of entrained gamma on detection (Fig. 5b). We also applied the same manipulation in a PV-Cre negative control mouse, and saw no significant relationship between baseline performance and the impact of the optical drive ($R^2 = 0.1842$, $P > 0.05$; Supplementary Fig. 3b).

As a summary metric across mice of the impact of entrained FS-gamma as a function of

stimulus detectability, we ranked the relative detectability of each of the 17 naturalistic stimuli for each animal, and then averaged the effect of this manipulation across these ordinal data (Fig. 5d). When considered as a function of relative detectability, 8 of 9 of the weakest stimuli showed increased detection with FS-gamma entrainment. We also found that 7 of 8 of the strongest stimuli were negatively impacted. This finding of impaired detection for salient stimuli is in apparent contrast to the lack of an association between endogenous gamma and hit or miss trials for highly salient natural stimuli, and a lack of an impact of driving FS-gamma on detection of highly salient periodic stimuli described below.

This result does not directly support the claim that FS gamma—rather than some other consequence of our optogenetic manipulation—is responsible for altering detection performance. Driving FS at other frequency bands could have provided a control condition to test the effect of stimulation frequency, but we rejected this approach because synchronous activation of FS at higher or lower frequencies would not have been a physiologically relevant manipulation. We also had to confirm that the effect of enhancement of less salient stimuli and impairment of more salient stimuli was not simply a "regression to the mean." We consider this possibility unlikely for two reasons: (1) shuffling the trial labels removes the strong correlation between stimulus detectability and degree of modulation by light and (2) we saw no such correlation in a PV– control mouse with virus injection and fiber–optic implant (Supplementary Fig. 3b). Although the behavioral effects of FS drive are genuine, they cannot be attributed specifically to FS gamma without an appropriate control. For this reason, we conducted what we consider to be a more powerful set of experiments by manipulating the timing of FS synchronization relative to incoming excitation on gamma timescales.

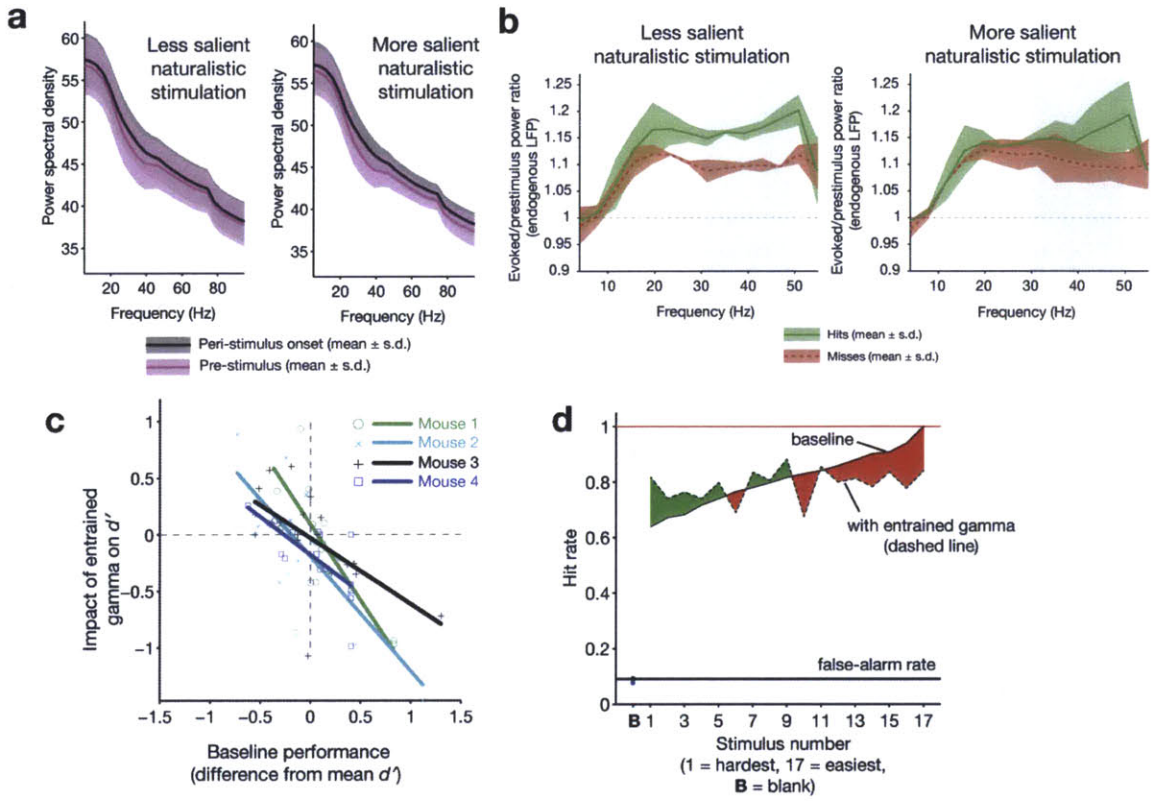

Figure 5 | Endogenous and entrained FS-gamma predict enhanced detection of less salient naturalistic stimuli. a, Baseline versus evoked power spectra, for both less salient and more salient naturalistic stimuli. Error bars indicate mean ± s.e.m. **b,** Endogenous peri-stimulus spectral power in the gamma range during detection of less salient (left, mean of naturalistic stimuli across mice with baseline detectability levels below the median) and most salient (right, above the median) naturalistic stimuli, shown for hit (green) and miss (red) trials. Only for less salient stimuli did power in the 30–50 Hz range (gray bar) differ between hits and misses (red and green shaded regions represent mean ± s.e.m). **c**, Relationship between baseline (laser off) performance and the effect of optogenetically entrained FS-gamma on d' ($N = 17$ stimuli presented to 4 mice). All mice displayed enhanced detection of less salient natural stimuli and a significant correlation between detectability and the entrained impact of FS-gamma (Mouse 1, $R^2 = 0.47$, $P = 0.0024$; Mouse 2, $R^2 = 0.66$, $P = 0.001$; Mouse 3, $R^2 = 0.38$, $P = 0.008$; Mouse 4, $R^2 = 0.54$, $P = 0.0008$). **d**, Average impact of entrained gamma on detection of stimuli ordered by detectability. For each mouse, the 17 natural stimuli were sorted as a function of detectability for that animal and then the impact of FS-gamma on these ordinal rankings were summed across animals. The difference between detection performance with and without entrained gamma is indicated in green (enhanced performance) or red (impaired performance).

2.3.7 Entrained FS-gamma and periodic stimulation

The detection of less salient naturalistic stimuli was enhanced by entrainment of FS-gamma. This enhancement could have arisen from non-specific effects due to selective recruitment of FS, such as generalized arousal or enhanced signal-to-noise through the presence of increased inhibition. Alternatively, these effects could depend on enhancement of the specific temporal pattern of activity associated with FS-gamma. Specifically, prior computational studies (Tiesinga et al., 2004; Knoblich et al., 2010) and experimental work (Fellous et al., 2003; Hasenstaub et al., 2005; Cardin et al., 2009) predict that excitatory input arriving ~20–25 ms after an event of FS synchrony will be preferentially transmitted downstream, as a result of increases in synchrony across pyramidal neurons (Börgers et al., 2008; Knoblich et al., 2010) and potentially by increases in their firing rate (Tiesinga et al., 2004; Knoblich et al., 2010). This optimal temporal window is predicted to be created by a variety of mechanisms, including but not limited to the relaxation of inhibition, as described further in the Discussion.

To test whether arrival of input during this window was associated with improved psychophysical performance, we combined 40 Hz FS activation with 40 Hz vibrissae deflections at a range of temporal alignments. Such precise alignment of an external stimulus with an internal gamma rhythm is unlikely to occur under natural conditions for several reasons, most importantly the fact that natural micro-motion patterns of vibrissae themselves created during surface contact will typically occur at much higher frequencies (Ritt et al., 2008; Wolfe et al., 2008). While artificial, this manipulation allowed us to directly test whether the timing of peripheral input relative to the optimal window created by FS-gamma shaped behavior.

We shifted the temporal lag between optical stimulus and sensory drive at 40 Hz by 5 ms steps, as in a prior study in anesthetized mice (Cardin et al., 2009). The latency from peripheral deflection to arrival of sensory-driven excitatory input to neocortex is ~8–10 ms (Supplementary Fig. 5). As such, neocortical spikes evoked by vibrissal stimuli presented 12.5 msec after a light pulse would be predicted to arrive at the optimal window 20–25 ms after a prior bout of FS synchrony. For this temporal offset, the propagation delay between the periphery and neocortex positions the excitatory barrage from sensory drive to arrive after the decay of FS-driven inhibition, but just before the onset of the subsequent bout of inhibition (Fig. 6a).

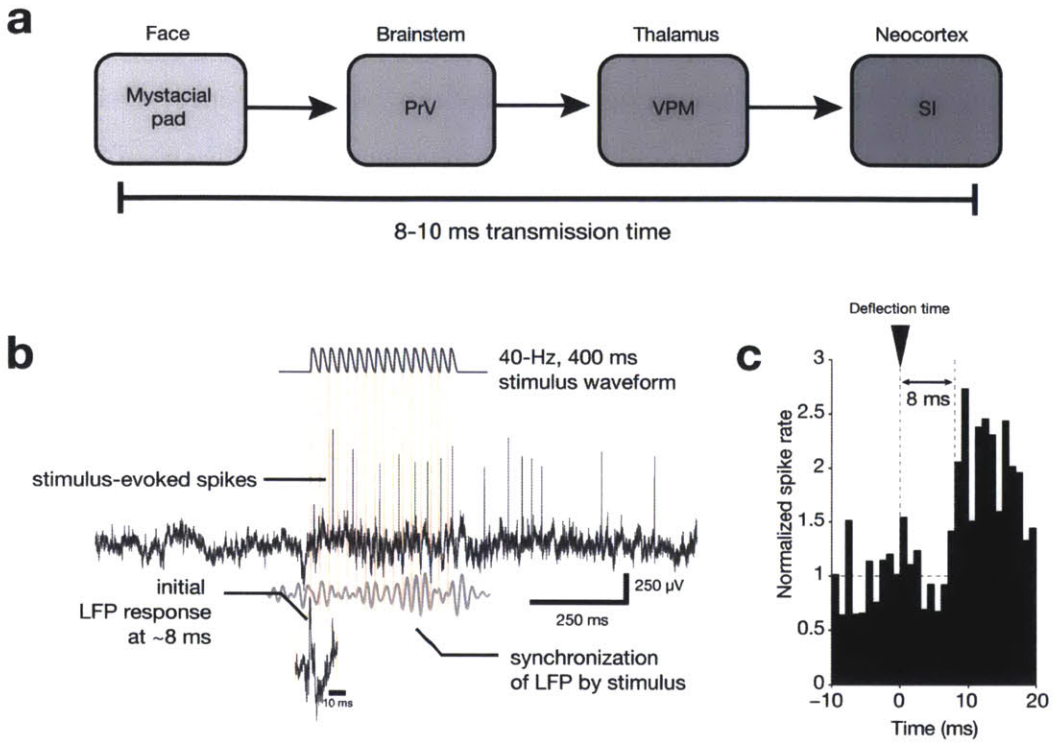

Supplementary Figure 5 | Propagation of signals from the periphery to barrel cortex. a, Schematic showing three synapses between the periphery and the cortex. PrV = principal trigeminal nucleus; VPM = ventral posterior medial nucleus; SI = primary somatosensory cortex. Evoked responses initiated in the periphery take, on average, 8–10 ms to propagate to barrel cortex. **b**, Raw voltage trace recorded during a single trial. A sharp transient in LFP traces from single trials is apparent ~8 ms after the onset of stimulation, indicating the arrival of excitatory EPSPs from the thalamus. **c**, Average peri-stimulus spike histogram for N = 48 RS and FS cells, showing the ~8 ms delay between deflecting the vibrissae (T = 0 ms) and the onset of spike activity in barrel cortex.

In the first two mice trained, we presented stimuli of 9 different amplitudes to construct psychophysical response curves (Supplementary Fig. 4a). Overall performance was relatively poor, with false-alarm rates of around 0.2 and hit rates for maximum-amplitude stimuli peaking around 0.8. There are a number of explanations for this, with the most obvious being that the degree of engagement with the task fluctuates across days and within each session. Supplementary Fig. 4c depicts an example of lick times (and corresponding d') for one mouse across 11 sessions. During periods in which the mouse licks to the stimulus, d' is high (2.0 or greater). But there are some days in which the mouse refuses to lick or performs at chance. These variations are likely due to differences in handling and positioning of the vibrissae stimulator across days. Trial-to-trial variation could result from differences in whisking behavior, which were not measured for this task. If the mouse is actively whisking during stimulus delivery, the movement of the vibrissae within the stimulator could affect detection performance.

Based on the psychophysical curves from the first two animals, we selected for intensive study a less salient amplitude of stimulation that generated intermediate detection rates. In these 2 subjects and in 6 subsequent mice, we randomly interleaved these threshold stimuli (~800°/s peak velocity) with maximum-amplitude stimuli (~2000°/s peak velocity). Across all mice (N = 8), less salient stimuli were harder to detect than maximum-amplitude stimuli (hit rate of 0.74 ± 0.13 vs. 0.90 ± 0.03) and took longer to detect (247 ± 5 ms vs. 224 ± 5 ms), although this difference was not significant. As in the naturalistic stimulation condition, we only analyzed trials from blocks in which d' was above 1.25 (minimum of 631 trials per mouse).

For less salient stimuli, detection was enhanced only for the 12.5 ms temporal offset condition (Fig. 6a). On these trials (represented in green), d' was higher than baseline for 7/8 mice and in the average (Fig. 6a mean difference from baseline = 0.23 ± 0.18, P = 0.0391; Wilcoxon signed-rank test with Bonferroni correction). Hit rate for the 12.5 ms offset was also enhanced and, as indicated above, no change in the false alarm rate on catch trials was observed (Fig. 1c; Supplementary Fig. 6). To determine the likelihood that a single offset would show these effects under the hypothesis that entrained FS-gamma has no impact on performance, we applied a permutation test ($P < 0.01$, see Methods section). In contrast, for all temporal offsets other than 12.5 ms, d' did not differ from baseline or was significantly lower (7.5 ms offset, mean

difference from baseline = –0.30 ± 0.21, $P = 0.0195$; Wilcoxon signed-rank test with Bonferroni correction). Thus, shifting entrained gamma oscillations by as little as 5 ms relative to sensory input led to significant variation in task performance. These data indicate that behavioral enhancement by optogenetic stimulation requires temporal processes specific to FS synchronization, rather than non-specific effects such as heightened arousal or general signal-to-noise transformations resulting from the manipulation.

We also quantified the impact of entrained FS-gamma on the detectability of maximum-amplitude, perceptually salient periodic stimuli, which had a baseline hit rate of 0.90 ± 0.03. In contrast to the enhancement observed with the less salient stimulus, we saw a non-significant trend towards enhancement at the 12.5 ms condition, and no significant changes for any other temporal offset with our optogenetic manipulation (Fig. 6b; $P = 0.711$, Friedman test). Further, when the performance of individual mice was analyzed, we found a range of detection levels for the less salient stimulus, resulting from the fact that stimulus intensity chosen based on performance of our first 2 mice was the same for all animals, and not scaled according to individual psychophysical abilities in each (Fig. 6c). The relative enhancement of detection by entrainment of FS-gamma at the 12.5 ms offset was predicted from the innate performance levels of each mouse. Those that performed inherently better in detecting the less salient stimulus, in the absence of the optogenetic manipulation, showed less benefit of entrained gamma ($N = 8$ mice; $R^2 = 0.64$, $P = 0.017$; Fig. 6d). The best-performing mouse, which had a baseline hit rate of 0.90 for less salient stimuli (identical to the mean detection rate for the "maximal" stimulus), had its performance negatively impacted by laser presentation. These inter-subject analyses further reinforce the finding that less salient stimuli were selectively benefitted by FS-gamma.

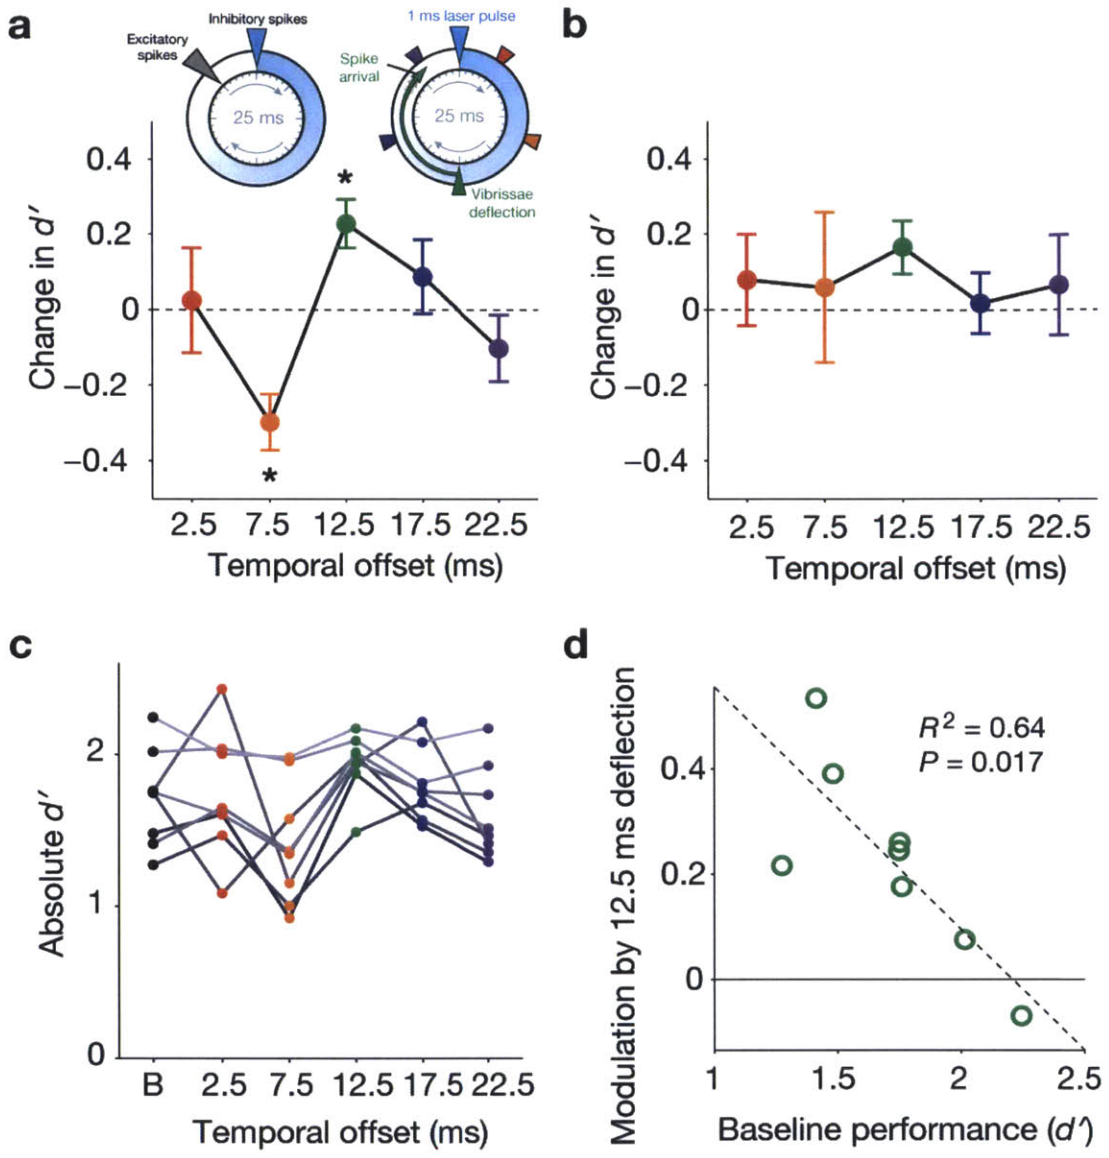

Figure 6 | Entrained FS-gamma enhances detection of less salient periodic stimuli at a specific optimal temporal offset between FS synchronization and sensory drive. a, The inset on the upper left summarizes the findings of prior studies[8-12,17,19,22] showing that FS-gamma creates a window of opportunity for optimal excitatory firing immediately prior to a subsequent FS synchronization event. For a 40 Hz gamma rhythm, this cycle repeats every 25 ms. The inset on the upper right illustrates the temporal offsets between optogenetic light pulses and vibrissae deflections used in this study. Vibrissa stimulation 12.5 ms after FS synchronization (green triangle) will cause stimulus-evoked neocortical excitation to occur in this window of opportunity, reinforcing the endogenous pattern of FS-gamma expression. The main plot shows the relationship between temporal offset of vibrissal stimulation and difference in *d′* from

baseline for $N = 8$ mice (threshold stimulus intensity trials only, see Figure 1). The temporal offset was a significant modulator of response ($P = 0.0016$, Friedman test), with a significant change relative to the "no laser" condition for 7.5 and 12.5 ms offsets ($P < 0.05$, one-sided Wilcoxon signed-rank test with Bonferroni correction). **b**, Same as in b, but for trials on which the maximal, highly salient amplitude stimulus was presented. There was no significant effect of changing the temporal offset for detection of these stimuli ($P = 0.711$, Friedman test). **c**, Absolute d' values for individual mice (B = baseline "no laser" condition; different shades of gray are used to distinguish individuals). **d**, Relationship between baseline performance and d' modulation by the 12.5 ms temporal offset condition for individual mice.

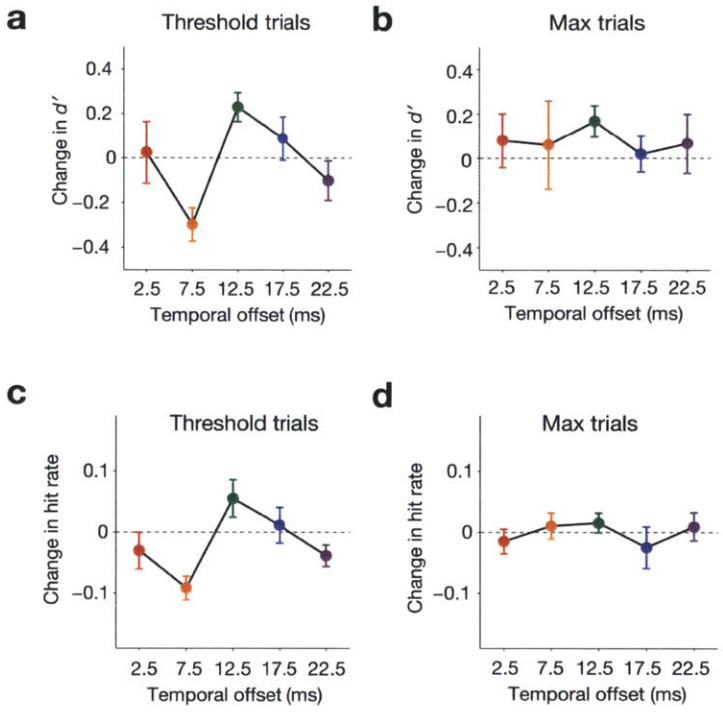

Supplementary Figure 6 | Comparing *d′* and hit rate. a, Average d' as a function of the temporal offset between sensory stimulus and entrained gamma for "threshold" trials with periodic stimulus presentation (same as Figure 6a in the main text). **b**, Average d' as a function of temporal offset between sensory stimulus and entrained gamma for "max" trials (same as Figure 6b in the main text). **c**, Average hit rate as a function of temporal offset between sensory stimulus and entrained gamma for "threshold" trials. **d**, Average hit rate as a function of temporal offset between sensory stimulus and entrained gamma for "max" trials. All panels are for the same $N = 8$ animals as in Figure 6. Error bars indicate s.e.m.

2.3.8 Physiological effects of vibrissa stimulation and FS-gamma

We analyzed 35 well-isolated regular-spiking cells, classified by waveform shape (Supplementary Fig. 7), or multi-unit activity (15 different electrodes, 3 mice). We only included units obtained from stereotrodes on which a rapid sensory-driven response was observed. Entrained FS-gamma transformed the spiking response in two ways that replicated prior findings in anesthetized mice (Cardin et al., 2009). First, for naturalistic and periodic stimulation, optogenetic manipulation suppressed baseline firing rates and decreased the average sensory-driven spike rate during the peri-stimulus period, independent of the relative alignment of sensory stimulation and light (Fig. 7a–e). Second, across all laser-to-vibrissae stimulus alignments, there was a significant increase in the vector strength (a measure of response periodicity of single units) during the peri-stimulus period (Fig. 7a–c,e), indicative of enhanced phase-locking to entrained FS-gamma. Third, in agreement with prior findings (Cardin et al., 2009), delivering the sensory stimulus 12.5 ms after each laser pulse resulted in a selective enhancement in spike precision (a measure of the width of the multiunit evoked response to the first vibrissae deflection) and the strongest enhancement of gamma power in spike time distributions (Fig. 7e, Supplementary Figure 8). This finding is consistent with the arrival of sensory-driven excitatory input ~8–10 ms post-stimulation, which aligns the excitatory spikes to the optimal phase of post-inhibitory decay (Fig. 6a; Supplementary Fig. 5).

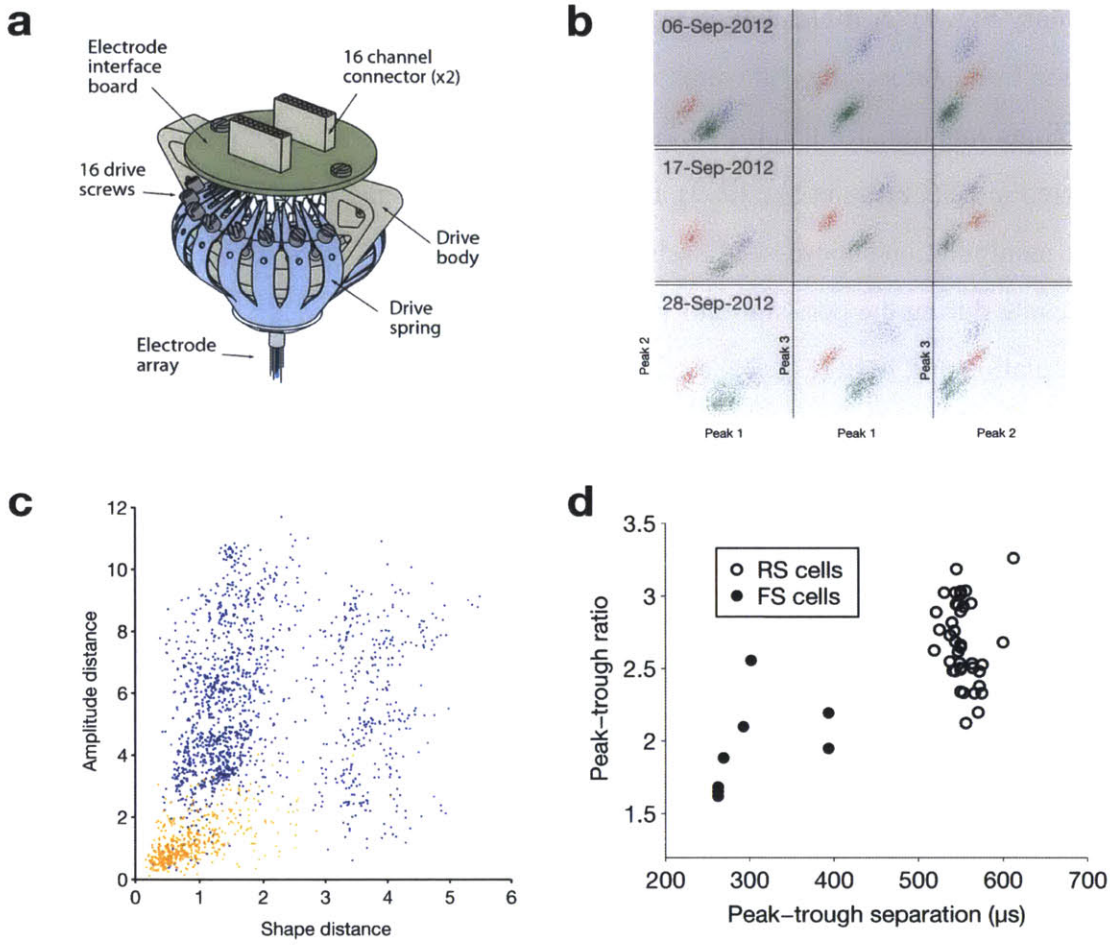

Supplementary Figure 7 | Electrophysiological methods. a, Diagram of the fiber-optic–electrode implant used to record single unit activity in barrel cortex. **b**, Example projection plots for waveforms recorded on the same electrode over the course of a month. Spikes from the same cell are color-coded across days. **c**, Assertions of cell identity, used to minimize chances of double-counting the same cells across days, were analyzed with a waveform similarity metric. Orange dots indicate distance between waveforms identified as coming from the same cell; blue dots indicate distance between waveforms identified as coming from different cells. Distance from the origin indicates degree of waveform similarity. **d**, Measures used to identify cells as FS (fast-spiking) or RS (regular spiking). Cells in the present sample were separated into two groups based on peak–trough ratio and peak–trough separation.

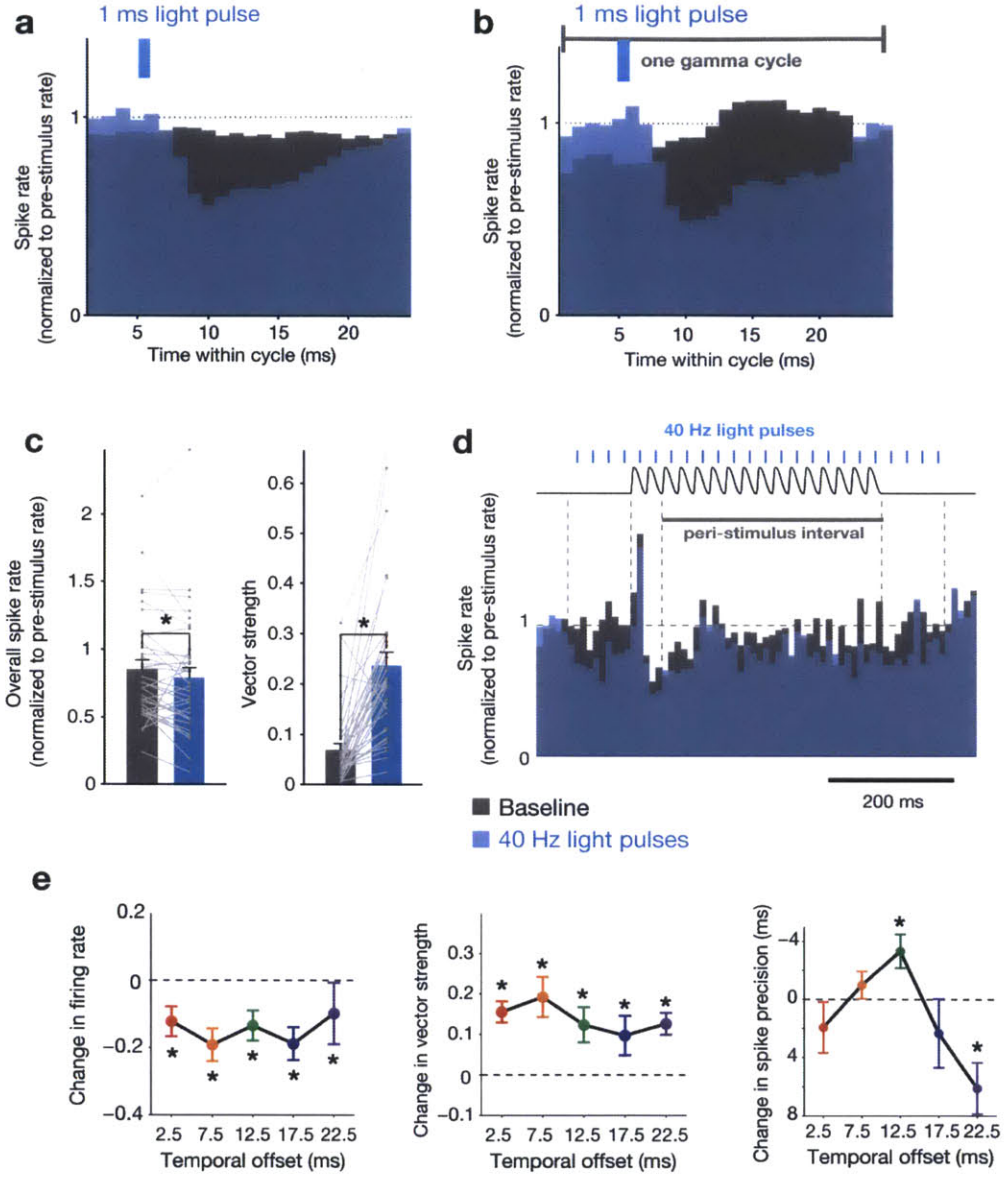

Figure 7 | Interaction between optogenetically entrained gamma and the response to vibrissal deflections. **a**, Mean peri-stimulus (naturalistic stimulation) spike histogram averaged over individual cycles of optogenetically entrained gamma for $N = 35$ regular-spiking single units under two conditions: laser off (black bars) and laser on (blue bars). **b**, Mean spike histogram for periodic stimulation, averaged over individual cycles of gamma, time-locked to laser onset (same units as in a). **c**, Mean changes in firing rate and vector strength for the same units as in a. Firing rate is normalized to pre-stimulus rate. Values for individual cells are overlaid in gray. Error bars indicate mean ± s.e.m. (* = $P < 0.05$, Wilcoxon signed-rank test; $P = 0.0420$ for firing rate, $P = 4.06 \times 10^{-6}$ for vector strength). **d**, Mean peri-stimulus spike histogram for the periodic stimulation condition. Time course of the vibrissae stimulus (black trace, top) and laser pulses (blue lines, top) are superimposed. **e,** Left panel: Change in mean magnitude of

the peak evoked response relative to the interval 100 ms prior to stimulus onset with laser stimulation, normalized to the no laser condition, for the five temporal offsets diagrammed in Fig. 6a ($N = 35$ single units). Center panel: Change in vector strength for the same interval. Right panel: Mean spike precision within the first 25 ms of stimulus onset (multi-unit activity for $N = 15$ electrodes) relative to the no laser condition (* = $P < 0.05$, Wilcoxon signed-rank test with Bonferroni correction; error bars indicate mean ± s.e.m.). The first 25 ms (time between the first and second vibrissae deflections) was the only interval in which firing rates were high enough for spike precision to be reliably calculated.

Supplementary Figure 8 | Additional spike quantification. a, Mean PSTH of multi-unit

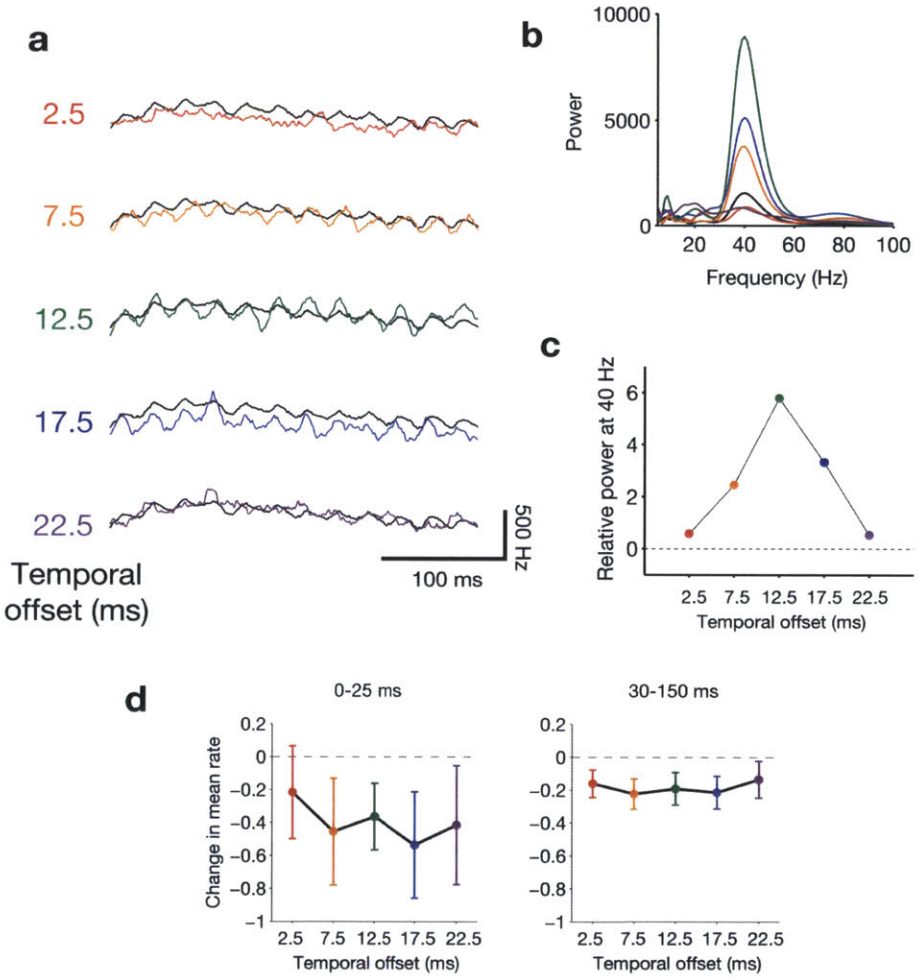

activity from 15 electrodes during the peri-stimulus period (black lines = baseline condition). The condition with a 12.5 ms temporal offset (Figure 6a) is visibly more entrained to the 40 Hz stimulus. **b**, Wavelet transform of the mean PSTH for each temporal offset (colors map to plots in a). **c**, Power at 40 Hz for each of 5 temporal offsets, taken from the plot in panel b. d, Mean rate and peak rate for 0–25 ms and 30–150 ms, relative to baseline ($N = 35$ well-isolated RS cells).

The absence of mean increases in firing rate across our sample—we observed, in fact, significant *decreases* in mean rate—directly indicate that population-level gain cannot explain our behavioral findings. We note that a small subset of RS did show rate increases, indicating that firing enhancement in a select group could have provided sensory benefit. In contrast, enhanced temporal sharpening did predict conditions that led to behavioral benefit for the natural and artificial stimuli. Neither rate nor temporal transformations predicted the decreased detectability observed with the 7.5 ms latency artificial alignment condition (Fig. 6a), underscoring that population-level transformations of the layer II/III neurons we sampled cannot provide a sufficient explanation for the full range of perceptual results.

2.4 Discussion

The current experiments tested the behavioral impact of mechanistically realistic FS-gamma in a local neocortical circuit. Endogenous expression of higher gamma power predicted the detection of less salient sensory stimuli, and driving FS at 40 Hz enhanced the detection of less salient sensory stimuli. This finding was replicated across naturalistic and periodic stimuli. Stimuli that were highly salient were not benefitted by FS-gamma. Detection of highly salient periodic stimuli was neither aided nor penalized by entrained FS-gamma, and endogenous gamma did not predict enhanced or decreased detection of highly salient naturalistic stimuli, though FS-gamma did impair detection of these naturalistic inputs, potentially due to ceiling effects. Studies in rodents (Cohen and Castro-Alamancos, 2010; Miyashita and Feldman, 2013; Sachidhanandam et al., 2013) have shown that transient silencing or lesioning of SI impairs the detection of less salient tactile stimuli. The present findings are consistent with the view that FS-gamma beneficially impacts processing of hard-to-perceive stimuli that require neocortical circuitry, and especially those that may require the allocation of attention to be detected consistently.

2.4.1 FS-gamma generates windows of opportunity for spiking

A simple conceptual framework can explain how emulated FS-gamma enhanced detection of

the naturalistic and periodic stimuli used here (Supplementary Fig. 9). Prior studies have predicted that synchronous activation of FS should create a "window of opportunity" 20-25 ms later, and that inputs arriving from the periphery during this window should show increased gain and/or synchrony. Below, the motivations for these mechanistic predictions are discussed, but considering the basic conceptual framework requires only the a-mechanistic prediction that such a window exists.

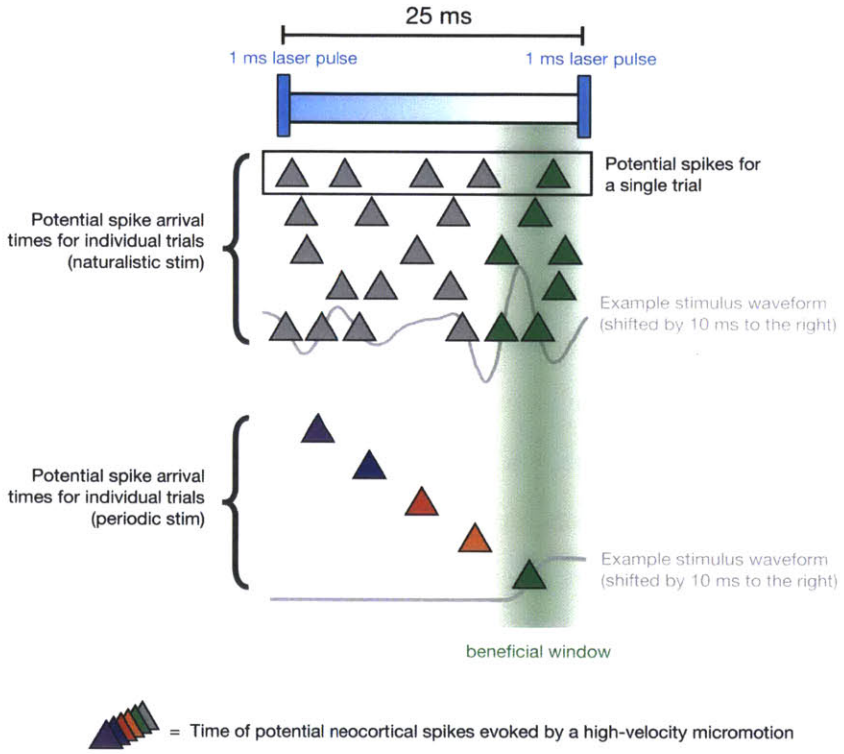

Supplementary Figure 9 | FS-gamma enhances perception when peripheral sensory drive arrives in neocortex during a beneficial window 20-25 milliseconds after FS synchronization. Prior computational and experimental studies have predicted that FS-gamma may create an optimal window 20-25 milliseconds after FS synchronization, benefitting sensory relay when inputs arrive during this period. The conceptual diagram shows the timing of potential spikes generated by high-velocity vibrissae micro-motions (filled triangles) relative to the laser pulses that define the FS-gamma cycle (blue rectangles). In the naturalistic stimulation condition (top), stochastic vibrissal deflections at 180–200 Hz generated many high velocity micro-motions within a 25 ms cycle. As shown in Figure 4, these events had no systematic relationship to the FS-gamma cycle, and afferent spikes generated by these stimuli should arrive during the beneficial window. In the periodic stimulation condition (bottom), only one temporal offset (green, 12.5 ms) was predicted to generate stimulus-driven spikes inside the beneficial window, given the predicted lag of 8–10 milliseconds between peripheral drive and arrival in neocortex.

High-velocity micro-motions are widely regarded as a key driver of sensory responses and behavioral detection in the vibrissa sensory system (Jadhav et al., 2009; Gerdjikov et al., 2010). During the periodic stimulation condition in the present experiments, high-velocity micro-motions repeatedly occurred at specific latencies relative to FS synchrony. Given the well-documented lag of ~8–10 ms for transmission of high-velocity vibrissal motions from periphery to SI (Supplementary Fig. 5), the 12.5 ms latency (green) drove peripheral inputs arriving in neocortex 20–25 ms after FS synchronization, during the predicted optimal window. Under the "artificial" conditions of alignment in the periodic stimulation task, only sensory stimulation arriving at this latency should access this window. Therefore, it is the only latency that would be predicted to drive enhanced relay and, in turn, improved behavioral performance.

In contrast, naturalistic stimuli have a much higher rate of high-velocity micro-motions. In the present experiment, mean vibrissal frequencies fell in the range of 180–200 Hz, corresponding to ~10 such events per 25 ms FS-gamma cycle. As diagrammed in Figure 4d and Supplementary Figure 9, these high-velocity micro-motions occurred at random times relative to the beneficial window created by FS synchronization. However, because micro-motions occur at a high rate, they can consistently drive peripheral input that arrives during the beneficial window and yield enhanced sensory processing in the presence of FS-gamma. The enhanced detection observed when combing FS-gamma with the least salient naturalistic stimuli indicates that there is not an apparent a penalty for also having inputs arrive outside the beneficial window. This finding is in contrast to the more artificial periodic stimulation condition with a 7.5 ms latency, in which all inputs arrived just before the window and resulted in impaired performance. Put another way, the conceptual framework predicts that naturalistic stimuli received the benefit of FS-gamma presence by having high-velocity micro-motions that drove activity in the optimal window, but there was no perceptual penalty for extra high-velocity micro-motions at other latencies.

An open question is why detection of more salient naturalistic stimuli was suppressed by entrained FS-gamma. In apparent contrast, highly salient periodic stimuli were not impaired, and endogenous expression of higher gamma power did not predict impaired detection of highly salient naturalistic stimuli. Baseline detection rates for highly salient periodic and naturalistic

stimuli were around 90% (Fig. 5d Supplementary Fig. 6), so ceiling effects alone cannot explain this discrepancy. Further experiments will be required to directly address this issue.

Our findings and this conceptual framework indicate that FS-gamma will benefit perception if sensory inputs are less salient, and consist of a high rate of peripheral events (as in many vibrissal or auditory natural scenes), or a consistent, tonic presence of input generating high-frequency drive (such as during sustained visual input while fixating for hundreds of milliseconds between saccades). Stochastic more sparse afferent input could also access this window. The present results suggest that only stimuli occurring at the same frequency as FS-gamma, and with the "wrong" latency offset, would be penalized by the emergence of this dynamic. In concept, however, such stimuli should be relatively rare in natural sensation.

2.4.2 Potential mechanisms underlying the perceptual benefits

Generation of FS-gamma has been hypothesized to benefit processing by producing enhanced firing rate in response to sensory stimuli (increased gain), and/or synchronization of local pyramidal activity. Several mechanisms have been proposed for how FS-gamma could lead to firing rate gain, including cellular-level induction of enhanced excitability resulting from a prior deep inhibition (Fellous et al., 2003; Tiesinga et al., 2004; Rudolph et al., 2007). Circuit-level mechanisms have also been proposed, including increased pyramidal excitability through decreased feed-forward recruitment of FS in the period immediately following synchronous FS activation (Knoblich et al., 2010). Synchronization of pyramidal activity—with or without concomitant rate changes—has also been a proposed consequence of FS-gamma expression, and could occur through the enhanced temporal precision enforced by multiple cycles of inhibition. Firing probability will be reduced during the peak of deep inhibition generated by synchronous FS activation, and again by the following cycle of inhibition, leaving the "window of opportunity" for pyramidal activation (Hasenstaub et al., 2005). Limiting the temporal occurrence of pyramidal firing to this window should increase synchrony, and synchronous activation could benefit downstream relay by an increase in the probability of driving spikes in downstream targets. Supplementary Fig. 10 provides a graphical representation of the impact of these predictions and further description of them.

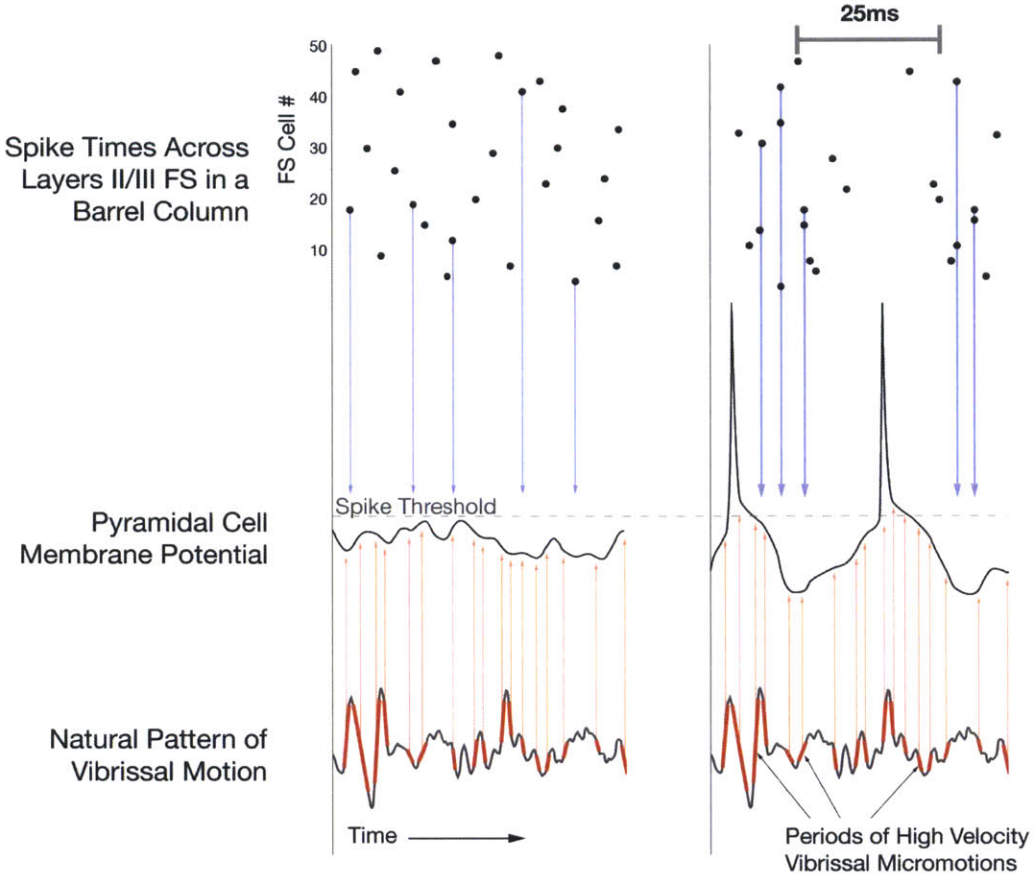

Supplementary Figure 10 | Proposed schema of gamma's effects on evoked spiking in pyramidal cells. In this example, high-velocity deflections arrive while the cortex is in one of two "states," one in which low synchrony among FS leads to generalized inhibition of firing and desynchronized firing in pyramidal cells (left), and one in which high FS synchrony creates a localized gamma oscillation (right). In the former state, disorganized but high rates of FS activity create a persistent state of inhibitory tone, diminishing the probability of spiking. In contrast, FS synchronization creates a greater transient depth of inhibition, but recovery from this hyperpolarization creates a window in which a host of excitatory currents are more readily recruited. This relief from synchronized inhibition concentrates firing of local neurons to this period. Motion patterns (bottom) measured from an *ex vivo* vibrissa; FS spike times and pyramidal cell membrane potential are simulated.

2.4.3 Entrained FS-gamma did not increase net firing rate

In the limited sample of single units from layer II/III recorded in the present study, we did not observe widespread increases in rate in the presence of entrained FS-gamma, and instead observed significant decreases in mean rate (Figure 7c,e). This finding is inconsistent with the view that rate increases across the local population (e.g., layer II/III of the relevant neocortical column) are required to drive the perceptual enhancement we observed. Our limited sample did reveal gain increases in a few RS, posing the possibility that FS-gamma driven rate enhancement in a sparse, maximally informative subset of neurons could be crucial to behavior, even though the net effect was suppression.

In contrast to the absence of population rate effects, we observed enhanced temporal precision in sensory evoked responses that predicted enhanced performance in the naturalistic and periodic stimulus conditions. These findings imply that local synchrony could underlie the perceptual benefit observed. However, neither rate changes nor temporal precision changes predicted the decreased detection probability observed at the 7.5 ms latency. This failure of local neural transformations to predict positive and negative impacts on behavior underscores the limitations of the current data set, and more generally the likelihood that the effect of local FS-gamma on behavior is a network-level phenomenon, and not simply the altered readout of a single local neural-population–level metric.

2.4.4 Interpretations relative to other recent reports

Recent studies examining the neural correlates of sensory performance in rodents have come to a variety of conclusions regarding the relative importance of rate transformations for behavioral success. In a recent study, O'Connor and colleagues concluded that the rate of layer IV firing of barrel neocortex neurons predicted detection of vibrissal contact with a pole, independent of the phase of vibrissal position at the time of the rate increase (O'Connor et al., 2013). Similarly, Sachidhanandam and colleagues found that detection of highly salient vibrissal deflections was predicted by increased firing in layer II/III neurons that was not temporally locked to the stimulus, occurring in a window 50–400 ms after initial passive vibrissa motion. High-frequency (200 Hz) optogenetic activation of FS during this epoch suppressed behavioral

performance (Sachidhanandam et al., 2013). Detection of direct optogenetic stimulation by mice has also suggested that rate, and not temporal pattern or synchronization, are exclusively predictive of performance. Histed and colleagues (Histed and Maunsell, 2014) recently reported that the rate of activation resulting from optogenetic drive of excitatory neurons across neocortical layers in V1 was the sole predictor of detection of these stimuli. They failed to observe a benefit of patterning excitatory drive in the beta and gamma frequency ranges.
In potential contrast to these recent results, our findings do not suggest that mean rate increases across a local population predict perceptual success. However, the present findings are supported by a recent study from Doron and colleagues (Doron et al., 2014). This study found that behavioral detection of single neuron stimulation was not simplistically predicted by rate increases, with evoked firing rates of 30 Hz better detected than rates of 90 Hz. Further, they observed that direct stimulation of FS was more effective in driving detection than direct stimulation of pyramidal neurons. This result agrees with our finding that FS drive at 40 Hz can systematically enhance detection probability. In further agreement with the present results, detection of the change in the orientation of visual input in mice was also enhanced by optogenetic FS drive of area V1, although this study did not measure whether gamma LFP increases were induced by their optogenetic stimuli (Lee et al., 2012).

Reconciliation of these apparently disparate results could arise from experimental differences. The current study probed layer II/III in a passive detection task, in contrast to manipulation of layer IV dynamics in a sensorimotor task by O'Connor et al. (O'Connor et al., 2013) and the stimulation of excitatory neurons across layers in Histed and Maunsell (Histed and Maunsell, 2014). If population-level rate increases in deeper layers predict enhanced detection, while they do not for layer II/III, the present findings may be consistent. While similar in paradigm, Sachidhanandam et al. (Sachidhanandam et al., 2013) probed only supramaximal stimuli in their experiment, and did not attempt to emulate FS-gamma, instead applying FS stimulation at 200 Hz, differences that could readily account for their dissimilar findings to our own.

Although these potential explanations could account for the different conclusions reached across studies, the similarity in the findings of Doron et al. (2014) and the present results suggest

that, at least under the experimental conditions in these studies, mean increases in rate across larger populations of neurons are not a requisite predictor of detection success, and that the dynamics driven by FS recruitment are an effective mechanism for enhancing relay.

2.4.5 A role for FS-gamma can be distinct from a role in binding

Correlative studies have shown that local neocortical gamma oscillations are often present at an appropriate time to enhance perceptual processing. The allocation of attention to a specific position in space leads to increased gamma expression during stimulus presentation in the associated region of area V4 (Fries et al., 2001; Womelsdorf et al., 2006). These findings, and theoretical and empirical studies implicating fine-timescale synchronization across dispersed representations, coincide with a hypothesis that gamma-mediated temporal organization underlies the binding of images into a coherent percept (Engel and Singer, 2001). This hypothesis, as well as related predictions linking gamma timescale activity to the encoding of stimulus features, has been criticized as inconsistent with the fluctuations in the precise frequency of gamma observed in area V1 with changes in stimulus salience (Ray and Maunsell, 2010).

These recent results showing gamma fluctuations do not contradict a different view about the potential benefit of gamma, supported by the present findings, that increased FS-gamma in a local neocortical circuit could enhance the relay of stimuli to benefit detection (Fellous et al., 2003; Tiesinga et al., 2004; Knoblich et al., 2010). Under this view, transformations in pyramidal synchrony or increases in local firing rate associated with FS-gamma can increase the probability of downstream relay. In agreement with this prediction, higher levels of gamma over human occipital lobe predict enhanced detection of change (Hoogenboom et al., 2010), and similar increases in human somatosensory neocortex predict heightened detection of innocuous input (Meador et al., 2002) and the salience of painful stimuli (Gross et al., 2007). Further, local gamma expression in V1 predicts increased firing in synaptic recipients in V2 with aligned receptive fields (Jia et al., 2013). The present study supports this predicted role for FS-gamma, while the psychophysical demands of the present experiments do not directly test a role in binding or encoding by this dynamic.

2.4.6 Implications for further study

The current study was designed to test the behavioral impact of emulating a natural dynamic, the repeated cyclical synchronization of FS in the gamma frequency range. While the present experiments are conclusive in showing enhancement of sensory detection by entrained FS-gamma, and the specific requirement for accessing a window of opportunity to achieve this enhancement, they are not conclusive as to the specific neural mechanisms mediating this benefit. Systematic examination of this question will require the local measurement of synchronization and gain of well-isolated single units localized to additional specific cell layers in SI. Simultaneous interrogation of downstream activity in cortical and subcortical targets—including correlative recording and causal control—will be necessary to determine what impact these transformations have on relay.

Another key question and topic for further study is whether the periodicity and frequency range of FS-gamma matters for behavioral benefit, or whether individual events of FS synchronization, potentially occurring stochastically and with greater or lower frequency, can drive the same effects. Our conceptual model does not require a subsequent cycle of synchronized FS activity after the optimal temporal window for relay of signals to benefit (Supplementary Figure 9). Further, as discussed above, recent work by Histed and colleagues (Histed and Maunsell, 2014) suggests that rates of activity in the gamma range are not preferentially relayed, though work by Doron and colleages (Doron et al., 2014) suggests that such patterning may be beneficial. Our own computational modeling (Knoblich et al., 2010) has predicted that the timing of the next event of FS synchrony—e.g., 25 ms after the prior event—may be optimal in enhancing the efficiency of neocortical relay, by preventing the expression of spikes that would not have contributed to relay occurring outside the optimal temporal window. From this perspective, the rate of FS-gamma at 30–80 Hz may be optimal because it is the fastest rate at which cycles of a beneficial individual temporal transformation can be realized, not because the frequency per se carries information about the stimulus.

2.4.7 Conclusion

The current results directly address a long-standing debate in neocortical function, FS-gamma can impact sensory processing (either positively or negatively). While theoretical studies have long suggested a role for this dynamic in processes such as attention, the evidence in support of this view has so far been correlative. Arguments against a possible role for FS-gamma have included mechanistic propositions, such as the view that the biophysical properties of neocortical neurons are not suited to benefit from this dynamic (Shadlen and Movshon, 1999), and information-processing perspectives, such as the view that neocortical transformations on such short timescales would not benefit read-out of perceptual signals (Shadlen and Movshon, 1999; Ray and Maunsell, 2010). Here, we provide causal evidence that FS-gamma can improve psychophysical performance, and that such enhancement requires transformations with temporal precision on the ~5 ms timescale. The degree to which these findings generalize to other tasks and brain regions requires further study, as does detailed examination of the local and downstream transformations in neural activity driven by FS-gamma. Nevertheless, these results provide a unique existence proof that this dynamic can improve the relay of sensory signals through the neocortex.

2.5 Experimental procedures

2.5.1 Animals

All mice were male parvalbumin-Cre (PV-Cre) heterozygotes, derived from PV-Cre BAC transgenics back-crossed into a C57BL/6J line (gift from S. Arber, now available as Jackson Laboratory strain B6;129P2-*Pvalb*$^{tm1(cre)Arbr}$). Mice were 8–14 weeks old at the time of initial surgery (mean age = 11.5 ± 2.6 weeks). Animals were individually housed and maintained on a 12-hour reversed light/dark cycle (lights out at 9:00 AM or 12:00 PM). All experimental procedures and animal care protocols were approved by the Massachusetts Institute of Technology and Brown University Institutional Animal Care and Use Committees and were in accordance with NIH guidelines.

2.5.2 AAV vectors

ChR2 (Boyden et al., 2005) fused to the fluorescent protein mCherry was cloned in the antisense direction into pAAV-MCS (Stratagene) to create AAV DIO *ChR2-mCherry* (Supplementary Fig. 2). *ChR2-mCherry* was flanked by a pair of canonical loxP sites (Tsien et al., 1996) and a pair of mutated lox2272 sites. A woodchuck hepatitis B virus post-transcriptional element was placed in sense direction 5′ of the poly(A). Adeno-associated viral particles of serotype 2 were produced by the Vector Core Facility at UNC Chapel Hill.

2.5.3 Fiber-optic–electrode implants

Implants were constructed around a custom plastic base, designed in SolidWorks (Dassault Systems, Waltham, MA) and printed via stereolithography in Accura 55 plastic (American Precision Prototyping, Tulsa, OK). The plastic base held 8 electrodes around a central fiber optic cable (Doric Lenses, Quebec, Canada). Electrodes were made from 12.5 μm polyimide-coated nichrome wire (Kanthal, Halstahammar, Sweden), twisted and heated to form tetrodes or stereotrodes (Nguyen et al., 2009). In three animals, the electrodes were stationary. In five animals, electrodes were fixed to laser-cut plastic springs (Pololu, Las Vegas, NV) individually driven by miniature screws (Voigts et al., 2013). Electrodes were attached to a custom electrode interface board (Sunstone, Mulino, OR) with gold pins (Neuralynx, Bozeman, MT). Individual electrodes were gold plated to an impedance of 200–400 kΩ. The fiber optic cable had an inner diameter of 200 μm and a cladding diameter of 260 μm, and was terminated with a 1.25 mm metal ferrule (Supplementary Fig. 7).

2.5.4 Surgical procedure

Naïve mice were anesthetized with isofluorane gas anesthesia (0.75–1.25% in 1 L/min oxygen) and secured in a stereotaxic apparatus. The scalp was shaved, wiped with hair removal cream, and cleaned with iodine solution and alcohol. Following IP injection of Buprenex (0.1 mg /kg, as an analgesic) and dexamethasone (4 mg/kg, to prevent tissue swelling) and local injection of lidocaine, the skull was exposed with an incision along the midline. After the skull was cleaned, two small stainless steel watch screws were implanted

in the skull plates anterior to bregma, one of which served as ground. Next, a ~1.5 mm–diameter craniotomy was drilled over barrel cortex of the left hemisphere (1.5 mm posterior to bregma and 3.5 mm lateral to the midline). We did not target a specific barrel, but post-hoc analysis revealed that our implants were located near the rear of the barrel map, centered around the C2 column. Virus was delivered through a glass micropipette attached to a Quintessential Stereotaxic Injector (Stoelting). The glass micropipette was lowered through the dura to a depth of 450 μm below the cortical surface. A bolus of 1 μl of virus (AAV DIO ChR2-mCherry; 2×10^{12} viral molecules per ml) was injected into the cortex at 0.05 μl/min. After the injection, the pipette was held in place for 10 min at the injection depth and 10 min at a depth of 200 μm before being fully retracted from the brain. The fiber-optic–electrode implant (Supplementary Fig. 7a) was aligned with the craniotomy at an angle of 15° from vertical and lowered to the surface of the cortex, centered approximately 200 μm anterior to the injection site. If notable bleeding occurred during implantation, the surgery was aborted. Once the implant was stable, a small ring of dental acrylic was placed around its base. A drop of surgical lubricant (Surgilube) prevented dental acrylic from contacting the cortical surface. A custom head post made from durable plastic (APProto) was then affixed to the skull with adhesive luting cement (C&B Metabond). Once the cement was dry, the scalp incision was closed with VetBond (3M), and mice were removed from isoflurane. Mice were given 3–10 days to recover prior to the start of water restriction.

2.5.5 Stimulus delivery and behavioral control

Vibrissae were stimulated by computer-controlled movements of piezoelectric wafers (Noliac, Kvistgård, Denmark). Stimulations consisted of high-speed deflections in the dorsal direction with a raised cosine velocity profile (6 ms rising phase) or naturalistic velocity profiles based on actual vibrissa movements as described above and in the text. Vibrissae were held with a silk suture loop fed through a 4 mm, 21-gauge stainless steel cannula, which was attached to the piezoelectric wafer via a glass capillary tube (0.8 mm outer diameter). Several vibrissae were secured, centered around the C2 vibrissa, and gripped approximately 5 mm from the mystacial pad. Mice were photographed from above at the start of each session

to confirm which vibrissae were stimulated. For maximum-amplitude deflections, vibrissae moved approximately 1 mm at the point of contact. All stimulation amplitudes were calibrated using a linear CCD array.

Water delivery was based on gravitational flow controlled by a solenoid valve (Lee Co., Westbrook, CT) connected with Tygon tubing. Mice received distilled water through a plastic tube mounted on a piezoelectric wafer. This "lick tube" was positioned near the animal's mouth using a micromanipulator. The water volume was controlled by the duration of valve opening (30–60 ms), calibrated to give ~8 µL per opening. Individual licks were detected by amplifying and thresholding the output of the piezoelectric wafer using a custom circuit board or Arduino microcontroller.

The light stimulus was delivered through a jacketed fiber-optic cable 200 µm in diameter and 2.5 m long with a numerical aperture of 0.22 (Doric Lenses, Quebec, Canada) connected to a 473 nm laser (Opto Engine, Salt Lake City, UT). Laser light passed through an adjustable neutral-density filter and a collimator (Thorlabs PAF-X-15-PC-A) before entering the fiber. The 2.5 m fiber was connected to the animal's head via two mating metal ferrules sheathed in a zirconia sleeve. Light loss for this connection was measured for each implant prior to surgery, and was around 50%. The amplitude of the light stimulus was calibrated daily with an optical power meter (Thorlabs PM100D with S120C sensor). Light power at the surface of the cortex was estimated to be ~1 mW, or 30 mW/mm^2 for a 200 µm fiber, a lower level than that used in a previous study from our lab (Cardin et al., 2009). We wanted our optogenetic manipulation to modulate FS synchrony, rather than flood the cortex with inhibition. Because the fiber was not implanted into the brain, this irradiance was roughly the same throughout the craniotomy (1.5 mm diameter). To calculate the depth of light penetration, we used a model based on direct measurements in mammalian brain tissue (Yizhar et al., 2011) (http://www.stanford.edu/group/dlab/cgi-bin/graph/chart.php). Light power was strongest in layer II/III, falling off to 1.75 mW/mm^2 at a depth of 500 microns (~layer IV).

All behavioral events, piezoelectric control, reward delivery, laser stimulation, and lick measurements were monitored and controlled by custom software written in LabVIEW or Matlab and interfaced with a PCI DIO board (National Instruments, Austin, TX).

2.5.6 Trial structure

On each trial with 40 Hz stimuli, vibrissae were stimulated for 400 ms at a single amplitude, varying between 0 ("blank" trials) and 1 mm ("maximal" trials). Amplitude did not vary within a given trial. For naturalistic stimuli, vibrissae were stimulated with one of seventeen 400 ms velocity profiles (main text and Supplementary Fig. 3). On laser-stimulation trials, 1 ms light pulses were delivered at 40 Hz for 600 ms. Vibrissae deflections began after the fourth light pulse. The precise timing of the vibrissae deflections relative to the light pulses varied across five temporal offsets in the periodic stimulus trials, but remained consistent within individual trials. For all sessions, inter-trial intervals were uniformly distributed between 4 and 6 s.

If mice licked the reward spout at any point up to 500 ms after the onset of the vibrissae stimulus, a drop of water was delivered. After a slight delay, any remaining water was removed via vacuum suction. This prevented mice from receiving reward that was not immediately preceded by a vibrissae stimulus. There was no punishment or time-out period for false alarms.

When the vibrissae were removed from the lasso, but the stimulator remained near the face, performance immediately dropped to chance levels. This indicates that mice were not licking to some other aspect of the stimulus, such as auditory cues or vibrations transmitted through the table.

2.5.7 Behavioral training

Training began after at least 3 days of post-operative recovery and at least 7 days of water restriction (1 ml /d). All training sessions took place near the beginning of the animals' dark cycle. At the beginning of each session, mice were secured to the head post apparatus with their bodies placed inside a Falcon tube. For the first five training sessions, mice were head

fixed for periods of 5 to 15 minutes and given water ad libitum, to allow them to accommodate to the experimental apparatus. Vibrissae stimulation was added during the next week of training. At first, post-stimulus reward was delivered regardless of the animal's response, to establish a contingency between vibrissae deflections and drops of water. During this week, training time was gradually increased to 45 minutes. As mice became accustomed to performing the task (at approximately day 10 of training), water was only delivered if the reward spout was licked within 500 ms of stimulus onset. Initially, all trials included maximum-amplitude stimuli, with no blank trials. On successive days, blank stimuli (up to 10%) were randomly interleaved and the probability of non-maximal stimuli was increased.

After reaching criterion on the training task, laser light stimulation was added on a subset of trials. Laser stimulation was given ~100 ms preceding tactile stimulation for a duration of 600 ms (100 ms post stimulation) at ~1 mW power on the cortical surface. 50% of trials included delivery of stimuli at a single "threshold" amplitude ranging from 20–40% of the maximum stimulus amplitude, with the same duration to peak amplitude (6 ms). 30% of trials were blank, with the remaining 20% of trials at maximum amplitude to keep the animals engaged in the task. On half of all trials, laser stimulation was presented at one of 5 temporal offsets (Fig. 6a). All trial types were randomly interleaved.

Mice that did not consume 1 ml of water during the training sessions were supplemented with water in their home cage several hours after the experiment finished.

2.5.8 Electrophysiology

Spikes and local field potentials were recorded using twisted-wire nichrome stereotrodes or tetrodes integrated into a custom implant (Voigts et al., 2013; Siegle, 2012) (Supplementary Fig. 7a). Recording electrodes were either implanted 300–400 μm from the surface of the cortex during surgery, or lowered to the same depth over the course of several days. Once they reached their final depth, electrodes were not moved for the remainder of the experiment. Throughout each recording session, broadband signals referenced to ground were digitized at 40 kHz (Recorder/64, Plexon, Dallas, TX). Electrophysiology data was synchronized with behavioral data via TTL pulses at the start of each trial.

We measured the delay between the onset of the first deflection and the appearance of the evoked response on our recording electrodes. We found the delay to be 8–10 ms, similar to previous reports (Carvell and Simons, 1988; Armstrong-James et al., 1992; Moore and Nelson, 1998; Brecht et al., 2003) (Supplementary Fig. 5).

2.5.9 Histology

At the end of training, electrodes sites were lesioned with 15 μA of current for 10 s. Mice were transcardially perfused with 100 mM PBS followed by 4% formaldehyde in PBS. Brains were post-fixed for at least 18 h at 4 °C. 60 μm sections were mounted on glass slides with Vectashield (Vector Laboratories), coverslipped, and imaged with an upright fluorescent microscope (Supplementary Fig. 2b). Viral expression was confirmed in all animals ($N = 8$). Expression was generally limited to the upper layers of barrel cortex (II–IV). In one animal, expression was also observed in the lower layers (V and VI). Ventral/medial and anterior/ posterior spread of the virus ranged between 0.3 and 1.0 mm, encompassing 1–3 barrel columns.

2.5.10 Behavioral analysis

Data analysis was performed in Matlab (Mathworks, Natick, MA). Raw data from LabVIEW or Matlab was converted to event times for behavioral analysis. Trials were first selected based on d' according to the following procedure: (1) hit rate and false alarm rates were calculated for blocks of 50 trials, slid in 1-trial intervals. A negative offset was added to the hit rate and an equal-but-opposite offset was added to false alarms, to prevent d' saturation. An offset of 0.04 was used for all analyses. Hit rate was capped at a minimum value equivalent to the offset and false alarm rate was capped at a maximum value equivalent to 1 minus the offset. (2) d' was calculated as Z(hit rate) – Z(false alarm rate), where Z stands for the inverse of the Gaussian distribution function (Supplementary Fig. 5). (3) For any blocks with d' above a value of 1.25, the middle trial was included in the analysis. If fewer than 25% of trials remained after this procedure was carried out, the subject was excluded

from analysis (1/9 mice for periodic stimuli, 1/5 mice for naturalistic stimuli). This criteria was established prior to the start of experimentation.

After d' trial rejection, trials with pre-stimulus licks within 1.5 s of trial onset were also eliminated. Trials were selected from the remaining subset based on the parameters of the vibrissae and laser stimuli. Hit rate and d' were calculated for each animal for each set of parameters.

Because all mice were of the same genotype, and all experienced the same experimental conditions (except the four mice that did not receive the naturalistic stimulation), data collection and analysis were not performed blind to the conditions of the experiments.

2.5.11 Statistical analysis

P-values were obtained using nonparametric statistical tests, unless otherwise stated. Where there were more than two treatments per animal (such as the 5 stimulus-to-laser temporal offsets plus baseline conditions), the Friedman test was used (nonparametric version of a repeated-measures ANOVA). For post-hoc comparisons and any tests where we collapsed across temporal offsets, we used the one-sided Wilcoxon signed-rank test. Whenever multiple tests were performed on the same dataset, we adjusted our *P*-values using the Bonferroni correction procedure. In total, we trained enough animals (≥ 8) such that we could observe $P < 0.05$ significance levels via a Wilcoxon signed-rank test after Bonferroni correction for our five phase offset conditions.

To analyze the significance of the results from Figure 6, we also employed a permutation procedure to determine the likelihood that at least 7/8 animals would improve relative to baseline at any one temporal offset *and* at least 5/8 animals would have their best performance at that temporal offset, under the null hypothesis that our manipulation had no effect. For each round of the simulation, 5 temporal offsets + 1 baseline condition were randomly ordered for each of 8 mice. After 10,000 repetitions, we counted the total number of rounds in which the requirements were satisfied, and found the probability was always less than 0.01.

In one case (Fig. 5b), we used a *t*-test with Bonferroni correction to test significance. We used the Kolmogorov–Smirnov test to test for normality before applying this test ($P < 0.0005$).

2.5.12 Offline electrophysiological analysis

Continuous electrophysiological data from the Plexon system was either low-pass filtered (3rd-order Butterworth) and downsampled to 2 kHz or filtered between 600 and 6000 Hz (acausal FIR filter) and thresholded to extract spikes. Spikes were sorted offline using custom software written in Matlab (Simple Clust v0.5, available online at https://github.com/moorelab/simpleclust). Well-isolated single units were readily matched across days by eye; afterward, cell identity was confirmed using a waveform similarity metric (Tolias et al., 2007). In total, we isolated 43 units from 3 mice, which persisted an average of 3.29 ± 3.2 d each (range: 1 to 13 d). The distribution of units was as follows: 26 units from Mouse 1, 1 from Mouse 2, and 16 from Mouse 3. Cluster quality was quantified using L-ratio and isolation distance of the four peak heights of bandpass-filtered waveforms for each tetrode (Schmitzer-Torbert et al., 2005). If putative single units had an L-ratio greater than 1 or an isolation distance less than 5, they were rejected. On average, our units had an L-ratio of 0.19 ± 0.20 and an isolation distance of 19.68 ± 15.42, similar to previously reported values for peak heights (Davidson et al., 2009). We used the peak-to-trough ratio and peak-to-trough separation on 600 to 6000 Hz filtered waveforms to classify cells as fast spiking or regular spiking. We found two distinct clusters of units, one with 35 regular-spiking cells and one with 8 fast-spiking cells (Supplementary Fig. 7d). A subset of well-isolated FS units (2 of 8) were activated by light.

Peristimulus firing rate was calculated as the mean firing rate between 30 ms and 400 ms after stimulus onset. Vector strength was calculated for the same period using the procedure described in (Goldberg and Brown, 1969). Spike precision was calculated as the inter-quartile range of spike times within the first 25 ms of stimulus onset (Kumbhani et al., 2007). Spike precision was calculated only on multi-unit activity, as single-unit spike rates were not high enough to observe consistent peri-stimulus time histograms after the first vibrissae

deflection. This is the same way that spike precision was measured in a previous paper from our lab (Cardin et al., 2009).

2.6 Acknowledgments

We thank R. Clary, J. Feather, H. Farrow, R. Lichtin, S. Bechek, J. Klee, N. Padilla, and C. Burley for help running experiments. We are grateful to J. Cardin, U. Knoblich, M. Halassa, J. Ritt, J. Voigts, C. Deister, B. Higashikuibo, D. Meletis, and M. Carlén for technical assistance. We thank members of the Moore lab, M. Wilson, M. Andermann, R. Haslinger, N. Kopell, C. Börgers, R. Sekuler, and D. Sheinberg for their comments on the manuscript. This study was supported by a grant from the NIH to C.I.M., an NRSA Fellowship to D.L.P., and NDSEG and NRSA Fellowships to J.H.S.

2.7 References

Adesnik H, Scanziani M (2010) Lateral competition for cortical space by layer-specific horizontal circuits. *Nature* **464**: 1155-1160.

Armstrong-James M, Fox K, Das-Gupta A (1992) Flow of excitation within rat barrel cortex on striking a single vibrissa. *J Neurophysiol* **68**: 1345-1358.

Baranauskas G, Maggiolini E, Vato A, Angotzi G, Bonfanti A, Zambra G, Spinelli A, Fadiga L (2012) Origins of 1/f2 scaling in the power spectrum of intracortical local field potential. *J Neurophysiol* **107**: 984-994.

Bosman CA, Schoffelen JM, Brunet N, Oostenveld R, Bastos AM, Womelsdorf T, Rubehn B, Stieglitz T, De Weerd P, Fries P (2012) Attentional stimulus selection through selective synchronization between monkey visual areas. *Neuron* **75**: 875-888.

Boyden ES, Zhang F, Bamberg E, Nagel G, Deisseroth K (2005) Millisecond-timescale, genetically targeted optical control of neural activity. *Nat Neurosci* **8**: 1263-1268.

Börgers C, Epstein S, Kopell NJ (2005) Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc Natl Acad Sci U S A* **102**: 7002-7007.

Börgers C, Epstein S, Kopell NJ (2008) Gamma oscillations mediate stimulus competition and attentional selection in a cortical network model. *Proc Natl Acad Sci U S A* **105**: 18023-18028.

Börgers C, Kopell N (2003) Synchronization in networks of excitatory and inhibitory neurons

with sparse, random connectivity. *Neural Comput* **15**: 509-538.

Brecht M, Roth A, Sakmann B (2003) Dynamic receptive fields of reconstructed pyramidal cells in layers 3 and 2 of rat somatosensory barrel cortex. *J Physiol* **553**: 243-265.

Buzsáki G, Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci.*

Cardin JA, Carlén M, Meletis K, Knoblich U, Zhang F, Deisseroth K, Tsai LH, Moore CI (2009) Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* **459**: 663-667.

Cardin JA, Carlén M, Meletis K, Knoblich U, Zhang F, Deisseroth K, Tsai LH, Moore CI (2010) Targeted optogenetic stimulation and recording of neurons in vivo using cell-type-specific expression of Channelrhodopsin-2. *Nat Protoc* **5**: 247-254.

Carlén M, Meletis K, Siegle JH, Cardin JA, Futai K, Vierling-Claassen D, Rühlmann C, Jones SR, Deisseroth K, Sheng M, Moore CI, Tsai LH (2011) A critical role for NMDA receptors in parvalbumin interneurons for gamma rhythm induction and behavior. *Mol Psychiatry.*

Carvell GE, Simons DJ (1988) Membrane potential changes in rat SmI cortical neurons evoked by controlled stimulation of mystacial vibrissae. *Brain Res* **448**: 186-191.

Cohen JD, Castro-Alamancos MA (2010) Detection of low salience whisker stimuli requires synergy of tectal and thalamic sensory relays. *J Neurosci* **30**: 2245-2256.

Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* **63**: 497-507.

Doron G, von Heimendahl M, Schlattmann P, Houweling A, Brecht M (2014) Spiking Irregularity and Frequency Modulate the Behavioral Report of Single-Neuron Stimulation. *Neuron* **81**: 653-663.

Engel AK, Singer W (2001) Temporal binding and the neural correlates of sensory awareness. *Trends Cogn Sci* **5**: 16-25.

Fellous -M, Rudolph M, Destexhe A, Sejnowski J (2003) Synaptic background noise controls the input/output characteristics of single cells in an in vitro model of in vivo activity. *Neuroscience* **122**: 811-829.

Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**: 1560-1563.

Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* **28**: 4823-4835.

Gerdjikov TV, Bergner CG, Stüttgen MC, Waiblinger C, Schwarz C (2010) Discrimination of vibrotactile stimuli in the rat whisker system: behavior and neurometrics. *Neuron* **65**: 530-540.

Goldberg JM, Brown PB (1969) Response of binaural neurons of dog superior olivary complex to dichotic tonal stimuli: some physiological mechanisms of sound localization. *J Neurophysiol* **32**: 613-636.

Gross J, Schnitzler A, Timmermann L, Ploner M (2007) Gamma oscillations in human primary somatosensory cortex reflect pain perception. *PLoS Biol* **5**:e133.

Hamada Y, Miyashita E, Tanaka H (1999) Gamma-band oscillations in the "barrel cortex" precede rat's exploratory whisking. *Neuroscience* **88**: 667-671.

Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* **47**: 423-435.

Histed MH, Maunsell JH (2014) Cortical neural populations can guide behavior by integrating inputs linearly, independent of synchrony. *Proc Natl Acad Sci U S A* **111**: E178-E187.

Hoogenboom N, Schoffelen JM, Oostenveld R, Fries P (2010) Visually induced gamma-band activity predicts speed of change detection in humans. *Neuroimage* **51**: 1162-1167.

Jadhav SP, Wolfe J, Feldman DE (2009) Sparse temporal coding of elementary tactile features during active whisker sensation. *Nat Neurosci.*

Jia X, Tanabe S, Kohn A (2013) Gamma and the coordination of spiking activity in early visual cortex. *Neuron* **77**: 762-774.

Jones MS, Barth DS (1997) Sensory-evoked high-frequency (gamma-band) oscillating potentials in somatosensory cortex of the unanesthetized rat. *Brain Res* **768**: 167-176.

Knoblich U, Siegle JH, Pritchett DL, Moore CI (2010) What do we gain from gamma? Local dynamic gain modulation drives enhanced efficacy and efficiency of signal transmission. *Front Hum Neurosci* **4**: 1-12.

Kumbhani RD, Nolt MJ, Palmer LA (2007) Precision, reliability, and information-theoretic analysis of visual thalamocortical neurons. *J Neurophysiol* **98**: 2647-2663.

Lee SH, Kwan AC, Zhang S, Phoumthipphavong V, Flannery JG, Masmanidis SC, Taniguchi H, Huang ZJ, Zhang F, Boyden ES, Deisseroth K, Dan Y (2012) Activation of specific interneurons improves V1 feature selectivity and visual perception. *Nature* **488**: 379-383.

Meador KJ, Ray PG, Echauz JR, Loring DW, Vachtsevanos GJ (2002) Gamma coherence and conscious perception. *Neurology* **59**: 847-854.

Miyashita T, Feldman DE (2013) Behavioral detection of passive whisker stimuli requires somatosensory cortex. *Cereb Cortex* **23**: 1655-1662.

Moore CI, Carlen M, Knoblich U, Cardin JA (2010) Neocortical interneurons: from diversity, strength. *Cell* **142**: 189-193.

Moore CI, Nelson SB (1998) Spatio-temporal subthreshold receptive fields in the vibrissa

representation of rat primary somatosensory cortex. *J Neurophysiol* **80**: 2882-2892.

Nguyen DP, Layton SP, Hale G, Gomperts SN, Davidson TJ, Kloosterman F, Wilson MA (2009) Micro-drive array for chronic in vivo recording: tetrode assembly. *J Vis Exp.*

O'Connor DH, Hires SA, Guo ZV, Li N, Yu J, Sun QQ, Huber D, Svoboda K (2013) Neural coding during active somatosensation revealed using illusory touch. *Nat Neurosci.*

Penttonen M, Kamondi A, Acsady L, Buzsaki G (1998) Gamma frequency oscillation in the hippocampus of the rat: intracellular analysis in vivo. *Eur J Neurosci* **10**: 718-728.

Pinto DJ, Brumberg JC, Simons DJ (2000) Circuit dynamics and coding strategies in rodent somatosensory cortex. *J Neurophysiol* **83**: 1158-1166.

Ray S, Maunsell JH (2010) Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron* **67**: 885-896.

Ray S, Ni AM, Maunsell JH (2013) Strength of gamma rhythm depends on normalization. *PLoS Biol* **11**:e1001477.

Ritt JT, Andermann ML, Moore CI (2008) Embodied information processing: vibrissa mechanics and texture features shape micromotions in actively sensing rats. *Neuron* **57**: 599-613.

Rudolph M, Pospischil M, Timofeev I, Destexhe A (2007) Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *J Neurosci* **27**: 5280-5290.

Sachidhanandam S, Sreenivasan V, Kyriakatos A, Kremer Y, Petersen CC (2013) Membrane potential correlates of sensory perception in mouse barrel cortex. *Nat Neurosci* **16**: 1671-1677.

Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish AD (2005) Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* **131**: 1-11.

Shadlen MN, Movshon JA (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**: 67-77, 111-25.

Shao YR, Isett BR, Miyashita T, Chung J, Pourzia O, Gasperini RJ, Feldman DE (2013) Plasticity of recurrent l2/3 inhibition and gamma oscillations by whisker experience. *Neuron* **80**: 210-222.

Shaw FZ, Chew JH (2003) Dynamic changes of gamma activities of somatic cortical evoked potentials during wake-sleep states in rats. *Brain Res* **983**: 152-161.

Siegle JH (2012) Combining optical stimulation with extracellular electrophysiology in behaving mice. *In: Neuronal Network Analysis* (Halassa MM, Fellin T, eds). New York: Humana Press.

Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, Buzsáki G (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. *Neuron* **60**: 683-697.

Sohal VS, Zhang F, Yizhar O, Deisseroth K (2009) Parvalbumin neurons and gamma rhythms enhance cortical circuit performance. *Nature* **459**: 698-702.

Tiesinga PH, Fellous JM, Salinas E, José JV, Sejnowski TJ (2004) Inhibitory synchrony as a mechanism for attentional gain modulation. *J Physiol Paris* **98**: 296-314.

Tolias AS, Ecker AS, Siapas AG, Hoenselaar A, Keliris GA, Logothetis NK (2007) Recording chronically from the same neurons in awake, behaving primates. *J Neurophysiol* **98**: 3780-3790.

Traub RD, Whittington MA, Colling SB, Buzsáki G, Jefferys JG (1996) Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *J Physiol* **493 (Pt 2)**: 471-484.

Voigts J, Siegle JH, Pritchett DL, Moore CI (2013) The flexDrive: An ultra-light implant for optical control and highly parallel chronic recording of neuronal ensembles in freely moving mice. *Front Sys Neurosci* **7**.

Wang XJ, Buzsáki G (1996) Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *J Neurosci* **16**: 6402-6413.

Whittington MA, Traub RD, Kopell N, Ermentrout B, Buhl EH (2000) Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *Int J Psychophysiol* **38**: 315-336.

Wolfe J, Hill DN, Pahlavan S, Drew PJ, Kleinfeld D, Feldman DE (2008) Texture coding in the rat whisker system: slip-stick versus differential resonance. *PLoS Biol* **6**:e215.

Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**: 733-736.

Yizhar O, Fenno LE, Davidson TJ, Mogri M, Deisseroth K (2011) Optogenetics in neural systems. *Neuron* **71**: 9-34.

Chapter 3: Enhancement of encoding and retrieval functions through theta phase–specific manipulation of hippocampus[2]

3.1 Abstract

Assessing the behavioral relevance of the hippocampal theta rhythm has proven difficult, due to a shortage of experiments that selectively manipulate phase-specific information processing. Using a closed-loop stimulation protocol, we triggered inhibition of dorsal CA1 at specific phases of the endogenous theta rhythm in freely behaving mice. This intervention enhanced performance on a spatial navigation task that requires the encoding and retrieval of information related to reward location on every trial. In agreement with prior models of hippocampal function, the behavioral effects depended on both the phase of theta and the task segment at which we stimulated. Stimulation in the encoding segment enhanced performance when inhibition was triggered by the peak of theta. Conversely, stimulation in the retrieval segment enhanced performance when inhibition was triggered by the trough of theta. These results suggest that processes related to the encoding and retrieval of task-relevant information are preferentially active at distinct phases of theta.

[2]The findings in this chapter were in press at the time of thesis submission. (Siegle and Wilson, *eLife*, 2014)

3.2 Introduction

Theta oscillations (4-12 Hz) are one of the most prominent rhythms in the mammalian brain (Vanderwolf, 1969; Buzsáki et al., 1983; Colgin, 2013). Theta is a distributed oscillation, which is broadly coherent between the left and right hippocampi, the entorhinal cortex, the medial septum, and various other cortical and subcortical recipients of hippocampal projections (Buzsáki, 2002). Neural activity is highly structured within each cycle of theta, with the firing rates of genetically defined cell types peaking at different phases (Klausberger et al., 2003; Klausberger et al., 2004; Klausberger et al., 2005). Spatially selective principal cells in the hippocampus fire at progressively earlier phases of theta as animals traverse their individual place fields, a phenomenon known as phase precession (O'Keefe and Recce, 1993; Schmidt and Lipson, 2009). Thus, the information content of hippocampal outputs changes throughout each cycle (Mehta et al., 2000; Mehta et al., 2002).

The organization of activity relative to theta appears to be important for behavior, since the degree to which other regions synchronize to the hippocampal theta rhythm is correlated with spatial decision-making performance (Jones and Wilson, 2005; Sigurdsson et al., 2010). But the specific role of theta in guiding behavior remains unclear, due to a lack of studies employing causal interventions with adequate temporal precision to selectively disrupt or enhance activity within this rhythm. Here, we employed a closed-loop approach to target an optogenetic manipulation to particular phases of endogenously generated theta oscillations. Closed-loop control is an under-utilized strategy for interrogating neural circuits, as it facilitates the testing of hypotheses that would be difficult or impossible to address through correlative methods (Fetz, 1969; Rolston et al., 2010; Newman et al., 2012; Wallach, 2013; Paz et al., 2013; Ngo et al., 2013).

One specific hypothesis about the mnemonic role of theta is that it partitions processes related to the encoding of new information and the retrieval of stored information (Hasselmo et al., 2002). In order to carry out their roles in spatial navigation, the hippocampus and related structures must be able to distinguish activity that tracks the current state of the world from activity that reflects prior experience. Theta could serve to coordinate cell assemblies such that

encoded and retrieved information are less likely to interfere (Hasselmo et al., 2002; Hasselmo, 2005; Colgin and Moser, 2010).

There is abundant correlative evidence that inputs to the hippocampus vary as a function of theta phase. Input from the entorhinal cortex (EC), the major source of cortical projections to the hippocampus, is highest at the trough of theta waves recorded at the hippocampal fissure (Brankack et al., 1993; Kamondi et al., 1998). Because it conveys information about the outside world, this input is likely associated with encoding of the current state of the environment (Hasselmo, 2005). At this same phase, the hippocampus is more susceptible to long-term potentiation (Hyman et al., 2003; Kwag and Paulsen, 2009), consistent with the idea that this phase is optimized for encoding new information. At the 180° phase offset (the peak of fissure theta), CA1 cells receive greater input from upstream cells in CA3 (Hasselmo, 2005). At this phase, stimulation of Schaffer collateral or temporoammonic inputs induces long-term depression (Hyman et al., 2003; Kwag and Paulsen, 2009), which could suppress information storage during memory retrieval. As a result of these phase-specific physiological changes, hippocampal networks can regulate the behavioral impact of different types of information as a function of task context.

Simultaneously recording in CA1, CA3, and EC reveals that oscillations in the high gamma range (60-100 Hz) are a signature of enhanced coordination between CA1 and EC, whereas low gamma (25-50 Hz) indicates enhanced coordination between CA3 and EC. Furthermore, these oscillations occur at different phases of theta, and typically on different cycles (Colgin et al., 2009). Taken together, these results indicate that the balance between the relative influence of the outside world (via EC) and internal states (via CA3) on hippocampal outputs is strongly modulated as a function of theta phase. Other studies have observed task-dependent modulation of hippocampal firing that is consistent with encoding of novel stimuli and retrieval of stored memories being biased to different phases (Manns et al., 2007; Lever et al., 2010; Newman et al., 2013; Douchamps et al., 2013). Computational modeling studies provide further support for this hypothesis (Hasselmo et al., 2002; Kunec et al., 2005; Hasselmo and Eichenbaum, 2005).

These results do not imply that new information is encoded and stored information is retrieved on every cycle of theta (~8 Hz). When encoding and retrieval do occur, though, they

may be preferentially active at different phases, to take advantage of the temporal structure of activity within the hippocampal–EC loop (Mizuseki et al., 2009; Colgin and Moser, 2010). Thus, a manipulation that targets a specific phase of theta could, on average, selectively modulate one process or the other.

If encoding and retrieval processes are most active at different times within the theta cycle, the consequences of a phase-specific intervention should depend on the behavioral context. Manipulations that alter hippocampal outputs at one phase of theta may have a strong impact on behavior if they occur while information is being encoded, but no effect (or the opposite effect) if they occur while information is being retrieved. Conversely, manipulations that occur with a 180° phase offset may have their behavioral impact limited to intervals in which retrieved information is used to guide behavior, but have no effect (or the opposite effect) at times when task-relevant information is being encoded.

In this study, we used millisecond-timescale optogenetic control of intrinsic inhibition to gate hippocampal outputs at specific phases of the ongoing theta cycle. Although previous studies have used optogenetic interventions to highlight the role of inhibition in phase precession (Royer et al., 2012) and theta resonance (Stark et al., 2013), here the goal was to suppress firing of CA1 in a phase-specific manner. We directly activated parvalbumin-positive interneurons, which deliver fast and powerful inhibition to the cell bodies of pyramidal neurons in the hippocampus (Bartos et al., 2007), at either the falling phase or rising phase of theta recorded in the local field potential (LFP). We performed LFP-phase-triggered optogenetic feedback in the context of a spatial navigation task, in which mice must encode and retrieve location information on every trial (Jones and Wilson, 2005). Our stimulation occurred relative to the phase of locally recorded theta on the trigger electrodes, rather than the phase at the hippocampal fissure, to which much of the previous literature uses as a landmark (Hasselmo et al., 2002; Brankack et al., 1993; Kamondi et al., 1998). However, post-hoc analysis revealed that light pulses were delivered at similar absolute phases across animals.

Using closed-loop optogenetics to intervene on the timescale of theta oscillations is a powerful approach. It allows us to alter hippocampal outputs relative to ongoing theta rhythms on a trial-by-trial basis, providing within-animal controls for all stimulation conditions. We

found that triggering inhibition on the peak of theta improved performance when it occurred in the encoding segment of the task, but had no effect in the retrieval segment. Triggering inhibition on the trough of theta had the opposite effects: it enhanced retrieval performance, but did not affect encoding processes.

3.3 Results

3.3.1 Mice learn to perform a spatial navigation task

We trained mice on a navigation task that required encoding and retrieval of reward location on individual trials. Mice were placed on an H-shaped track, which consisted of two choice points separated by a central arm (Jones and Wilson, 2005) (Figure 1a). At one junction, a movable barrier forced mice to make a left or right turn in order to arrive at the start location. At the other junction, mice were free to turn in either direction. A food reward was delivered only if mice chose the arm closest to the most recent start location.

In order to perform the task above chance, mice must update their knowledge of reward location on a trial-by-trial basis. During the encoding segment of the task (start arms), environmental cues signal the location of the upcoming reward. During the retrieval segment of the task (central arm), information about the start arm is no longer present, and thus activity that drives decision-making must be generated internally. This task makes it possible to dissociate the effects of theta phase–specific inhibition on encoding and retrieval processes by separating the cues to reward location from the time and location of the mouse's decision.

Four mice were trained on this task over the course of two to four weeks. All mice expressed the gene for Cre-recombinase in parvalbumin-positive cells, to allow us to target expression of channelrhodopsin to these cells later in the experiment. In the last five days of training, all mice performed at levels significantly above chance ($P < 0.05$, based on p.d.f. of binomial distribution with chance level of 0.5), with an average probability of correct response of 0.61 ± 0.05 (Figure 1b). In addition to improving their accuracy, mice also increased the speed at which they performed the task, to 1.49 ± 0.69 trials per minute during the last 5 days of training (Figure 1c).

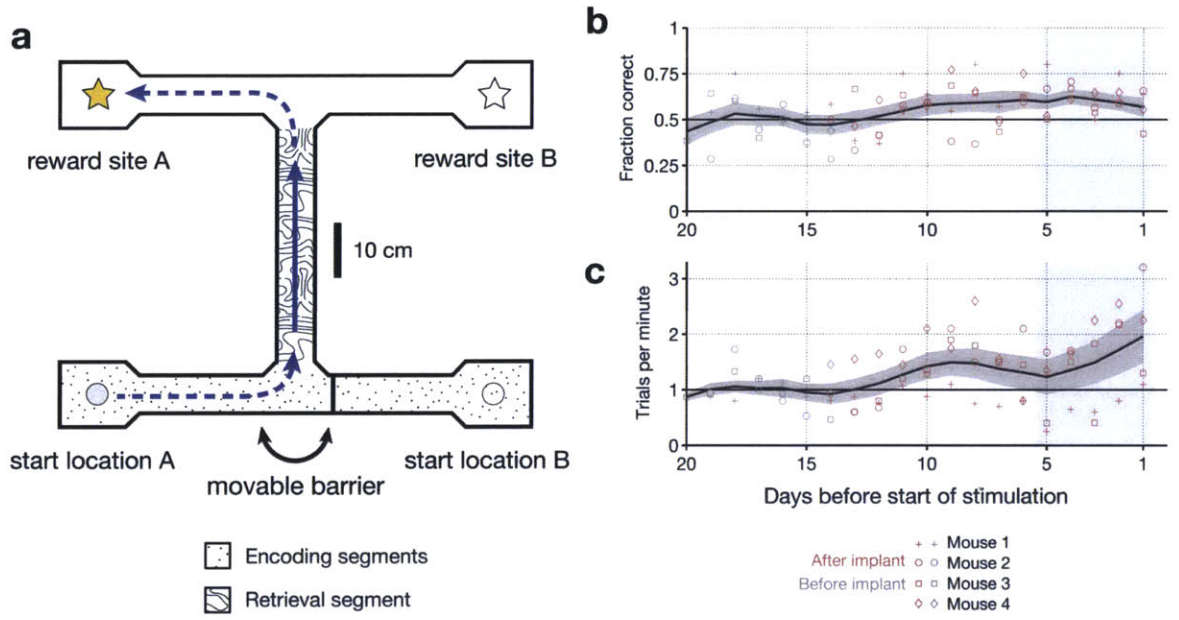

Figure 1 | Overview of the behavioral task. a, Scale drawing of the end-to-end T-maze used in all experiments. On each trial, mice navigate through the "retrieval segment" in the direction of the solid arrow and must choose between one of two reward sites. Reward is delivered for trajectories that involve two turns in the same direction (e.g., the "left/left" trajectory shown). Once the reward site is reached, mice must travel back to one of two start locations in order to initiate the next trial. A movable barrier determines the start location for that trial, and hence which reward site will contain the food pellet. The barrier is repositioned randomly after each visit to a reward site (whether correct or incorrect). A second barrier (not shown) prevents mice from navigating between reward sites after a decision has been made. The maze is surrounded by 10 cm walls made of clear acrylic, through which distal cues are visible. **b**, Fraction of correct trials in each session leading up to the start of optogenetic stimulation for $N = 4$ individual mice (open shapes) and the mean ± s.e.m. across all subjects (5-day running average). In the 5 days before the start of optogenetic stimulation (shaded region), all mice perform significantly above chance ($P < 0.05$, based on p.d.f. of the binomial distribution with probability of 0.5). **c**, Trials per minute for the same sessions as in c. Mean for last 5 days (shaded region) is 36.1 ± 18.3 trials per session per mouse.

3.3.2 Recruiting fast inhibition as a function of ongoing theta phase

After at least 8 days of training, mice were implanted with a multielectrode array that targeted movable tetrodes and stationary fiber optic cables to hippocampus bilaterally. Two fiber optic cables (one per hemisphere) were implanted to a depth of 0.9 mm at the time of surgery. In the same procedure, we injected 1.0 μL of an adeno-associated virus carrying the gene for channelrhodopsin-2 (Nagel et al., 2003) into both sides of the brain, centered on CA1 approximately 1 mm posterior to the septal pole of hippocampus. Expression spread at least 2 mm along the septotemporal axis, covering most of dorsal CA1 as well as overlying cortex (Figure 2a).

We waited at least two weeks for ChR2 expression levels to increase, during which we lowered electrodes toward the hippocampus and continued to train animals to criterion. During test sessions, we used a 465 nm LED light to drive parvalbumin-positive interneurons, which are primarily fast-spiking, soma-targeting basket and chandelier cells in the hippocampus (Pawelzik et al., 2002). All light pulses lasted 10 ms and had an irradiance of 50 mW/mm^2 (~2.5 mW from a 250 micron fiber optic cable). Individual pulses reliably elicited up to 4 spikes from well-isolated fast-spiking units (peak rate of 400 Hz, Figure 2b-e). Nearby regular-spiking units were inhibited for a period of 25 ms following light onset (Figure 2f), consistent with the known time constant of fast-spiking inhibition (Bartos et al., 2007).

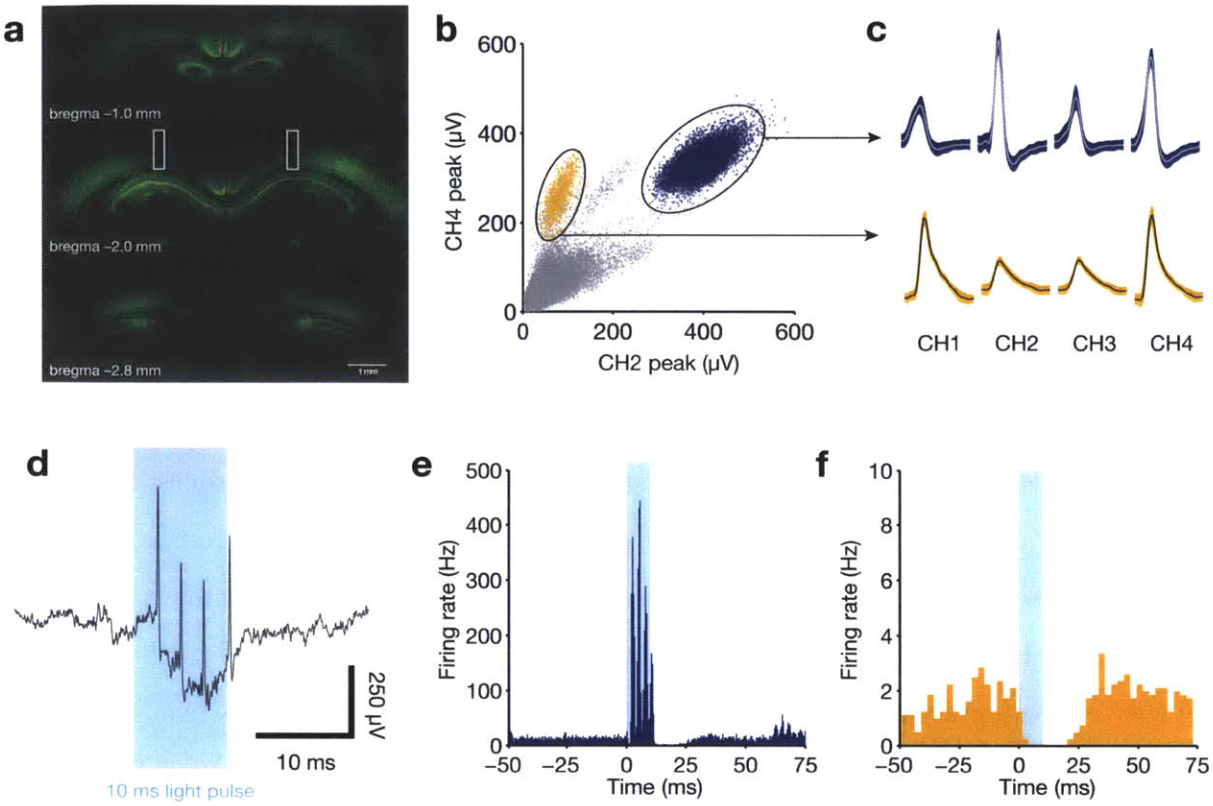

Figure 2 | Direct recruitment of fast-spiking inhibition with light. a, Expression of ChR2-EYFP throughout the dorsal hippocampus. Note the strong labeling in stratum pyramidale, indicative of dense PV+ projections in this layer. Bilateral fiber optic lesions are marked with white rectangles, centered at ~2 mm posterior to bregma and ~1.75 mm lateral to the midline. **b**, Projection plot of peak heights from a CA1 electrode containing a well-isolated fast-spiking unit (blue) and a well-isolated regular-spiking unit (yellow). **c**, Mean waveforms (with s.d.) for each tetrode channel for the same units as in panel b. **d**, Raw, broadband trace for a single trial, aligned to the 10 ms light pulse. Four light-evoked spikes from the fast-spiking unit are clearly identifiable. **e**, Peri-stimulus time histogram for the fast-spiking unit in b, c, and d, aligned to the start of each light pulse (N = 1106 pulses from one session). This unit responds with 3-4 spikes per stimulus, then remains silent for a period of ~15 ms following light offset. **f**, Peri-stimulus time histogram for the regular-spiking unit in b and c, aligned to the start of each light pulse (N = 1106 pulses from one session). This unit is silenced for ~25 ms following light onset.

We combined optogenetic stimulation with closed-loop feedback in order to trigger inhibition at specific phases of theta. In each mouse ($N = 4$), an electrode with high theta power in the local field potential was chosen as the "trigger" channel. The signals from these electrodes were filtered between 4 and 12 Hz in software. When the signal reached a local maximum or minimum, the software triggered a 10 ms light pulse delivered simultaneously to both implanted fiber optic cables (Figure 3a-b). A light pulse was delivered once per theta cycle as long as the mouse remained in the stimulation zone. The same trigger channel was used throughout the course of the experiment.

Within an individual session, stimulation was confined to the retrieval (middle arm) or encoding (sample arms) segments of the track (Figure 1a). In the *retrieval* segment, mice run toward the choice point. Stimulation in this region may affect the retrieval of stored information about reward location, but not encoding of information directly relevant for task performance. In the *encoding* segments, mice explore one of two sample arms. In these regions, stimulation could affect the encoding of available information about reward location. In both cases, however, the behavioral readout is the same: whether or not the mouse turned in the correct direction to retrieve the reward for that trial.

Individual trials were classified as one of three types: baseline (no stimulation), peak-triggered stimulation, or trough-triggered stimulation (Figure 3b). All trials types were randomly interleaved and occurred with equal probability. Three mice experienced the retrieval stimulation condition first, followed by the encoding stimulation condition. One mouse experienced the conditions in the opposite order. Analysis was limited to the first 150 trials for each condition (encoding or retrieval).

The properties of our closed-loop stimulation were as follows: the mean delay between the trigger event (peak or trough of theta reached) and the onset of the light pulse was 21.7 ± 7.2 ms for peak-triggered stimulation and 21.3 ± 7.4 ms for trough-triggered stimulation (Figure 3c). This is equivalent to approximately 1/6 of a 125 ms theta cycle. The mean phase of stimulation (based on offline-filtered theta with no phase delay) was 96 ± 54° for peak-triggered stimulation and –131 ± 63° for trough-triggered stimulation (Figure 3e, 0° = peak). The phase targeting for trough-triggered stimulation was less precise, as the rising phase of theta is shorter than the

falling phase (Belluscio et al., 2012). The mean number of pulses per trial in the retrieval-stimulation condition (middle arm) was 8.8 ± 3.3 for peak-triggered stimulation and 8.3 ± 8.3 for trough-triggered stimulation. For the encoding-stimulation condition, the mean number of pulses was 46.1 ± 57.2 for peak-triggered stimulation and 39.6 ± 39.6 for trough-triggered stimulation (Figure 3d-f).

Stimulation did not generally alter occupancy time in different segments of the track (Figure 3g-h). On each trial, mice spent the majority of time in the encoding segment (average of 2-3 s for inbound trajectories and 5-8 s for outbound trajectories). Once they left the sample arm, they ran quickly toward the goal, spending 1-2 s in the retrieval segment and a similar amount of time running toward the reward location after making their decision. The addition of optogenetic feedback only changed occupancy times significantly for one mouse in the retrieval and outbound encoding segments, and a second mouse in the outbound encoding segment. Otherwise, all occupancy times were similar ($P > 0.05$, Wilcoxon rank sum test with Bonferroni correction for 2 tests per segment, $N \geq 44$ trials per segment per mouse).

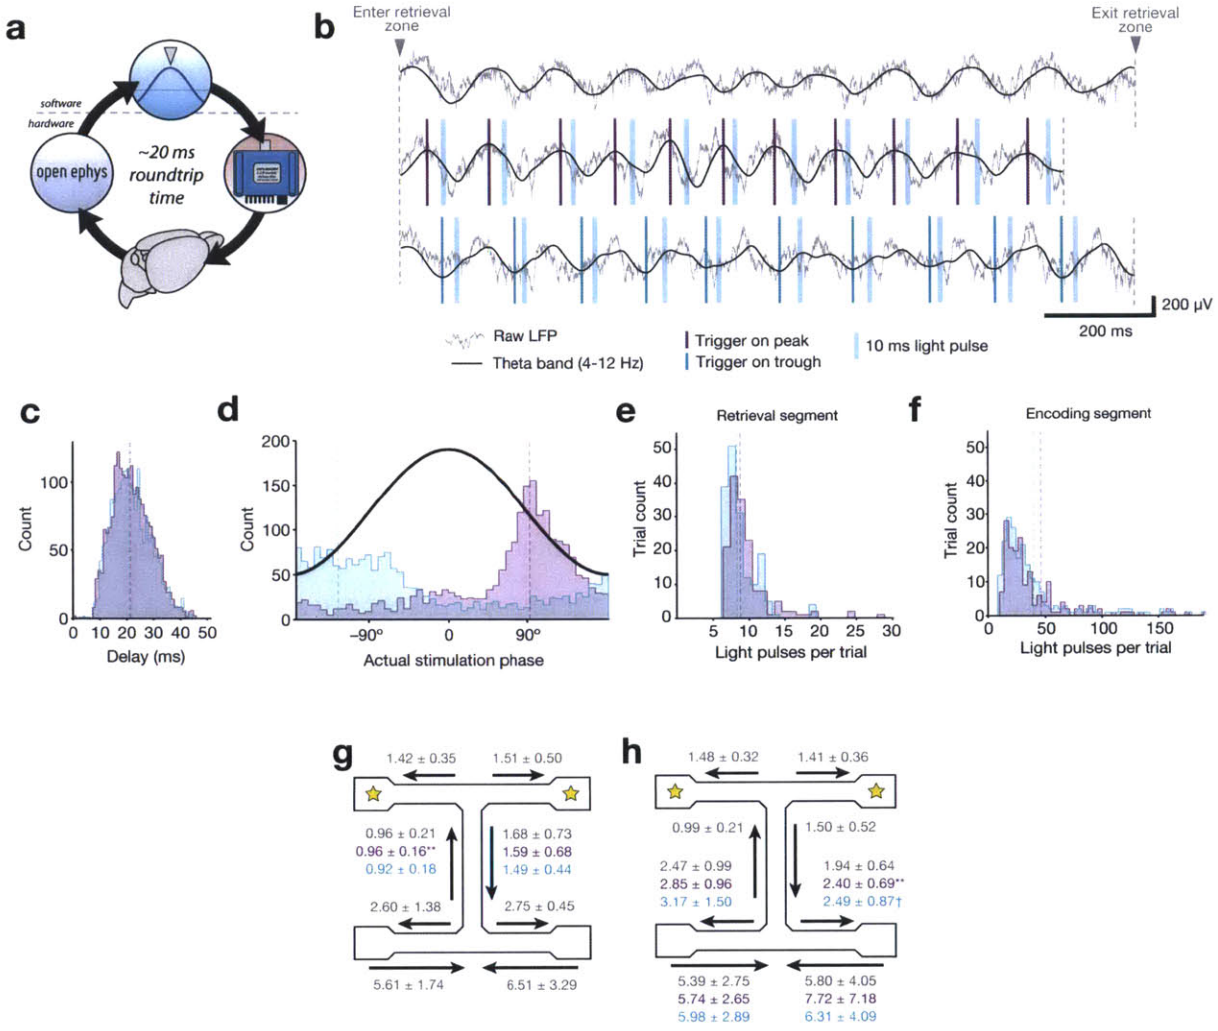

Figure 3 | Properties of theta-triggered stimulation. a, Schematic of steps involved in delivering closed-loop feedback. An event occurs in the brain (bottom), which is detected and digitized by the Open Ephys recording hardware (left), and sent to software for analysis (top). When the target event is detected, the software activates an LED (right) which delivers light to brain via implanted fiber optic cables (bottom). **b**, Examples of raw and theta-bandpassed LFP during baseline trials (top), peak-triggered stimulation trials (middle), and trough-triggered stimulation trials (bottom). Vertical blue bars indicate the time at which 10 ms light pulses occur on each cycle. **c**, Distribution of delays between detection of the actual theta peak (purple) or trough (teal) and the time of stimulus delivery. **d**, Distribution of actual theta phases at which stimulation occurred, for both peak (purple) and trough (teal) trials. Peak-triggered stimulation tends to occur during the falling phase of theta, whereas trough-triggered stimulation occurs around the actual trough and rising phase. Phase was calculated for data filtered offline between

4 and 12 Hz, to eliminate the phase delays inherent in online filtering. **e**, Distribution of pulses per trial for the retrieval segment of the track. **f**, Same as e, but for encoding segments of the track. **g**, Occupancy times in different segments of the track for trials with retrieval-segment stimulation. Values for peak-triggered and trough-triggered stimulation are shown in purple and teal, respectively (mean ± s.d. for $N = 4$ mice; ** = occupancy time decreased significantly for one mouse, $P < 0.005$, Wilcoxon rank sum test with Bonferroni correction). **h**, Occupancy times in different segments of the track for trials with encoding-segment stimulation. Values for peak and trough-triggered stimulation are shown in purple and teal, respectively (mean ± s.d. for $N = 4$ mice; ** = occupancy time decreased significantly for one mouse, $P < 0.005$; † = occupancy time increased significantly for one mouse, and decreased significantly for a different mouse, $P < 0.05$; Wilcoxon rank sum test with Bonferroni correction).

Optogenetic feedback altered the average power spectrum across trigger channels, for example by increasing the peak frequency and amplitude of theta during the peak-triggered stimulation condition (Figure 4a). There was also an increase in power in the low-gamma band (25-35 Hz) for both peak and trough stimulation, but this was associated with a much stronger peak in the beta band (16-25 Hz), which may have affected the low-gamma band via spectral leakage. Based on the shape of the evoked response to each optogenetic stimulus, it appears that these effects are due to the frequency content of the average waveform, rather than non-phase-aligned induced power in different frequency bands (Figure 4b). Aligning the local field potential to the start of each light pulse revealed a large deflection, 200-400 μV in amplitude. The shape of the average response accounts for both the shifts in theta frequency (based on the location of the subsequent peak), and the beta-range power increases (due to ~50 ms deflections). Individual pulses affected the amplitude of subsequent cycles of theta, as evidenced by the difference in the mean LFP between -100 and -75 ms for actual (purple and teal) versus dummy (gray) stimulation conditions.

Optogenetic stimulation was always aligned to the relative peak or trough of the 4-12 Hz bandpassed signals on each trigger electrode (Figure 3d). To permit meaningful interpretation of the analysis of our behavioral results, it was necessary to measure the time of stimulation relative to an absolute indicator of theta phase. We chose high gamma (60-80 Hz) power, which showed strong phasic modulation across all hippocampal electrodes, and has been previously shown to occur at a consistent phase of theta (Colgin et al., 2009). Therefore, the peak of high gamma on

baseline trials served as a landmark within each cycle of theta. In 3/4 mice, we measured stimulation times relative to the peak of high gamma for both the trigger electrode and a neighboring electrode that was passively recording signals (Figure 4c). Although post-mortem analysis of electrolytic lesions revealed different locations for each electrode, all electrodes indicated that peak-triggered stimulation occurred just after the trough of high gamma, whereas trough-triggered stimulation occurred around or after the high gamma peak. In one mouse, we could not measure absolute stimulation phase, due to the trigger electrode's location in L5/6 of cortex overlying hippocampus. This electrode expressed high theta power, presumably volume-conducted from hippocampus, which was used to trigger stimulation. However, it lacked associated high gamma power, which occurs more locally.

The consistency of this result indicates that, despite variations in electrode location, absolute stimulation phase was similar across animals. Although we did not measure CA1–MEC synchronization directly, previous studies have shown high gamma power to be a reliable indicator of enhanced coordination between these regions (Colgin et al., 2009; Yamamoto et al., 2014). Therefore, we hypothesize that trough-triggered stimulation resulted in optogenetic stimulation occurring during phases of theta in which CA1–MEC coordination was high, thereby providing CA1 with access to information about the current state of the world (Hasselmo et al., 2002; Hasselmo and Eichenbaum, 2005; Colgin et al., 2009; Colgin and Moser, 2010). Peak stimulation, on the other hand, targeted stimulation to phases in which CA1 and CA3 are most active (Mizuseki et al., 2009), during which information from the hippocampus can drive downstream structures.

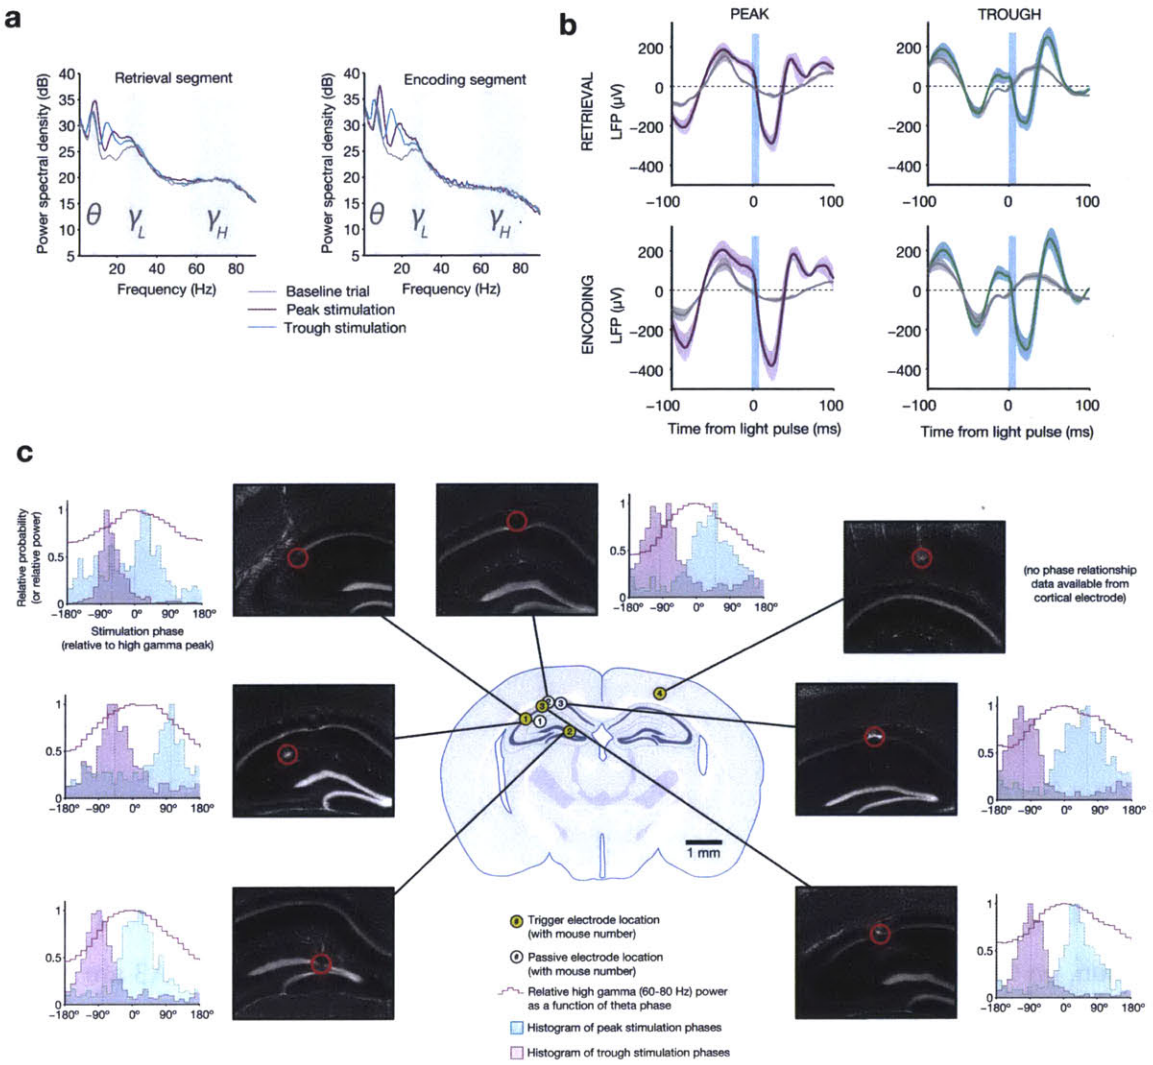

Figure 4 | Electrophysiological changes induced by theta-triggered stimulation. a, Mean power spectra for baseline, peak-triggered stimulation, and trough-triggered stimulation trials, while mice were in the retrieval segment heading toward the reward arm (left) or the encoding segment prior to entering the trial start location at the end of the sample arm (right) ($N = 4$ electrodes from 4 mice used for triggering online feedback). Theta, low gamma, and high gamma frequency bands are highlighted. **b,** Average light-evoked LFP response from $N = 3$ hippocampal electrodes for peak and trough-triggered stimulation trials (purple and teal traces, respectively), for both encoding and retrieval epochs (mean ± s.e.m.). Gray traces indicate the average theta waveforms for baseline trials, aligned to the time a stimulus would have occurred, but for which no actual light pulse was present. **c**, Locations of trigger electrodes (yellow) and passive recording electrodes (white) for four mice used in this experiment. The location of each lesion is indicated by red circles superimposed over histological sections (DAPI stain, grayscale image of

blue channel). Next to each of the images is a histogram of peak and trough stimulation phases, relative to the peak of high gamma power on that electrode for baseline (no stimulation) trials (indicated by 0°). High gamma power (a signature of synchronization between hippocampus and medial entorhinal cortex (Colgin et al., 2009), provides an absolute indication of theta phase, against which the time of our optogenetic stimulation can be compared. In all electrodes (except for the one trigger electrode in cortex, where high gamma was not measured), trough stimulation occurs after the peak of high gamma power, while peak stimulation occurs before the peak of high gamma.

3.3.3 Impact of closed-loop inhibition on behavior depends on both theta phase and task segment

The effects of closed-loop optogenetic feedback on behavior depended on both the phase of theta used to trigger stimulation and the region of the track in which the stimulation occurred. On individual trials, 10 ms light pulses were triggered on either the peak or trough of theta (Figure 5a, phase relative to theta at the hippocampal fissure). When stimulation occurred in the retrieval segment, performance did not differ between baseline and peak-triggered stimulation for 4/4 mice (mean of 57.3 ± 10.0% correct for baseline versus 57.8 ± 10.5% correct for peak, individual results in Table 1). For trough-triggered stimulation, however, performance improved significantly in 4/4 mice (71.0 ± 8.2% correct for trough; significance determined by the p.d.f. of the binomial distribution, with baseline accuracy for each mouse used as the "chance" level). The opposite effects were observed for stimulation in the encoding segment. In this condition, performance during trials with trough-triggered stimulation did not differ from baseline in 3/4 mice (mean of 59.1 ± 2.4% correct for baseline versus 58.7 ± 10.5% correct for trough; 1 mouse had significantly impaired performance in the trough-stimulation condition). For peak-triggered stimulation, performance improved significantly in 3/4 mice (mean of 69.6 ± 14.7% correct; 1 mouse showed no difference from baseline).

On average, trough-triggered stimulation resulted in a 13.7% improvement in accuracy for the retrieval condition, while peak-triggered stimulation resulted in a 10.5% improvement in accuracy for the encoding condition (Figure 5b). The effects were consistent across individual mice, with trough-triggered stimulation improving performance more for the retrieval segment than the encoding segment in 4/4 mice, and peak-triggered stimulation improving performance

more for the encoding segment than the retrieval segment for 3/4 mice (Figure 5c). Such effects represent a double-dissociation, as phase-specific optogenetically recruited inhibition reversed its behavioral impact depending on the region of stimulation.

To estimate the probability that these results could have occurred by chance, we had to consider all outcomes in which a double dissociation was present. Our initial hypothesis was only that the effects of stimulation would depend on both task phase and theta phase, not that performance would be specifically impaired or improved. We used a bootstrap procedure with 10,000 repetitions to determine the probability that any of the possible double dissociations shown in Figure 5d could have occurred by chance. We randomized the labels for all trials (baseline, peak-triggered, and trough-triggered) and looked for the presence of significant changes relative to baseline in any of the conditions in the 2x2 square. If 3/4 mice showed the same behavior (enhancement, impairment, or no change), we considered that a "consistent" quadrant. The probability that the same effects would be seen along any diagonal was 0.0013. The probability that the same effects were seen along *both* diagonals (what we observed in the actual data) was 0.0001.

To better understand the source of these behavioral effects, we analyzed the types of mistakes made by the mice, and how activating inhibition at the appropriate phase serves to correct them. It is clear that the outcome of the previous trial has a strong effect on decision-making: mice are much more likely to make a correct choice if they are cued to switch arms after failing to receive reward (82 ± 14% correct) or return to the same arm after receiving reward (72 ± 18% correct; mean response across encoding and retrieval conditions, baseline trials only). They are less likely to switch arms after a correct decision (32 ± 9% correct) or to return to the same arm after an incorrect choice (40 ± 21% correct). Thus, mice exhibit a bias toward a "win-stay, lose–switch" strategy. They favor returning to arms in which they just received reward, or switching to the opposite arm if there was no reward on the previous trial (Figure 5e).

Adding optogenetic stimulation on certain trials allows mice to overcome this inherent bias. The strongest effects were seen for trials in which mice were cued to enter the opposite reward arm after making a correct choice. On average, trough-triggered stimulation in the retrieval segment improved performance on these trials by 19.5 ± 7.3%, with improvement seen in 4/4

mice (versus 5.1 ± 14.7% for peak-triggered stimulation). Peak-triggered stimulation in the encoding segment improved performance on these trials by 25.0 ± 10.3%, again with improvement in 4/4 mice (versus 8.7 ± 14.6% for trough-triggered stimulation). The effects of phase-specific stimulation on other types of errors were less pronounced, but there was no evidence for reduced performance by the "optimal" stimulation phase for any trial type (Figure 5f).

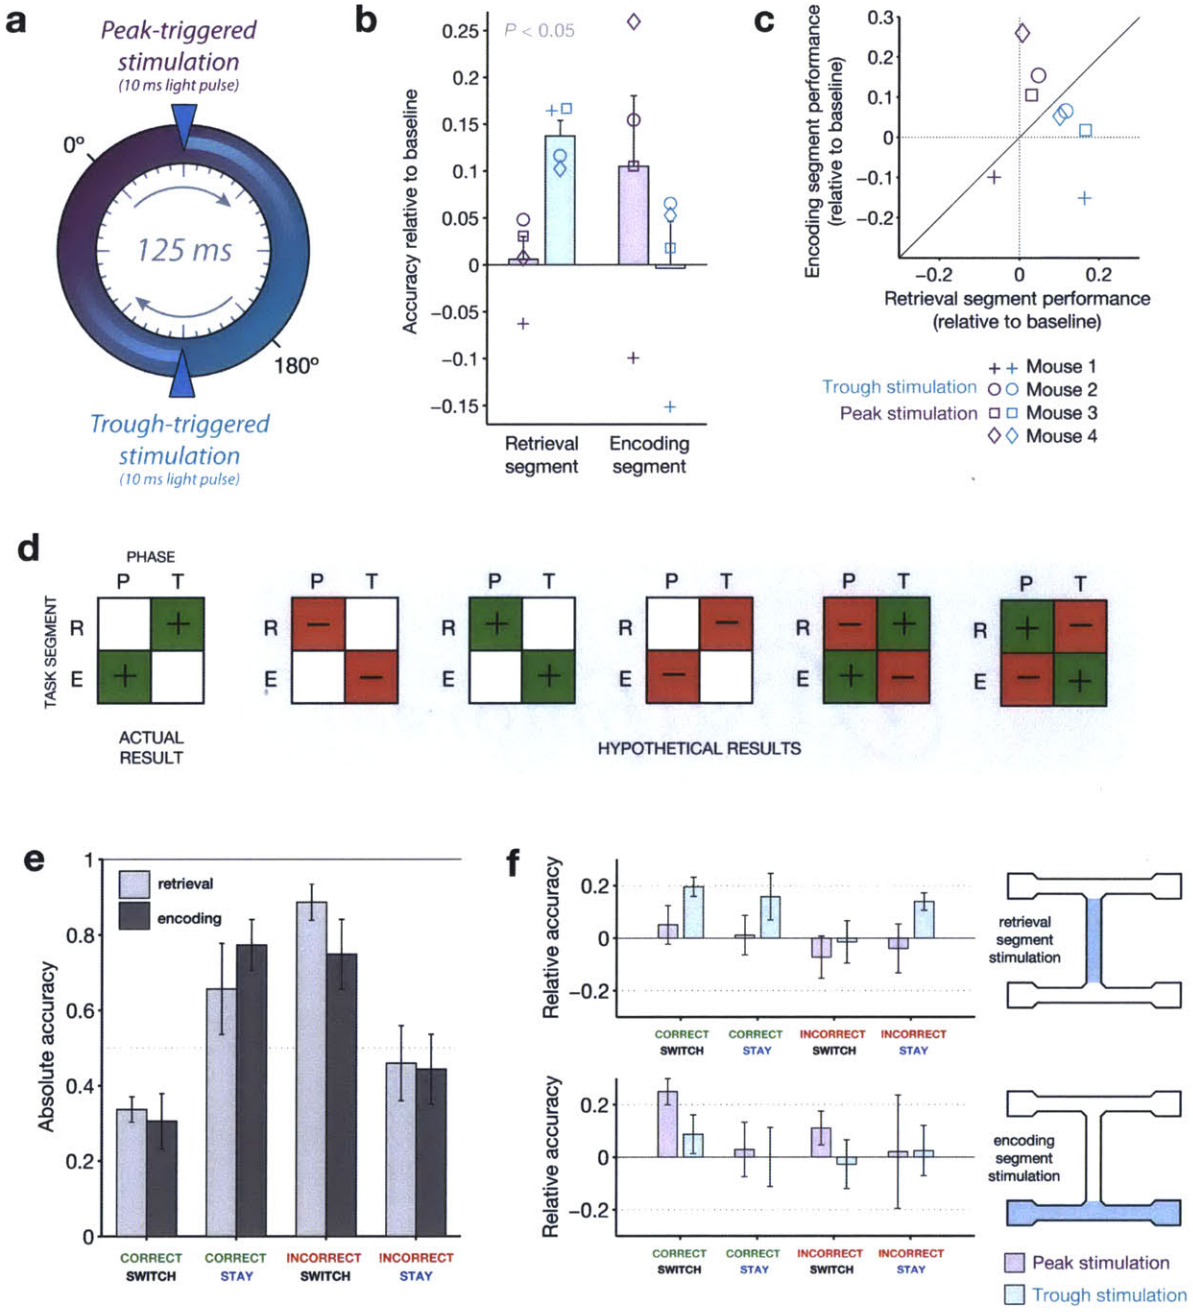

Figure 5 | Behavioral modulation depends on both theta phase and task segment. a, Illustration of the two manipulations performed in this experiment. On any given "non-baseline" trial, stimulation was triggered by the peak (purple phase) or trough (teal phase) of the 4-12 Hz theta rhythm. The resulting light pulses recruited inhibition for ~25 ms, or approximately 1/5 of the 125 ms theta cycle. **b**, Accuracy relative to baseline for four mice in four conditions:

optogenetic stimulation triggered at the peak (purple) or trough (teal) of theta, in either the retrieval (left) or encoding (right) segments of the track. Mean ± s.e.m., with results for each mouse overlaid. Individual results in the gray regions are significantly different from baseline ($P < 0.05$, p.d.f. of binomial distribution with probability equal to baseline accuracy). **c**, Same data as in b, but represented on the same axes. Note that peak-triggered stimulation in the encoding segment consistently improves performance more than the same type of stimulation in the retrieval segment (points above diagonal line). The opposite effects are seen for trough-triggered stimulation. **d**, Schematic of all possible "double-dissociation" scenarios used for establishing bootstrap significance levels of the actual result. **e**, Performance on baseline (no stimulation) trials for four different trial types: (1) mice are cued to switch arms after a correct choice (correct/switch), (2) mice are cued to return to the same arm after a correct choice (correct/stay), (3) mice are cued to switch arms after an incorrect choice (incorrect/switch), and (4) mice are cued to return to the same arm after an incorrect choice. Trials are grouped by retrieval stimulation or encoding stimulation conditions. For both conditions, changing trial type has a significant effect on performance: retrieval stimulation, $\chi^2 = 8.4$, $P = 0.038$; encoding stimulation, $\chi^2 = 8.1$, $P = 0.044$; Friedman test (nonparametric, repeated-measures ANOVA). **f**, Change in performance with the addition of closed-loop optogenetic stimulation for the four trial types in e.

	RETRIEVAL			ENCODING		
	Baseline	**Peak**	**Trough**	**Baseline**	**Peak**	**Trough**
Mouse 1	0.55	0.49	0.71	0.58	0.48	0.43
		$P = 0.09$	$\mathbf{P = 0.02}$		$P = 0.07$	$\mathbf{P = 0.02}$
Mouse 2	0.52	0.57	0.64	0.60	0.75	0.66
		$P = 0.09$	$\mathbf{P = 0.02}$		$\mathbf{P = 0.01}$	$P = 0.06$
Mouse 3	0.50	0.53	0.67	0.62	0.73	0.64
		$P = 0.10$	$\mathbf{P = 0.01}$		$\mathbf{P = 0.03}$	$P = 0.11$
Mouse 4	0.72	0.73	0.82	0.56	0.82	0.62
		$P = 0.12$	$\mathbf{P = 0.04}$		$\mathbf{P = 0.0001}$	$P = 0.08$

Table 1 | Results for individual mice. Probability of a correct response for four mice under six conditions: baseline (no stimulation), peak-triggered stimulation, and trough-triggered stimulation in both the retrieval and encoding segments. *P*-values computed from the p.d.f. of the binomial distribution with chance levels equal to the baseline performance for that mouse. Green cells indicate significant improvement relative to baseline, while red cells indicate significant impairment under that condition ($P < 0.05$).

3.4 Discussion

These results provide new evidence for a hypothesis that was previously supported by correlational studies and computational models: processes related to encoding new information and retrieving stored information occur preferentially at different phases of the theta oscillation (Hasselmo et al., 2002; Hasselmo, 2005; Colgin et al., 2009; Colgin and Moser, 2010). We have shown that interventions targeting the falling or rising phases of theta have different effects depending on the behavioral context. When environmental cues to reward location are available (as in the encoding segment of the task), triggering hippocampal inhibition on the peak of theta enhanced navigational accuracy. When behavioral guidance must be based on internal signals alone (as in the retrieval segment of the task), triggering hippocampal inhibition on the trough of theta increased the probability of a correct choice.

What is the neural basis these effects? Our favored explanation is that phase-specific inhibition serves to reduce the response to task-irrelevant inputs. Figure 6, which was inspired by a similar diagram in Hasselmo et al. (2002), illustrates the mechanism by which this could occur. On average, the influence of CA3 and entorhinal cortex inputs to CA1 changes as a function of theta phase (Hasselmo et al., 2002). Under baseline conditions, the relative influence of CA3 and EC is "balanced." With the addition of closed-loop optogenetic feedback, excess inhibition reduces spike activity either during EC-dominant or CA3-dominant periods of the theta cycle. Although the parvalbumin-positive interneurons recruited by this manipulation are typically active during the CA3-dominant phases of theta (Lasztóczi and Klausberger, 2014), causal control allows us to activate them synchronously at arbitrary times during the theta cycle. Under the proposed mechanism, inhibiting CA1 during EC-dominant cycles in the retrieval segment improves task performance by increasing the relative influence of CA3. Conversely, inhibiting CA1 during CA3-dominant cycles in the encoding segment improves task performance by increasing the relative influence of EC or by suppressing retrieval of interfering cross-trial information. In both cases, enhanced navigational accuracy could result from suppression of task-irrelevant information, rather than the enhancement of task-relevant information.

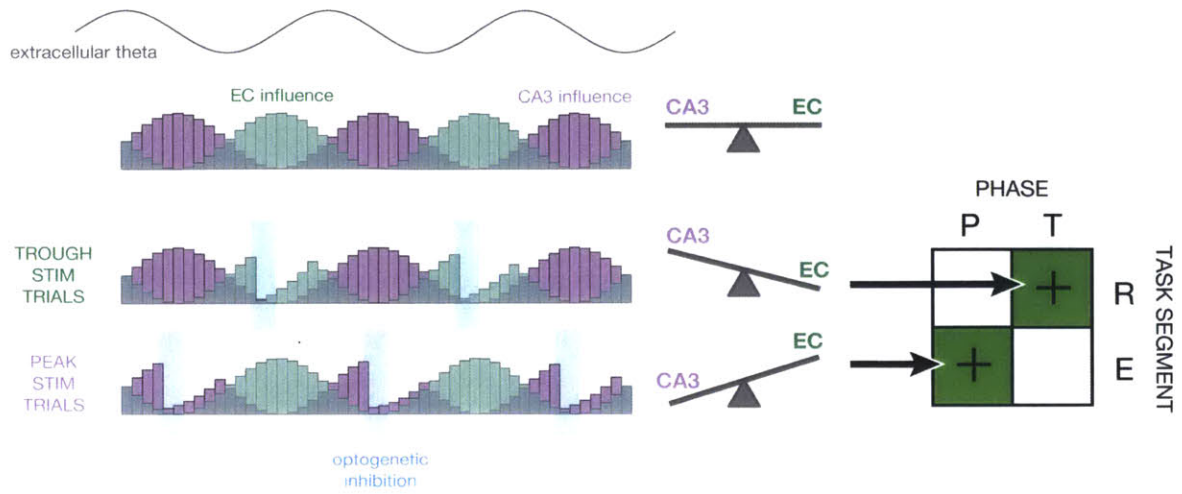

Figure 6 | Proposed mechanism. This diagram illustrates how the relative influence of CA3 versus entorhinal cortex (EC) inputs to CA1 could explain the experimental results. At the top, a sine wave indicates the phase of theta. Below, the purple and green histograms show the fluctuating influence of CA3 and EC on CA1 on each cycle. Levels represent averages; on individual cycles, one or the other may dominate (Colgin et al., 2009). When optogenetic inhibition is triggered on the trough of theta (T), it tends to reduce firing rates in CA1 during periods of high EC influence. This tips the balance in favor of CA3, thereby improving performance during periods of retrieval (R). When optogenetic inhibition is triggered on the peak of theta (P), it tends to reduce firing rates in CA1 during period of high CA3 influence. This tips the balance in favor of EC, thereby improving performance during periods of encoding (E). On the whole, our closed-loop manipulation may improve performance by reducing the influence task-irrelevant inputs as a function of both theta phase and maze region.

Our data supports the presence of strong local inhibition in CA1 at specific phases of theta (Figure 2f, 3d). The duration of this inhibition (~25 ms) is equivalent to approximately 1/5 of a 125 ms (8 Hz) theta cycle, long enough to impact encoding or retrieval functions, but precise enough to avoid disrupting the entire cycle. This suggests a simple, CA1-specific mechanism could be sufficient to explain our behavioral results. According to our proposed mechanism, suppression of inputs carrying information about the current state of the world improves performance in the retrieval segment, whereas suppression of inputs carrying information about past states improves performance in the encoding segment.

Although a local mechanism can provide the simplest explanation of our results, we cannot rule out a mechanism that involves changes in inter-regional coupling strength without simultaneous recordings from the relevant downstream areas. The navigation task used in this study must engage a wide network of brain regions, including those that receive monosynaptic projections from CA1. Besides projections back to the entorhinal cortex, the major cortical output of dorsal CA1 is retrosplenial cortex (Cenquizca and Swanson, 2007). Retrosplenial cortex is known to be important for spatial working memory (Vann and Aggleton, 2004; Vann et al., 2009; Keene and Bucci, 2009), and changes in coupling between hippocampus and this region could affect performance on the present task. There is also a projection from CA1 to prefrontal cortex, although this originates primarily from the ventral regions (Swanson, 1981; Cenquizca and Swanson, 2007). It is possible that our optogenetic manipulation is affecting prefrontal-dependent decision-making via ventral CA1 (Colgin, 2011; Schmidt et al., 2013; O'Neill et al., 2013), subicular projections (Jay and Witter, 1991), or multisynaptic pathways. Our manipulation could also be exerting long-range effects through the actions of projecting interneurons, a small fraction of which are parvalbumin-positive (Jinno, 2009).

Coupling strength between CA1 and medial prefrontal cortex is known to depend on task phase (Jones and Wilson, 2005), and it is therefore possible that phase-specific inhibition of CA1 could have a differential impact on CA1–mPFC synchrony during the encoding and retrieval segments. However, if a change in inter-regional coupling strength is at play, we might expect the impact of our optogenetic manipulation to be most pronounced during the retrieval segment, when cortical regions involved in behavioral guidance are likely to access the hippocampal

representation of space. In this case, it is not clear that phase-specific stimulation of the hippocampus in the encoding segment should also affect performance. Although our observation of a double-dissociation suggests that the main mechanism is local to the hippocampus, disruption of CA1–mPFC coupling during the encoding segment could minimize cross-trial interference, also leading to enhanced performance.

The influence of alternative behavioral strategies must also be considered. Mice could employ multiple strategies for completing the task, such as an egocentric, hippocampus-independent strategy based on turn direction or an allocentric, hippocampal-dependent strategy based on an internal map (Packard et al., 1989; Floresco et al., 1997). It is possible that inhibiting the hippocampus biases the mouse toward an egocentric strategy, which happens to improve performance. If this were the case, again, we would not expect to see such a striking double-dissociation effect in our results. We would instead predict that the same phase of stimulation would enhance performance in both the encoding and retrieval segments. Further evidence against a "strategy switching" mechanism could come from an experiment in which inhibition was activated at the optimal phase for *both* the encoding and retrieval segments on individual trials. If stimulation is merely invoking a change in strategy, we would not expect to see additive effects; that is, stimulation in the encoding phase would be sufficient to reach peak performance. However, if combined encoding and retrieval stimulation improved performance more than one or the other in isolation, it would suggest that specific effects on encoding and retrieval operations are at play.

The analysis of types of errors in Figure 5e makes it clear that encoding and retrieval are continuous processes. They are not necessarily confined to the "encoding" and "retrieval" segments specific to this task. The mice are actually performing at least two tasks concurrently, one which involves entry into the cued reward arm (the trained task) and one which involves acting based on the outcome of the previous trial (an untrained task). Could optogenetic stimulation merely serve to "reset" the system, reducing interference between trials? We consider this possibility unlikely, due to the fact that the effects of stimulation at different phases depended on the track segment in which it occurred. This double-dissociation indicates that stimulation at the appropriate phase allowed mice to more accurately update and retrieve

knowledge of the upcoming reward location, rather than simply suppress the influence of the previous trial. In addition, optogenetic stimulation did not impair performance for "easy" trials, in which the cued reward location was consistent with animals' tendency toward a "win–stay, lose–switch" strategy (Figure 5f). If stimulation brought mice back to a naïve "baseline" state, we would expect them to make more errors when inhibition is recruited on these trials. Overall, this analysis highlights the fact that encoding and retrieval cannot be considered discrete states that depend on the task at hand, and are instead occurring continuously as animals explore their environment.

Our results indicate that optogenetic inhibition of CA1 serves to *improve* performance across all mice tested. Our initial hypothesis was that the effects of stimulation would depend on both the task segment and phase, but we were unsure if they would be beneficial or punitive. Given that we are recruiting inhibition, and thereby suppressing CA1 output, one might expect the behavioral impact on a hippocampal-dependent task to be negative. Recruiting inhibition during the "retrieval" phase should impair performance in the retrieval segment, whereas recruiting inhibition during the "encoding" phase should impair performance during the encoding segment. The fact that we instead observed only enhanced performance (rather than enhancement for one phase of stimulation and impairment for the other) may be explained by a floor effect. Baseline performance was modest at the start of testing (Figure 1a, Table 1), which makes it unlikely that we could observe a significant impairment, even if it did exist. Mice were strongly influenced by the outcome of the previous trial (Figure 5e), which explains why their accuracy on the trained task is only slightly (but significantly) above chance. Our phase-specific optogenetic intervention helps them overcome this bias, especially in the case of trials in which they are required to switch arms after receiving reward (Figure 5f). However, even for trials in which the reward location was consistent with animals' intrinsic biases, stimulation did not interfere with performance. It is possible that higher light intensities, alternate fiber placements, or a different target phase could have created the conditions necessary to negatively impact behavior.

Revealing a convincing mechanistic explanation for the behavioral effects seen in this study will require further investigation. The present results justify more extensive inquiry along these lines by providing evidence that processes related to encoding new information and retrieving

stored information are most active at different phases of theta. This hypothesis was originally based on a computational model (Hasselmo et al., 2002), which was later supported by correlative evidence (Manns et al., 2007; Colgin et al., 2009; Newman et al., 2013; Douchamps et al., 2013). We advance this line of investigation through the use of a closed-loop optogenetic intervention that allowed us to interact with the hippocampus at specific phases of theta on a trial-by-trial basis. As the tools for closed-loop control become more accessible, experiments that couple precise stimulation to internal state variables have the potential to enhance our understanding of a wide range of topics related to the study of neural systems.

3.5 Experimental Procedures

3.5.1 Animals

All mice were male parvalbumin-Cre (PV-Cre) heterozygotes, derived from PV-Cre BAC transgenics back-crossed into a C57BL/6J line (Jackson Laboratory strain B6;129P2-Pvalbtm1(cre)Arbr). Mice were 8-12 weeks old at the start of training (mean age = 10.8 ± 1.5 weeks) and 10-15 weeks old at the time of surgery (mean age = 13.5 ± 2.1 weeks). Animals were individually housed and maintained on a 12-hour light/dark cycle (lights on at 7:00 AM). All experimental procedures and animal care protocols were approved by the Massachusetts Institute of Technology Institutional Animal Care and Use Committees and were in accordance with NIH guidelines.

3.5.2 Task structure

The task was adapted from that used in a previous study (Jones and Wilson, 2005). The track consisted of two T-mazes placed end-to-end to form an "H" shape, with movable gates at both choice points. When running toward the sample arms, the location of the gate forced the mouse in one direction or the other. No reward was delivered in the sample arm, but mice were required to reach the end of it in order to initiate a new trial. When running in the opposite direction, mice could choose between one of two "free choice" arms. Reward was only delivered if the mouse entered the free choice arm closest to the most recent sample arm it had visited. Rewards

consisted of one 14 mg sugar pellet (Bio-Serv product #F05684), and were always preceded by a 2 kHz tone lasting 250 ms, triggered by entry into the reward zone.

The track was made from laser-cut acrylic, with transparent walls and a black floor. Distal cues were provided by three large black curtains with high-contrast patterns in the center and the experimenter's body, which remained in a consistent location across days. IR sensors were used to monitor entry and exit from different regions of the track. An Arduino sent information about IR beam breaks to a computer running custom software written in Processing. The experimenter manually moved the "forced choice" gates at the start of each trial according to a sequence generated randomly by the behavior computer. If mice were biased toward one reward arm, the probability of reward appearing in the opposite arm increased according to the following equation: P(reward in left arm) = P(mouse chose right arm during last 12 trials).

3.5.3 Behavioral training

Prior to the start of training, mice were restricted to 2-3 g of dry food per day, with unlimited access to water. Training began with 4-6 days of habituation, during which mice freely explored the track while receiving reward in both the choice and sample arms. Next, a period of "forced choice" training began, in which a gate always forced the mice in the correct direction at each choice point. After 5-6 days of forced choice training, a "free choice" condition was added, in which mice were allowed to make incorrect decisions. Subsequent sessions typically consisted of 10-15 minutes of forced choice training, followed by 15-20 minutes of free choice training. Mice received free choice training for 0-10 days before surgery, and 14-26 days prior to the start of behavioral testing.

3.5.4 AAV vectors

We used AAV-5 viral vectors containing double-floxed, inverted, open-reading-frame ChR2 (H134R variant) coupled to EYFP and driven by the EF1α promoter.

3.5.5 Fiber optic–electrode implants

Implants were constructed according to a previously described procedure (Voigts et al., 2013). Design files can be found on GitHub (https://github.com/open-ephys/flexdrive), and assembly instructions are hosted on the Open Ephys wiki (https://open-ephys.atlassian.net/wiki/display/OEW/flexDrive). The base of the drive consisted of two stainless steel cannulae with their centers 3.6 mm apart. Each cannula held four electrodes spaced in a ring around a central fiber optic cable (240 micron core diameter, 0.51 NA, Edmund Optics part #02-531). The fiber optic cables protruded 0.9 mm past the end of each cannula. Electrodes were made from 12.5 μm polyimide-coated nichrome wire (Kanthal), twisted and heated to form tetrodes (Nguyen et al., 2009). Individual electrodes were gold plated to an impedance of 200-400 kΩ.

3.5.6 Surgical procedure

Mice were anesthetized with isofluorane gas anesthesia (0.75-1.25% in 1 L/min oxygen) and secured in a stereotaxic apparatus. The scalp was shaved, wiped with hair removal cream, and cleaned with iodine solution and alcohol. Following IP injection of Buprenex (0.1 mg/kg, as an analgesic), the skull was exposed with an incision along the midline. After the skull was cleaned, six steel watch screws were implanted in the skull, one of which served as ground.

Next, a ~1.5 mm-diameter craniotomy was drilled over left hippocampus (2.0 mm posterior to bregma and 1.8 mm lateral to the midline) and the dura was removed. Virus was delivered through a glass micropipette attached to a Quintessential Stereotaxic Injector (Stoelting). The glass micropipette was lowered through the center of the craniotomy to a depth of 1.2 mm below the cortical surface. A bolus of 1 μl of virus (see details above) was injected at a rate of 0.05 μl/min. After the injection, the pipette was held in place for 10 min at the injection depth before being fully retracted from the brain. The same procedure was then repeated for the opposite hemisphere.

The fiber optic–electrode implant (see details above) was aligned with the two craniotomies, and lowered until the cannulae were flush with the cortical surface. This placed the two fiber optic cables just above the CA1 region of hippocampus (depth of ~0.9 mm). Once the implant was stable, a small ring of black dental acrylic was placed around its base. A drop of surgical

lubricant (Surgilube) prevented dental acrylic from contacting the cortical surface. Adhesive luting cement (C&B Metabond) was used to further affix the implant to the skull. Once the cement was dry, the scalp incision was closed with VetBond (3M), and mice were removed from isoflurane.

Following 2-4 days of recovery, electrodes were lowered to their final location over the course of 2-3 weeks. Once stimulation began, electrodes were not adjusted.

3.5.7 Electrophysiology

On testing days, the track was wiped with an anti-static liquid (Staticide) and cleared of all debris. Electrophysiological data was recorded with the Open Ephys platform (http://open-ephys.org), an open-source data acquisition system based on Intan amplifier chips (http://www.intantech.com). Tetrode signals were referenced to ground, filtered between 1 and 7500 Hz, multiplexed, and digitized at 30 kHz on the headstage (design files available at https://github.com/open-ephys/headstage/tree/master/1x32_Omnetics_Standard). Digital signals were transmitted over a 12-wire cable counter-balanced with a system of pulleys and weights. Mouse location was determined via IR gates at behaviorally relevant points along the track and an overhead camera monitoring a red LED mounted on the headstage.

3.5.8 Stimulation protocol

Online feedback was delivered using the Open Ephys GUI (full source code available at https://github.com/open-ephys/GUI). The trigger channel was filtered between 4 and 12 Hz (2nd-order Butterworth) and sent to a "Phase Detector" module. When the mouse entered the stimulation segment (either one of two sample arms for "encoding" sessions or the central arm on forward and reverse trajectories for "retrieval" sessions), the Phase Detector emitted trigger events when the signal reached a local maximum ("peak") or local minimum ("trough"). Trials of each type ("peak," "trough," or "blank") were randomly interleaved with equal probability. Stimulation was triggered via a USB connection to a Pulse Pal (https://sites.google.com/site/pulsepalwiki/), and consisted of 10 ms light pulses from a Plexon PlexBright LED (465 nm, ~50 mW/mm^2).

3.5.9 Histology

At the end of training, electrodes sites were lesioned with 15 μA of current for 10 s. Mice were transcardially perfused with 100 mM PBS followed by 4% formaldehyde in PBS. Brains were post-fixed for at least 18 h at 4° C. 60 μm sections were mounted on glass slides with Vectashield (Vector Laboratories), coverslipped, and imaged with an upright fluorescent microscope. Viral expression was confirmed by observing EYFP expression beneath the fiber optic lesions in CA1 of all animals. Expression spread ~2 mm along the length of the dorsal hippocampus, primarily in CA1, but also in the lateral portion of CA3. Labeling was strongest in the hippocampal cell layer, where parvalbumin-positive cells have the densest projections. There was also expression in overlying cortex. No expression was observed in the dentate gyrus. In all animals, the lesion corresponding to the electrode used to trigger stimulation was identified.

3.5.10 Data analysis

All data analysis was performed using custom Matlab scripts. Spike activity was extracted offline by thresholding the 300-6000 Hz bandpassed signal. Units were clustered with Simple Clust software (https://github.com/moorelab/simpleclust), based on peak heights and regression coefficients for individual waveforms. Spikes were aligned to light pulses using event timestamps, or to the phase of LFP theta.

To determine the actual phase of theta without the phase shift associated with online filtering, we filtered the wideband, full-sample-rate data offline using the Matlab "filtfilt" function (2nd-order Butterworth, 4-12 Hz bandpass). We used the angle of the Hilbert-transformed signal to compute the phase in degrees (–180° to 180°, peak at 0°). Spectral analysis was performing using the Chronux toolbox (http://www.chronux.org), using multitaper methods (time–bandwidth product = 2, number of tapers = 3).

Behavioral analysis was limited to the first 150 trials performed in each condition (encoding stimulation versus retrieval stimulation). Trials were grouped by stimulation type (blank, peak-triggered, or trough-triggered) and the responses (0 = correct choice, 1 = incorrect choice) were averaged. *P*-values for individual mice were computed using the probability density function of the binomial distribution, with N = the number of trials of a given type and p = baseline accuracy.

The probability that a double-dissociation would occur by chance was computing using a bootstrap method with randomized trial labels (10,000 iterations).

3.6 Acknowledgments

We thank K. Kim and R. LeCoultre for their help running experiments. We are grateful for G. Hale, J. Sanders, J. Voigts, D. Meletis, M. Carlén, and R. Harrison for technical assistance. We thank members of the Wilson Lab for their comments on the figures and manuscript. This study was supported by a grant from the NIH to M. Wilson and NDSEG and NRSA fellowships to J. Siegle.

3.7 References

Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nat Rev Neurosci* **8**: 45-56.

Belluscio MA, Mizuseki K, Schmidt R, Kempter R, Buzsáki G (2012) Cross-frequency phase–phase coupling between theta and gamma oscillations in the hippocampus. *J Neurosci* **32**: 423-435.

Brankack J, Stewart M, Fox SE (1993) Current source density analysis of the hippocampal theta rhythm: associated sustained potentials and candidate synaptic generators. *Brain Res* **615**: 310-327.

Buzsáki G (2002) Theta oscillations in the hippocampus. *Neuron* **33**: 325-340.

Buzsáki G, Leung LW, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. *Brain Res* **287**: 139-171.

Cenquizca LA, Swanson LW (2007) Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Res Rev* **56**: 1-26.

Colgin LL (2011) Oscillations and hippocampal-prefrontal synchrony. *Curr Opin Neurobiol* **21**: 467-474.

Colgin LL (2013) Mechanisms and functions of theta rhythms. *Annu Rev Neurosci* **36**: 295-312.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**: 353-357.

Colgin LL, Moser EI (2010) Gamma oscillations in the hippocampus. *Physiology (Bethesda)* **25**: 319-329.

Douchamps V, Jeewajee A, Blundell P, Burgess N, Lever C (2013) Evidence for encoding versus retrieval scheduling in the hippocampus by theta phase and acetylcholine. *J Neurosci* **33**: 8689-8704.

Fetz EE (1969) Operant conditioning of cortical unit activity. *Science* **163**: 955-958.

Floresco SB, Seamans JK, Phillips AG (1997) Selective roles for hippocampal, prefrontal cortical, and ventral striatal circuits in radial-arm maze tasks with or without a delay. *The Journal of neuroscience* **17**: 1880-1890.

Hasselmo ME (2005) What is the function of hippocampal theta rhythm?--Linking behavioral data to phasic properties of field potential and unit recording data. *Hippocampus* **15**: 936-949.

Hasselmo ME, Bodelón C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput* **14**: 793-817.

Hasselmo ME, Eichenbaum H (2005) Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw* **18**: 1172-1190.

Hyman JM, Wyble BP, Goyal V, Rossi CA, Hasselmo ME (2003) Stimulation in hippocampal region CA1 in behaving rats yields long-term potentiation when delivered to the peak of theta and long-term depression when delivered to the trough. *J Neurosci* **23**: 11725-11731.

Jay TM, Witter MP (1991) Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *J Comp Neurol* **313**: 574-586.

Jinno S (2009) Structural organization of long-range GABAergic projection system of the hippocampus. *Front Neuroanat* **3**:13.

Jones MW, Wilson MA (2005) Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol* **3**:e402.

Kamondi A, Acsády L, Wang XJ, Buzsáki G (1998) Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. *Hippocampus* **8**: 244-261.

Keene CS, Bucci DJ (2009) Damage to the retrosplenial cortex produces specific impairments in spatial working memory. *Neurobiol Learn Mem* **91**: 408-414.

Klausberger T, Magill PJ, Márton LF, Roberts JD, Cobden PM, Buzsáki G, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* **421**: 844-848.

Klausberger T, Marton LF, O'Neill J, Huck JH, Dalezios Y, Fuentealba P, Suen WY, Papp E, Kaneko T, Watanabe M, Csicsvari J, Somogyi P (2005) Complementary roles of cholecystokinin- and parvalbumin-expressing GABAergic neurons in hippocampal network oscillations. *J Neurosci* **25**: 9782-9793.

Klausberger T, Márton LF, Baude A, Roberts JD, Magill PJ, Somogyi P (2004) Spike timing of dendrite-targeting bistratified cells during hippocampal network oscillations in vivo. *Nat Neurosci* **7**: 41-47.

Kunec S, Hasselmo ME, Kopell N (2005) Encoding and retrieval in the CA3 region of the hippocampus: a model of theta-phase separation. *J Neurophysiol* **94**: 70-82.

Kwag J, Paulsen O (2009) The timing of external input controls the sign of plasticity at local synapses. *Nat Neurosci* **12**: 1219-1221.

Lasztóczi B, Klausberger T (2014) Layer-specific GABAergic control of distinct gamma oscillations in the CA1 hippocampus. *Neuron* **81**: 1126-1139.

Lever C, Burton S, Jeewajee A, Wills TJ, Cacucci F, Burgess N, O'Keefe J (2010) Environmental novelty elicits a later theta phase of firing in CA1 but not subiculum. *Hippocampus* **20**: 229-234.

Manns JR, Zilli EA, Ong KC, Hasselmo ME, Eichenbaum H (2007) Hippocampal CA1 spiking during encoding and retrieval: relation to theta phase. *Neurobiol Learn Mem* **87**: 9-20.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* **417**: 741-746.

Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* **25**: 707-715.

Mizuseki K, Sirota A, Pastalkova E, Buzsáki G (2009) Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64**: 267-280.

Nagel G, Szellas T, Huhn W, Kateriya S, Adeishvili N, Berthold P, Ollig D, Hegemann P, Bamberg E (2003) Channelrhodopsin-2, a directly light-gated cation-selective membrane channel. *Proc Natl Acad Sci U S A* **100**: 13940-13945.

Newman EL, Gillet SN, Climer JR, Hasselmo ME (2013) Cholinergic blockade reduces theta-gamma phase amplitude coupling and speed modulation of theta frequency consistent with behavioral effects on encoding. *J Neurosci* **33**: 19635-19646.

Newman JP, Zeller-Townson R, Fong MF, Arcot Desai S, Gross RE, Potter SM (2012) Closed-loop, multichannel experimentation using the open-source NeuroRighter electrophysiology platform. *Front Neural Circuits* **6**:98.

Ngo HV, Martinetz T, Born J, Mölle M (2013) Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron* **78**: 545-553.

Nguyen DP, Layton SP, Hale G, Gomperts SN, Davidson TJ, Kloosterman F, Wilson MA (2009) Micro-drive array for chronic in vivo recording: tetrode assembly. *J Vis Exp.*

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

O'Neill PK, Gordon JA, Sigurdsson T (2013) Theta oscillations in the medial prefrontal cortex are modulated by spatial working memory and synchronize with the hippocampus through its ventral subregion. *J Neurosci* **33**: 14211-14224.

Packard MG, Hirsh R, White NM (1989) Differential effects of fornix and caudate nucleus lesions on two radial maze tasks: evidence for multiple memory systems. *J Neurosci* **9**: 1465-1472.

Pawelzik H, Hughes DI, Thomson AM (2002) Physiological and morphological diversity of immunocytochemically defined parvalbumin- and cholecystokinin-positive interneurones in CA1 of the adult rat hippocampus. *J Comp Neurol* **443**: 346-367.

Paz JT, Davidson TJ, Frechette ES, Delord B, Parada I, Peng K, Deisseroth K, Huguenard JR (2013) Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat Neurosci* **16**: 64-70.

Rolston JD, Gross RE, Potter SM (2010) Closed-loop, open-source electrophysiology. *Front Neurosci* **4**: 31.

Royer S, Zemelman BV, Losonczy A, Kim J, Chance F, Magee JC, Buzsáki G (2012) Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat Neurosci* **15**: 769-775.

Schmidt B, Hinman JR, Jacobson TK, Szkudlarek E, Argraves M, Escabí MA, Markus EJ (2013) Dissociation between dorsal and ventral hippocampal theta oscillations during decision-making. *J Neurosci* **33**: 6212-6224.

Schmidt M, Lipson H (2009) Distilling free-form natural laws from experimental data. *Science* **324**: 81-85.

Sigurdsson T, Stark KL, Karayiorgou M, Gogos JA, Gordon JA (2010) Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. *Nature* **464**: 763-767.

Stark E, Eichler R, Roux L, Fujisawa S, Rotstein H, Buzsáki G (2013) Inhibition-induced theta resonance in cortical circuits. *Neuron* **80**: 1263-1276.

Swanson LW (1981) A direct projection from Ammon's horn to prefrontal cortex in the rat. *Brain Res* **217**: 150-154.

Vanderwolf CH (1969) Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol* **26**: 407-418.

Vann SD, Aggleton JP (2004) Testing the importance of the retrosplenial guidance system: effects of different sized retrosplenial cortex lesions on heading direction and spatial working memory. *Behav Brain Res* **155**: 97-108.

Vann SD, Aggleton JP, Maguire EA (2009) What does the retrosplenial cortex do? *Nat Rev Neurosci* **10**: 792-802.

Voigts J, Siegle JH, Pritchett DL, Moore CI (2013) The flexDrive: An ultra-light implant for optical control and highly parallel chronic recording of neuronal ensembles in freely moving mice. *Front Sys Neurosci* **7**: 8.

Wallach A (2013) The response clamp: functional characterization of neural systems using closed-loop control. *Front Neural Circuits* **7**: 5.

Yamamoto J, Suh J, Takeuchi D, Tonegawa S (2014) Successful execution of working memory linked to synchronized high-frequency gamma oscillations. *Cell* **157**: 845-857.

Chapter 4: Open Ephys: A flexible and affordable platform for high-channel-count extracellular electrophysiology[3]

4.1 Abstract

Open Ephys is the first widely adopted, open-source, community-driven platform for multichannel extracellular electrophysiology based on Intan amplifier chips. It represents a complete data acquisition solution, with a feature set comparable to that of leading commercial systems. Our headstages use a single chip to amplify, filter, and digitize neural signals, making the hardware compact and affordable. Our acquisition board is capable of interfacing with as many as 8 such chips, for a total of 256 simultaneously recorded channels. Our software includes functionality for visualizing, recording, and analyzing both local field potentials and spike waveforms. The Open Ephys platform was designed to make setting up new experiments as convenient as possible, with hardware built around open standards and a software architecture based on modular processing units. We hope the flexibility, transparency, and low cost of our system will make it appeal to researchers who might otherwise be skeptical of relying on open-source tools.

[3]The findings in this chapter were in preparation at the time of thesis submission. (Siegle, Voigts, et al., *in preparation*)

4.2 Introduction

Extracellular electrodes have provided countless insights into the workings of the nervous system, from the discovery of orientation tuning in the visual cortex (Hubel, 1957; Hubel, 1959) to the characterization of spatial selectivity in the hippocampus (O'Keefe and Dostrovsky, 1971; O'Keefe, 1976). As the field of systems neuroscience shifts its focus from individual cells to the interactions of distributed networks, the need to record neural activity with single-spike, single-neuron resolution across hundreds or thousands of electrodes will become increasingly vital (Stevenson and Kording, 2011; Marblestone et al., 2013).

Currently, there are many impediments to setting up extracellular electrophysiology experiments. For one, the commercial recording systems that dominate the field are often prohibitively expensive for labs that do not specialize in electrophysiology, and these systems cannot scale affordably to higher channel counts. Their closed nature means that the component parts of different commercial systems are not interoperable, and changes to their functionality cannot be implemented without the assistance of the original vendor. Therefore, hardware and software innovations made by individual researchers are unlikely to propagate to the larger community, resulting in an exorbitant amount of redundant effort across laboratories.

We have developed a platform for multichannel extracellular electrophysiology that aims to overcome these barriers. Our designs are completely open source, meaning anyone is free to build the system or update its functionality to suit their needs. Our system was designed from the ground up with modularity in mind, with open interface standards adopted wherever possible. The hardware is easy to mass-produce, making it easier for those without a background in engineering to obtain it. In order to highlight the accessibility and transparency of our platform, we have chosen to name it "Open Ephys."

The main difference between Open Ephys and previous open-source recording systems is our use of industry-standard Intan amplifier chips (Harrison and Charles, 2003; Harrison, 2008). These chips, which have now been adopted by most of the major commercial vendors, allowed us to create a compact, affordable system without compromising on data quality. On the software side, we chose to build our application around a cross-platform, plugin-based architecture, in

order to simplify the process of extending its functionality. Throughout the design process, we strove to maximize flexibility for the end-user without sacrificing polish.

The target application of Open Ephys is multichannel recordings from extracellular tetrodes or silicon probes. Our hardware has been optimized for use in mammalian model systems (mice, rats, and nonhuman primates). However, it has already been successfully adapted for use with other organisms, such as zebra finches. The Open Ephys data processing pipeline is general enough that it can be applied to any experiment that requires the digitization of biopotentials ranging from tens to thousands of microvolts.

In this paper, we describe the components of the Open Ephys data acquisition system, which provides a complete, open-source alternative to commercial recording platforms. Our system includes headstages for digitizing neural signals, an acquisition board for interfacing with a computer, software for real-time visualization and analysis, and I/O boards for synchronizing with external devices (Figure 1). A single acquisition board can record up to 256 channels simultaneously (via eight Intan chips), and multiple systems can be daisy-chained for higher channels counts. Since our tools are "plug-and-play," their low cost and high flexibility should make them an attractive option for those looking to upgrade or extend the recording capabilities of their laboratory. Our system is already being used by over 30 research groups around the world, with that number expected to grow substantially over the coming years.

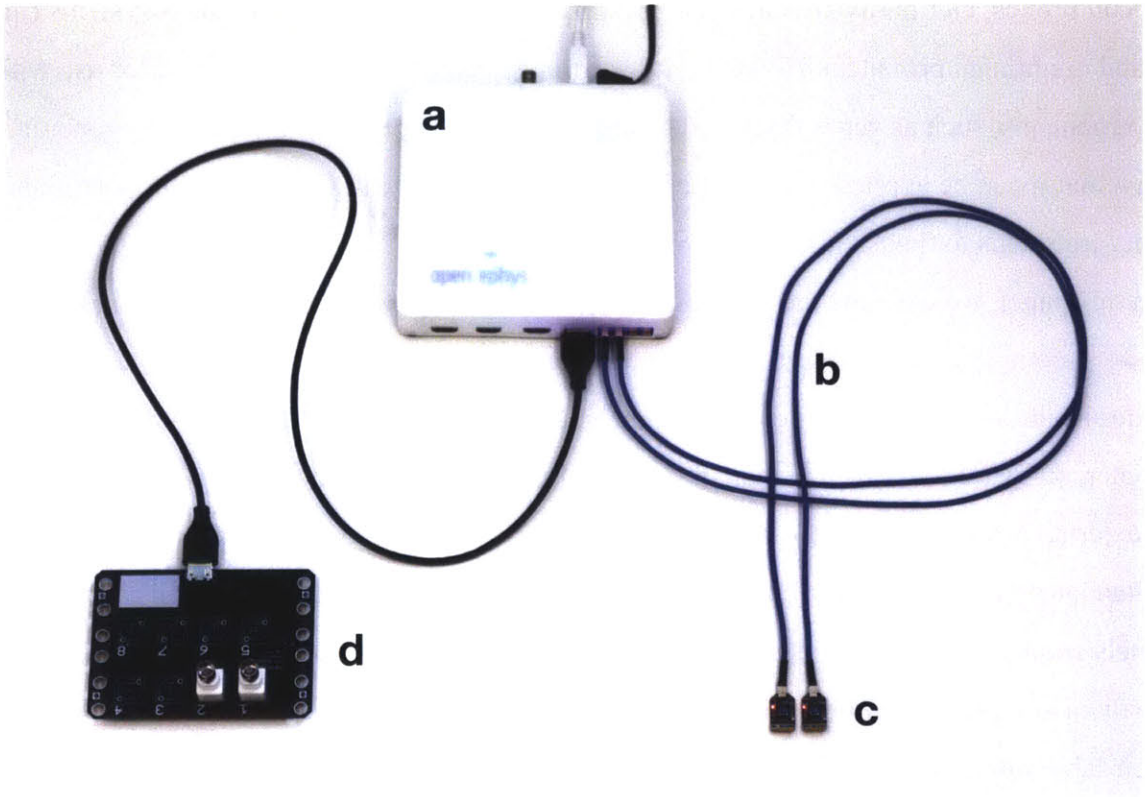

Figure 1 | Overview of the Open Ephys data acquisition platform. The basic Open Ephys hardware setup consists of four main components: the acquisition board (**a**), which sends data to a computer via USB; tethers (**b**), which connect headstages (**c**) to the acquisition board; and I/O boards (**d**), which allow peripheral devices to interface with the acquisition board via analog and digital inputs and outputs. The example setup in this image is capable of acquiring 64 channels of neural data (via two 32-channel headstages) and 8 channels of auxiliary digital input (via one I/O board). With more headstages, the acquisition board can support up to 256 channels of neural data. Open Ephys also includes free software for recording, visualization, and real-time feedback.

4.3 Results

4.3.1 Headstages and tethers

The first step in the Open Ephys data processing pipeline is the transmission of neural signals from implanted electrodes to a headstage. These signals typically consist of voltage differences between a high-impedance electrode placed in the neuropil and a low-impedance electrode embedded in the skull or cerebrospinal fluid. This configuration is capable of detecting fast voltage fluctuations caused by action potentials and lower-frequency changes in extracellular voltage known as the local field potential (LFP).

Each electrode must interface with a single input channel on the headstage. Unlike a traditional "buffer" headstage, which uses operational amplifiers to boost the current for each channel, the Open Ephys headstage transmits a digital representation of each channel's voltage as a serial data stream. This eliminates artifacts induced by movement of the tether and reduces the number of conductors needed for data transmission. We accomplish this via integrated amplifier chips from Intan Technologies.

4.3.1.1 Intan chips

Intan chips encapsulate much of the functionality of traditional data acquisition systems inside a 8 x 8 mm package (Figure 2a). Open Ephys uses Intan's RHD2132 chips, first released in 2012, which include a bank of analog filters and amplifiers for each of 32 channels (Figure 2b). The filtered and amplified signals are sent to a multiplexer, which connects them one by one to an analog-to-digital converter. Samples of each channel are represented as 16-bit integers, which are transmitted serially over the tether. The use of low-voltage differential signaling (LVDS) facilitates reliable data transmission over long, thin conductors.

The Intan chips are controlled by a standard SPI bus, which can be used to change the settings of the analog filters (low cut between 0.1 and 500 Hz, high cut between 100 and 20,000 Hz), the status of an on-chip digital high-pass filter for amplifier offset removal (on or off), and the rate of digitization (up to 30 kHz per channel). The chip also includes 3 auxiliary analog inputs for external sensors and integrated impedance measurement circuitry for all 32 channels. As of early 2014, the price of each chip is $390 or less, depending on the quantity ordered.

We characterized the noise levels and frequency response of the Intan chips in the context of the Open Ephys system. Although the Intan chips are listed as having a root-mean-squared (RMS) noise of 2.4 µV, RMS noise levels are measured at 8.2 µV in a real experimental setup (Figure 2c). The internal filters on the Intan chips have a minimum low-cut of 0.1 Hz and a maximum high-cut of 20 kHz, although in practice the rolloff starts around 1 Hz and 5000 Hz, respectively (Figure 2d).

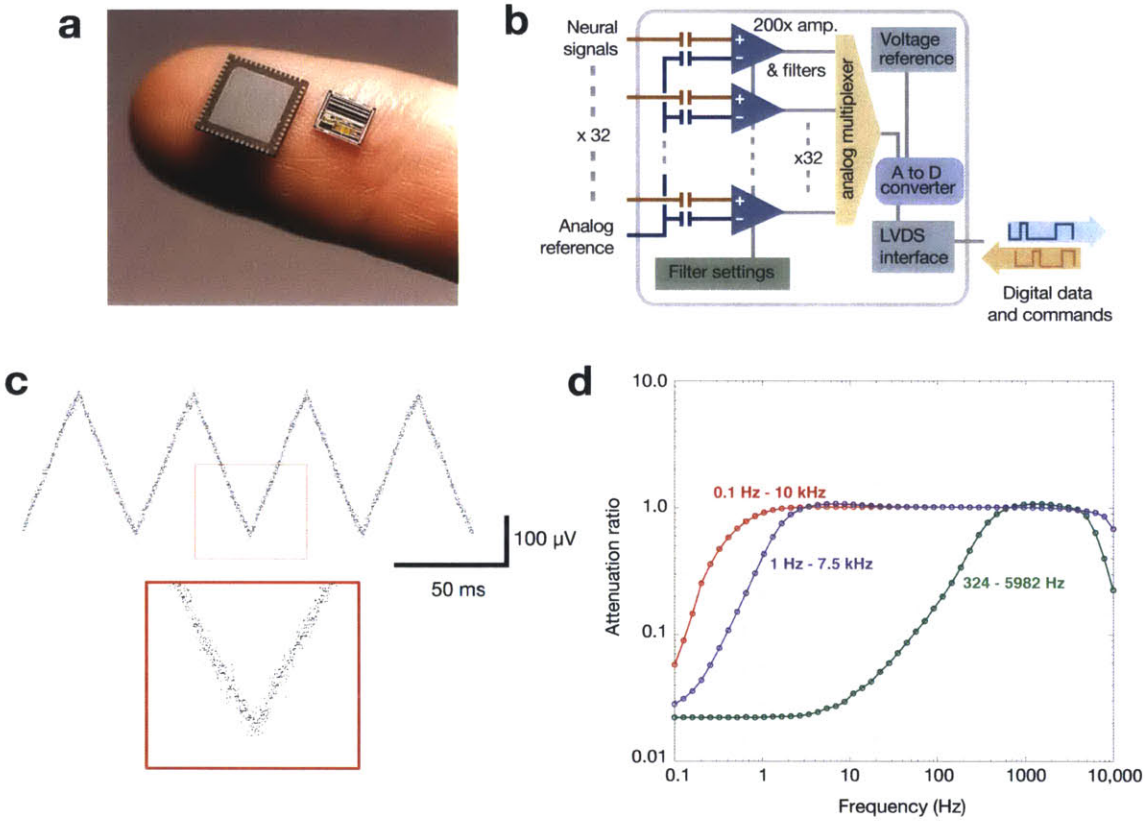

Figure 2 | Amplifier chips from Intan Technologies. a, Photograph of the Intan RHD2132 32-channel digital-output amplifier chip in a QFN package (left) and a bare die (right). Photo credit, Intan Technologies (www.intantech.com/products.html). **b**, Schematic of an Intan chip showing 32 channels of neural signals on the input side (left) and digital data on the output side (left). Inside, each channel is filtered, amplified, multiplexed, and converted to a 16-bit digital sample at a rate of up to 30 kHz per channel. **c**, Analog noise visualized for a 20 Hz, 300 µV triangle wave sampled at 30 kHz. **d**, Frequency response of the Intan chip for three different internal filter settings: 0.1 Hz – 10 kHz (completely open), 1 Hz – 7.5 kHz (default settings), and 324 – 5982 Hz (spike band). Attention ratio was measured for a sine wave with an amplitude of 1000 µV.

4.3.1.2 Headstages

Headstages compatible with the Open Ephys system consist of one or more Intan chips, an "electrode-facing" connector, and a "tether-facing" connector. The electrode-facing connector must include one conductor for each electrode, which is electrically continuous with one of the inputs on the Intan chip. The tether-facing connector must be a 12-pin Omnetics connector (product #A79623-001) that interfaces with tethers conforming to the Intan SPI cable standard.

Most of the existing headstage designs (either from Open Ephys or Intan) use 16 or 32-channel Omnetics connectors with 0.025" pitch on the electrode-facing end (Figure 3a-c,e). These connectors are already widely adopted in neuroscience, due to the reliability of their connection over a high number of mating cycles. Because they use the same reference and ground configuration, these headstages can be immediately swapped in for headstages made by Plexon, Neuralynx, Blackrock, Ripple, and Triangle Biosystems. The headstage designs are open source, which makes it is possible to create versions that interface with other types of connectors. We already have a headstage that includes a 0.05" pitch Samtec connector, used by many laboratories performing tetrode recordings in rats.

The Open Ephys headstage includes an onboard 3-axis accelerometer, which interfaces with the auxiliary analog inputs of the Intan chip. These signals can be used to monitor overall activity levels of freely moving subjects. Our headstage also include a red indicator LED, to facilitate tracking with an overhead camera.

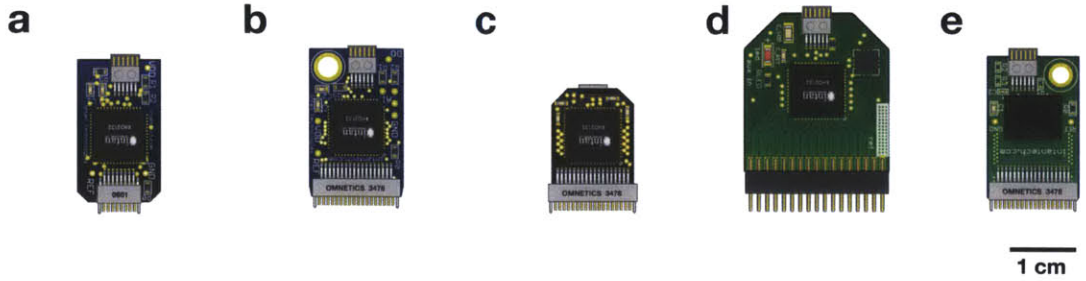

Figure 3 | Examples of headstages compatible with the Open Ephys acquisition board. a, 16-channel Omnetics headstage. **b**, 32-channel Omnetics headstage. **c**, 32-channel Omnetics "mini" headstage, with 3-axis accelerometer and LED. **d**, 32-channel Samtec headstage, with 3-axis accelerometer. **e**, 2x32-channel Omnetics headstage (64 channels total). Headstages a, b, and e can be purchased from Intan Technologies (as of 4/19/2014), and all designs are available on GitHub (github.com/open-ephys/headstage).

4.3.1.3 Tethers

Open Ephys uses tethers that conform to the SPI cable specification developed by Intan Technologies. This specification allows us to send high channel counts over only 12 conductors (Figure 4a). Tethers can be purchased from Intan in "standard" (Figure 4b) or "ultra-thin" (Figure 4c) thicknesses. Standard tethers have a diameter of 2.9 mm and a mass of 8.20 g/m. Ultra-thin tethers have a diameter of 1.8 mm and a mass of 4.07 g/m. We recommend using the ultra-thin tethers for experiments involving freely moving animals.

The tethers have the same 12-pin Omnetics connector on each end, and are designed to be daisy-chained together. The use of LVDS by the Intan chips allows tethers to reach up to 10 m in length, equivalent to five 3-foot cables placed end-to-end. A simple Y-adaptor (Figure 4d) can be used to transmit data from two headstages through a single tether, further reducing the number of necessary wires.

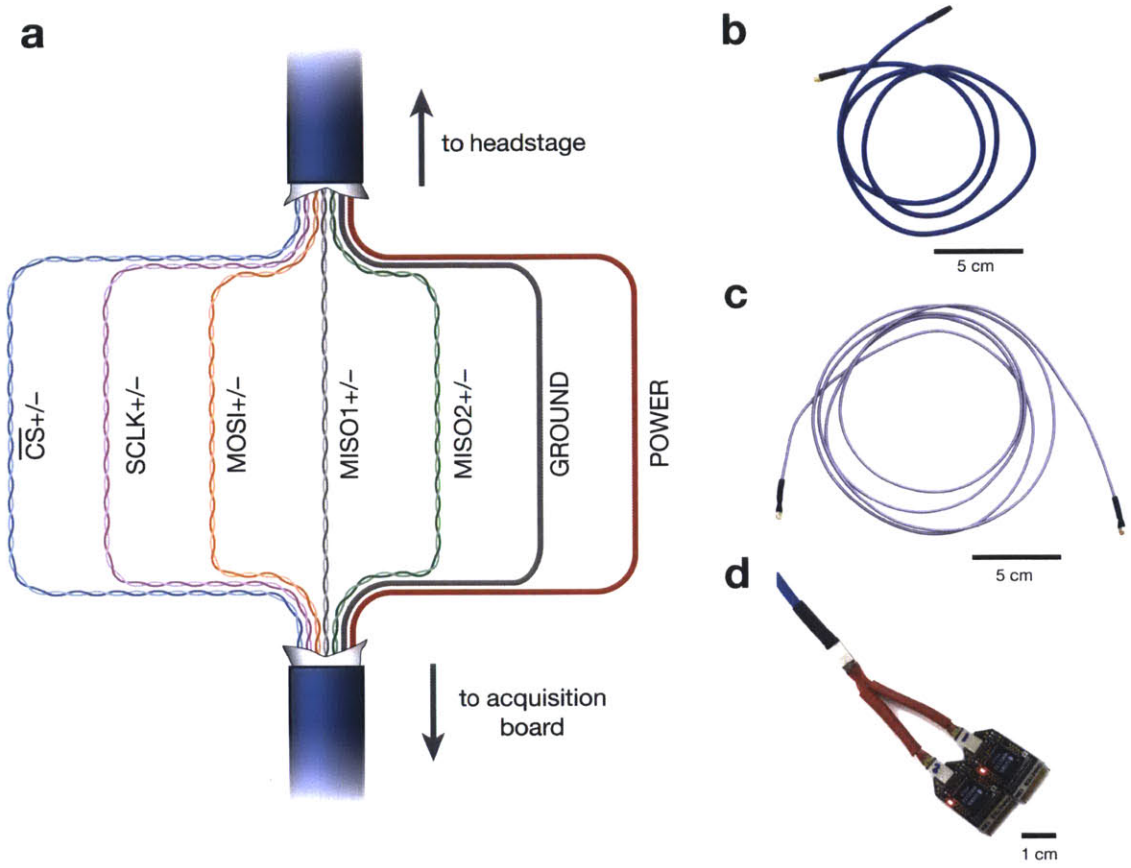

Figure 4 | Tethers used with the Open Ephys system. a, Tethers connecting headstages and acquisition boards are capable of transmitting up to 128 channels of data using only 12 conductors (see the RHD2000 SPI cable/connector specification at intantech.com/downloads.html for a complete specification). Briefly, one CS+/– twisted pair selects the Intan chip to communicate with, one SCLK+/– twisted pair provides the clock signal for data input and output, one MOSI+/– twisted pair sends control signals to the Intan chips, and two MISO+/– twisted pairs enable neural data to stream simultaneously from two headstages. The use of both MISO pairs is optional, but in combination they can transmit data from two 32-channel or two 64-channel channels Intan chips (maximum of 128 channels). Twisted-pairs allow data to be sent using low-voltage differential signaling (LVDS), a standard that greatly increases transfer rates compared to single-wire transmission. In addition to the five LVDS pairs, one power and one ground line provide the necessary supply voltage for the Intan chip to run. **b**, Standard 3-foot rubbed-coated tether, available for purchase from Intan Technologies. **c**, Ultra-thin 6-foot tether made from Cooner wire, available for purchase from Intan Technologies. **d**, A custom Y-adapter (red) allows two headstages to interface with one tether.

4.3.2 Acquisition board

The Open Ephys acquisition provides an interface between up to 8 headstages and a computer. Its main function is to send control signals to each of the headstages, buffer the incoming data, and stream it to the computer via USB 2.0. Because most of the heavy lifting is accomplished by the Intan chips on each headstage, the acquisition board is more compact than traditional recording hardware. Data transfer is orchestrated by a field-programmable gate array (FPGA) running a custom variant of Intan's Rhythm firmware. We chose to adopt the Rhythm standard so that our acquisition board could use the same firmware as the RHD2000 Evaluation System sold by Intan Technologies. If Intan adds new features to its chips, these will be easy to integrate into our platform.

4.3.2.1 Inputs and outputs

The acquisition board has the following ports for input and output (Figure 5):

- **Four 12-pin Omnetics connectors for Intan SPI cables**, each capable of supporting two headstages. These ports send power lines and control signals to the headstages, while receiving two lines of serial data (one per Intan chip). All eight headstages can be run simultaneously at 30 kHz.
- **One HDMI connector for digital inputs.** This port allows digital devices with 5V logic levels to communicate with the acquisition board. The 8 digital input lines are sampled at the same rate as the neural data.
- **One HDMI connector for digital outputs.** This port allows the acquisition board to send digital signals to devices with 5V logic levels. The status of the 8 digital output lines can be updated via software, or directly from the FPGA.
- **One HDMI connector for analog inputs.** This port allows the acquisition board to acquire analog signals ranging between 0-5V or ±5V. The range of each of the 8 analog input channels is set via hardware jumpers inside the acquisition board. The analog inputs are sampled at the same rate as the neural data.

- **One HDMI connector for analog outputs.** This port allows signals in the ±5V range to be generated by the acquisition board and sent to external devices. The first two analog outputs (out of 8) can be programmed to follow any of the headstage channels, creating a replica of the amplified signal with a delay of 1 sample. These signals can be accessed through the built-in audio jack to create a low-latency audio monitor of any of the neural data channels.

The digital inputs, digital outputs, analog inputs, and analog outputs must be used in conjunction with Open Ephys I/O boards (Figure 1). Signals are transmitted between the I/O boards and the acquisition boards via standard passive HDMI cables. We chose these cables because they provide 8 shielded conductors capable of transmitting digital and analog signals bidirectionally and are readily available at a low price point.

The parallel incoming data streams are synchronized by the FPGA, which is part of an Opal Kelly XEM6010-LX45 development board. Opal Kelly provides a convenient interface for software-to-FPGA communication. The data transmission protocol can be upgraded to USB 3.0 or PCI-express by swapping in a different development board (Opal Kelly XEM6310 or XEM6110v2) and making minor changes to the firmware.

4.3.2.2 Additional features

In addition to acting as a data pipe, the acquisition board includes several features to simplify the process of setting up new experiments:

- **"Daisy chain" connector for synchronizing multiple acquisition boards.** If an experiment requires more channels than can be acquired by a single acquisition board, multiple boards can be synchronized via this port. One board will act as the "master," controlling the timestamps of the other "slave" boards. Each board will interface with its own computer, but all of the recorded data will be referenced to the same hardware clock. This allows channel counts to be scaled up indefinitely.
- **Isolation boards to prevent ground loops.** When recording with multiple headstages, the same ground and power lines travel along all the tethers in parallel. This can lead to increased

noise as a result of ground loops. By inserting an isolation board between the acquisition board and the headstage, each headstage can have its own independent power supply.

- **TTL-triggered amplifier reset.** The acquisition board can be programmed to trigger a hardware reset of the headstage amplifiers in response to a 5V signal on any of the digital input lines. Experiments that involve electrical stimulation can take advantage of this feature in order to avoid saturating the amplifiers in response to currents near the electrode.

The acquisition board also includes a "prototyping area" for testing new hardware features. There are 18 unused input/output pins on the FPGA, which can be configured for a variety of applications. Although we have included features that will satisfy the majority of users, we wanted to make it as simple as possible for others to add custom functionality.

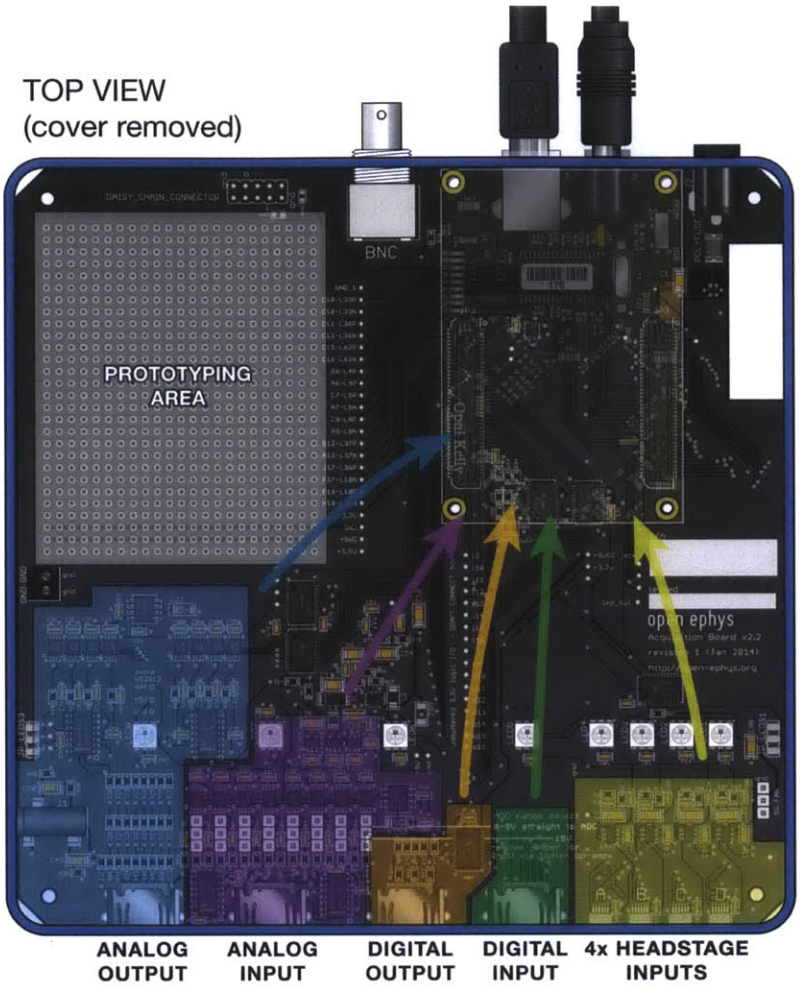

Figure 5 | The Open Ephys acquisition board. Illustration of the inside of the acquisition board, with important regions overlaid. The purpose of the board is to guide digital signals to and from an Opal Kelly XEM6010 FPGA (upper right), which transmits data to the host computer via USB. Analog outputs (blue) generate ±5V waveforms and provide audio monitoring of neural signals with sub-millisecond latency. Analog inputs (purple) digitize continuous data in the ±5V or 0–5V range, synchronized to the headstage inputs. Digital outputs (orange) generate binary signals for delivering closed-loop feedback. Digital inputs (green) can be used to capture synchronization pulses from other devices. All auxiliary inputs and outputs use standard HDMI cables to interface with Open Ephys I/O boards (Fig. 1). Headstage inputs (yellow) interface with RHD2000 SPI cables (Fig. 4) connected to compatible headstages (Fig. 3). Design files and assembly information are available on GitHub (github.com/open-ephys/acquisition-board).

4.3.3 Software

The Open Ephys Graphical User Interface (GUI) enables data from the acquisition board to be recorded and visualized. It is also capable of performing real-time data analysis, in order to trigger feedback based on neural events. The GUI's modular architecture simplifies the process of configuring new experiments, while making it easier to extend its functionality with custom code. It is written in C++ using the JUCE open-source toolkit, and it runs on Windows, Mac OS X, and Linux. The complete source code is available on GitHub (https://github.com/open-ephys/GUI). Precompiled binaries can be downloaded from the Open Ephys website (http://open-ephys.org).

4.3.3.1 General features

The GUI was designed to facilitate access to all controls and visualizations necessary for running and monitoring an experiment. If desired, the entire application can be run inside a single window, making it straightforward to use on computers with smaller displays. If more screen real estate is available, however, visualizations can be expanded into separate windows. The user interface consists of four main components (Figure 6):

- **The control panel**, at the top of the main application window, holds buttons to toggle acquisition and recording, a clock to track the amount of time spent acquiring data, controls for audio output and recording parameters, and displays that indicate CPU usage and available disk space.
- **The processor list**, at the left of the main application window, contains a list of data processing modules available for constructing the signal chain. These are organized as "Sources," "Filters," "Sinks," and "Utilities."
- **The editor viewport**, at the bottom of the main application window, provides access to the parameter editors for every processing module in the signal chain. Here, the user can control which channels are recorded and monitored, as well as alter custom settings for individual modules.

- **The data viewport**, which fills the remaining space in the main application window, holds tabs for visualizers, as well as a graph that displays the layout of the signal chain at a glance.

For most electrophysiology experiments, the user will need to visualize the continuous signals from all channels in order to assess electrode placement and recording quality. The GUI includes an "LFP Viewer" module that displays signals in real time (Figure 6). The LFP Viewer streams data from left to right, like an oscilloscope, for windows of 1-10 seconds. Events are overlaid as translucent color bars.

Many experiments also require the detection and display of spike waveforms in real time. The GUI separates this functionality across three modules: a Filter Node, a Spike Detector, and a Spike Viewer. The Filter Node streams the incoming data through a bandpass filter; the default settings are 300 Hz for the low cut and 6000 Hz for the high cut. This allows the data to be thresholded by the Spike Detector, in order to extract the waveforms of candidate spikes. The Spike Detector can detect spikes on single electrodes, stereotrodes, or tetrodes. These spikes are sent as events (parallel to the continuous data stream) to the Spike Viewer, which displays waveforms and peak-height projection plots for stereotrodes and tetrodes.

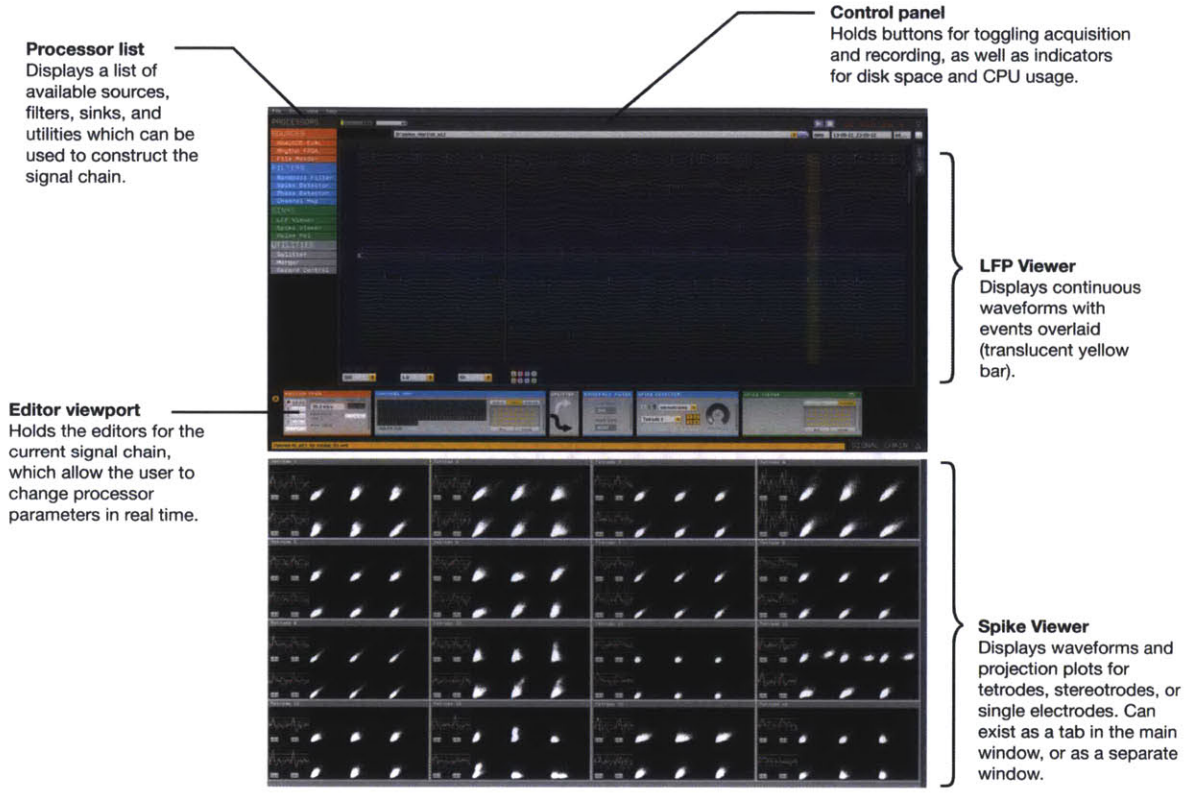

Figure 6 | Layout of the Open Ephys graphical user interface (GUI). The major components of the software are labeled. The layout is very flexible; the Processor List and Editor Viewport can be hidden when not in use, and any visualizers can either inhabit a tab within the main window or their own floating window. The software runs on Windows, Linux, and Mac OS X. Pre-compiled executable files are available at open-ephys.org, and the source code can be viewed on GitHub (github.com/open-ephys/GUI).

4.3.3.2 Plugin architecture

We chose to build the GUI around a collection of processing modules, rather than a hard-coded signal chain, in order to improve flexibility while encouraging code reuse. This "plugin" architecture was inspired by audio processing software such as Ableton Live (DeSantis et al., 2013), which makes it easy to load custom instruments and filters while the application is running. In the same vein, the GUI treats data sources as "instruments" that bring new data into the signal chain. Each subsequent step in the processing pipeline modifies the incoming data stream or sends it to an external device (such as a stimulus generator or a display). In order to use the same signal chain with a different data stream, one can remove the current "source" processor and replace it with another source from the processor list. The GUI is intended to be used with any source of neural data, not just the Open Ephys acquisition board.

We hope that the GUI's plugin architecture will encourage code reuse. Unnecessarily rewritten code is omnipresent in neuroscience, and is a huge time sink for researchers everywhere. We carefully considered the minimum functionality required for each module, to ensure that it is specialized enough to be useful on its own, yet general enough to fit into a variety of signal chains. For example, we chose to place the three steps involved in displaying spikes inside their own modules. As a result, the Filter Node can be used as a general bandpass filter with configurable cutoffs, the Spike Detector can be used to detect spikes from signals that have been filtered in hardware, and the Spike Viewer can be used to display spikes that originate on a separate machine and are streamed over a network. A user might wish to swap out the Spike Detector with a custom module for detecting spikes on high-density silicon probes (which may not be amenable to static channel groups), without having to reimplement functions for filtering or spike display.

Adding a new processing module is as simple as creating a new C++ class (Figure 7). At a bare minimum, the new module need only implement the "process" method, which determines how it reacts to incoming data. Developers can also create classes that specify the module's editor (visible in the editor viewport) or a visualizer. Example classes are available online, and a tutorial on our wiki describes the process of creating new modules in detail.

The GUI is available as both as a stable, standalone application and an experimental host/plugin architecture. We hope to soon transition entirely to the host/plugin version, which will make it easier for researchers to share new modules without the need to recompile code.

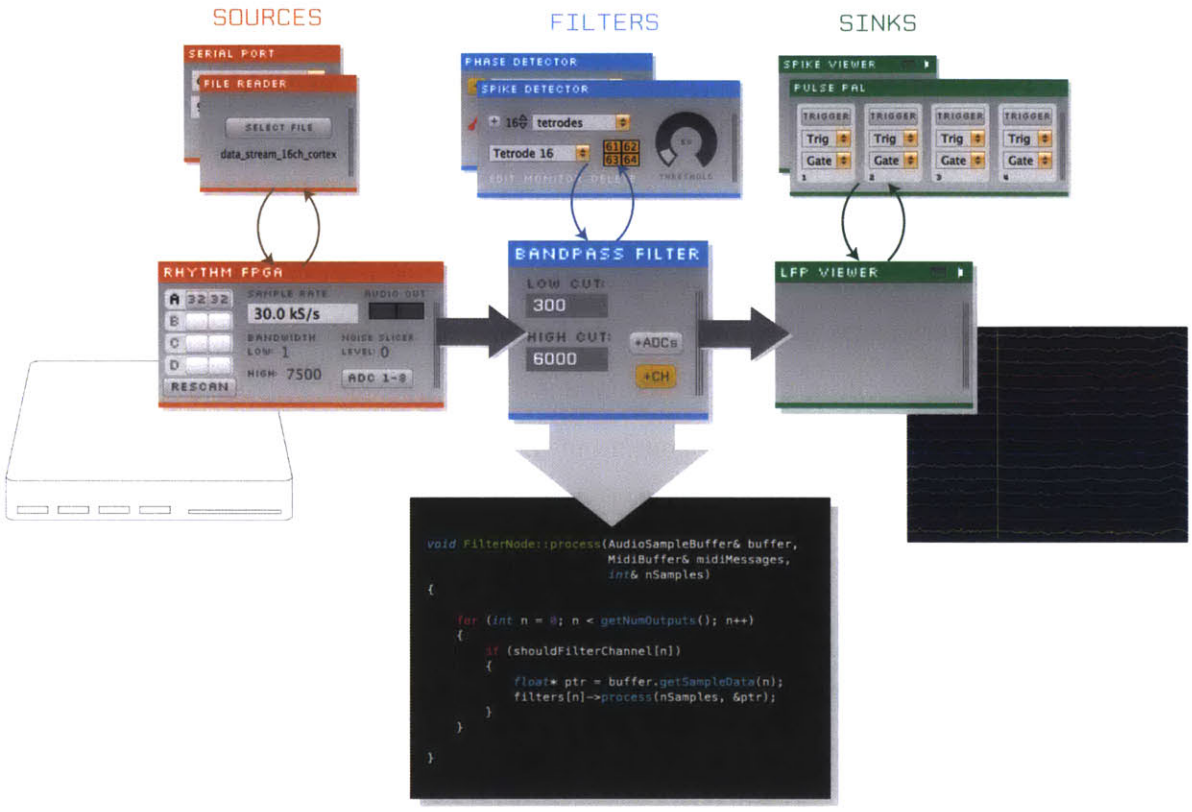

Figure 7 | The GUI's modular architecture makes it easy to add new functionality. Sources (left), filters, middle, and sinks (left) can be mixed and matched to change the GUI's processing pipeline on the fly. In this example, a "Rhythm FPGA" module communicates with the Open Ephys acquisition board to bring 64 channels of neural data into the processing pipeline. A "Bandpass Filter" module filters the data between 300 and 6000 Hz. Finally, an "LFP Viewer" module send the data stream to the display for visualization. The example code shows the Bandpass Filter's "process" method, written in C++, which determines how the module responds to or transforms the incoming data stream. All modules must implement this method, at a minimum, in order to be integrated into the signal processing pipeline.

4.3.3.3 Closed-loop feedback

The ability to trigger stimulation based on brain states opens up a large class experiments that are not accessible to the typical neuroscientist. These include protocols for implementing brain machine interfaces (Schwartz, 2004; Hatsopoulos and Donoghue, 2009), adaptive sampling (Paninski, 2005; Benda et al., 2007), and entrainment or disruption of intrinsic oscillations (Berényi et al., 2012; Ngo et al., 2013; Paz et al., 2013; Buetfering et al., 2014). Making closed-loop experiments feasible for a wider range of researchers was one of the primary goals of developing the Open Ephys GUI. Although implementing closed-loop algorithms in software involves inherent and uncertain delays (~20 ms), these latencies are acceptable for a wide range of applications. Having real-time access to data in software also makes it easier to prototype feedback algorithms that can later be transferred to hardware. And for certain types of experiments, such as real-time decoding of spike trains (Lawhern et al., 2010; Kloosterman et al., 2014), hardware implementations are impractical.

The GUI's real-time feedback engine is built on top of JUCE's audio processing library, which can handle complex floating point data processing steps in real time. The requirements for audio processing—in terms of sample rate, bit depth, and channel count—are in a similar range as those for neural data. We were therefore able to use the existing classes with minimal modification.

A basic closed-loop feedback experiment with the GUI requires an input module, an event-detector module, and an output module (Figure 8). The input module receives signals from a data source, such as the Open Ephys acquisition board, and transmits these to the next step in the processing pipeline. The event-detector searches the incoming data stream for specific patterns of activity. These can be as simple as a threshold-crossing or a peak, or as complex as a decoded spike train. The events from this processor are received by an output module, which communicates with some external piece of hardware. We favor using a Pulse Pal, an open-source stimulus generator developed by Josh Sanders at Cold Spring Harbor Laboratory (https://sites.google.com/site/pulsepalwiki). The Pulse Pal triggers a pulse of light from an LED, which activates light-sensitive ion channels in the brain.

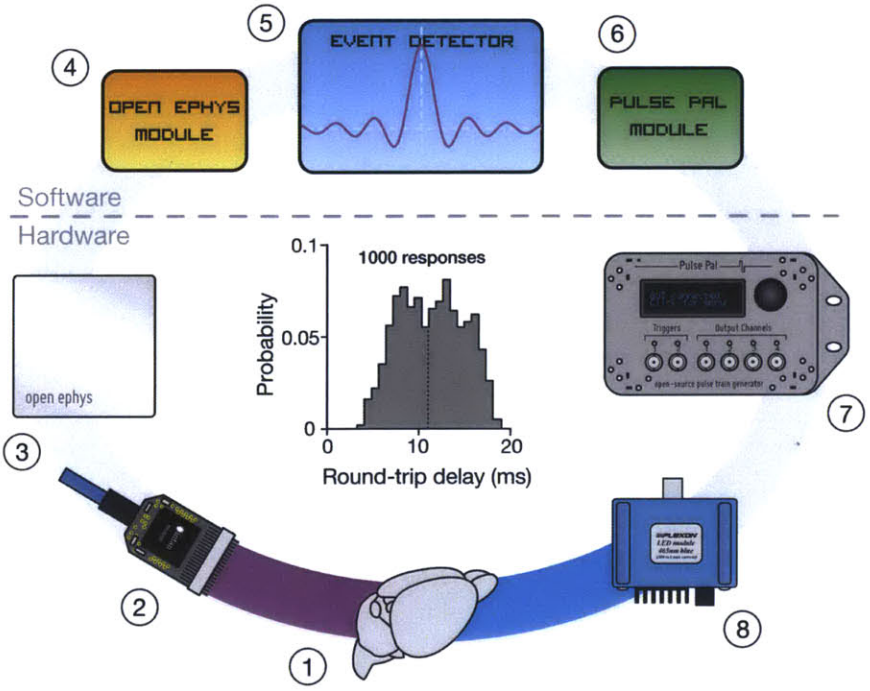

Figure 8 | Closed-loop feedback with the Open Ephys platform. One of the key advantages of the Open Ephys acquisition system is its prioritization of closed-loop feedback. Open Ephys software makes it easy to mix and match modules for acquiring data, detecting events, and sending triggers to external devices. This figure depicts one example of a configuration used for closed-loop experiments. The histogram in the middle shows the amount of time elapsed between an event occurring and the feedback being delivered (measured for 1000 events). The elements of the loop are as follows: **1.** Electrical potentials are generated by the brain and detected by an array of electrodes. **2.** Analog electrical signals are digitized by an Open Ephys headstage and sent to the acquisition board. **3.** The acquisition board packages samples into a buffer, synchronizes them with auxiliary digital and analog inputs, and sends them to a computer via USB. **4.** A dedicated module in our software unpacks the incoming data, converts it to floating-point values, and places it into a separate buffer for analysis. **5.** Another software module analyzes the incoming data stream and looks for specific types of events. These can be simple, first-order characteristics of a continuous signal, such as phase or amplitude, or more complex features that are matched to a template. Other modules could be created to analyze the attributes of spike activity, such as multiunit firing rate or decoded sequences. When the event is detected, this module generates a signal that's passed to all the modules farther down the signal chain. **6.** In the final step in the software processing, we have a module that communicates with a Pulse Pal, an open-source stimulator designed by Josh Sanders in the Kepecs Lab at CSHL. This module receives the signal sent by the event generator and triggers the Pulse Pal via USB. **7.** The

Pulse Pal is pre-programmed to generate pulses of a precise frequency and duration. Any of its four output channels can be triggered independently via software. **8.** Lastly, a PlexBright LED (Plexon, Inc., Dallas, TX) activated by the Pulse Pal, generating a brief pulse of light that activates channelrhodopsin in the brain. At every step of the way, it's possible to swap out the hardware or software modules for ones with different functionality. For example, someone proficient in C++ could write a module in Step 4 that communicates with National Instruments hardware. The module in Step 6 could be programmed to deliver feedback via an Arduino. Feedback can also come directly from the Open Ephys acquisition board. For example, the hardware in Step 8 could be a laser, a current source, or a mechanical actuator.

4.3.4 Example data

The Open Ephys system has been used to record hundreds of hours of data from a variety of model organisms. Figure 9 shows examples of raw data from mouse hippocampus (CA1 region), monkey neocortex (middle face patch), and zebra finch anterior forebrain (LMAN). Spikes and low-frequency fluctuations are clearly visible in all traces. Open Ephys is general enough to be used in any preparation that includes extracellular recording electrodes.

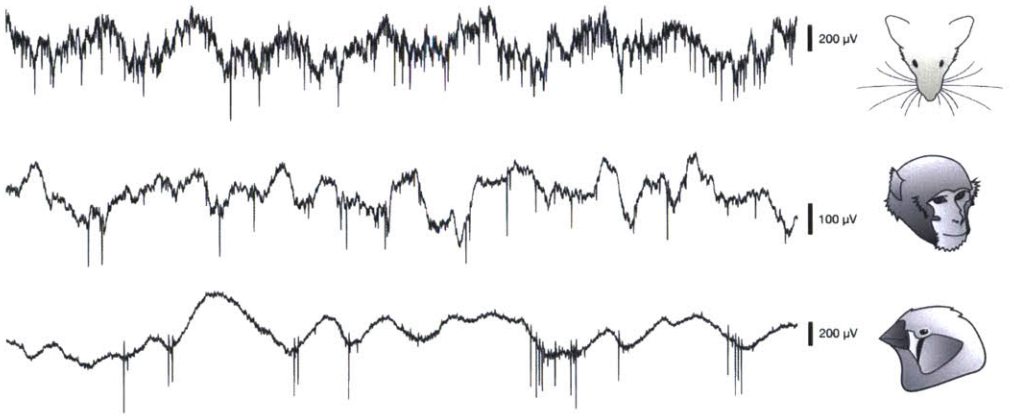

Figure 9 | Data collected with the Open Ephys platform. Three traces showing 1 s of 30 kHz raw data from three species: awake mouse hippocampus (top), awake macaque neocortex (middle), and anesthetized zebra finch LMAN (bottom). Spike waveforms (vertical lines) are clearly visible in all examples.

4.3.5 Distribution and development

One of the hurdles any open-source hardware initiative must overcome is the distribution of its tools to users. There are many examples of open-source software gaining widespread adoption in neuroscience, such as ImageJ for image processing (Abràmoff et al., 2004), NEURON for simulating neural networks (Hines and Carnevale, 1997), and ScanImage for data acquisition with two-photon microscopes (Pologruto et al., 2003). But open-source hardware has yet to catch on in the same way, for obvious reasons. Software can be distributed for free, but the infrastructure needed to manufacture and ship hardware can be costly. The need to generate revenue to support hardware distribution discourages tool developers from keeping their designs open.

Our top priority is maintaining an ecosystem built around open-source data acquisition hardware. If we can find a way to do this sustainably, Open Ephys may succeed in gaining market share in a field currently dominated by closed-source, commercial systems. We have the advantage of being able to leverage recent advances in rapid prototyping and circuit board manufacturing. It is now more cost effective than ever before to obtain high-quality hardware in small quantities. There are companies that can manufacture the custom components of the Open Ephys acquisition board in batches of one for far less than the cost of an equivalent commercial system. When the number of units increases to 100 or more, a fully functional acquisition board can be acquired for under $700 per unit (Table 1). Just as we are able to distribute the software development across multiple labs, we hope to distribute the manufacturing efforts across the broader community, by making it as simple as possible for anyone to source the hardware we've designed.

Item	Notes	Price for 5	Price for 100
Standard components	Purchased from DigiKey.com	$656	$3284
Texas Instruments	Donated through TI University partnership	$1058	$14,059
LEDs	WS2812s, purchased through Alibaba.com	$20	$130
Connectors	Omnetics connectors for SPI cables	$460	$5456
PCBs	Manufactured by Advanced Circuits	$1976	$4950
Assembly	Assembled by Advanced Circuits	$2586	$7000
FPGAs	Opal Kelly XEM6010-LX45	$2000	$24,995
Heatsinks	DigiKey part #294-1097-ND	$13	$207
Cases	3D printed (5) or cast urethane (100)	$569	$8944
Top	Laser cut by Ponoko.com	$45	$600
Screws	10 mm M3, stainless steel	$3	$52
Rubber feet	1/2" diameter	$6	$24
	TOTAL FOR ALL BOARDS	$9392	$69,701
	PRICE PER BOARD	$1878.40	$697.01

Table 1 | Manufacturing costs. Approximate costs for outsourcing board manufacturing, for quantities of 5 or 100 boards per run. This does not include the cost of purchasing headstages and cables, which adds approximately $1000 per block of 32 channels.

4.3.5.1 Obtaining the hardware

The design files for Open Ephys hardware (headstage, acquisition board, and I/O board) are available on our GitHub page (https://github.com/open-ephys). These include circuit board designs (in Eagle format), Gerber files, bills of materials, and 3D CAD files (for the acquisition board case). These files can be downloaded individually, or as part of an entire "repository" for each component.

Building the hardware from scratch involves three steps, described in more detail on our wiki (http://open-ephys.atlassian.net):

- **Manufacturing the circuit board.** There are a number of companies that can produce bare printed circuit boards based on Gerber files. A solder stencil is required as well.
- **Ordering the standard components.** Most of the parts for the acquisition board can be purchased from DigiKey, but a few components (such as connectors from Samtec and Omnetics) need to be ordered directly from the manufacturer. See Table 1 for approximate costs.
- **Populating the circuit board.** First, use the stencil to apply solder paste to the exposed metal on the circuit board. Next, match the components in the bill of materials to the reference designators printed on the board. This step may take a few hours. Finally, use a hot air station or reflow oven to solidify the solder paste.

If this seems beyond the reach of your capabilities, the manufacturing process can be outsourced. Open Ephys has partnered with CircuitHub (https://circuithub.com) to distribute acquisition boards and I/O boards as needed. The cost will be higher than building it from scratch, but the time investment will be considerably less.

4.3.5.2 Contributing to the software

The Open Ephys GUI is freely available as pre-compiled binaries from our website (http://open-ephys.org). In order to contribute new code, however, would-be developers must create a GitHub account and become familiar with the "git" software package. Git makes it much easier

for a large network of programmers to collaboratively develop an application. Although it's not necessary, we encourage all users to learn about the source code and familiarize themselves with git protocols.

Several features of GitHub and git have helped streamline the development of the Open Ephys GUI:

- **Version control.** Git tracks the changes made to the code, who made them, and when. This makes it easy to "roll back" accidental changes, and to understand what updates have taken place after each development cycle. In addition, the use of "branches" facilitates the testing of experimental features in parallel.
- **Pull requests.** Under GitHub's development model, each lab can maintain their own copy of the source code, and add changes to the main repository via "pull requests." This makes collaborative development run much more smoothly.
- **Issue tracker.** Bug reports and feature requests are centralized through GitHub's "Issues" page. All contributors can view the open issues and coordinate their development efforts.

Without these features, it would have been much harder to get the software to the point where it is today. As our tools spread to more labs, GitHub will be essential for tracking feature requests and integrating contributions from around the world.

4.4 Discussion

Open Ephys offers a flexible, affordable platform for collecting data from extracellular electrodes. Our main target audience is researchers performing tetrode and silicon probe recordings in mammalian model systems, but the hardware and software are general enough to be used for other applications. The system has been tested in a variety of experimental settings, and has been shown to be robust and reliable. As additional users adopt the system, the functionality of the software will continue to expand.

4.4.1 Advantages of Open Ephys

For labs looking to purchase a new multichannel electrophysiology system, Open Ephys offers three advantages over its closed-source commercial counterparts:

- **Low cost.** A complete Open Ephys system can be obtained for less than $100 per channel, compared to commercial systems costing $1000 per channel or more. For labs on a tight budget, Open Ephys may be the only option for setting up high-channel-count experiments. For labs that want to add recording capabilities to multiple rigs in parallel, our system is an attractive choice. As the throughput of systems neuroscience research continues to expand—both in terms of the number of simultaneously recorded channels, and the number of subjects per experiment—Open Ephys may be the best solution for growing a lab's electrophysiology resources.
- **Transparency.** Because the designs are freely available, Open Ephys encourages scientists to look "under the hood" and understand the details of its implementation. This not only makes for a better-educated user base, but also alleviates the dependency on a particular company to upgrade functionality or fix bugs. The one caveat to this is the Intan amplifier chips in the headstages, which do contain proprietary technology. But this is also a problem for the numerous commercial companies that are using Intan chips in their hardware.
- **Flexibility.** Not only is our hardware and software completely open source, but it was designed from the start with modularity in mind. The acquisition board is compatible with many types of headstages, and the software is built around processor modules that can be swapped in and out independently. Hardware and software modules designed by various labs can be made immediately accessible to the broader community, reducing the amount of time spent on redundant development. It is highly possible that, in the future, Open Ephys hardware will be widely used with a different piece of software. Or, the hardware could be made obsolete by new technologies, while the software lives on as a separate entity.

4.4.2 Drawbacks of Open Ephys

There are several drawbacks of our platform that must be taken into consideration:

- **Hardware distributed in limited quantities.** There is no guarantee that Open Ephys acquisition boards will be available fully assembled at any given time. There's no company to coordinate the distribution efforts, so the timing and location of production runs is currently unpredictable. However, we are working to mitigate this issue by partnering with CircuitHub to provide one-click ordering of individual acquisition boards.
- **Small user base.** Open Ephys hardware has only been under development for 3 years, and only became available to a wider audience in the past year. Therefore, it has not had the opportunity to grow its user base to sustainable levels. That said, there are already users at many major universities and neuroscience research institutes. Since users are connected via GitHub and the Open Ephys wiki, technical support efforts can be distributed throughout the community.
- **Developed by amateur engineers.** Open Ephys was developed by neuroscientists for neuroscientists, which means some of the design decisions may not be optimal from an engineering perspective. Anyone who chooses to use our system must accept responsibility for ensuring that everything works the way they expect it to. This should not be a substantial burden, but labs that aren't willing to put in the extra consideration may be better off buying a commercial system.

4.4.3 Comparison to other open-source systems

How does Open Ephys compare to other open-source recording platforms? The most similar is NeuroRighter, implemented by the Potter Lab at Georgia Tech (Rolston et al., 2009; Newman et al., 2012). Like Open Ephys, NeuroRighter offers open-source hardware and software optimized for multielectrode recordings and closed-loop stimulation. NeuroRighter is more mature, having been in use for over 5 years. However, there are two aspects of NeuroRighter that make it less flexible than Open Ephys: its reliance on National Instruments digitization hardware, and its use of the C# programming language, which is Windows-specific. The former also makes the system more costly: NeuroRighter costs around $10,000 for 64 channels, whereas a 64-channel Open Ephys system can be built for less than $3000. The use of Intan chips in our

headstages makes Open Ephys more compact and more affordable. Nevertheless, the success of NeuroRighter was an inspiration during the early days of developing Open Ephys, especially their commitment to making tools for delivering closed-loop feedback more accessible.

Besides NeuroRighter, there are no other published open-source multichannel electrophysiology platforms that include designs for both software and hardware. There are Matlab-based toolboxes for neurophysiology, such as PLDAPS (Eastman and Huk, 2012) and MANTA (Englitz et al., 2013), but these are intended to work with closed-source commercial hardware from Plexon, Blackrock, and TBSI. They also require a Matlab license to run. Another promising approach is to use real-time Linux to achieve lower closed-loop latencies than are possible with Open Ephys. This approach has been taken by RTXI (http://www.rtxi.org), but their main focus is on intracellular electrophysiology, rather than multichannel extracellular electrophysiology.

A key advantage of open-source tools is that it is possible to make them interoperate. If someone had the need for it, they could make Open Ephys interface with any of the platforms listed above. Our hardware is meant to be used with other software, so we have made the necessary details of the communication protocols available to all. Likewise, the Open Ephys GUI was intended to be used with a variety of data sources. Creating a module to interface with a different type of hardware is a straightforward process.

4.4.4 Conclusion

Open Ephys is a relatively young initiative, but it has already been adopted by labs around the world. Our beta testing program sent acquisition boards to researchers in the U.S., Germany, Canada, the UK, Sweden, Spain, Portugal, France, Israel, and Brazil. In our second round of manufacturing, we added China, Korea, Finland, Belgium, and the Netherlands to that list. We hope our platform continues to spread, making it easier for advanced users to customize their experiments, and easier for those new to electrophysiology to access the necessary tools.

The biggest challenge will be gaining enough momentum to turn Open Ephys into a viable alternative to commercial recording platforms. We believe something along the lines of Open Ephys is long overdue in neuroscience research, and it would be a shame if it didn't gain traction

in the field. There is currently a huge amount of development occurring inside individual labs, and most of this goes unshared and unnoticed. Neuroscientists would greatly benefit from a centralized, standardized tool base which they can iteratively improve. We believe Open Ephys is poised to fill this role for the extracellular electrophysiology research. Whether or not it succeeds will depend on the community's enthusiasm for a more collaborative tool development model.

4.5 Acknowledgments

We are grateful to R. Harrison of Intan Technologies for advice and encouragement. We thank Texas Instruments for their generous donation of components. We thank D. Meletis, M. Carlén, M. Jazayeri, J. Goldberg, O. Yizhar, D. Huber, and D. Moorman for donations. Special thanks to all beta testers for their feedback on our tools.

4.6 References

Abràmoff MD, Magalhães PJ, Ram SJ (2004) Image processing with ImageJ. *Biophotonics international* **11**: 36-43.

Benda J, Gollisch T, Machens CK, Herz AV (2007) From response to stimulus: adaptive sampling in sensory physiology. *Curr Opin Neurobiol* **17**: 430-436.

Berényi A, Belluscio M, Mao D, Buzsáki G (2012) Closed-loop control of epilepsy by transcranial electrical stimulation. *Science* **337**: 735-737.

Buetfering C, Allen K, Monyer H (2014) Parvalbumin interneurons provide grid cell-driven recurrent inhibition in the medial entorhinal cortex. *Nat Neurosci.* **17**: 710-718.

DeSantis D, Gallagher I, Haywood K, Knudsen R, Behles G, Rang J, Henke R, Slama T (2013) Ableton Reference Manual Version 9. Berlin: Ableton AG.

Eastman KM, Huk AC (2012) PLDAPS: A Hardware Architecture and Software Toolbox for Neurophysiology Requiring Complex Visual Stimuli and Online Behavioral Control. *Front Neuroinformatics* **6**:1.

Englitz B, David SV, Sorenson MD, Shamma SA (2013) MANTA—an open-source, high density electrophysiology recording suite for MATLAB. *Front Neural Circuits* **7**.

Harrison RR (2008) A versatile integrated circuit for the acquisition of biopotentials. *Custom Integrated Circuits Conference, 2007. CICC'07. IEEE*: 115-122.

Harrison RR, Charles C (2003) A low-power low-noise CMOS amplifier for neural recording applications. *IEEE Journal of Solid-State Circuits* **38**: 958-965.

Hatsopoulos NG, Donoghue JP (2009) The science of neural interface systems. *Annu Rev Neurosci* **32**: 249-266.

Hines ML, Carnevale NT (1997) The NEURON simulation environment. *Neural Comput* **9**: 1179-1209.

Hubel DH (1957) Tungsten microelectrode for recording from single units. *Science* **125**: 549-550.

Hubel DH (1959) Single unit activity in striate cortex of unrestrained cats. *J Physiol* **147**: 226-238.

Kloosterman F, Layton SP, Chen Z, Wilson MA (2014) Bayesian decoding using unsorted spikes in the rat hippocampus. *J Neurophysiol* **111**: 217-227.

Lawhern V, Wu W, Hatsopoulos N, Paninski L (2010) Population decoding of motor cortical activity using a generalized linear model with hidden states. *J Neurosci Methods* **189**: 267-280.

Marblestone AH, Zamft BM, Maguire YG, Shapiro MG, Cybulski TR, Glaser JI, Amodei D, Stranges PB, Kalhor R, Dalrymple DA, Seo D, Alon E, Maharbiz MM, Carmena JM, Rabaey JM, Boyden ES, Church GM, Kording KP (2013) Physical principles for scalable neural recording. *Front Comput Neurosci* **7**:137.

Newman JP, Zeller-Townson R, Fong MF, Arcot Desai S, Gross RE, Potter SM (2012) Closed-loop, multichannel experimentation using the open-source NeuroRighter electrophysiology platform. *Front Neural Circuits* **6**:98.

Ngo HV, Martinetz T, Born J, Mölle M (2013) Auditory closed-loop stimulation of the sleep slow oscillation enhances memory. *Neuron* **78**: 545-553.

O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. *Exp Neurology* **51**: 78-109.

O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* **34**: 171-175.

Paninski L (2005) Asymptotic theory of information-theoretic experimental design. *Neural Comput* **17**: 1480-1507.

Paz JT, Davidson TJ, Frechette ES, Delord B, Parada I, Peng K, Deisseroth K, Huguenard JR (2013) Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat Neurosci* **16**: 64-70.

Pologruto TA, Sabatini BL, Svoboda K (2003) ScanImage: flexible software for operating laser scanning microscopes. *Biomed Eng Online* **2**:13.

Rolston JD, Gross RE, Potter SM (2009) A low-cost multielectrode system for data acquisition enabling real-time closed-loop processing with rapid recovery from stimulation artifacts. *Front Neuroengineering* **2**:12.

Schwartz AB (2004) Cortical neural prosthetics. *Annu Rev Neurosci* **27**: 487-507.

Stevenson IH, Kording KP (2011) How advances in neural recording affect data analysis. *Nat Neurosci* **14**: 139-142.

Chapter 5: Conclusion

5.1 Contributions of the thesis

What have we learned from the experiments in Chapters 2 and 3? Both studies provide evidence that manipulating oscillations—defined as repeated patterns of temporally organized spikes—can have a *positive* impact on behavior. The fact that emulating gamma rhythms *enhanced* detection abilities and theta phase–specific inhibition *improved* spatial decision-making is crucial. It shows that these interventions were precise enough and subtle enough to augment the functioning of behaviorally relevant network states. This allows us to draw stronger conclusions than if we had merely disrupted behavior through rhythmic stimulation.

In the case of emulating gamma in somatosensory neocortex, the findings represent a compelling "proof of principle": we have shown that precisely aligned excitation and inhibition in the gamma frequency range can, in practice, benefit detection in an awake, behaving animal. This is not entirely unexpected, given prior modeling work that predicts gamma can be used as a form of gain control (Tiesinga, 2005; Börgers et al., 2005; Börgers and Kopell, 2008; Knoblich et al., 2010). But it's one thing to demonstrate the benefits of gamma *in silico*, and quite another to show that it works in a real brain. Given the long line of arguments that gamma has *no* effect *in vivo* (Shadlen and Movshon, 1999; Burns et al., 2011; Xing et al., 2012), procuring evidence that gamma can benefit a behavior as fundamental as stimulus detection is a major advance.

That said, the precise mechanism by which our optogenetically controlled gamma rhythms impacted behavior remains unclear. We have evidence that presenting vibrissae deflections at the optimal temporal offset (12.5 ms) elevated the synchrony of the evoked response, which could enhance the downstream impact of spikes originating in SI. However, we did not see a corresponding *decrease* in synchrony at the 7.5 ms temporal offset, the alignment that impaired performance. Therefore, we cannot claim that gamma enhances behavior by enhancing synchrony, as measures of synchrony were not correlated with behavior. This should not come as a surprise, given our small sample size of regular-spiking units ($N = 35$) limited to layer 2/3. It is possible that the synchrony effects are most pronounced in a subset of L2/3 cells, or are more

apparent in another layer. We cannot rule out—or rule in—any mechanism based on the currently available physiological data.

We are also not at liberty to claim that changes in spike alignment in the gamma range are the basis for attentional modulation—only that they *could* be. We have correlative evidence that gamma is preferentially enhanced on hit trials in mouse SI. But in order to study gamma correlatively in a mouse model, we need much denser sampling of the networks involved. Ideally, we would want simultaneous recordings from as many fast-spiking interneurons as possible, in order to see how changes in gamma-range inhibitory synchrony impact detection on a trial-by-trial basis. In our present data set, we have, at most, two such cells monitored at any given time. Advances in recording technology may make it possible to one day observe gamma at the appropriate level.

Despite the lack of evidence for a precise mechanism by which our intervention enhanced behavior, our fundamental claim is a mechanistic one: changing the alignment between excitation and inhibition in the gamma frequency range can have a significant impact on stimulus detectability. The only alignment that improved performance was the one that reinforced the intrinsic tendency of cortical networks to fire a barrage of excitatory spikes followed by a barrage of inhibitory spikes (Hasenstaub et al., 2005). Rather than allow oscillatory spike timing to fluctuate naturally, we have brought it under causal control. The findings from Chapter 2 imply that, in behaving animals, gamma *can* enhance signal transmission by aligning excitation and inhibition in this way. Whether or not non-rhythmic inhibition would have a similar effect is less relevant; our goal was to explore the putative behavioral impact of spike patterns that manifest themselves under natural conditions. It is clear that gamma needs to be conceived not as a frequency band, but as repeated bursts of excitatory and inhibitory spikes. Future experiments that use optogenetic interventions to induce and impair gamma rhythms on a trial-by-trial—or even cycle-by-cycle—basis should reflect this fact.

The interventions from Chapter 3 represent a different approach to studying oscillations: rather than enforcing spike timing in a given frequency range, we leveraged intrinsic inhibition to remove spikes at specific phases of an ongoing rhythm. This would not have been a feasible approach to studying gamma, since the timing of local inhibition defines the rhythm itself.

Recruiting fast-spiking inhibition would reset the phase of the next cycle, rather than reduce the firing rate at specific phases of an ongoing rhythm. Because the cycle length of theta is much longer than the impact of fast-spiking inhibition (125 ms vs. 25 ms), and theta is coherent across a distributed network that is less susceptible to local perturbations, phase-specific inhibition is viable. Studying theta in the same way we have studied gamma—by shifting the alignment of the spikes that define the rhythm—would have required multi-site, multi-wavelength optogenetic control that is beyond our current technological capabilities.

The implication of the Chapter 3 results, then, is not simply that "theta can impact behavior," but that changes in hippocampal information content at different phases of theta are used to guide decision-making. This has been predicted by various lines of computational (Hasselmo et al., 2002; Hasselmo and Eichenbaum, 2005; Kunec et al., 2005) and correlative (Colgin et al., 2009; Lever et al., 2010; Douchamps et al., 2013) evidence, but it has yet to be demonstrated via a causal intervention on the timescale of theta itself. Our findings also represent the first use of a closed-loop optogenetic manipulation to enhance a behavior of *any* kind, an achievement that will hopefully be replicated by many future studies.

As with the gamma results, we do not know the mechanism by which our intervention enhanced performance. Several possibilities are discussed in section 3.4, with the most likely being that precisely timed inhibition shifts the relative of influence of EC and CA3 on CA1, reducing task-irrelevant information in different segments of the track. Verifying that this is, in fact, the case, will require more extensive sampling of the networks involved. Ideally, this will be combined with optogenetic interventions that target all three areas. If, for example, inactivating the projections from EC has the same effect as recruiting CA1 inhibition at a specific phase, it will provide further evidence for the claim that our effects stem from changing the relative influence of the inputs that drive outputs from CA1.

Overall, the studies presented in this thesis make the case for the use of optogenetics as the primary means for studying oscillatory phenomena in the brain. All other "causal" interventions that claim to affect rhythms—lesions, pharmacology, genetics, entrainment—seem hopelessly crude in comparison. Once we conceive of oscillations in terms of their underlying mechanism, it becomes obvious that a spike-centric intervention is needed to study them directly. Even

subthreshold oscillations, which underlie many important physiological states (Harvey et al., 2009; Domnisoru et al., 2013), stem from rhythmic spike activity in upstream populations. Upgrading the tools (and the vocabulary) we use to study oscillations has the potential to resolve the confusion that has plagued the field for decades. T. H. Bullock's original question was: do oscillations "represent a novel slow, rhythmical change of state not necessarily involving impulses?" (Bullock, 1945). We should now be able to answer with a definitive "no": spikes both define oscillations and constrain their function.

5.2 Comparing gamma and theta

When it comes to biological systems in motion, rhythms are the rule, rather than the exception. Rhythms are an essential part of respiration, circulation, digestion, locomotion, and a host of other activities at every level of the animal kingdom. Given that nervous systems evolved to guide action, it should come as no surprise that rhythms are ubiquitous in the brain as well. Nervous tissue, as a physical object with its own intrinsic dynamics, has properties that are equivalent to the "momentum" of bones and muscles. Any mechanism for coordinating networks of cells must take these dynamics into account.

Gamma and theta represent examples of two broad classes of neural rhythms. Gamma can be conceived of as a "cellular" rhythm, whereby two or more reciprocally connected cells fire in repeating patterns (Skinner et al., 1994). Cellular rhythms have been widely characterized in invertebrates, for example in networks controlling leech swimming (Stent et al., 1978), locust flight (Pearson and Wolf, 1987), and crustacean digestion (Marder and Bucher, 2001). On each cycle, cells may fire bursts or single spikes, but they must become active in a particular order. The important feature of these rhythms is that they are defined on the scale of individual cells, rather than networks. The similarity between gamma and invertebrate locomotor rhythms has led some to call the cortex itself a "central pattern generator" (Yuste et al., 2005). But this simplistic view does not capture the properties of the more complex class of rhythms to which theta belongs.

Theta represents a "network" rhythm, whereby two or more reciprocally connected *areas* are activated sequentially. Other examples of network rhythms include slow waves, which mediate inter-areal interactions during sleep in mammalian brains (Hahn et al., 2012). The activity of individual cells is modulated by such rhythms, but does not define them. Instead, they depend on the distributed coordination of multiple regions, such as medial septum, hippocampus, and entorhinal cortex in the case of theta. By defining theta in this way, we limit the labeling of rhythms as "theta" to those that explicitly involve septo-hippocampal networks. Oscillations in the same frequency range, but with a distinct mechanism, should not be called "theta."

Cellular rhythms can be nested inside network rhythms, as evidenced by the widely studied relationship between theta and gamma. On each cycle of theta, CA1 gamma oscillations occur at well-defined phases, with low gamma occurring earlier in the cycle and high gamma occurring later (Colgin et al., 2009). Low gamma is a signature of CA3 influence on CA1, whereas high gamma is a signature of EC influence on CA1. Mechanistically speaking, these two rhythms are mediated by two different classes of local interneurons (Lasztóczi and Klausberger, 2014). Thus, hippocampal theta is defined by the sequential activation of CA3, CA1, and EC on each cycle (Mizuseki et al., 2009), while hippocampal gamma is defined by local barrages of synchronous inhibition (Buzsáki and Wang, 2012).

It should be clear from this discussion that gamma and theta are very different phenomena. Gamma is a cellular rhythm, which only needs two cells to occur. Theta is a network rhythm, which requires distributed interactions among participating brain regions. It would be inaccurate, bordering on uncouth, to label rhythms with different underlying mechanisms, but falling in the same frequency range, "gamma" or "theta." Standardizing the nomenclature we use to describe rhythmic activity would go a long way to ameliorating the confusion surrounding the purpose rhythms serve in neural networks.

The primary purpose of gamma is to synchronize the spikes of principal cells in order to enhance their downstream efficacy, as well as to make them more excitable following recovery from inhibition. The results from Chapter 2 show, for the first time, that shifts in gamma-range spike timing can enhance the behavioral impact of a sensory stimulus. Presumably, gamma serves a similar purpose in the hippocampus, allowing neurons in CA1 to maximize their

influence on cortical targets (Cenquizca and Swanson, 2007). Parvalbumin-positive interneurons seem optimized to mediate this synchronization, as a result of their fast inhibition, widely divergent connections, and electrical synapses (Bartos et al., 2007). Over the course of evolution, less specialized interneurons were likely co-opted to serve this purpose. The degree to which they are suited to enforce synchronization suggests that gamma conferred an evolutionary advantage.

Local synchrony in the gamma range, in turn, could have subsequently been co-opted to mediate inter-areal coordination—not just as a byproduct of upstream synchrony, but as an active mechanism for routing signals. This role for gamma has been proposed (Fries, 2009), and supported by correlative evidence (Gregoriou et al., 2009; Gregoriou et al., 2012; Roberts et al., 2013), but it has yet to be shown causally. More sophisticated closed-loop feedback experiments have the ability to test these hypotheses more directly (see next section).

The primary purpose of theta is to ensure that the various participants in the EC–HPC loop do not interfere with one another. Information retrieved from CA3 should not overlap with novel information arriving from EC, lest old memories be overwritten or new memories be ignored. The experiments in Chapter 3 suggest that this segregation of activity plays a role in behavioral guidance, as phase-specific inhibition impacts decision-making in a manner that depends on task segment. Similarly to local gamma synchrony being used for signal routing, theta may also serve a "higher-level" function in the context of a wider network. Given that the inputs to CA1 are active in a well-defined sequence, regions that receive CA1 projections could entrain to different phases of theta in order to access different types of information. CA1 represents retrospective information early in the cycle and prospective information late in the cycle (Mehta et al., 2002; Lubenov and Siapas, 2009). In a study that recorded simultaneously from CA1 and mPFC, the authors found evidence for enhanced mPFC phase-locking to theta prior to decision making, but no preferred phase for the population (Jones and Wilson, 2005). However, a recent study with recordings from anterior cingulate and CA1 found a consistent 80 ms phase lag between the two regions prior to decision-making, with the hippocampus leading cortex (Remondes and Wilson, 2013). This provides correlative evidence for the utility of phase-specific communication at theta

frequencies. Just as with a gamma, more advanced closed-loop experiments will provide the opportunity to test this hypothesis causally.

5.3 The future of closed-loop feedback

Chapter 4 of this thesis presented a new set of tools for implementing experiments involving closed-loop feedback. In order to take oscillations research to the next level, neuroscience needs a widely adopted, affordable, and extensible platform for configuring real-time analysis and stimulation. The experiments in Chapter 3 offer a glimpse of what is possible through this approach, but these barely scratch the surface. The tools for implementing closed-loop experiments are well within our reach; it would be wise to apply them wherever it is practical to do so.

Closed-loop feedback is currently used for the purpose of disrupting certain types of neural activity, such as seizures (Berényi et al., 2012; Paz et al., 2013) or hippocampal ripples (Girardeau et al., 2009; Ego-Stengel and Wilson, 2010; Jadhav et al., 2012). In this case, software or hardware is used to detect large-amplitude oscillations in the LFP, which triggers nonspecific electrical or optogenetic stimulation. While such approaches are useful in certain instances, they do not begin to take advantage of the computational resources at our disposal.

One simple addition to the neuroscientist's toolbox that would be tremendously beneficial is the optimization of stimulation parameters in real time. This has been explored extensively by theoreticians (Paninski, 2005; Huys et al., 2006; Lewi et al., 2009; Lewi et al., 2011; Ahmadian et al., 2011), but is rarely applied in experimental settings. Such optimization is crudely applied when searching for a neuron's receptive field while monitoring its spike output. But more elegant application of real-time parameter estimation could be used to titrate light power for optogenetic manipulations or search for sensory stimulus patterns that bring a population into a particular state (such as gamma). Such experiments would be straightforward to implement with new processing modules in the Open Ephys GUI.

Interrogation of hippocampal circuits would be improved by the use of real-time decoding, which would allow the content of spike activity to influence experimental parameters. This could

be used to condition reward on different patterns of ripple-associated replay (Foster and Wilson, 2006; Davidson et al., 2009), exposing the degree to which replay is under intentional control. Or, it could be used to implement a much more sophisticated version of the experiments presented in Chapter 3, whereby stimulation is triggered on the *information content* of CA1, rather than the theta phase. This would be a more direct test of the hypothesis that segregation of information by theta phase is useful for guiding behavior.

What about the "higher-order" hypothesis that differential phase-locking to the hippocampal theta rhythm is employed by cortical recipients of CA1 projections? This could be tested by creating a phase-dependent "dynamic clamp" of population activity through optogenetic inhibition or excitation. With recording electrodes placed in both hippocampus and prefrontal, retrosplenial, or anterior cingulate cortices, optogenetic stimulation of the cortex could be used to shift the preferred firing phase earlier or later in the theta cycle, or to limit firing to a single phase. Active monitoring of cortical response would be necessary to titrate light powers at different phases, to account for the natural tendency of cortical neurons to phase-lock to hippocampal theta. Thus, stimulation would be conditioned simultaneously on both hippocampal and cortical activity. A null effect would be hard to interpret, but enhancement or disruption of decision-making performance would indicate that theta phase–specificity matters outside the hippocampus as well.

Closed-loop experiments would also benefit research on gamma. The "communication through coherence" hypothesis could be tested directly by entraining fast-spiking interneurons in higher-order areas. This entrained gamma would be locked at a particular phase offsets to gamma in a primary sensory region with an overlapping receptive field. Ideally, the sensory-evoked gamma would arise naturally, and would be measured in terms of synchrony of fast-spiking interneurons, rather than just extracellular oscillations. Then, the downstream impact of spikes could be measured for different delays between the evoked and entrained gamma. There should be one delay that leads to optimal signal transfer between the two regions, which should be the same delay seen in synchronization under natural conditions. Combined with behavior, this would be the first direct test of communication through coherence, one of the most widely cited reasons that gamma is important.

All of these proposed experiments are feasible, given today's technology. This is fortunate, since causal tests will provide much more convincing evidence for the behavioral impact of oscillations than any correlative measures. The experiments in Chapter 2 and 3 offer a blueprint for studying rhythms in a causal manner, while the tools described in Chapter 4 will put closed-loop experiments within reach of more researchers. But there's much more work to be done. If neuroscientists are serious about demonstrating that oscillations have a behavioral impact, we need to become equally serious about engineering. That entails enforcing standards for software and hardware, coordinating projects across laboratories, and not placing development efforts in the hands of scientists alone. We are finally on the cusp of being able to interact with neural circuits at the right level of detail to show that oscillations are not just a side-effect of anatomy, but that they can actually do useful work in the brain.

5.4 References

Ahmadian Y, Packer AM, Yuste R, Paninski L (2011) Designing optimal stimuli to control neuronal spike timing. *J Neurophysiol* **106**: 1038-1053.

Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nat Rev Neurosci* **8**: 45-56.

Berényi A, Belluscio M, Mao D, Buzsáki G (2012) Closed-loop control of epilepsy by transcranial electrical stimulation. *Science* **337**: 735-737.

Börgers C, Epstein S, Kopell NJ (2005) Background gamma rhythmicity and attention in cortical local circuits: a computational study. *Proc Natl Acad Sci U S A* **102**: 7002-7007.

Börgers C, Kopell NJ (2008) Gamma oscillations and stimulus selection. *Neural Comput* **20**: 383-414.

Bullock TH (1945) Problems in the comparative study of brain waves. *Yale J Biol Med* **17**: 657-679.

Burns SP, Xing D, Shapley RM (2011) Is gamma-band activity in the local field potential of V1 cortex a "clock" or filtered noise? *J Neurosci* **31**: 9658-9664.

Buzsáki G, Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci.* **35**: 203-25.

Cenquizca LA, Swanson LW (2007) Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Res Rev* **56**: 1-26.

Colgin LL, Denninger T, Fyhn M, Hafting T, Bonnevie T, Jensen O, Moser MB, Moser EI (2009) Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature* **462**: 353-357.

Davidson TJ, Kloosterman F, Wilson MA (2009) Hippocampal replay of extended experience. *Neuron* **63**: 497-507.

Domnisoru C, Kinkhabwala AA, Tank DW (2013) Membrane potential dynamics of grid cells. *Nature* **495**: 199-204.

Douchamps V, Jeewajee A, Blundell P, Burgess N, Lever C (2013) Evidence for encoding versus retrieval scheduling in the hippocampus by theta phase and acetylcholine. *J Neurosci* **33**: 8689-8704.

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

Foster DJ, Wilson MA (2006) Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**: 680-683.

Fries P (2009) Neuronal gamma-band synchronization as a fundamental process in cortical computation. *Annu Rev Neurosci* **32**: 209-224.

Girardeau G, Benchenane K, Wiener SI, Buzsáki G, Zugaro MB (2009) Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* **12**: 1222-1223.

Gregoriou GG, Gotts SJ, Desimone R (2012) Cell-type-specific synchronization of neural activity in FEF with V4 during attention. *Neuron* **73**: 581-594.

Gregoriou GG, Gotts SJ, Zhou H, Desimone R (2009) High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* **324**: 1207-1210.

Hahn TT, McFarland JM, Berberich S, Sakmann B, Mehta MR (2012) Spontaneous persistent activity in entorhinal cortex modulates cortico-hippocampal interaction in vivo. *Nat Neurosci.* **15**: 1531-38.

Harvey CD, Collman F, Dombeck DA, Tank DW (2009) Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* **461**: 941-946.

Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* **47**: 423-435.

Hasselmo ME, Bodelón C, Wyble BP (2002) A proposed function for hippocampal theta rhythm: separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Comput* **14**: 793-817.

Hasselmo ME, Eichenbaum H (2005) Hippocampal mechanisms for the context-dependent retrieval of episodes. *Neural Netw* **18**: 1172-1190.

Huys QJ, Ahrens MB, Paninski L (2006) Efficient estimation of detailed single-neuron models. *J Neurophysiol* **96**: 872-890.

Jadhav SP, Kemere C, German PW, Frank LM (2012) Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**: 1454-1458.

Jones MW, Wilson MA (2005) Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol* **3**:e402.

Knoblich U, Siegle JH, Pritchett DL, Moore CI (2010) What do we gain from gamma? Local dynamic gain modulation drives enhanced efficacy and efficiency of signal transmission. *Front Hum Neurosci* **4**: 1-12.

Kunec S, Hasselmo ME, Kopell N (2005) Encoding and retrieval in the CA3 region of the hippocampus: a model of theta-phase separation. *J Neurophysiol* **94**: 70-82.

Lasztóczi B, Klausberger T (2014) Layer-specific GABAergic control of distinct gamma oscillations in the CA1 hippocampus. *Neuron* **81**: 1126-1139.

Lever C, Burton S, Jeewajee A, Wills TJ, Cacucci F, Burgess N, O'Keefe J (2010) Environmental novelty elicits a later theta phase of firing in CA1 but not subiculum. *Hippocampus* **20**: 229-234.

Lewi J, Butera R, Paninski L (2009) Sequential optimal design of neurophysiology experiments. *Neural Comput* **21**: 619-687.

Lewi J, Schneider DM, Woolley SM, Paninski L (2011) Automating the design of informative sequences of sensory stimuli. *J Comput Neurosci* **30**: 181-200.

Lubenov EV, Siapas AG (2009) Hippocampal theta oscillations are travelling waves. *Nature* **459**: 534-539.

Marder E, Bucher D (2001) Central pattern generators and the control of rhythmic movements. *Current biology* **11**: R986-R996.

Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* **417**: 741-746.

Mizuseki K, Sirota A, Pastalkova E, Buzsáki G (2009) Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron* **64**: 267-280.

Paninski L (2005) Asymptotic theory of information-theoretic experimental design. *Neural Comput* **17**: 1480-1507.

Paz JT, Davidson TJ, Frechette ES, Delord B, Parada I, Peng K, Deisseroth K, Huguenard JR (2013) Closed-loop optogenetic control of thalamus as a tool for interrupting seizures after cortical injury. *Nat Neurosci* **16**: 64-70.

Pearson KG, Wolf H (1987) Comparison of motor patterns in the intact and deafferented flight system of the locust: I. Electromyographic analysis. *J Comp Physiol A* **160**: 259-268.

Remondes M, Wilson MA (2013) Cingulate-Hippocampus Coherence and Trajectory Coding in a Sequential Choice Task. *Neuron* **80**: 1277-89.

Roberts MJ, Lowet E, Brunet NM, Ter Wal M, Tiesinga P, Fries P, De Weerd P (2013) Robust gamma coherence between macaque V1 and V2 by dynamic frequency matching. *Neuron* **78**: 523-536.

Shadlen MN, Movshon JA (1999) Synchrony unbound: a critical evaluation of the temporal binding hypothesis. *Neuron* **24**: 67-77, 111-25.

Skinner FK, Kopell N, Marder E (1994) Mechanisms for oscillation and frequency control in reciprocally inhibitory model neural networks. *J Comput Neurosci* **1**: 69-87.

Stent GS, Kristan WB, Friesen WO, Ort CA, Poon M, Calabrese RL (1978) Neuronal generation of the leech swimming movement. *Science* **200**: 1348-1357.

Tiesinga PH (2005) Stimulus competition by inhibitory interference. *Neural Comput* **17**: 2421-2453.

Xing D, Shen Y, Burns S, Yeh CI, Shapley R, Li W (2012) Stochastic generation of gamma-band activity in primary visual cortex of awake and anesthetized monkeys. *J Neurosci* **32**: 13873-13880.

Yuste R, MacLean JN, Smith J, Lansner A (2005) The cortex as a central pattern generator. *Nat Rev Neurosci* **6**: 477-483.