

Distribution Agreement

In presenting this thesis or dissertation as a partial fulfillment of the requirements for an advanced degree from Emory University, I hereby grant to Emory University and its agents the non-exclusive license to archive, make accessible, and display my thesis or dissertation in whole or in part in all forms of media, now or hereafter known, including display on the world wide web. I understand that I may select some access restrictions as part of the online submission of this thesis or dissertation. I retain all ownership rights to the copyright of the thesis or dissertation. I also retain the right to use in future works (such as articles or books) all or part of this thesis or dissertation.

Signature:

Nuri Jeong

Date

Goals Unhindered: How spatially selective disinhibition drives new important memories.

By

Nuri Jeong
Doctor of Philosophy

Graduate Division of Biological and Biomedical Science
Neuroscience

Annabelle C. Singer
Advisor

Joseph R. Manns
Committee Member

Matthew J. Rowan
Committee Member

Samuel J. Sober
Committee Member

Lary C. Walker
Committee Member

Accepted:

Kimberly Jacob Arriola, Ph.D, MPH
Dean of the James T. Laney School of Graduate Studies

Date

Goals Unhindered: How spatially selective disinhibition drives new important memories.

By

Nuri Jeong
B.S., Thomas Jefferson University, 2014
B.S., Immaculata University, 2014

Advisor: Annabelle C. Singer, Ph.D.

An abstract of
a dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Graduate Division of Biological and Biomedical Science
Neuroscience
2023

Abstract

Goals Unhindered: How spatially selective disinhibition drives new important memories.

By Nuri Jeong

Animals rely on quickly identifying and remembering the most important locations for successful goal-directed spatial navigation. Developing new spatial memories rapidly requires orchestrated firing activity of excitatory pyramidal cells in the hippocampus, but the mechanisms underlying new goal learning at precise locations are not fully understood. Inhibitory interneurons, including those that directly inhibit the somas of pyramidal cells, are known to modulate excitatory firing but this inhibitory modulation is typically thought to follow and depend on excitatory inputs. According to this view, inhibition supports excitatory reorganization during spatial learning by preventing hyperexcitation and suppressing low-firing activity to increase the signal-to-noise ratio of spatial coding. Distinct from this view, we tested a novel hypothesis that spatially selective reduction in inhibition would drive learning-associated excitatory reorganization at specific locations to serve goal-directed navigation. Spatially selective disinhibition, could drive excitatory reorganization for new learning and therefore gate enhanced information transfer at specific locations most pertinent to learning. We reasoned that such gating would occur at important locations such as rewarded areas. To test our hypothesis, we simultaneously recorded from many single neurons in mouse hippocampal subregion CA3, a region known to be important for developing new spatial memories, during rapid learning of new reward locations in virtual reality environments. We found a spatially selective reduction in firing rates of most interneurons when mice approached learned reward zones. This inhibitory reduction could not be explained by position-related changes in speed or licking behavior, nor was the timing or magnitude of the reduction consistent with simple balancing of changes in excitatory activity. To test the causal role of reductions in interneuron activity in learning, we optogenetically disrupted the normal reduction in inhibition at goal locations in new environments. Consistent with our hypothesis, goal location-specific stimulation of a small subset of CA3 parvalbumin interneurons, which provide perisomatic inhibition onto pyramidal cells, disrupted new goal learning without affecting performance in the familiar environment. Learning impairment was accompanied by deficits in the goal-relevant spatial information and sharp wave ripple activity. These results highlight a novel inhibitory gating mechanism for new goal-specific spatial learning.

Goals Unhindered: How spatially selective disinhibition drives new important memories.

By

Nuri Jeong
B.S., Thomas Jefferson University, 2014
B.S., Immaculata University, 2014

Advisor: Annabelle C. Singer, Ph.D.

A dissertation submitted to the Faculty of the
James T. Laney School of Graduate Studies of Emory University
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy
in Graduate Division of Biological and Biomedical Science
Neuroscience
2023

Acknowledgements

I would like to thank the incredible people who helped me make my goals a reality. First and foremost, I am immensely grateful for the exceptional mentorship provided by my advisor, Dr. Annabelle Singer. Her dedication, enthusiasm, measured risk-taking, and genuine concern for my well-being have been instrumental in the successful completion of this study. Special appreciation goes to my dream team committee members, Drs. Joe Manns, Sam Sober, Matt Rowan, and Lary Walker. Together, they fostered stimulating discussions and provided thoughtful feedback that significantly strengthened our work.

In addition, I would like to acknowledge my fellow lab members, with whom I have had the pleasure of getting to know and work alongside over the years. They have made a positive lab culture that encourages collaboration and diversity of thought, all while celebrating small victories along the way. I want to extend a special thank you to Steph Prince, Abby Paulson, Lu Zhang, and Matty Attokaren for their support and kindness since the beginning of my graduate career.

My time at Emory University has been enriched by the friendships I have formed with my cohort. They have been a constant source of laughter, inspiration, and cherished memories. I am truly grateful to call them my friends and trusted colleagues. Special thanks go to Mallika Halder, Fu Hung (aka Thomas) Shiu, and Esra Sefik for their authenticity and sense of humor, which have made my graduate school experience all the more enjoyable.

My family's sacrifices, particularly my mom's unwavering support, have enabled me to live and study abroad and have been the cornerstone of my achievements. From a young age, my dad instilled in me the value of always giving my best, and I carry his motto with me to this day. Without these life lessons, I would not have been able to complete the project I started. Furthermore, I would like to thank my American parents, Denise and Larry Coleman, who have unconditionally loved me as their own and believed in me since I became a part of their lives.

Finally, I would like to acknowledge my partner, Mark Snyder, who has consistently responded to my moments of vulnerability with love, kindness, and emotional intelligence. His rock-solid support has helped me overcome numerous setbacks and achieve my goals. I am forever grateful for his calm patience and belief in me.

Table of Contents

Introduction	1
Chapter 1	4
Abstract.....	4
Main Text.....	4
Figures.....	16
Chapter 2	20
Abstract.....	20
Introduction.....	20
Results.....	23
Discussion.....	36
Materials and Methods.....	39
Figures.....	52
Supplementary Figures	63
Supplementary Text	83
Chapter 3	92
Figure	97
Appendix A BrainWAVE: A Flexible Method for Noninvasive Stimulation of Brain Rhythms across Species	98
Abstract.....	98
Introduction.....	99
Results.....	102
Discussion.....	106
Materials and Methods.....	110
Figures.....	119
Extended Data.....	128
Appendix B A rapid and generalizable goal-directed spatial learning paradigm.....	136
Abstract.....	136
Introduction.....	136
Results.....	137
Discussion.....	141
Materials and Methods.....	143

Figures.....	149
References.....	154

Introduction

We tend to remember the things that matter the most. The ability to store specific important information into our memory thus helps us use our mental resources efficiently to serve our goals. Our memories of everyday experiences develop much like how we learn about new environments when we move to a new city. To navigate within an environment, we often use remembered cues (e.g., red brick house around the corner) to tell us where we are in relation to places of importance (e.g., where we can get the best tacos in town). As such, novel spatial learning, or learning in a new environment, has been studied as a model with which we can understand the formation of new important memories. In goal-directed spatial learning, the cues that predict target destinations take on important meaning over repeated exposure. Not surprisingly, our brains have developed specialized teams of neurons to quickly identify and prioritize remembering specific locations to serve goal-directed navigation.

Much of what we know about the neural mechanisms of spatial learning comes from electrophysiological studies conducted in the hippocampus, where specialized excitatory neurons that fire at specific locations of an environment were discovered in rodents (O'Keefe and Dostrovsky, 1971). These excitatory neurons, or "place cells," have been studied as an elegant example of how the brain develops an internal model or "cognitive map" of external environments with which we interact (O'Keefe and Nadel, 1978; Tolman, 1948). Place cells are thought to collectively represent the entire environment being learned but they tend to represent more of the important locations like rewarded areas over less important ones. Recently, goal-representing place cells, or pyramidal cells that fire specifically at goal locations, were found to play a causal role in behavior around goal locations during goal-directed navigation (Robinson et al., 2020). Likewise, groups of coactively firing excitatory neurons during the high-frequency oscillatory period called sharp-wave ripples observed in the hippocampus are also biased to represent trajectories toward goals (Pfeiffer and Foster, 2013). The importance of sharp-wave ripples for learning and memory is illustrated in

several studies that either perturbed or enhanced sharp-wave ripples artificially and found significant impairment or improvement, respectively, of spatial memory performance (Fernández-Ruiz et al., 2019; Jadhav et al., 2012).

While the current hippocampal literature is dominated by findings from excitatory neurons and excitatory population activity during sharp-wave ripples, they are an incomplete picture of the story, which continues to be written and updated. Inhibitory interneurons, particularly those that express the calcium-binding protein parvalbumin, are known to be important for regulating excitatory pyramidal cell firing and generating sharp-wave ripples (Evangelista et al., 2020; Schlingloff et al., 2014). Previous models of inhibitory roles in spatial learning posit that interneurons follow and respond to spatially modulated excitatory inputs with little location selectivity themselves. It has been traditionally thought that inhibition balances excitation and increases the signal-to-noise ratio of spatial coding by suppressing low responding excitatory cells.

These traditional models are based on the assumptions that changes in inhibitory activity depend on altered excitatory drive onto interneurons. These models have failed to address the mechanisms of location-selective learning that may not necessarily be driven by major changes in excitatory activity. A potential mechanism for new spatial learning may instead involve location-selective changes in inhibitory activity that drive enhanced excitation at specific locations. This potential function of inhibitory interneurons in spatial learning has not been examined in previous studies. As a result, it is not clear whether inhibitory interneurons reorganize to show spatially selective firing patterns over learning. It is also unclear whether inhibitory organization is necessary for learning-dependent excitatory reorganization and ultimately, successful goal-directed navigational performance. Most electrophysiological studies of hippocampal functions have been conducted in CA1, even though plasticity in CA3, the subregion upstream of CA1, is known to be a critical driver of rapid learning of new information. These gaps in knowledge are an important precursor to understanding how new specific information is learned to best serve our goals.

Recent studies with cell-type-specific targeting methods during awake behavior argue for an update to our models of inhibitory functions in spatial learning and memory. For example, hippocampal interneurons show spatial modulation and even refine place fields after initial formation, contrasting previous studies of uniform spatial modulation in interneurons. In Chapter 1, we will review these recent studies using cell-type-specific tools to examine inhibitory roles in local circuit computations during active spatial navigation. In this Chapter, a peer-reviewed article published in *Current Opinion in Neurobiology*, we focus on the hippocampal subregion CA1 in which most prior work was conducted and discuss limitations and unanswered questions. In Chapter 2, a manuscript under review, we will propose a novel function of inhibitory interneurons, especially those that directly suppress excitatory firing, in learning specific important locations in a new environment. In this Chapter, we focus on the subregion CA3, one synapse away from CA1, because plasticity events in this region are known to be critical for rapid learning of new information. Finally, in Chapter 3, we will summarize the new findings and discuss important caveats of this work.

Chapter 1

This chapter is a published review article as it appears in Current Opinion in Neurobiology.

Learning from inhibition: Functional roles of hippocampal CA1 inhibition in spatial learning and memory

Abstract

Hippocampal inhibitory interneurons exert a powerful influence on learning and memory. Inhibitory interneurons are known to play a major role in many diseases that affect memory, and to strongly influence brain functions required for memory-related tasks. While previous studies involving genetic, optogenetic, and pharmacological manipulations have shown that hippocampal interneurons play essential roles in spatial and episodic learning and memory, exactly how interneurons affect local circuit computations during spatial navigation is not well understood. Given the significant anatomical, morphological, and functional heterogeneity in hippocampal interneurons, one may suspect cell-type specific roles in circuit computations. Here we review emerging evidence of CA1 hippocampal interneurons' role in local circuit computations that support spatial learning and memory and discuss open questions about CA1 interneurons in spatial learning.

Main Text

Learning is the process by which an animal acquires knowledge or skills through experience. This process includes forming neural representations of novel information which are initially stored for short-term access and then stabilized into long-term memory through the process of consolidation. The hippocampus is essential for rapid learning and consolidation of spatial and episodic memory, or memories of events. Hippocampal excitatory pyramidal cell activity is thought to represent the internal perception of the external environment such as one's position in space-time dimensions and associations between external stimuli

relevant for task performance (Gauthier and Tank, 2018a; Góis and Tort, 2018; Knierim, 2002; Lee et al., 2012; MacDonald et al., 2011; O’Keefe and Dostrovsky, 1971; Singer and Frank, 2009; Tanaka et al., 2018; Wang et al., 2020a). Inhibitory interneurons strongly influence brain functions required for memory-related tasks and are known to play a major role in many diseases with common symptoms of memory impairment (Allen et al., 2011; Andrews-Zwilling et al., 2012; Artinian et al., 2019; Bissonette et al., 2015; Delorme et al., 2021; Lovett-Barron et al., 2014; Martinez-Losa et al., 2018). While genetic, optogenetic, and pharmacological interventions have shown that hippocampal interneurons play essential roles in spatial and episodic learning and memory (Andrews-Zwilling et al., 2012; Artinian et al., 2019; Oliveira da Cruz et al., 2020; de Salas-Quiroga et al., 2020), exactly how interneurons affect local circuit computations during behavior is an active area of investigation and debate. Recent studies have revealed significant functional heterogeneity in spatial and contextual selectivity, as well as heterogeneity in temporal dynamics of various interneurons, suggesting cell-type specific roles in circuit computations (Geiller et al., 2020; Klausberger and Somogyi, 2008; Pelkey et al., 2017). Indeed, interneuron subclasses in CA1 not only exhibit transcriptomic heterogeneity, but also target different cell-types and cellular compartments in CA1 (**Fig. 1-1A**) (Francavilla et al., 2019; Harris et al., 2018; Sik et al., 1995).

Here we review the emerging role of hippocampal interneurons in circuit computations for spatial learning and memory and open questions about interneurons in learning. In this review, we focus on recent papers investigating the main hippocampal output region CA1 using recording and manipulation of select neuronal populations during hippocampus-dependent spatial navigation tasks. While one of the primary established roles of interneurons is the generation of neural oscillations, this has been reviewed elsewhere and therefore will not be the focus here (Allen and Monyer, 2015; Antonoudiou et al., 2020; Butler and Paulsen, 2015; Buzsáki, 2015; Colgin, 2016; Dvorak et al., 2018; Mann et al., 2005; Strüber et al., 2022). Recent studies have shown that CA1 interneurons exhibit previously underappreciated feature selectivity that is relevant for spatial learning and memory (Ego-Stengel and Wilson, 2007; Geiller et al., 2021). Growing evidence suggests that this selectivity for features—whether it be a specific location, type of

information, or context—may be attributed to specific subclasses of interneurons (Geiller et al., 2020). Thus, we review how interneuron roles in computations for learning and memory are actively being rewritten by recent findings investigating cell-type-specific inhibitory populations during behavior.

Formation and Maintenance of Place Codes

When an animal navigates an environment, hippocampal pyramidal cells develop location-specific receptive fields, or place fields, where they preferentially fire action potentials when the animal is in a specific part of the environment (O'Keefe and Dostrovsky, 1971). Place cells have been studied extensively as an elegant example of how neural firing patterns develop to form internal representations of the external world (Best and White, 1998; Wilson et al., 1993). Some studies have shown that place fields tend to cluster around reward zones (Dupret et al., 2010; Gauthier and Tank, 2018a; Hollup et al., 2001; Lee et al., 2020; Singer and Frank, 2009), consistent with the idea that they prioritize behaviorally relevant locations over less important ones. Recent work has shown that place cells are causally involved in goal-directed behaviors (Robinson et al., 2020). For example, optogenetically stimulating groups of place cells that typically fired when animals were at a rewarded location led mice to lick at an unrewarded location, behaving as if they were at the reward location (Robinson et al., 2020). Thus, it is important to understand how interneurons affect place field formation and maintenance because place fields are relevant for memory-guided navigation.

The role of inhibitory interneurons in place field formation and maintenance is currently debated in the field. On one hand, one study showed near uniform inhibitory firing across spatial positions in CA1 of head-fixed navigating mice (Grienberger et al., 2017). Under uniform inhibitory firing, spatial modulation of excitatory cells is thought to arise primarily from spatially specific excitatory inputs or strengthening of excitatory synapses, requiring little spatially-specific inhibitory input (Bittner et al., 2015, 2017). On the other hand, spatial selectivity in CA1 interneurons has been previously reported in varying degrees (Ego-Stengel and Wilson, 2007; Hangya et al., 2010; Nitz and McNaughton, 2004; Wilent and

Nitz, 2007), and some have speculated that location-specific release of inhibition at least partly contributes to the spatial tuning of place cells (Hangya et al., 2010). Other studies that have examined CA1 excitatory and inhibitory responses to spontaneously developed or optogenetically induced place fields in head-fixed behaving mice have reported seemingly conflicting findings (Dudok et al., 2021; Geiller et al., 2020, 2021; McKenzie et al., 2021; Rolotti et al., 2022). Below we highlight recent empirical data investigating the role of interneurons in place field formation and maintenance.

A recent study reports a role for CA1 interneurons in place field maintenance (Geiller et al., 2021). This study by Geiller *et al.* (2021) elucidated interneuron effects on place fields in mice by using *in vivo* single-cell electroporation combined with monosynaptic retrograde tracing and optogenetics to identify synaptically connected cells *in vivo* (Geiller et al., 2021). The authors retrogradely labeled neurons presynaptic to an electroporated starter pyramidal cell and found that over 90% of the presynaptic inputs were from local inhibitory interneurons. Given this observation, the authors then expressed a genetically encoded Ca^{2+} indicator in all inhibitory interneurons using the VGAT-Cre mouse line, followed by electroporation of a starter pyramidal cell in CA1. This approach allowed for calcium imaging of functionally coupled pyramidal cells and interneurons in head-fixed mice running on a belt decorated with different tactile cues and licking for randomly delivered water rewards. Strikingly, while there was no detectable reduction in the activity of presynaptic interneurons during spontaneous place field formation, presynaptic interneurons showed “inverse” spatial selectivity once newly formed place fields stabilized, with strong reduction in firing at field locations of the starter pyramidal cell (**Fig. 1-1B**). These findings contrast with prior work that showed inhibitory inputs are spatially uniform and contrast to theories that reduced inhibition is required for place field formation (Grienberger et al., 2017). The discrepancy in spatial selectivity of interneurons between different studies may be due to the cells included. Geiller *et al.* specifically examined interneurons that were presynaptic to cells with place fields in the current environment and these cells may have higher spatial selectivity than other interneurons, e.g. those that synapse onto cells without a current place field. Additionally, differences in the behavior paradigm (like

randomly delivered reward versus reward delivered in specific parts of the track), may affect the animal's use of position-specific cues leading to differences in spatial selectivity of interneurons.

In contrast, another study suggests CA1 interneurons do play a key role in place field formation, but these effects may be interneuron subtype specific. Dudok *et al.* (2021) show that disinhibition by axo-axonic Chandelier cells, which target the axon initial segment of pyramidal cells, can induce place field formation (Dudok *et al.*, 2021). In this study, the authors created a new genetic mouse line to specifically target and optogenetically manipulate CA1 axo-axonic Chandelier cells (Dudok *et al.*, 2021). This work found axo-axonic cells inhibit firing activity of pyramidal cells in awake behaving mice. Both optogenetic activation and silencing of axo-axonic cells led to remapping of place fields. Interestingly, optogenetic silencing induced new place fields near the photostimulation site that persisted during post-photostimulation laps while optogenetic activation did not change the total number of place fields (**Fig. 1-1C**). Thus, location-specific reduction in axo-axonic cells' firing may be sufficient to generate place fields at that location. Together, the findings of Dudok *et al.* and Geiller *et al.* suggest that spatially organized activity of interneurons plays a key role in the development and maintenance of hippocampal representations of spatial experience with the exact role dependent on the interneuron subclass.

Indeed, spatial selectivity in CA1 interneurons differs between interneuron subclasses, or interneurons that have the same molecular markers, like parvalbumin, but differ in their morphological classification, like basket cells and bistratified cells (Geiller *et al.*, 2020). Geiller *et al.* (2020) used three-dimensional 2-photon calcium imaging and molecular verification by immunohistochemistry to survey several subclasses of CA1 interneurons simultaneously recorded in behaving head-fixed mice. This study found that many interneuron subclasses showed spatial selectivity, although the degree or stability of modulation varied by subclass. For example, spatial selectivity was more stable across sessions for parvalbumin-positive basket cells than for somatostatin-positive interneurons, although they both had similar spatial modulation. Given the diversity in interneuron responses, inhibitory populations likely contribute significantly to spatial representations that are typically thought to be represented by excitatory

populations with significant spatial modulation. Indeed, Geiller *et al.* (2020) found that the animal's current position was decoded above chance levels from CA1 interneuron activity alone although decoding performance was better with CA1 place cells than with interneurons only.

Beyond the development of individual place fields, Rolotti *et al.* (2022) suggest that local feedback inhibition may control the size of the excitatory population to represent a specific location (Rolotti *et al.*, 2022). This is analogous to the previously reported role of dentate somatostatin interneurons in controlling the size of fear memory ensembles (Stefanelli *et al.*, 2016). In this study, Rolotti *et al.* used a tamoxifen-dependent Cre virus and Cre-dependent excitatory opsin ChRmine to achieve sparse opsin expression across CA1 subpopulations in head-fixed mice running on a treadmill for randomly delivered water rewards. By varying the dose of tamoxifen injected in each mouse, the authors titrated the fraction of opsin-expressing CA1 pyramidal cells, while co-injection of GCaMP6f allowed for imaging calcium dynamics of the entire CA1 excitatory population. Rolotti *et al.* found that most single neurons could be optogenetically induced to develop place fields at the stimulation location that lasted at least 24 hours post-stimulation. Thus, new place fields may require strong enough excitatory drive above a certain threshold. Indeed, they showed that the percentage of pyramidal cells that were stimulated to fire together affected the efficacy of place field induction at the stimulation location (Rolotti *et al.*, 2022). Stimulation of fewer cells together due to lower opsin expression density led to a higher induction of place fields at the stimulation location compared to mice with higher opsin expression density. The authors then reasoned that a larger stimulated excitatory subpopulation might be more likely to recruit local interneurons, providing lateral inhibition onto nearby CA1 pyramidal cells. As a result, place fields could be more difficult to induce due to this lateral inhibition. Consistent with this hypothesis, the number of activated excitatory cells in a subpopulation was increased by chemogenetically suppressing local inhibition with inhibitory DREADDs specific to interneurons. Similarly, McKenzie *et al.* (2021), which used relatively low-power stimulation of CA1 pyramidal cells in head-fixed behaving mice, observed induction failures at the stimulation site perhaps due to larger recruitment of interneurons than place field induction via intracellular recordings

(McKenzie et al., 2021). These experiments suggest that interneurons play a key role in a competitive mechanism by which some pyramidal cells are selected to be part of a neuronal assembly while other pyramidal cells are suppressed and excluded (Rao-Ruiz et al., 2019; Roux and Buzsáki, 2015). Such groups of neurons firing together during an experience would then strengthen their connections which is thought to be the basis of memory formation (Buzsáki, 2010; Wallace and Kerr, 2010).

Interneurons have been hypothesized to play a key role in developing and refining spatial codes thus improving the signal-to-noise of such coding (McNaughton and Morris, 1987). These recent studies support this hypothesis and refine it by showing that specific functions vary by interneuron subclass. Together, these studies point to a model in which interneurons refine spatial codes in multiple ways. First, the activity of some spatially modulated interneurons can decode animal location (Geiller et al., 2020). Second, presynaptic interneuron activity exhibits an inverse relationship with new place field activity to maintain sharp and stable place fields (Geiller et al., 2021). Third, while interneurons disinhibit high firing cells at their place field locations, interneurons also suppress low firing pyramidal cells such that cells with inadequate excitatory drive do not participate in an assembly (Rolotti et al., 2022). Thus, these studies suggest that interneurons increase spatial coding signal-to-noise to stabilize new memories and suppress low or unstable excitatory activity with their exact roles differing by cell-type. The notion of cell-type-specific spatial selectivity in interneurons is further supported by cell-type-specific plasticity mechanisms discovered in CA1, including at glutamatergic synapses onto interneurons and at inhibitory synapses onto excitatory pyramidal cells (Bannon et al., 2020; Nissen et al., 2010; Szabo et al., 2012; Udakis et al., 2020). Spatially-selective interneurons may also play a role in behavior timescale synaptic plasticity (BTSP) in which excitatory cells exhibit extended periods of subthreshold depolarizations, or dendritic plateau potentials, over behavioral (seconds) timescales, although this has traditionally been attributed to excitatory inputs (Bittner et al., 2017; Milstein et al., 2021). Furthermore, spatial selectivity that seems to be driven by excitatory inputs may have an inhibitory component. Previous *in vitro* work has shown that the efficacy of inhibition depends on the distance to adjacent excitatory inputs on the same hippocampal dendritic

branch, while excitatory and inhibitory drives are also balanced across dendritic branches (Liu, 2004). This within-dendrite inhibitory influence over excitatory inputs may contribute to spatial selectivity of excitatory inputs. Future work is necessary to determine which plasticity mechanisms contribute to different phases of spatial learning.

Learning Goal-Directed Navigation

The hippocampus is essential for rapid learning, whether after a single experience or a few minutes (Kim and Frank, 2009; Moser et al., 1993; Nakazawa et al., 2003). While recent papers described above examined place field formation and maintenance as a model of memory formation, place field formation occurs even when animals explore open fields and does not necessarily require animals to express their learning via behavioral changes (Wilson et al., 1993). To study learning-related behaviors explicitly, a few recent studies have shed light on the role of interneurons in circuit function as animals learn to find new goal locations (Dupret et al., 2013; Turi et al., 2019). These studies have examined neural inputs and connections as well as place fields that tend to cluster around reward location (**Fig. 1-2A**).

Interneuron-targeting interneurons, which ultimately disinhibit pyramidal cells, were recently discovered to play a key role in goal-directed spatial learning in the work by Turi *et al.* (2019) (Turi et al., 2019). This study focused on CA1 vasoactive intestinal polypeptide (VIP) positive interneurons, a subset of interneurons that primarily target parvalbumin- or somatostatin-expressing interneurons (Francavilla et al., 2018; Tyan et al., 2014), and ultimately release pyramidal cells from perisomatic or dendritic inhibition. After verifying that these VIP interneurons were indeed disinhibitory on CA1 pyramidal cells *in vitro* and *in vivo*, the authors virally expressed GCaMP6f in the dorsal CA1 region of VIP-Cre mice to observe chronic Ca^{2+} activity in head-fixed mice running on a treadmill. The authors first found that VIP interneurons could be separated into two groups based on whether they were positively or negatively modulated by speed. Upon training in a behavioral task that required the mice to learn to lick at a specific location, the firing of both VIP interneuron groups was modulated by proximity to reward, but in opposite

directions. Interestingly, this reward modulation was not observed if the mice did not have to learn the task as in random foraging conditions. The authors then tested the hypothesis that disinhibition of CA1 pyramidal cells is necessary for overrepresentation of goal location by place cells during learning and subsequent improved performance. To do this, Turi *et al.* virally expressed either the opsin channelrhodopsin or the opsin archaerhodopsinT in CA1 of VIP-Cre mice to optogenetically activate or silence, respectively, CA1 VIP interneurons in mice learning a new reward location. In support of their hypothesis, the authors found that optogenetic silencing of VIP interneurons reduced the number of place cells near goal locations as well as the rate of learning measured by licking earlier as animals approached the reward location (**Fig. 1-2 B,C**). Mice learned significantly faster than controls when CA1 VIP interneurons were optogenetically stimulated near the reward location, even if it did not affect CA1 pyramidal cell reorganization. Based on these results, the authors concluded that disinhibition mediated by VIP interneurons is necessary, but not sufficient, for learning-dependent reorganization of CA1 pyramidal cells.

Prior work has shown that the strength of connections between pyramidal cells and interneurons rapidly reorganize during spatial learning (Dupret *et al.*, 2013). Dupret *et al.* (2013) found that rats rapidly learning new goal zones in an otherwise familiar environment had learning-dependent changes in the strengths of putative monosynaptic connections between pre-synaptic excitatory place cells and post-synaptic inhibitory cells in CA1 (Dupret *et al.*, 2013). This study showed that putative pyramidal-to-interneuron connections alter input weight distributions following learning, with the direction of change dependent on whether the assembly represented new learning. Pyramidal assemblies that represented newly learned goal locations strengthened their connections to postsynaptic interneurons that preferentially fired when the new map was expressed. Conversely, these same pyramidal assemblies weakened their coupling to interneurons whose firing was more correlated with the old map. Furthermore, the changes in connection strength were more likely to occur around rewarded locations compared to unrewarded locations, even though pairing events were observed in all locations (Dupret *et al.*, 2013). These findings indicate that local

pyramidal-to-inhibitory circuits reconfigure dynamically based on both learned information and their affiliation to the map of relevance, and thus may contribute to outputting learned behavior. While this study was unable to differentiate between different subclasses of interneurons, it would be important for future work to determine whether and how cell-type-specific changes in connection strengths drive new learning (Geiller et al., 2020).

Together these studies show that CA1 interneurons develop learning-dependent changes important for goal-directed behavior. First, interneurons alter firing responses to specific spatial cues that predict reward, with the magnitude and direction of response modulation varied by subclass. Second, interneurons rapidly reorganize over learning by redistributing their synaptic weights onto nearby neurons. Ultimately, interneuron changes in firing activity and synaptic strength lead to stable representation of newly learned locations directly relevant for task performance.

Current Gaps in Knowledge

Together these recent studies reveal previously unappreciated spatial selectivity in some interneuron types which in turn increases the signal-to-noise ratio of spatial codes and enables learning (Dudok et al., 2021; Geiller et al., 2020, 2021; Rolotti et al., 2022). We must also keep important experimental limitations in mind when considering these findings. First, several studies have made clever use of induced place fields (Geiller et al., 2021; McKenzie et al., 2021; Rolotti et al., 2022). However, these place fields are generated artificially so may differ in their inputs or structure than naturally occurring place fields. Second, many of these studies were performed in head-fixed mice. This approach enables complex imaging and electrophysiology, but limits vestibular inputs to animals among other limitations (Minderer et al., 2016). Third, optogenetic stimulation can have many unintended side effects (Allen et al., 2015). Fourth, imaging of hippocampus often includes extensive damage to overlying cortex which may affect hippocampal functions or animal behavior since interactions between hippocampus and cortex are involved in cognition

(Anderson and Floresco, 2021; Sigurdsson and Duvarci, 2015; Wang et al., 2020b). Even with these limitations, these studies have provided key insights into interneuron roles in spatial coding.

These findings raise several key questions. First, it remains to be examined whether inhibitory interneurons are crucial in selecting specific assemblies over others to represent or respond to particular experiences. Second, the specific plasticity mechanisms at play in the development of interneuron spatial selectivity and in changing connection strength between inhibitory and excitatory cells are not well understood. Third, while inhibitory deficits have been well-documented in diseases with learning and memory impairments, the diverse circuit mechanisms by which interneurons influence learning and memory have not been elucidated. Finally, multiple brain regions may have unifying inhibitory mechanisms for learning and disease susceptibility, which opens avenues for future research. Studies that bridge multiple scales, from single neurons to networks to behavior, will shine light on the inhibitory mechanisms of learning and memory in health and disease.

Functional Outcomes of Inhibitory Deficits in Disease

Interneurons are sensitive to disease states and deficits in inhibitory activity are likely to lead to network dysfunction due to their powerful influences on neural circuits (Marín, 2012; Paterno et al., 2020; Ruden et al., 2020). Multiple studies have shown abnormal inhibitory activity of multiple subclasses in several animal models of Alzheimer's disease (AD), revealing common pathophysiological mechanisms (Busche et al., 2008; Leung et al., 2012; Li et al., 2009; Martinez-Losa et al., 2018; Palop et al., 2007; Reid et al., 2021; Verret et al., 2012). Furthermore, abnormal GABAergic inhibition is observed in depression, autism spectrum disorder, and Down syndrome, suggesting the sensitivity of inhibitory neurons to disease and potential common inhibitory mechanisms across multiple pathological conditions (Contestabile et al., 2017; Filice et al., 2020; Umschweif et al., 2021). Recent studies using *in vivo* electrophysiology in animal models of disease have demonstrated the need for subclass-specific investigation of interneurons that may affect different aspects of spatial learning and memory. Paterno *et al.* (2021) using the *Cntnap2* KO mouse model

of autism showed reduced inhibitory transmission *in vitro* by parvalbumin-expressing interneurons onto CA1 pyramidal cells in KO mice, which also had impaired performance in spatial object recognition and object congruence tasks (Paterno et al., 2021). This model also had layer- and frequency-specific deficits in gamma phase-amplitude coupling, which could be directly related to deficits in subclass-specific transmission at inhibitory synapses. In another study, Chung *et al.* (2020) found that in mice injected with toxic amyloid beta oligomers, optogenetically activating CA1 parvalbumin- or somatostatin-expressing interneurons restored peak gamma or theta power, respectively, to normal levels *in vivo*, albeit during anesthesia (Chung et al., 2020). The optogenetic manipulation also led to resynchronization of pyramidal cell spiking relative to gamma or theta oscillation, respectively. While this study did not examine mice during task performance, it would be of interest to consider interneuron subclass-specific rescue of frequency-specific oscillations during spatial memory-guided behavior. Similar to these observations, our study using the 5XFAD genetic mouse model of AD and a virtual reality spatial navigation task found fewer and weaker sharp-wave ripples in AD mice, which correlated with a significant reduction in the strength of putative interneuron-onto-pyramidal neuron monosynaptic connections in CA1 (Prince et al., 2021). These studies suggest that interneurons are involved in memory impairment in multiple models of disease either by disrupting oscillations or driving network imbalances and hyperactivity. Considering recent studies revealing the diversity and specialization of inhibitory populations and their roles for spatial coding and memory formation described above, interneuron dysfunction likely plays more direct roles in memory impairment, however empirical data for the specific roles of interneurons in circuit dysfunction in disease models is lacking.

Figures




Figure 1-1. Inhibitory roles in place field formation and maintenance

A. Simplified diagram of CA1 interneuron subclasses and their inhibitory synaptic targets on a pyramidal cell (Pyr.) with circles and triangle illustrating the location of the cell body of the interneuron and pyramidal cell, respectively. Calretinin (CR)- and vasoactive intestinal polypeptide (VIP)-expressing interneurons inhibit dendritic-targeting O-LM cells that primarily express somatostatin (SST), as well as soma-targeting parvalbumin-expressing basket cells (PVBC), thereby disinhibiting pyramidal cells. Basket cells co-expressing cholecystokinin (CCK) and VIP target the soma of pyramidal cells, whereas CCK and calbindin (CB)-expressing interneurons typically target the dendrites. Axo-axonic Chandelier (AAC) cells primarily target the axon initial segment. Note that the diagram focuses on interneurons described in this review and is not a comprehensive overview of hippocampal inhibitory circuitry. **B.** Prior work shows that spatially selective reduction in inhibition is not required for new field formation but may be important for stabilization and maintenance of new fields. Top, several types of inhibitory interneurons that synapsed onto a pyramidal cell were identified and targeted by their expression of vesicular GABA transporter (VGAT). The inhibitory neurons included multiple cell-types with multiple possible synaptic target locations. Bottom, no change in presynaptic inhibition was found prior to the lap where a new place field formed, but as the place field stabilized, the presynaptic interneuron developed “inverse” spatial selectivity such that it fired less than baseline levels at the postsynaptic pyramidal cell’s place field location (Geiller et al., 2021). **C.** Spatial selectivity in place cells is inducible with cell-type-specific reduction in inhibition. Optogenetic activation or silencing of axo-axonic cells targeting the axon initial segment of the pyramidal cell led to place field remapping in vivo during behavior (Dudok et al., 2021). Specifically, optogenetic silencing led to induction of new, persistent place fields at the photostimulated location. In contrast optogenetic activation led to no significance change in the number of place fields at the stimulation site although rare pre-existing place fields were suppressed.




Figure 1-2. Inhibitory roles in learning goal-directed navigation

A. Place cells overrepresent behaviorally relevant locations such as reward zones and, while place fields tile the entire environment, more place fields cluster around these important locations. **B.** Vasoactive intestinal polypeptide (VIP)-expressing interneurons target other interneurons (Int.) that provides either perisomatic or dendritic inhibition onto pyramidal cells (Pyr.). In a recent study investigating a local CA1 disinhibitory circuit in goal-directed learning, optogenetically activating VIP interneurons (light blue) was found to induce faster learning of a new reward zone, as demonstrated by increased licking near the reward location (Turi et al., 2019). Optogenetic silencing (yellow) led to impaired learning in the same goal-directed navigation task. These findings suggest that new goal learning (and its behavioral expression) is mediated by local inhibition and disinhibition of pyramidal cells. **C. Left,** consistent with these behavioral changes in **B**, optogenetic silencing of VIP neurons led to a reduced proportion of place cells with fields near the goal (yellow, inhibitory opsin ArchT in VIP interneurons). **Right,** in contrast to the observed behavioral effects, VIP optogenetic activation (light blue, excitatory opsin ChR2 in VIP interneurons) did not lead to observable changes in the proportion of goal-representing place cells. These results suggest

that transient release of pyramidal cells from local inhibition is necessary, but not sufficient, to induce learning-dependent neuronal reorganization.

Declaration of interest: none

Acknowledgements: A.C. Singer acknowledges the Packard Foundation, NIH-NINDS R01 NS109226, McCamish Foundation, Friends and Alumni of Georgia Tech, and the Lane Family. N. Jeong acknowledges the Michael Kuhar Neuroscience Research Fellowship. We would like to thank members of the Singer lab and Lary Walker for feedback on the manuscript.

Chapter 2

Spatially selective inhibition gates new goal memory

Abstract

Goal-directed navigation in a new environment requires quickly identifying and exploiting crucial spatial information. Prevailing theories posit that interneurons play a supporting role in spatial learning by increasing the signal-to-noise ratio of excitatory place cells. We tested a novel hypothesis that spatially selective disinhibition drives plasticity at crucial locations during learning. We found location-selective, learning-dependent reduction in the firing of hippocampal interneurons around new goal locations. This goal-associated inhibitory reduction was required for spatial learning and excitatory reorganization to represent new goal locations. Sparse optogenetic stimulation to prevent goal-related decreases in interneuron firing impaired learning, reactivation of goal locations, and coding for goal locations in new but not familiar environments. These results reveal a novel role for interneurons in gating new spatial memories.

Introduction

Navigating to new goals is a fundamental behavior that requires rapidly identifying and remembering goal locations. Location-selective excitatory pyramidal cells, or place cells, in the hippocampus represent many positions in the environment including goal locations and are considered to drive goal-directed navigation (O'Keefe and Dostrovsky, 1971; Robinson et al., 2020). Place cells are intimately coupled with fast-spiking interneurons, including parvalbumin (PV) interneurons, that both respond to and control coordinated excitatory cell firing (English et al., 2017; Royer et al., 2012; Stark et al., 2014). Recent work reveals that hippocampal interneurons in CA1 exhibit spatially modulated firing, which is thought to result from place

cell inputs (Geiller et al., 2021). By following changes in excitatory input inhibitory interneurons are understood to increase the signal-to-noise ratio of place cell coding and balance excitation. In contrast to this view, we propose a new principle of circuit function in which spatially modulated interneuron firing acts as the driver of selecting important information, like goal locations, for rapid learning and consolidation. Prior work has not examined changes in inhibitory activity that precede or drive changes in excitatory firing during new spatial learning. As a result, it is unclear whether interneurons develop spatial selectivity preferentially for important locations and whether this interneuron selectivity is necessary for new learning.

Prior work posits two major hypotheses of inhibitory roles in spatial learning that support rapid learning of new goals. First, a long-standing view is that inhibitory recruitment balances excitation to prevent network instability (Hennequin et al., 2017; Higley and Contreras, 2006; Wehr and Zador, 2003). To balance excitation, inhibitory activity is expected to increase proportionally to increases in excitation to keep excitatory activity within an optimal dynamic range. It was previously shown in the hippocampus that the inhibitory balancing activity occurs after a time delay whose duration (tens of milliseconds) depends on the input amplitude (Bhatia et al., 2019) (**Fig. 2-1A**, left). A second prominent hypothesis is that weakly tuned inhibition improves the signal-to-noise ratio of spatial coding by suppressing low firing activity due to weak excitatory drive while strong excitatory drive overcomes this inhibition to produce high firing (Grienberger et al., 2017) (**Fig. 2-1A**, middle). This inhibitory model is analogous to the one proposed in sensory cortices in which broadly tuned inhibition sharpens the feature selectivity of excitatory neurons (Liu et al., 2011; Poo and Isaacson, 2009; Wu et al., 2008).

We propose a new model of inhibitory roles in rapid spatial learning. In this model, interneurons function as a disinhibitory gate permitting specific plasticity events to occur at locations of special relevance to learning (**Fig. 2-1A**, right). In gating, changes in inhibitory and excitatory activity occur in opposite directions (i.e., one increases and the other decreases in firing), with unequal magnitude of change, and

with inhibitory change preceding excitatory change. We can directly test these hypotheses by investigating the relationship between the direction, magnitude, and timing of excitatory and inhibitory activity. Prior work has primarily examined hippocampal interneurons in CA1, even though CA3, upstream of CA1, is the nexus of rapid spatial learning in a novel environment (Lee and Kesner, 2002; Nakazawa et al., 2003; Rolls, 2018). Thus, understanding the role of interneurons, especially those that directly inhibit pyramidal cells, in CA3 is crucial for distinguishing their role in rapid learning.

Because previous studies have not explicitly tested the inhibitory gating hypothesis, there are important questions to be addressed about the potential role of disinhibition in learning of precise locations. (1) Does disinhibition of excitatory firing occur in task-relevant locations during learning? (2) Does disinhibition increase learning-associated plasticity events at those locations? (3) Does disinhibition improve goal-directed navigational behavior? In the aforementioned models of inhibitory roles in new location-selective learning, the first two models assume reactively occurring inhibition, that is, inhibition that occurs in response to a change in excitation. In contrast, the gating hypothesis supports the model in which proactive inhibition occurs in response to specific locations and drives response selectivity of excitatory neurons by biasing them to represent specific locations over others.

Here, we addressed these questions by combining *in vivo* electrophysiology, optogenetics, and a novel virtual reality (VR) spatial learning task in head-fixed mice. We found that about half of fast-spiking CA3 interneurons, most likely PV interneurons as determined by opto-tagging, reduce their firing activity over a period of seconds as mice approach learned reward locations. This reward location-selective reduction in inhibitory firing occurred in a manner consistent with the gating hypothesis, and is inconsistent with the simple reactive balancing of excitatory change. To determine if this reduced inhibitory activity is required for new goal learning, we triggered optogenetic stimulation of a small subset of CA3 PV interneurons specifically at goal locations of novel environments to disrupt the reduction in PV activity during goal approach. Consistent with the gating hypothesis, optogenetic stimulation to disrupt PV firing

reduction at goals, but not sham stimulation at non-goals, impaired learning in the new environment. In contrast, optogenetic stimulation to disrupt PV firing reduction at goals did not impair performance in the familiar environment. Furthermore, PV stimulation at goal locations led to unstable goal representation by pyramidal cells and reduced reactivation of goal locations during sharp-wave ripples, both of which predict stable memory consolidation. These results demonstrate a mechanism of novel memory formation via disinhibition in selective locations that are most pertinent to new learning. These findings reveal new principles of neural circuit function in which prolonged, behaviorally relevant decreases in neural activity are crucial to learning, in contrast to transient increases in activity. Furthermore, we discover that spatially-tuned inhibition plays an active role in spatial learning by gating prioritized memories and plasticity processes during spatial navigation.

Results

Mice rapidly learn new goal locations in a novel virtual reality spatial learning task

To investigate hippocampal activity during new location-specific learning, we developed a new VR behavioral paradigm in which we could quantify progressive learning of goal locations in a controlled environment. We trained head-fixed mice to lick for rewards while running unidirectionally on treadmill (**Fig. 2-1B** and **fig. 2-S1A**). The track consisted of 36 equally sized areas around a circular track that had non-overlapping wall patterns and colors. Mice had to lick in three of these areas, or reward zones (RZ), to receive a reward. Licking outside the RZ or the zone just before the RZ, the anticipatory zone (AZ), more than 25-50 times resulted in a 4-second timeout. Prior to electrophysiological recordings, mice were exposed to one environment only (Track A or familiar track; **Fig. 2-1C** and **fig. 2-S1A**). After 7-14 days of training, mice reduced their movement speed and increased lick rate by an average of 28.95% and 21.36%, respectively, when they approached the AZ, indicating that they had learned the RZ locations (**fig. 2-S1A**, box). Successful task performance required visual cues because mice did not show a significant change in

movement speed or licking around well-learned goal areas if visual cues were removed (**fig. 2-S1B**). We used timeouts to discourage mice from licking indiscriminately and placed three RZs unevenly around the track to prevent mice from using a distance-based strategy.

On the first day of recording, mice alternated between sessions navigating on the familiar track (Track A) and a novel track (Track B or C) where they learned three new RZs through trial and error over three days (**Fig. 2-1C** and **fig. 2-S1A**). Consistent with knowing the specific reward locations, mouse movement speed in the familiar AZ decreased across all three days compared to a non-reward zone (NRZ), defined as the control zone that appeared 30 degrees after each RZ within the same environment (**Fig. 2-1D**, shades of grey; **fig. 2-S1C**, Days 1-3). On the first day of novel track exposure, mice showed no significant difference in movement speed between the novel AZ and NRZ. This lack of differences in movement speed shows that mice did not slow down, consistent with not knowing the reward locations (**Fig. 2-1D**, lightest shade of green; **fig. 2-S1D**, Day 1). By Day 2 or 3, mice exhibited slowing prior to arriving at the RZs, indicating that they had learned to differentiate rewarded areas from unrewarded areas (**Fig. 2-1D**, darker shades of green; **fig. 2-S1D**, Days 2-3). We quantified learning over days by animals' behavioral differentiation between goal and non-goal locations based on either movement speed, lick latency, or lick rate using receiver-operating characteristic curves in which larger areas under the curves (AUC) indicate better behavioral performance. We observed that, on average, AUC increased 20.05% on Day 2 and 22.65% on Day 3 from Day 1 in the novel environment, demonstrating that mice had rapidly learned to differentiate between AZs and NRZs (**Fig. 2-1E**, green; $F_{2,111}$ (Day x Environment) = 3.5, $P = 0.01$, linear mixed-effects model, LMM). Thus, mice learned to differentiate goals from non-goals over only a few days of novel track exposure. The performance metric using movement speed data tended to be better at position differentiation than licking-based metrics, likely because mice were not required to lick in the AZ to receive a reward and well-trained animals licked very sparsely (**fig. 2-S1E**). Overall, our novel virtual reality spatial navigation task enabled robust quantification of rapid learning of new goals.

Hippocampal inhibitory interneurons reduce firing activity around learned goal locations

A key prediction of the inhibitory gating hypothesis is that changes in inhibitory activity precede that of excitatory activity and occur in the opposite direction, while other models predict inhibitory activity follows or co-occurs with that of excitatory cells and proceeds in the same direction (**Fig. 2-1A**, right). To distinguish between the different hypothesized roles of interneurons in rapid goal learning (**Fig. 2-1A**), we first asked if CA3 interneurons show a spatially selective firing pattern around learned goal locations. We recorded 3,465 single units in CA3 during active spatial navigation in head-fixed mice and classified them based on their waveform features, inter-spike interval statistics, and autocorrelograms (Petersen et al., 2021) (**fig. 2-S2**). We identified putative pyramidal cells and two fast-spiking groups of interneurons with narrow or wide spike widths (**fig. 2-S2**). Opto-tagging and juxtacellular data indicate that narrow (NS) and wide spike-width (WS) interneurons most likely correspond to PV basket cells and somatostatin-positive interneurons, respectively (English et al., 2017; Royer et al., 2012; Stark et al., 2014). We confirmed that PV interneurons fall into the NS interneurons' classification via opto-tagging.

Strikingly, we found a goal-related reduction in the firing activity of the majority of NS and WS interneurons around three different RZs along the familiar track (**Fig. 2-1F**). Given the observed firing reduction across all goal locations despite distinct visual cues, we collapsed the responses to three RZs into a single firing map over the relative distance to RZ (**Fig. 2-1G**). We observed approximately a 15-20% reduction in firing during the approach to familiar RZs spanning approximately 30 degrees or 6 seconds prior to the RZ (**Fig. 2-1H**; $F_{25,4879.8}(\text{Position}) = 12.76$, $P < 2.2\text{e-}16$, LMM). We found significant firing reduction in 40-70% of all recorded interneurons over days in the familiar environment (**Fig. 2-1I**, left and middle, black). As expected, pyramidal cells tended to fire at one or a few locations that collectively represented the entire environment, with approximately 15-20% of units with increased activity around the RZ (**Fig. 2-1I**, right, black). In the novel environment, the proportion of NS interneurons with goal-associated activity reduction increased to comparable levels in the familiar environment by Day 2 (**Fig. 2-**

1I, left, green), whereas the proportion of pyramidal cells with increased goal-associated activity did not change significantly over days (**Fig. 2-1I**, right, green).

Because hippocampal neurons are known to be speed-modulated (Góis and Tort, 2018; Nitz and McNaughton, 2004; Wiener et al., 1989), we considered the possibility that the neuronal firing reduction is purely driven by reduced movement speed during reward approach behavior. To account for changes due to speed, we regressed out the estimated contributions of movement speed and lick rate to the raw firing rates (see *Materials and Methods*). Here we report the residual firing rates that cannot be attributed to position-related changes in movement speed or lick rate (**Fig. 2-1, G-I**). The observed raw firing patterns were preserved in residuals. Indeed, the amount of variance explained by movement speed and lick rate was less than 25% for the majority of units (**fig. 2-S3**). The statistically significant reduction in the residual firing activity of interneurons shows that these results could not be explained by position-dependent changes in movement speed or lick rate.

We performed several different control analyses that consistently support the finding that NS interneurons reduce firing rates around goals beyond what is expected by position-related behavioral changes. First, interneuron firing rates in the RZ during navigation in the familiar environment were significantly different from speed-matched firing rates during baseline no-VR periods at all speed quartiles except for the fastest speeds (**fig. 2-S4A**). Second, we separated out familiar reward approach trials by the animal's movement speed. When we compared the firing rates of the same neurons at the familiar AZ between speed quartiles, we found no significant speed effect on the AZ firing rates in all cell types (**fig. 2-S4B**). Finally, to further control movement speed and lick rate inherently associated with animal position on the track, we introduced a new trial type in the familiar environment in a small subset of animals ($n = 2$). Here, mice slowed down and licked more at the familiar RZ as they normally would, except that, upon RZ entry, the visual cues on the screen froze for 3-5 seconds even if the animal was moving on the treadmill, when the visual cues would normally progress forward through the track. With these screen-freeze trials,

we effectively dissociated position-related visual cues in VR from the animal's movement on the treadmill (**fig. 2-S4C**). We examined a period after RZ entry during which movement speed and lick rates were similar but the visual cues dissimilar. We found that the firing rates of NS interneurons were lower when reward-associated visual cues stayed fixed during screen-freeze trials than during normal trials. These results show that inhibitory firing is reduced for goal-related cues and not purely for speed changes (**fig. 2-S4C**).

Our data are consistent with the inhibitory gating hypothesis in which a spatially selective inhibitory reduction occurs prior to excitatory activity changes. Importantly, we find that this gating occurs at locations relevant for goal-directed navigation. To determine if this goal-specific inhibitory activity is correlated with behavior performance, we compared NS interneuron firing on correct and incorrect trials. NS interneurons reduced firing rates at the familiar AZ by a moderately greater amount on correct trials than on incorrect trials (**fig. S5**; familiar correct, $n = 196$ versus incorrect, $n = 75$, $F_{1,124.84}(\text{TrialOutcome}) = 3.93$, $P = 0.05$; LMM). These results show that successful task performance co-occurs with a goal-specific reduction in inhibitory activity that precedes excitatory activity changes and develops over days of new spatial learning.

Optogenetically disrupting inhibitory firing reduction at goal locations impairs new goal learning without affecting performance in the well-learned environment

We next asked whether the goal-associated reduction in inhibitory activity was necessary for learning new goal locations. To answer this question, we used optogenetic stimulation to disrupt the normal inhibitory reduction around goals during active spatial navigation. We focused our stimulation on PV interneurons for two reasons. First, we observed the earliest and largest pre-RZ activity reduction in NS interneurons, which are most likely PV interneurons as verified by opto-tagging (**Fig. 2-1H**). Second, PV interneurons are strong inhibitors of pyramidal cells at the somas readily found in *stratum pyramidale* and accessible in our recordings (Pelkey et al., 2017). For PV-specific targeting *in vivo*, we crossed the PV-Cre mouse line with

Ai32 mice to achieve Cre-dependent expression of the blue light-sensitive opsin channelrhodopsin-2 (ChR2) in PV interneurons (Madisen et al., 2012) (**fig. 2-S6A**). We inhibited PV cells in a small part of dorsal CA3, a circular truncated cone volume of approximately 1.7 mm³ and 4.3 mm³ in CA3 with the light intensities of 4.9 Mw/mm² and 12.7 Mw/mm², respectively (Stark et al., 2012). We confirmed the optical targeting of 2-3 PV units on average in a single recording, consistent with prior studies estimating between 2.4 and 32 PV cells per mm³ in CA3 (Deng et al., 2019; Li et al., 2022; Uchida et al., 2014). Consistent with light-induced activation, we observed an increase in PV firing activity within ~3 ms of blue LED light (470 nm) onset with the rate increase occurring in a dose-dependent manner (**fig. 2-S6, B and C**). As expected, strong PV stimulation increased firing rates of NS interneurons, while it suppressed firing in WS interneurons and pyramidal cells at the highest stimulation intensity (**fig. 2-S6C**). Low PV stimulation increased PV cell firing without significantly decreasing firing of WS interneurons or pyramidal cells (**fig. 2-S6C**). To disrupt inhibitory reduction around goals, LED light turned on from the time that the animal entered the AZ until it left the RZ or after a maximum of 10 seconds, whichever came first (**Fig. 2-2A**, “Goal stim,” blue). We compared the goal stimulation condition to a sham stimulation condition within the same animals in which PV interneurons were stimulated for up to 10 seconds but at the NRZ and the zone immediately after each NRZ in another novel environment (**Fig. 2-2A**, “Sham stim,” orange). The novel tracks used and the order of stimulation conditions were cross-balanced across animals.

Surprisingly, disrupting goal-associated firing reduction of a small subset of CA3 PV cells via goal stimulation profoundly impaired the learning of new goal locations. Mice that underwent PV goal stimulation did not show a significant speed difference between the AZ and NRZ over three days (**Fig. 2-2B**; goal stimulation using all trials: Day 1 versus Day 2, $P = 0.99$; Day 1 versus Day 3, $P = 0.95$; Day 2 versus Day 3, $P = 0.90$; $n = 4$ mice; LMM followed by Tukey’s correction). These same animals showed normal learning when they received sham stimulation in another novel environment (**Fig. 2-2B**; sham stimulation using all trials: Day 1 versus Day 2, $P = 0.01$; Day 1 versus Day 3, $P = 0.0003$; Day 2 versus Day 3, $P = 0.32$; $n = 5$ mice; LMM followed by Tukey’s correction). Goal-related PV stimulation

performance deficits over days were confirmed by significant fixed effects of day, stimulation condition, and interaction ($F_{2,17.091}$ (Day) = 6.5, $P = 0.01$; $F_{1,18.708}$ (Stimulation Condition) = 7.0, $P = 0.02$; $F_{2,17.091}$ (Day x Stimulation Condition) = 5.4, $P = 0.02$, LMM). Even after parsing out different trial types based on stimulation intensity and duration, the results were similar in that mice learned with sham stimulation, but not with stimulation that disrupted inhibitory reduction at goals (**Fig. 2-2C**).

It is Important to note that PV stimulation at goals did not result in the complete shut-down of excitatory activity that might explain the observed learning impairment. Because we used relatively low light intensities of a small subset of PV cells (0, 5, or 12 Mw/mm² used as no, low, or high stimulation, respectively, for each of the three goal locations), the normalized firing rate change of NS interneurons, WS interneurons, and pyramidal cells did not differ between no stimulation and low stimulation (**fig. 2-S6C**). Furthermore, the average reduction of normalized firing from baseline for pyramidal cells did not differ significantly between the two stimulation conditions with low or high stimulation (**Fig. 2-2D**; low stimulation: goal stimulation, -0.07 ± 0.18 , $n = 264$ versus sham stimulation, -0.03 ± 0.17 , $n = 314$ cells, $P = 0.84$; high stimulation: goal stimulation, -0.1021 ± 0.19 spikes/s, $n = 264$ cells versus sham stimulation, -0.06 ± 0.19 , $n = 314$ cells, $P = 0.66$; LMM followed by Tukey's correction). Likewise, PV firing did not differ significantly between goal and sham stimulation conditions at all intensities. At light onset, PV neurons increased their firing rates by similar amounts for both goal and sham stimulation conditions with low or high stimulation (**Fig. 2-2E**; low stimulation: goal stimulation, 0.31 ± 0.11 , $n = 12$ versus sham stimulation, 0.29 ± 0.17 , $n = 15$ cells, $P = 0.98$; high stimulation: goal stimulation, 0.45 ± 0.10 , $n = 12$ versus sham stimulation, 0.3770 ± 0.17 , $n = 15$ cells, $P = 0.96$; LMM followed by Tukey's correction).

Interestingly, mice failed to learn both the no-stimulation RZ that was never stimulated and the other stimulated RZs at the end of each session when they received no stimulation (**Fig. 2-2C**, "No stim"). This finding shows that there was no latent learning that emerges even when PV firing reduction was no longer disrupted. One explanation might be that the effects of PV stimulation on behavior are not spatially

confined and affect behaviors like licking or running speed outside stimulated zones. We found evidence against this hypothesis because there was no significant difference in the overall rate of licking outside the goal or stimulation locations between goal and sham conditions, ruling out the possibility that PV stimulation impaired the animals' ability to lick or that mice had different baseline levels of engagement in a task that requires licking for rewards (goal stimulation, 1.3 ± 0.14 licks/s versus sham stimulation, 1.5 ± 0.09 licks/s; Wilcoxon rank-sum test, $P = 0.47$). We found little to no differences in licking and movement speed between no stimulation and low stimulation in both familiar and novel environments (**fig. 2-S7 A and B**), which cannot explain the deficits in new goal learning even at low stimulation intensity. Thus, we conclude that disrupting interneuron firing reduction specifically at goal locations impairs new goal learning that cannot be due to potential differences in task engagement, and licking and running behavior at least at the lower intensity. Interestingly, disrupting goal-associated inhibitory reduction with PV stimulation at well-learned goals in the familiar environment did not compromise behavioral performance at any stimulation intensity (**fig. 2-S7C**; $F_{1,51.195}$ (Stimulation Condition) = 0.25, $P = 0.62$ for speed-based AUC; $F_{1,41.588}$ (Stimulation Condition) = 2.2, $P = 0.15$ for lick latency-based AUC; $F_{1,41.319}$ (Stimulation Condition) = 1.8, $P = 0.18$ for lick rate-based AUC; LMM). Taken together, these results show that goal-associated reduction in CA3 PV inhibition is required for learning new goal information but not for retrieving previously stored goal information.

Goal-associated reduction in inhibitory activity develops over learning and precedes goal-associated increase in excitatory activity

Because decreased inhibitory activity at goals is essential for learning and develops quickly, we wondered how this inhibitory reduction develops during learning. Prior work has shown that goal overrepresentation develops with learning (Danielson et al., 2016; Turi et al., 2019; Zaremba et al., 2017). We aimed to understand the time course of the development of inhibitory reduction at goal locations in wild-type (WT) animals to determine how it develops with learning (**Fig. 2-3, A and B**). When examining pooled data from

all days of novel exploration, we saw a significant ($P < 0.05$; one-sample t-test followed by Bonferroni's correction) reduction in firing activity at the goal location in NS interneurons (**Fig. 2-3, A and B**, top row), but not in WS interneurons (**Fig. 2-3, A and B**, middle row). As expected, pyramidal cells had a significant increase at the goal location (**Fig. 2-3, A and B**, bottom row).

The population activity in the novel environment tended to be more variable than in the familiar environment, with variable proportions of neurons that had goal-related firing patterns. Therefore, we opted to average across the units with similar goal-associated changes in firing activity only. To this end, we sub-selected units with a significant firing decrease or increase within 10 degrees of the novel RZ with the significance threshold established by each unit's shuffled rate distribution. Among the cell types, NS interneurons had the most consistent firing reduction during reward approach (**Fig. 2-3B**, top; Day 1, $-8.8 \pm 8.5\%$, $n = 8$; Day 2, $-14.35 \pm 5.3\%$, $n = 19$; Day 3, $-3.4 \pm 6.9\%$, $n = 22$). While there was no significant position effect on the residual firing rates of NS interneurons on Day 1 ($F_{25,525}(\text{Position}) = 1.2$, $P = 0.21$; LMM), NS interneurons did develop a small but significant reduction prior to the RZ even on the first session of Day 1. This rapid development preceded a significant increase in the pyramidal cell firing at the RZ that occurred by the second session (**fig. 2-S3D**). By Day 2, NS interneurons showed a significant position effect on the residual firing rates, which was also observed on Day 3 (Day 2, $F_{25,825}(\text{Position}) = 4.6$, $P = 1.5e-12$; Day 3, $F_{25,875}(\text{Position}) = 3.7$, $P = 2.8e-09$; LMM). On the other hand, few WS interneurons survived the significance threshold, making it difficult to see a reliable population activity change in the AZ (**Fig. 2-3B**, middle; Day 1, $11.48 \pm 17.4\%$, $n = 6$; Day 2, $-6.70 \pm 9.7\%$, $n = 7$; Day 3, $-34.94 \pm 1.4\%$, $n = 4$). Pyramidal cells did not have a consistent change in the firing rate prior to the RZ, with a slight increase on Day 1, a decrease on Day 2, and then an increase on Day 3 (**Fig. 2-3B**, bottom; Day 1, $3.4 \pm 2.2\%$, $n = 184$; Day 2, $-1.9 \pm 2.5\%$, $n = 168$; Day 3, 6.6 ± 2.4 , $n = 206$).

To further understand the functional consequences of excitatory firing changes over learning, we examined the spatial properties of goal-representing cells, or pyramidal cells that have at least one receptive

field with significant spatial modulation around goals. As expected, field stability was significantly higher in the familiar environment than the novel environment based on rate map correlations across trials (**Fig. 2-3C**; $F_{2,1091.6}$ (Environment) = 27.08, $P = 2.3\text{e-}07$; $F_{2,1175.0}$ (Day) = 11.25, $P = 1.5\text{e-}05$; LMM). The amount of spatial information about new goals carried by goal-representing cells did not increase significantly from Day 1 to Day 2, but it did by Day 3 in the novel environment (**Fig. 2-3D**; Day 1 versus Day 2, $P = 0.99$; Day 1 versus Day 3, $P = 0.04$; Day 2 versus Day 3, $P = 0.006$; $n = 133, 177, 192$ cells for days 1-3; LMM followed by Tukey's correction). The proportion of goal-representing cells did not change over days. We did not observe significant effects of day, environment, or interaction on the proportion of goal-representing cells (**Fig. 2-3E**; $F_{2,70.237}$ (Day) = 0.061, $P = 0.50$; $F_{2,70.003}$ (Environment) = 2.1, $P = 0.15$; $F_{2,70.003}$ (Day x Environment) = 2.6, $P = 0.09$; LMM). Thus, while the proportion of NS interneurons with a significant reduction in activity around the goal increased to comparable levels in the familiar environment by Day 2 (**Fig. 2-1I**), goal-related spatial information carried by excitatory cells did not significantly increase until day 3 in the novel environment and was lower in the novel than familiar environment. Together, these results show that learning-dependent changes in inhibitory activity occur rapidly during learning with evidence of inhibitory change before the enhancement of spatial information in excitatory cells.

Reduced inhibition coincides with the learning-dependent enhancement of new spatial representations by excitatory pyramidal cells

We then asked what coding is gated and causally unlocked by the reduction of inhibition around goals. Because goal stimulation resulted in similar patterns of behavioral deficits across stimulation intensities, we hypothesized that this stimulation affected goal-representing cells during learning as these cells can code for multiple goal locations in an environment (Gauthier and Tank, 2018b). To address this, we determined how optogenetic disruption of the goal-associated inhibitory reduction affects excitatory goal representations over learning in a new environment. During normal learning, the spatial information of goal-representing cells increases over days as animals perform better in the task (**Fig. 2-3D**). Goal

overrepresentation predicts successful goal-directed behavior and improves with learning (Danielson et al., 2016; Dupret et al., 2010; Gauthier and Tank, 2018a; Hollup et al., 2001; Turi et al., 2019; Zaremba et al., 2017). We found that goal stimulation disrupted the stability of new goal fields, as evidenced by a significant reduction in the rate map correlation over trials with goal stimulation compared to sham stimulation (**Fig. 2-4, A, B, and C**; $F_{2,267.89}$ (Day x Stimulation Condition) = 6.7, $P = 0.001$, LMM). Consistent with impaired goal representation due to disrupted PV reduction provided by goal stimulation, there was no significant increase in the spatial information of goal-representing cells over days (**Fig. 2-4D**; goal stimulation spatial information: Day 1 versus Day 2, $P = 0.99$; Day 1 versus Day 3, $P = 0.95$; Day 2 versus Day 3, $P = 0.73$; $n = 46, 36, 36$ cells for days 1-3; LMM followed by Tukey's correction). However, the same animals showed significant learning-dependent increase in goal information from Day 1 to Day 3, consistent with prior work, with sham stimulation in non-goal areas (**Fig. 2-4D**; sham stimulation spatial information: Day 1 versus Day 2, $P = 0.45$; Day 1 versus Day 3, $P = 0.0003$; Day 2 versus Day 3, $P = 0.09$; $n = 51, 60, 54$ cells for days 1-3; LMM followed by Tukey's correction). At the same time, there was no significant increase in the proportion of goal-representing cells over days for both goal and sham stimulation conditions (**Fig. 2-4E**; $F_{2,17.009}$ (Day) = 0.86, $P = 0.44$; $F_{2,17.677}$ (Stimulation Condition) = 1.5, $P = 0.24$; $F_{2,17.009}$ (Day x Stimulation Condition) = 0.52, $P = 0.60$; LMM).

Notably, we found that goal stimulation caused deficits in the stability of place cells in positions *away* from goals, outside the stimulated portion of the track (**Fig. 2-4F**; goal stimulation rate map correlations: Day 1 versus Day 2, $P = 0.97$; Day 1 versus Day 3, $P = 0.96$; Day 2 versus Day 3, $P = 1.0$; $n = 54, 38, 45$ cells for days 1-3; LMM followed by Tukey's correction). In contrast, when the same mice received sham stimulation in a different novel environment, it resulted in normal improvement over days (**Fig. 2-4F**; sham stimulation rate map correlations: Day 1 versus Day 2, $P = 1.0$; Day 1 versus Day 3, $P = 0.006$; Day 2 versus Day 3, $P = 0.04$; $n = 42, 41, 56$ cells for days 1-3; LMM followed by Tukey's correction). The amount of spatial information of place fields away from goals also did not increase over days with goal stimulation but did with sham stimulation (**Fig. 2-4G**; goal stimulation spatial information

in non-goal cells: Day 1 versus Day 2, $P = 0.97$; Day 1 versus Day 3, $P = 0.96$ Day 2 versus Day 3, $P = 1.0$; $n = 53, 37, 44$ cells for days 1-3; sham stimulation: Day 1 versus Day 2, $P = 1.0$; Day 1 versus Day 3, $P = 0.006$; Day 2 versus Day 3, $P = 0.04$; $n = 36, 46, 55$ cells for days 1-3; LMM followed by Tukey's correction). These spatial deficits occurred without significant changes in the proportion of non-goal-representing cells under both stimulation conditions (**Fig. 2-4H**; $F_{1,19.008}$ (Day) = 0.45, $P = 0.51$; $F_{1,19.116}$ (Stimulation Condition) = 0.99, $P = 0.33$; $F_{1,19.008}$ (Day x Stimulation Condition) = 0.43, $P = 0.52$; LMM). These results show that learning-dependent reduction in interneuron activity during goal approach is critical for both goal and non-goal representations. Disrupting PV reduction around goal locations resulted in deficits in both stimulated and unstimulated locations, whereas new spatial representations improved under sham stimulation. Taken together, our results support a goal-selective gating role of inhibition that recruits excitatory neurons to sharpen and stabilize new spatial representations that improve with performance.

Reduced inhibition increases goal-related reactivations during sharp-wave ripples

Goals are also represented at the population level during sharp-wave ripples (SWRs) when sequences of pyramidal cell activity represent paths toward goals (Pfeiffer and Foster, 2013; Singer and Frank, 2009; Xu et al., 2019). SWRs are required for rapid learning (**Fig. 2-5A**) and have been reported to affect synapses related to new spatial learning (Norimoto et al., 2018). We asked whether goal-associated inhibitory reduction is required for SWRs at goal locations. We hypothesized that disrupting the reduction in PV firing around goals would cause SWR activity deficits, in particular, deficits in goal-informative SWRs. We found that goal stimulation that disrupted PV firing reduction selectively at the novel goals resulted in a slightly lower rate of SWRs occurring around goal locations than sham stimulation in the same animals (**Fig. 2-5B**; goal stimulation, 0.018 ± 0.00088 Hz, $n = 20$ sessions versus sham stimulation, 0.034 ± 0.0014 Hz, $n = 20$ sessions, $P = 0.036$, LMM). Reduced SWR rate during goal stimulation was further corroborated by the observation that coactivation probability among simultaneously recorded pairs of goal-representing pyramidal cells during SWRs was lower with goal stimulation than sham stimulation (**Fig. 2-5C**; goal

stimulation, 0.043 ± 0.0080 , $n = 111$ pairs versus sham stimulation, 0.087 ± 0.0056 , $n = 516$ pairs, $P < 0.0001$, LMM). Furthermore, both SWR power and duration were significantly smaller in SWRs detected in sessions with goal stimulation than sham stimulation (**Fig. 2-5, D and E**; SWR power: goal stimulation, 3.5 ± 0.068 , $n = 54$ SWRs versus sham stimulation, 3.6 ± 0.10 , $n = 80$ SWRs, $P = 0.004$; ripple duration: goal stimulation, 0.076 ± 0.0036 seconds, $n = 54$ SWRs versus sham stimulation, 0.11 ± 0.015 seconds, $n = 80$ SWRs, $P < 0.0001$; LMM). These results show that disrupting goal-associated inhibitory reduction reduces SWR activity.

We then asked whether SWRs in the two stimulation conditions carried significantly different information about goal locations. To answer this question, we used sequenceless decoding of SWR content within a single time window to identify the most likely location about which individual SWRs carry information. To control for animal position and position-related variability in behavior, we included only SWRs that occurred while the animal was in the AZ or RZ in this analysis. We found significant differences in the information bias of SWRs between goal and sham stimulation. SWRs in sham stimulation sessions were much more likely to represent locations around goals than locations far from goals, similar to what we observed with no stimulation (**Fig. 2-5, F to I**; sham stimulation: near-goal versus far-goal proportions, $n = 12$ sessions, $P = 0.009$, Wilcoxon signed-rank test). On the other hand, when PV firing reduction was disrupted in the same animals, there was no significant difference in the proportion of SWRs representing locations near or far from goals, indicating a deficit in goal-related content during SWRs (**Fig. 2-5, F and G**; goal stimulation: near-goal versus far-goal proportions, $n = 8$ sessions, $P = 0.67$, Wilcoxon signed-rank test). These findings show that inhibitory reduction during reward approach is required for preferential reactivation of goal-related information. Consistent with this idea, we observed that SWRs in WT mice were much more likely to code for near-goal locations than far-goal locations in both familiar and novel environments (**Fig. 2-5, H and I**; familiar near-goal versus far-goal proportions, $n = 36$ sessions, $P = 3.1 \times 10^{-7}$; novel near-goal versus far-goal proportions, $n = 26$ sessions, $P = 0.001$; Wilcoxon signed-rank test). This bias toward reactivation of goal locations in the WT mice occurred without significant environmental

differences in SWR rate, power, duration, or generation (**fig. 2-S8**). Neither SWR rate nor duration in the familiar environment differed significantly with or without goal stimulation in the same PvxAi32 animals (**fig. 2-S9, C and D**; $F_{1,193.02}$ (Stimulation Condition) = 0.0034, $P = 0.49$ for rate; $F_{1,539.77}$ (Stimulation Condition) = 0.017, $P = 0.90$ for duration; LMM). These results are consistent with our previous finding that stimulation itself does not induce a complete shutdown of excitatory activity (**Fig. 2-2E**). While the rate of SWR rate was lower during goal stimulation, for SWRs that did occur during goal stimulation appeared similar to SWRs in the sham stimulation condition in terms of firing around the SWR, SWR power, and duration (**fig. 2-S9**). Overall, our results show that reduced activity of PV interneurons during reward approach gates goal-informative SWR activity for learning locations most pertinent to task performance.

Discussion

Here, we find a new role for PV inhibitory activity in gating information about goals that is essential for rapid learning of goal locations in new environments. This hippocampal inhibitory activity is learning-dependent, spatially selective, and specific to new goals. We found that CA3 interneurons reduced firing rates as mice approached learned goal locations. The observed inhibitory reduction was inconsistent with the simple balancing of increased excitatory. The inhibitory reduction started about 30 degrees or 6 seconds prior to RZs and preceded excitatory changes, in contrast to reactive inhibition. This goal-specific reduction in inhibitory activity developed quickly by the second day with hints of it on Day 1 while mice learned new goal locations in a novel environment. Optogenetically disrupting this inhibitory firing reduction around goals profoundly impaired new goal learning but did not impair the retrieval of previously learned goal memories. In line with this, disrupting this inhibitory firing reduction with goal stimulation resulted in less stable and less informative spatial representations by pyramidal cells, shorter and fewer goal-informative SWRs, and reduced co-activation of pyramidal cells during SWRs compared to sham stimulation. Together,

our data support the hypothesis that hippocampal inhibition gates new location-selective learning in service of goal-directed navigation.

Our findings provide a direct inhibitory link to previously reported hippocampal signatures of new learning: enhanced goal representation by pyramidal cells and sharp-wave ripple activity. We show that these signatures are under the influence of inhibition during new goal learning. Others have also reported goal-related changes in hippocampal activity as animals approach reward locations, including pre-reward bi-directional ramping activity of the retrohippocampal cortex (Tennant et al., 2022), calcium activity of CA1 astrocytes (Doron et al., 2022), reward-predictive signals of ventral tegmental area dopaminergic inputs to CA1 (Krishnan et al., 2022), and a dedicated reward-coding population found in CA1 and the subiculum (Gauthier and Tank, 2018a). However, the causal role of these previously goal-related hippocampal changes in rapid goal learning is unclear. Given its relevance to behavioral performance, new goal learning likely engages multiple plasticity mechanisms within or outside the hippocampus that may involve inhibitory gating.

One such plasticity mechanism worth noting is the recently proposed behavioral time scale synaptic plasticity (BTSP) rule (Bittner et al., 2015, 2017; Milstein et al., 2021). BTSP was illustrated in CA1 place field development with naturally occurring or artificially induced dendritic calcium spikes or “plateau potentials.” In our study, we found that changes in interneuron firing activity occurred over long, behaviorally relevant (seconds) timescales, similar to BTSP. While BTSP appears to be present without goal enrichment, reduced PV activity may underlie BTSP, with PV interneurons suppressing the activity of dendrite-targeting somatostatin-expressing (SST) interneurons preferentially at newly learned goal locations. Most prior work on the mechanisms of spatial learning has examined CA1, whereas this study focused on area CA3, directly upstream of CA1, because plasticity in this region is thought to drive rapid learning. Examining PV firing reduction in intra- and extrahippocampal regions will be an important next step in addressing a potential common inhibitory mechanism in new learning.

Our observation that goal stimulation to disrupt decreases in PV firing leads to deficits at single-unit, population, and behavioral levels raises the ultimate question about which deficit comes first. Prior work showed that temporally precise activation of PV interneurons, albeit on a much shorter timescale than in our task, is necessary and sufficient to generate SWRs (Schlingloff et al., 2014). However, it is unclear whether PV stimulation directly affects goal coding that causes learning deficits, or whether stimulation directly impairs learning, which in turn leads to disrupted goal coding. We hypothesize that, for PV interneurons in spatial coding, location-specific learning occurs because disinhibition triggers the induction of plasticity events at the key locations. Indeed, while disinhibition might occur anywhere around the track, coordinated disinhibition of about half of interneurons around the goal location unlocks plasticity selectively at the goal location. In examining the effects of stimulation intensities on learning in the novel environment, we found that mice failed to learn the no-stimulation RZs in which there was no disruption of PV firing. These results suggest that disrupting PV activity in one RZ may reduce an animal's learning capacity at another RZ, perhaps due to the inability to recognize goals, the inability to reduce inhibition, or both. Our observation that disrupted PV reduction resulted in less informative and less stable place cells across the environment could be due to disrupting SWRs which are thought to stabilize place fields. However, we found that goal-informative SWRs were particularly vulnerable to disrupting PV reduction. These findings suggest that PV reduction biases plasticity-related events to occur at and to represent important locations. Thus, disinhibition as a means of gating excitatory reorganization occurs preferentially at important locations over less relevant locations.

We initially found it surprising that stimulating PV interneurons in only a small subset of CA3 would have such a powerful effect on learning but others have also shown that stimulating as few as ~15 hippocampal pyramidal cells can profoundly alter learned behavior during spatial navigation (Robinson et al., 2020). Considering that fast-spiking interneurons are mostly coupled via gap junctions (Fukuda and Kosaka, 2000), it is possible that we stimulated more interneurons than we could detect at the site of recording. Importantly, our stimulation did not result in a “temporal lesion” of excitatory cell activity. The

effect of low-power stimulation on the firing activity of pyramidal cells was relatively small. Furthermore, we did not see signs of tissue damage or significant change in running and licking behavior. Future studies could use recently developed μ LED optogenetic tools to target single units in a more spatially confined area (Wu et al., 2015) to investigate the minimum number of interneurons required to generate the observed learning deficits.

In sum, this study reveals physiological disinhibition as a necessary means to endow excitatory activity with appropriate selectivity for behaviorally relevant information during spatial learning. Previous studies have described synaptic inhibition as a spatially uniform or temporally regular gain controller of excitatory activity without feature selectivity of its own. Others have suggested some feature selectivity of inhibition that followed changes in excitatory activity. In contrast to these studies, our study reveals that inhibition, preceding excitation, shows input selectivity and mediates subsequent response selectivity by tuning excitatory events to the most relevant information that meets the task demands.

Materials and Methods

Animals

We used C57BL/6J adult male wild-type (WT) mice (N = 7 mice, 25-35 g) at 10-12 weeks of age (3.5-5 months at the time of recordings) in our study. For optogenetics experiments, we crossed male PV-Cre knockin (The Jackson Laboratory #017320) with female homozygous Ai32 mice (The Jackson Laboratory #024109) to generate PvxAi32 mice (N = 10 mice) that express channelrhodopsin-2 specifically in parvalbumin-positive interneurons. All procedures involving animals were performed in accordance with the guidelines provided by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology. Animals were housed in a reverse dark-light cycle room (07:00 light off, 19:00 light on) with *ad libitum* access to food and water. We performed all behavioral training during the dark cycle.

Surgery

Mice were handled for at least three days prior to stereotaxic surgery. For head-plate implantations, mice were deeply anesthetized with isoflurane, and head-plates were affixed to the skull with dental cement (Parkell C&B Metabond, Edgewood, NY). Animals were single-housed and monitored for at least 3 days post-surgery. Mice received a subcutaneous injection of ketoprofen (0.5 mg/kg) for the first two days and saline for the first three days post-surgery. For acute extracellular recordings, craniotomies of 600-900 μ m in diameter were made on either hemisphere of well-trained mice one day before the first recording. We used the following coordinates to target CA3 *stratum pyramidale* (from bregma): 0.18 mm A/P, 0.24 mm M/L, and ~2.5 mm D/V.

Virtual reality spatial learning task

We began habituation of head-fixed mice on a treadmill at least 7 days of recovery after head-plate implantation. We trained mice to run on either a linear treadmill (N = 4 WT mice; PhenoSys SpeedBelt, Berlin, Germany) or a spherical Styrofoam treadmill floating on air (N = 3 WT mice, 8 PvxAi32 mice). To increase motivation for running, animals were food-deprived gradually to reach 85-90% of their original body weight. Mice were head-fixed and trained daily to run unidirectionally on the virtual track for progressively longer periods. We delivered drops of sweetened condensed milk as a reward, delivered through a plastic needle. Licks were detected using either a piezoelectric sensor attached to the reward needle or a custom-built photointerruptor-based system. All virtual tracks were designed and interfaced with animals using the open-source software Virtual Reality MATLAB Engine (ViRMEn) as previously described (Aronov and Tank, 2014a). Proximal and distal cues were projected onto a cylindrical screen, creating a rich and immersive environment for virtual navigation. Voluntary movement of a mouse on the linear or spherical treadmill automatically advanced movement in the virtual track environment.

We used three annular tracks (Tracks A, B, and C; **2-Fig. 1C**) of the same size but with distinct visual cues, reward locations, and distances between reward locations. All mice were initially trained on

Track A as the familiar environment and later (during recordings) introduced to Tracks B and C as the novel condition. Speed gain was set such that 1° advancement through the virtual environment was equivalent to the movement of about 1.5 cm on the belt for a total distance of approximately 540 cm around each track. Gain on the spherical treadmill was adjusted to ensure similar virtual experience as with the linear treadmill. Each of the tracks had 36 non-overlapping, equally sized wall cues (10° per cue), and 3 of them were associated with reward (reward zones). The reward zones were irregularly placed to prevent animals from solely using the inter-reward distances, as opposed to using distinct spatial cues, to perform the task across multiple environments.

For the first 3-5 days of virtual spatial experience, mice received rewards automatically when they arrived at the three reward zones on the track. After this first phase of training, mice showed anticipatory licking immediately before the reward delivery. In the subsequent training and recording sessions, animals were required to lick in the reward zones to trigger a reward delivery. To prevent generalized licking irrespective of position on training days, licking more than 25-50 times cumulatively outside the anticipatory and reward zones triggered a 4-second time-out period. During a time-out period, the animal was teleported to a grey box and received no reward. After 4 seconds, the mouse resumed running from the location it had left off prior to the time-out. We introduced mice to novel tracks on days of recording only after a mouse was deemed to have learned and reached behavioral criteria indicative of good performance in the familiar environment.

Behavioral data analysis

Mice were required to have 85% correct or higher performance in the familiar environment for at least 2 consecutive sessions before the recording commenced. Behavioral performance criteria were set based on an animal's propensity to slow down and lick more in the anticipatory zones.

Raw behavioral data were divided based on an animal's position (in degrees) into either 360-degree "laps" around the whole environment or RZ-centered "trials" that spanned the area 60 degrees before and

after each 10-degree RZ for analyses. Mean speed (in degrees/s) was calculated by dividing the total distance traveled in each 2-degree position bin by the total time spent in that bin. Lick probability was calculated by taking the ratio of the number of licks per position bin over the total number of licks within the lap or trial. Lick rate (in licks/s) was calculated by dividing the number of total licks per position by the total time spent in that bin over all position bins. Speed and lick behavior for each lap or trial was smoothed with a Gaussian-weighted moving average filter (SD = 2 bins, MATLAB function `smoothdata.m`). Once learned, mice tended to show stereotyped behavior in all three reward zones. Therefore, speed and lick behavioral data around reward zones during all trials were concatenated and averaged for each session regardless of the absolute position on the track. We used receiver-operating characteristics curves (MATLAB function `perfcurve.m`) to quantify behavioral performance based on how well mice differentiated between the pre-reward zone (anticipatory zone) and the non-reward control zone that appeared 30 degrees after the end of each reward zone. For lick latency-based performance, we quantified the temporal distance (in seconds) between the time an animal entered the reward zone and the first lick time. This lick latency at the reward zone was compared to the control latency between the time an animal arrived at the non-reward zone and the first lick time regardless of position.

Electrophysiology

All extracellular electrophysiology recordings were performed using a poly 5 two-shank 64-channel silicone probe or a 64-channel optoelectrode of the same channel geometry (NeuroNexus, Ann Arbor, MI). The probe was affixed to a 3-axis micromanipulator (Luigs & Neumann, Ratingen, Germany) and vertically advanced to the target region. Neural data were acquired using either two 32-channel Upright Headstages (SpikeGadgets, San Francisco, CA) or two RHD 32-Channel Recording Headstages (Intan RHD2000, Los Angeles, CA) at a sampling rate of 30 kHz with a ground pellet used as reference. Location was determined by stereotaxic coordinates, depth, and electrophysiological signatures. Electrophysiological features

indicative of CA3 were closely monitored: high-amplitude theta activity during running, prominent sharp-wave ripples during stillness, and high-amplitude (100+ μ V) action potentials appearing on many channels.

On each recording day, the animal had at least two 15-to-30-minute sessions in the familiar (Track A) and novel (Track B or C) environments per day (always starting with the familiar environment). Two behavioral sessions were separated by a shorter (5-10 minute) baseline recording period in the dark. The movement direction was consistent (clockwise) across environments. We performed recordings from the same craniotomy (typically from the right hemisphere first) for three consecutive days, using the same novel track during this period. After the first set of novel learning days, a second craniotomy was made on the contralateral hemisphere on the fourth or fifth day, and mice learned a brand-new track for three consecutive days starting on the following day. The same familiar track was used for each mouse for the entire duration of the study.

Optogenetic stimulation during behavior

We recorded and stimulated neural activity simultaneously using a customized two-shank 64-channel optoelectrode with two optical fibers terminating 100-200 μ m above the top recording site on each shank. Each optical fiber's (inner/outer) core diameter was 105/125 μ m and the numerical aperture 0.22. Ferrule patch cables ($\varnothing$ 105 μ m Core, 0.22 NA SMA905 to $\varnothing$ 1.25 mm; Thorlabs, M63L01) were connected to optical fibers on one end via ceramic split mating sleeves (Thorlabs, ADAL1-5) and 470 nm fiber-coupled LEDs (Thorlabs, M470F3) with T-Cube LED driver (Thorlabs, LEDD1B) with a 15-V power supply unit (Thorlabs, KPS101) on the other end. A custom MATLAB script and National Instruments data acquisition system were used for detection and triggering of on/offset and intensity of blue light stimulation. Animal position in virtual environments was detected in real time to trigger position-specific stimulation. In the goal location-specific condition, light turned on when an animal entered the anticipatory zone, the zone immediately before the reward zone, and stayed on for up to 10 seconds or until the animal left the reward zone. Similarly, in the sham stimulation condition, the light turned on at the non-reward zone for up to 10

seconds or until the animal left the zone immediately after the non-reward zone. A subset of animals was stimulated at the anticipatory or reward zone only, or for a fixed duration of 3 seconds regardless of position. For each of the three goal zones, we stimulated at the light intensity of either 0, 5, or 12 Mw/mm² from the fiber tip. The same stimulation intensity was used for the same goal location throughout the session and in the same environment over three days. Each stimulation session consisted of 20 minutes of stimulation trials followed by 5 minutes of stimulation-free trials to assess the potentially lasting effects of stimulation at the end of a session. All PvxAi32 mice had a shorter (~10 minutes) session with goal location-specific stimulation trials in the familiar environment (Track A) as their final session of the last day of recording.

Histological verification of probe location

On the last day of recording, the neural probe was dipped in fluorescent dye diI (0.9 mg/ml) prior to use. Following recording, mice were deeply anesthetized with isoflurane and perfused with ice-cold 4% paraformaldehyde. Brains were quickly extracted and stored at 4°C in 4% paraformaldehyde for 24 hr. Brains were then transferred to either 1X phosphate-buffered saline (PBS) until sectioning by a vibratome or to 20% sucrose solution overnight before being frozen at -80°C for sectioning on a cryostat. Fixed brain tissues were cut coronally in 100 µm thickness on a vibratome or in 60 µm thickness on a cryostat. Tissue sections were stained for nuclei with DAPI, mounted on slides (Vectashield Antifade Mounting Media), and cover-slipped for confocal imaging. Images were taken at 10X using an LSM 700 laser scanning confocal microscope (Zeiss).

Virtual reality manipulation

A small subset of animals (n = 2) was exposed to a behavioral manipulation in virtual reality on the last recording day to control for position-dependent changes in speed and licking activity. In this session, we introduced automatic screen freeze at the three reward zones in the familiar environment for a fixed duration of 3 to 5 seconds regardless of the animal's movement on the ball. The goal of this manipulation was to dissociate movement speed from position-specific visual cues.

Preprocessing of local field potentials and detection of sharp-wave ripples and theta periods

Raw neural signals were down-sampled to 2 kHz and bandpass filtered between 1-300 Hz to obtain local field potentials. Multi-unit spikes were extracted by bandpass filtering the raw signal between 300-6,000 Hz and thresholding the filtered signal above 5 standard deviations from the mean. We interpolated over outliers (noise) defined as 15 standard deviations above the mean of the pre-filtered signal. For analyses of theta, gamma, and sharp-wave ripple periods, the filtered signal free of outliers was bandpass-filtered based on frequency bands (4-12 Hz for theta, 1-4 Hz for delta, and 12-30 for beta) using a finite impulse response (FIR) equiripple filter. We then applied the following criteria to extract specific periods of interest and used the same thresholds for all sessions recorded from the same animal on the same day. To detect theta and non-theta periods, the envelope amplitude of the filtered theta signal was divided by the sum of the envelope amplitudes of the delta and beta signals. A theta period was defined as a period during which this theta to delta and beta ratio was 2 standard deviations above the mean for at least 2 seconds (Csicsvari et al., 1999; Iaccarino et al., 2016; Jackson et al., 2006). A non-theta period was defined as a period during which this ratio threshold was less than 1.1 for at least 2 s. Sharp-wave ripples were detected when the envelope amplitude of the filtered sharp-wave ripple trace was greater than 3 standard deviations above the mean for at least 20 ms (Karlsson and Frank, 2009; Singer and Frank, 2009; Singer et al., 2013). We excluded any events with a power ratio (power from 100 to 250 Hz / power from 250 to 400 Hz) less than 4 based on the typically observed frequency range of sharp-wave ripples (Ylinen et al., 1995). We only included in our ripple analyses the periods with at least 1 multi-unit spike and excluded periods during which movement speed was above 5 deg/s for 1 second before and after the mid-point of each period (2 seconds in total). In all of our local field potential analyses we used the channel with the highest envelope amplitude for the ripple band signal as the proxy site of *stratum pyramidale* (Gordon et al., 2005). Extracted periods of interest were visually inspected to ensure accurate detection. Duration of sharp-wave ripples was defined as the length of time the sharp-wave ripple envelope was greater than the threshold of 3 standard deviations above the mean. Sharp-wave ripple power was quantified as standard deviations above the mean ripple power for

the entire recording session. Ripple rate was quantified for ripple events detected during stopped periods defined as periods where movement speed fell below 2 deg/s. Ripple rate was expressed as the total number of ripple events divided by the total duration of stopped periods for individual recording sessions.

Single unit isolation

We identified and sorted single units using the automatic clustering software Kilosort2 (Pachitariu et al., 2016), followed by manual curation. Only well-isolated units with the signal-to-noise ratio of greater than 1 and < 0.01% inter-spike interval noise were included in the study. We visually verified the firing rate stability of single units by ensuring the firing rates during periods in between VR sessions did not fall below the threshold set at 10% of the peak firing rate of the entire duration of the recording.

Cell-type classification

We classified single units as pyramidal cells and interneurons with narrow- and wide- waveform interneurons based on the spike waveform's trough-to-peak distance in time and the autocorrelogram (ACG) fitted with a triple-exponential equation via CellExplorer software (Petersen et al., 2021):

$$ACG_{fit} = \max \left(c \left(\exp \left(\frac{-(x - t_{refrac})}{\tau_{rise}} \right) - d \cdot \exp \left(\frac{-(x - t_{refrac})}{\tau_{rise}} \right) \right) + h \cdot \exp \left(\frac{-(x - t_{refrac})}{\tau_{rise}} \right) + rate_{asymptote}, 0 \right)$$

where c is ACG τ decay amplitude, d is ACG τ rise amplitude, h is burst amplitude, t_{refrac} is ACG refractory period (ms). The CellExplorer software automatically identifies putative narrow spike-width (NS) interneurons with trough-to-peak latency ≤ 0.425 ms, wide spike-width (WS) interneurons with trough-to-peak latency > 0.425 ms and ACG τ rise amplitude > 6 ms, and the remaining units assigned as pyramidal cells. We recorded a total of 3,465 well-isolated single units recorded in this study (**2-fig. S2**). Among these, we identified a total of 463 NS interneurons, 258 WS interneurons, and 2,743 pyramidal cells. These criteria for fast-spiking NS interneurons were confirmed via optogenetics. To identify optically

tagged parvalbumin-positive interneurons, we used the stimulus-associated spike latency test (SALT) as previously described (Kvitsiani et al., 2013) and manually verified by visualizing light-evoked firing activity within 1-3 ms of light onset.

Place field analyses

To construct a one-dimensional, occupancy-normalized firing rate map for each cell, we first removed all spikes that occurred during stillness (movement speed < 2 deg/s or ~3 cm/s) from the spike trains. We then binned the remaining spikes into 5-degree position bins. We counted the number of spikes and time spent (in seconds) for each position bin and smoothed the spike count and occupancy map separately with a Gaussian kernel (SD = 2 bins). Finally, the rate map was constructed by dividing the smoothed spike count by the smoothed occupancy for each bin. The following criteria were used to identify place cells among putative pyramidal cells: (1) mean firing rate less than 10 spikes/s to exclude potential interneurons, (2) peak firing rate of least 1 spike/s, and (3) spatial information content greater than the 95th percentile of the information content generated from shuffled (repeated 1,000 times) data. Spatial information, expressed in bits/spike, was computed using the formula (Skaggs et al., 1996):

$$I = \sum_{i=1}^N p_i \frac{\lambda_i}{\lambda} \log_2 \frac{\lambda_i}{\lambda}$$

where p_i is the probability of the animal occupying the position bin i , λ_i is the mean firing rate of the cell in bin i , and λ is the mean firing rate across all position bins.

Multiple linear regression

To control for the well-documented linear and usually positive relationship between movement speed and neural activity, we used multiple linear regression. We also accounted for potential effects of position-specific lick rate on firing activity. Using MATLAB function regress.m, we fit for each unit a multiple linear regression model with position-binned trial-by-trial firing rates as the response and position-binned

speed, lick rate, and the interaction between speed and lick rate as predictors. We then subtracted the expected effects of speed and lick behavior from the observed firing rates. The remaining data (residuals) were used in further analyses as estimated neural activity that could not be attributed to position-related changes in speed or lick rate across trials. For visualizing population activity, we first normalized residual firing rates over binned position by dividing each bin by peak firing rate per unit and averaged over rewarded trials only per unit. To illustrate the change in residual firing activity from baseline, we subtracted the baseline firing rate defined as the mean normalized firing rate of the first two bins (in the position range of [-60, -50) degrees. We then sorted units based on the binned position with the peak or trough activity, for pyramidal cells or interneurons, respectively.

Determining units with a significant increase or decrease in firing activity at goal locations

To quantify the proportion of units with significantly increased or decreased activity in the familiar environment, we first created a position-binned firing rate map using a 2-degree bin size for each RZ-centered trial (from -60 to 70 degrees with zero being the start of the RZ) per unit. After removing all bins during which animal speed fell below 2 deg/s, we averaged binned firing rates over trials to result in a single rate map per unit. We then used a generalized linear model (MATLAB function `fitglm.m`) to find the best linear fit for the rate map over the pre-RZ position bins (-60 to 0 degrees) and obtain a *P*-value and estimated coefficient (slope) per unit. Only units with significant adjusted *P*-values after applying the false discovery rate of 0.01 were identified as units with significant activity change. For NS and WS interneurons, we identified units with significant negative coefficients as interneurons with significantly reduced activity. For pyramidal cells, we only included units with significant positive coefficients as units with significantly increased activity.

The firing rate changes around the new reward zones in the novel environment tended to be more variable than in the familiar environment. Therefore, we chose a different approach to identify units with a significant increase or decrease in firing in the novel environment to average over in **Fig. 2-3B**. We

randomly shuffled both the position-binned trial-by-trial maps of firing rates, movement speed, and lick rates used for multiple linear regression. This process was repeated 1,000 times for each unit to generate a shuffled distribution. We compared these shuffled distributions in each position bin against the observed session-averaged residual firing rate in the same position bin. We identified units with a significant decrease or increase at each bin if the observed residual was either less than the 10th percentile (for interneurons) or greater than 90th percentile (for pyramidal cells), respectively, of the shuffled distribution at each position bin. The population average in **Fig. 2-3B** are the average of units that had at least one bin determined to be significant 10 deg before and after the reward zone.

Ripple content decoding

To determine the content of individual ripple events, we performed sequenceless decoding to decode each ripple as a single time window (Carey et al., 2019). For each ripple, we calculated the spatial probability distribution using a simple Bayesian decoder as previously described (Karlsson and Frank, 2009). Briefly, the probability of particular positions given the spiking activity at each timepoint (expressed in $n\text{PositionBins} \times n\text{Time}$) was calculated using the formula:

$$P(X|N_1^C) = \frac{(\prod_{i=1}^C P(N_i|X) \cdot P(N_i|X)) \cdot P(X)}{P(N_1^C)}$$

where X is the set of all distance relative to the RZ (using 5-degree bins) and N_1^C is a vector of spike counts for all C cells recorded simultaneously within the specified time window used as testing spike counts. For all ripple events, we used a single 250-ms time window centered at the midpoint of each event regardless of its duration. For the training spike counts, we used the population firing rate map averaged across trials (a $n\text{Units} \times n\text{Bins}$ matrix using 5-degree position bins as a function of distance to RZ, from -60 to 70 degrees) multiplied by the decoding time window. $P(X|N_1^C)$ was then normalized across position bins to sum to 1. From the single probability estimate, we identified the decoded position bin with the highest spatial probability as the most likely spatial information (expressed as the relative distance to RZ) carried

by the population activity during each event. Only sharp-wave ripples that occurred when the animal was in either AZ or RZ were included for this analysis.

Immunohistochemistry

We used additional six PvxAi32 mice (3 females) to verify the expression of channelrhodopsin-2 localized in PV-expressing interneurons. We made 30 μ m coronal sections and used 5% goat serum in 1X PBS as the blocking solution for the first 30-minute incubation, followed by an additional 30-minute incubation in 5% goat serum with 0.3% Triton-X. Following blocking, we incubated the sections in primary antibodies at 4°C overnight, washed them in 1X PBS three times for 5 minutes each, and incubated them with secondary antibodies at room temperature for 1 hr. After three washes in 1X PBS, the sections were mounted on slides and imaged at 10X or 20X with an LSM 700 laser scanning confocal microscope (Zeiss). We used the following primary and secondary antibodies: rabbit anti-parvalbumin (Swant, PV27, 1:5,000), goat anti-rabbit IgG Alexa Fluor Plus 647 (A55055, 1:2,000).

Statistical analysis

We examined differences for statistical significance using a non-parametric two-sided Wilcoxon rank-sum test for non-uniformly distributed data. For non-parametric paired comparisons, we used the Wilcoxon signed-rank test with the Bonferroni method used to correct for multiple corrections. We used one-sample permutation t-test (5,000 times) identifying the position bins with population residual firing rate change that are significantly different from zero. Notations in figures, *, **, ***, and **** indicate $P < 0.05$, $P < 0.01$, $P < 0.005$, and $P < 0.001$, respectively, and ns indicates not significant. Data with error bars were reported as mean $\pm$ SEM. Details on statistical analyses have been described in the supplementary text.

To account for dependency in our data, we used linear mixed-effects model (LMM) analysis in R (version 4.2.2) and lme4 package (Bates et al., 2015) for most of the data. Significance was calculated using the lmerTest package (Kuznetsova et al., 2017) with Kenward-Roger's method to estimate degrees of freedom and generate p-values for mixed models. The emmeans package (<https://cran.r->

project.org/web/packages/emmeans/index.html) was used to adjust *P*-values for multiple comparisons. To determine statistically significant learning over days in WT mice, we included the areas under the behavioral receiver-operating characteristics curves as the dependent variable, day, environment, and day by environment interaction terms as fixed effects, and animal IDs as a random effect. The model specification for WT mice was as follows: $AUC \sim \text{day} + \text{environment} + \text{day} * \text{environment} + (1 | \text{AnimalID})$. For PvxAi32 mice, our main comparison was between goal stimulation and sham stimulation conditions that occurred in the novel environment only, therefore, we added fixed effects of day, stimulation condition (instead of environment), and the interaction between the two. We included animal number as a random effect. For PvxAi32 mice, we used the following model specification: $AUC \sim \text{day} + \text{stim_condition} + \text{day} * \text{stim_condition} + (1 | \text{AnimalID})$. Spatial information distributions between groups were compared using the following model specification: for WT mice, $\text{spatial information} \sim \text{day} * \text{environment} + (1 | \text{AnimalID}/\text{CellID})$; for PvxAi32 mice with goal and sham stimulation comparisons, $\text{spatial information} \sim \text{day} * \text{stim_condition} + (1 | \text{AnimalID}/\text{CellID})$.

Figures




Figure 2-1. Interneurons reduce firing as mice approach learned goals. **A.** Hypothesized roles of inhibitory interneurons in spatial learning. Left, inhibitory activity (blue) balances heightened excitatory activity (red) by increasing and decreasing firing around the location of importance (e.g. reward zone, pink) in about the same magnitude and timing of excitatory activity change. Middle, inhibitory activity increases the signal-to-noise ratio of spatial coding by suppressing low-firing excitatory activity without significant spatial modulation around the location of importance (pink). Right, inhibitory gating occurs with changes in excitatory following decreases in inhibitory activity around the area of importance (pink).

B. Head-fixed virtual navigation setup with the location of the probe. **C.** Experimental timeline. Visual cues that predict reward are unique in the three different environments, Tracks A, B, and C. During recordings, animals ran in both novel (green) and familiar (black) tracks each day. **D.** Example receiver operating characteristic curves based on speed from a single animal with larger area under the curve indicating higher behavioral differentiation between rewarded and unrewarded areas (familiar days 1-3, or Fam D1-3, in shades of grey; novel days 1-3, or Nov D1-3, in shades of green, lightest for day 1). Dotted diagonal line indicates chance-level performance. **E.** Behavioral performance based on speed in AZ and NRZ improves over days in the novel environment (green) and reaches performance levels in the familiar environment (black) by day 3. Data represent mean $\pm$ SEM area under the receiver operating characteristic curves for WT mice (linear mixed-effects model with Tukey's correction, $n = 14$ sessions from 7 mice in two different novel environments). **F.** Heatmap of residual firing rates of all recorded NS interneurons as a function of animal position on the familiar track. Pink areas indicate the three reward zones, or RZ1-3. **G.** Residual firing rates as a function of distance to reward zone for NS interneurons (dark blue, left), WS interneurons (light blue, middle), and pyramidal cells (red, right). **H.** Population average percent change in normalized residual firing as a function of distance (left) or time (right) to reward zone (RZ, pink) for each cell type (NS interneurons ("NS Int."), dark blue, $n = 196$; WS interneurons ("WS Int."), light blue, $n = 95$; pyramidal cells ("Pyr."), red, $n = 1157$). Color-coded horizontal bar indicates position bins that are significantly ($P < 0.05$) different from zero or baseline using data pooled from all days (t-test followed

by Bonferroni's correction). Data represent mean $\pm$ SEM. **I.** Proportion of NS (left) and WS interneurons (middle) with a significant reduction in activity over days in the novel environment (green) and familiar environment (black) and proportion of pyramidal cells (right) with significant increase in activity in either environment. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text.




Figure 2-2. Goal-associated inhibitory reduction is required for new goal learning. **A.** Left, schematic of optogenetic experiments targeting a few parvalbumin neurons (blue) in the stratum pyramidale (str.pyr.). Right, PVxAi32 mice received goal stimulation (blue) in one novel environment and sham stimulation (orange) in another novel environment for three days each. Both stimulation conditions were the same except for the location in which blue light turned on. **B.** Disrupting normal goal-associated inhibitory firing reduction impairs learning new goal locations. Data represent mean $\pm$ SEM area under the receiver operating characteristic curves based on speed differentiation over three days of learning (goal stimulation,

blue, $n = 4$ mice; sham stimulation, orange, $n = 5$ mice; linear mixed-effects model followed by Tukey's correction). **C.** Performance for each stimulation condition: no-stimulation trials only ("No stim," far left), low-intensity stimulation trials only ("Low stim," center left), high-intensity stimulation trials only ("High stim," center), all stimulation trials ("L+H," center right), or trials with the stimulation duration of 5 or less seconds ("Short stim," far right). Linear mixed-effects model followed by Tukey's correction. **D.** Top row, normalized change in pyramidal cell firing from baseline over time with goal stimulation (blue) and sham stimulation (orange) across stimulation intensities. Data represent mean $\pm$ SEM. Baseline firing rates were defined as the average firing rate in the 1-second window prior to stimulation onset. Bottom row, mean firing rates in the 2 second period after AZ for no stimulation or light onset across stimulation intensities. Goal stimulation, $n = 12$ cells; sham stimulation, $n = 15$ cells. Baseline firing rates were defined as the average firing rate in the 1-second window prior to stimulation onset. Goal stimulation normalized pyramidal cell firing rate change percentiles: no stimulation = [-0.50, -0.12, -0.0050, 0.080, 0.46], $n = 264$ cells; low stimulation = [-0.55, -0.15, -0.038, 0.055, 0.58], $n = 264$ cells; high stimulation = [-0.63, -0.21, -0.075, 0.031, 0.53], $n = 264$ cells. Sham stimulation percentiles: no stimulation = [-0.51, -0.093, 0.0085, 0.095, 0.50], $n = 313$ cells; low stimulation = [-0.60, -0.18, -0.059, 0.037, 0.54], $n = 302$ cells; high stimulation = [-0.72, -0.18, -0.048, 0.031, 0.50], $n = 313$ cells. **E.** Same as in D, but for PV interneurons. Goal stimulation normalized firing rate change percentiles: no stimulation = [-0.13, -0.094, -0.063, -0.017, 0.073], $n = 12$ cells; low stimulation = [0.18, 0.25, 0.31, 0.43, 0.52], $n = 12$ cells; high stimulation = [0.28, 0.37, 0.51, 0.57, 0.66], $n = 12$ cells. Sham stimulation percentiles: no stimulation = [-0.098, -0.0615, -0.014, 0.017, 0.11], $n = 15$ cells; low stimulation = [0.056, 0.24, 0.34, 0.46, 0.59], $n = 15$ cells; high stimulation = [-0.1271, 0.3345, 0.4020, 0.5524, 0.6759], $n = 15$ cells. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-3. Learning-dependent reduction in inhibitory activity precedes excitatory reorganization to overrepresent new goals. **A.** Heatmap of normalized residual firing rates as a function of distance to the novel RZ across three days of learning (from left to right) for NS interneurons (top row), WS interneurons (middle row), and pyramidal cells (bottom row). **B.** Percent change in normalized residual firing rates from

baseline for NS interneurons (top, dark blue), WS interneurons (middle, light blue), and pyramidal cells (bottom, red) as a function of distance to RZ (pink dotted lines) in the novel environment over three days. Shades of color indicate days 1-3. Data represent mean $\pm$ SEM. Color-coded horizontal bar indicates position bins that are significantly ($P < 0.05$) different from zero or baseline using data pooled from all days (t-test followed by Bonferroni's correction). **C.** Pearson's correlation of ratemaps across trials for goal-representing cells in the ("Fam," black) and novel ("Nov," green) environments; linear mixed-effects model followed by Tukey's correction. **D.** Spatial information carried by goal-representing pyramidal cells is significantly increased from day 1 to day 3 and from day 2 to day 3 in both novel (shades of green) and familiar (shades of black) environments. Novel spatial information percentiles on day 1, [0.028, 0.16, 0.38, 0.78, 2.4], $n = 133$; day 2, [0.0217, 0.20, 0.42, 0.81, 3.0], $n = 177$; day 3, [0.08, 0.32, 0.57, 1.1, 9.7], $n = 192$. Familiar spatial information percentiles on day 1, [0.032, 0.39, 1.1, 2.1, 13.8], $n = 179$; day 2, [0.056, 0.39, 0.79, 2.4, 8.5], $n = 156$; day 3, [0.086, 0.62, 1.5, 2.9, 7.3], $n = 202$; linear mixed-effects model followed by Tukey's correction. **E.** Proportion of goal-representing pyramidal cells does not significantly change over days in the familiar (black, $n = 13, 15, 14$ sessions for days 1-3) or novel (green, $n = 13, 14, 14$ sessions for days 1-3) environment; linear mixed-effects model. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-4. Inhibitory reduction co-occurs with sharper and more stable new spatial representation by pyramidal cells. A-B. Examples of goal-representing cells (“Goal cell”) or non-goal cells during sessions with goal stimulation (A, “Goal stim,” blue) and sham stimulation (B, “Sham stim,” orange) in a novel environment. Vertical bars in pink on the x-axis indicate the three reward zones (“RZ”). C. Goal stimulation disrupts stabilization of new place fields. Pearson’s correlation coefficients for rate map correlations across trials for all goal-representing pyramidal cells (“goal cells”) in goal stimulation

(“Goal stim,” blue) and sham stimulation (“Sham stim,” orange). Goal stimulation novel rate map correlation percentiles, [0.068, 0.18, 0.22, 0.28, 0.59], $n = 121$; sham stimulation correlation percentiles, [0.076, 0.20, 0.28, 0.37, 0.65], $n = 169$; linear mixed-effects model with Tukey’s correction. **D.** Goal stimulation disrupts new goal representation over days. Cumulative distribution of spatial information for all goal-representing pyramidal cells recorded in sessions with goal stimulation over three days (“Goal stim D1-3”, shades of blue) or sham stimulation (“Sham stim D1-3”, shades of orange) in the novel environment. Goal stimulation spatial information percentiles on day 1, [0.064, 0.15, 0.22, 0.36, 0.96], $n = 45$; day 2, [0.046, 0.18, 0.27, 0.41, 0.91], $n = 36$; day 3, [0.025, 0.11, 0.20, 0.31, 1.2], $n = 36$ cells. Sham stimulation spatial information percentiles on day 1, [0.025, 0.097, 0.19, 0.45, 0.96], $n = 51$; day 2, [0.044, 0.098, 0.29, 0.45, 2.48], $n = 60$; day 3, [0.047, 0.20, 0.40, 0.71, 2.9], $n = 54$ cells; linear mixed-effects model with Tukey’s correction. **E.** Proportion of goal-representing pyramidal cells does not significantly change with goal stimulation (blue) or sham stimulation (orange). Goal stimulation novel rate map correlation percentiles, [-0.0066, 0.17, 0.21609, 0.30, 0.62], $n = 143$; sham stimulation correlation percentiles, [0.044, 0.21, 0.29, 0.40, 0.86], $n = 145$; linear mixed-effects model with Tukey’s correction. **F.** Same as C, but for pyramidal cells with significant spatial modulation outside goals (“non-goal cells”). Goal stimulation spatial information percentiles on day 1, [0.027, 0.11, 0.17, 0.31, 1.0], $n = 54$; day 2, [0.041, 0.14, 0.25, 0.49, 0.89], $n = 38$; day 3, [0.064, 0.19, 0.29, 0.44, 0.80], $n = 45$ cells. Sham stimulation spatial information percentiles on day 1, [0.032, 0.085, 0.17, 0.36, 3.4], $n = 42$; day 2, [0.038, 0.15, 0.32, 0.46, 1.6], $n = 41$; day 3, [0.054, 0.20, 0.41, 0.74, 2.5], $n = 56$ cells. **G.** Same as D, but for non-goal cells. **H.** Same as E, but for non-goal cells. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-5. Inhibitory reduction is required for goal-related reactivation during sharp-wave ripples. A.

Example filtered traces of SWRs (marked in stars) in the local field potentials with (blue) or without (black) optogenetic stimulation. **B.** Overall SWR rate slightly decreased in goal stimulation ("Goal stim," blue) than sham stimulation ("Sham stim," orange). Sessions with at least 10 ripples detected were included. Goal stimulation percentiles, [0, 0.0023, 0.014, 0.026, 0.062], $n = 20$ sessions; sham stimulation percentiles, [0, 0.0093, 0.022, 0.052, 0.11], $n = 20$ sessions; linear mixed-effects model. **C.** Coactivation probability of goal-representing pyramidal cells during SWR activity in sessions with goal (blue) or sham (orange) stimulation. Goal stimulation percentiles, [0, 0, 0, 0, 0.38], $n = 111$ pairs of goal-modulated

pyramidal cells; sham stimulation, [0, 0, 0.043, 0.10, 0.55], $n = 516$ pairs of goal-modulated pyramidal cells; linear mixed-effects model. **D.** SWR power expressed as the number of standard deviations from the mean in goal stimulation (blue, percentiles, [3.01, 3.20, 3.39, 3.61, 6.26], $n = 75$ ripples) and sham stimulation (orange, percentiles, [3.01, 3.19, 3.50, 3.77, 10.87], $n = 122$ ripples); linear mixed-effects model. Note that part of the y-axis is removed to show the distributions; no data point was removed. **E.** Duration of SWRs in goal stimulation (blue, percentiles, [0.033, 0.06, 0.075, 0.091, 0.15], $n = 54$ ripples) and sham stimulation (orange, percentiles = [0.037, 0.062, 0.075, 0.10, 1.00], $n = 80$); linear mixed-effects model. Note that part of the y-axis is removed to show the distributions; three data points from sham stimulation were removed. **F.** Goal stimulation ("Goal stim," blue) results in fewer SWRs that represent positions near goals than sham stimulation ("Sham stim," orange). Heatmap of spatial probability of decoded ripple content. Each row is a single ripple event. **G.** SWRs during sham stimulation are more likely to represent positions near goals rather than far from goals, whereas no such preference is observed with goal stimulation. Proportions of SWRs with near-goal versus far-goal content ("Goal stim," blue, $n = 8$ sessions); sham stimulation proportions ("Sham stim," orange, $n = 12$ sessions). Wilcoxon signed-ranked test for near versus far goal representation. **H.** Same as in F, but for WT mice in the familiar ("Fam," black) and novel ("Nov," green) environments. **I.** Same as in G, but for WT mice. Familiar proportions of SWRs with near-goal versus far-goal content (black, $n = 36$ sessions); novel proportions (green, $n = 26$ sessions). Wilcoxon signed-ranked test for near versus far goal representation. $*P < 0.05$, $**P < 0.01$, $***P < 0.005$, $****P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].

Supplementary Figures




Figure 2-S1. Mice rapidly learn new goal locations in a novel virtual reality spatial learning task. **A.** Experimental timeline. Box, typical behavior of a well-trained animal. Data represent mean $\pm$ SEM change in lick rate (purple) and movement speed (yellow) of an animal, averaged over trials ($n = 80$) in a single session in the familiar environment. **B-D.** Example animal's speed distributions in AZ (black or green) or NRZ (tan) and respective receiver operating characteristic curves over three days during no-VR baseline (B), familiar (C), or novel navigation (D). **E.** Area under the receiver operating characteristic curves based on either speed (left), lick latency (middle), or lick rate (right). Error bars indicate mean $\pm$ SEM. Each line indicates a single animal per environment (linear mixed-effects model followed by Tukey's correction, $n = 14$; 2 novel environments per animal).




Figure 2-S2. Cell type classification. **A.** Top, 64-channel probe was inserted in stratum pyramidale of CA3 in head-fixed mice. Bottom, an example histology image taken at 10X, probe tract marked with the fluorescent marker DiI (pink) and the nuclei marker DAPI (blue). White scale bar indicates 500 μm . **B.** Distribution of autocorrelogram tau rise times and spike waveform trough-to-peak times of all recorded units color-coded by cell type (399 NS interneurons, dark blue; 237 WS interneurons, light blue; 2559 pyramidal cells, red). Interneurons with narrow and wide spike widths most likely correspond to PV basket cells and somatostatin-positive interneurons, respectively (English et al., 2017; Royer et al., 2012; Stark et al., 2014). Classification is based in part by the location of somas primarily found in our site of recording

and (stratum pyramidale), although we cannot rule out the possibility of a subset of cholecystokinin interneurons with wide spike widths that also have somas located in stratum pyramidale (Pelkey et al., 2017). **C.** Spike waveforms color-coded by cell type. Each line indicates a single unit. **D.** Distribution of firing rates color-coded by cell type. **E.** Distribution of burst index as used by Royer et al. 2012.






Figure 2-S3. Residual firing after estimating and removing contributions of speed and licking activity on neuronal firing rates. **A.** Top, model construction for performing multiple linear regression. Matrices of the same size (in $n\text{Trials} \times n\text{PositionBins}$) were created for firing rates, speed, and lick rate to fit the linear model per neuron. Bottom, example unit's model (illustrated in a grid) plotted against the trial data (blue circles) with its r -squared value on top right. **B.** Heatmaps of raw (top row) and residual (bottom row) firing rates for all recorded NS interneurons (left, dark blue), WS interneurons (middle, light blue), and pyramidal cells (right, red). Each row is a single unit. **C.** Percent variance explained by the multiple linear regression model fit for each unit, separated by cell type (NS interneurons, left, dark blue; WS interneurons, middle, light blue; pyramidal cells, right, red). Less than 25% of variance in the neural firing rate data is explained by movement speed and lick rate. **D.** Top row, change in residual firing rate from baseline for NS interneurons (dark blue) across first and second halves of day 1 (sessions 1-2), day 2 (sessions 3-4), and day 3 (sessions 5-6). Data represent mean $\pm$ SEM. Vertical dotted line indicates the start of RZ (RZ was 0-10 degrees). Color-coded horizontal bar indicates position bins that are significantly ($P < 0.05$) different from zero or baseline (t-test followed by Bonferroni's correction). Middle row, same as top, but for WS interneurons. Bottom row, same as top, but for pyramidal cells. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text.

A Compare firing rates between no-task and during-task periods at similar speeds



B Compare firing rates during reward approach across speed quartiles



C Compare firing rates between goal vs non-goal cues at similar speeds




Figure 2-S4. Inhibitory firing reduction cannot be explained by position-related changes in movement speed or lick rate. A-C. Analyses to control for the effects of speed on interneuron firing rates. **A.** Comparison of firing rates of NS interneurons in speed-matched trials during no-VR and familiar

navigation across speed quartiles. Linear mixed-effects model. **B.** Dividing the periods of familiar reward approach by speed quartiles. Firing rates do not change significantly across speed quartiles. **C.** Two animals exposed to a virtual reality manipulation in which the visual cues stayed frozen for at least 3 seconds upon RZ entry. *C*₁₋₂, Comparison of trial-averaged lick rate (top), movement speed (middle), or NS inhibitory firing rate (bottom) as a function of time to RZ or freeze onset between normal (grey, *n* = 43) and screen-freeze trials (teal, *n* = 30). Tan shaded bar indicates the time period during which quantification was done. Data represent mean $\pm$ SEM. For firing rates, only NS interneurons (*n* = 8 and 7 for normal and screen-freeze trials, respectively) with stable firing rates were selected and averaged across trials in animal 1 (*C*₁) and animal 2 (*C*₂; *n* = 5 and 4 NS interneurons for 34 normal and 29 screen-freeze trials). *C*₃₋₄, Quantification of firing rates within a 1-second time window (tan shaded bar) after arriving at the RZ during which movement speed was similar in the animal as in *C*₁ and the animal as in *C*₂. Bar plots with error bars indicate mean $\pm$ SEM; two-sided *t*-test. **P* < 0.05, ***P* < 0.01, ****P* < 0.005, *****P* < 0.001. Statistical details are available in the supplementary text.




Figure 2-S5. Goal-associated inhibitory firing reductions are more pronounced during correct trials than incorrect trials. **A.** Top row, normalized change in residual firing rates from baseline activity as a function of distance to the familiar RZ averaged across at least 5 correct (black) or incorrect (burgundy) trials in the familiar environment for each of NS interneurons (“NS Int.,” left), WS interneurons (“WS Int.,” middle), and pyramidal cells (“Pyr.,” right) from 7 WT mice. Data represent mean $\pm$ SEM. Bottom row, quantification of average residual firing rate change at the AZ (from -10 to 0 degrees) during correct (black) and incorrect (burgundy) trials. NS interneurons (left): familiar correct percentiles = [-0.97, -0.38, -0.10, 0.20, 0.83], $n = 196$ cells; incorrect = [-0.78, -0.24, 0.10, 0.26, 0.83], $n = 75$ cells; WS interneurons (middle): correct = [-0.86, -0.36, -0.12, 0.19, 0.68], $n = 95$ cells; incorrect = [-0.73, -0.29, -0.063, 0.19, 0.56], $n = 39$ cells; pyramidal cells (right): correct = [-0.95, -0.25, -0.015, 0.20, 0.91], $n = 1168$ cells; incorrect = [-0.98, -0.19, -0.0069, 0.18, 0.88], $n = 375$ cells. Each dot represents a single unit. Black circles with error bars represent mean $\pm$ SEM. **B.** Same as A, but for the novel environment. NS interneurons (left): novel correct percentiles = [-0.77, -0.21, 0.0050, 0.26, 0.80], $n = 193$ cells; incorrect = [-0.82, -0.25, -

0.031, 0.20, 0.86], $n = 146$ cells; WS interneurons (middle): correct = [-0.79, -0.18, 0.093, 0.34, 0.78], $n = 96$ cells; incorrect = [-0.77, -0.22, 0.010, 0.16, 0.67], $n = 84$ cells; pyramidal cells (right): correct = [-0.89, -0.24, -0.012, 0.20, 0.85], $n = 1144$ cells; incorrect = [-0.94, -0.21, -0.0065, 0.21, 0.84], $n = 898$ cells. $*P < 0.05$, $**P < 0.01$, $***P < 0.005$, $****P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-S6. Validation of optogenetic experiments in vivo. **A.** Confocal images taken at 20X showing CA3 stratum pyramidale in center. Fluorescence markers indicate nuclei (DAPI, blue), channelrhodopsin-2 (eYFP, green), and parvalbumin (red). White scale bar indicates 20 μm . We found $73.06 \pm 27.17\%$ of ChR2+ cells to be co-expressing PV, and $87.92 \pm 18.87\%$ of PV+ cells co-expressing channelrhodopsin-2. **B.** Left, schematic of in vivo recording with an optoelectrode inserted in CA3 stratum pyramidale. Right, spike raster plot of an example optically tagged unit as a function of time to light onset (blue shaded area).

C. Average change in the normalized firing rates from baseline across stimulation intensities (light to dark shades of purple, from 0 to 12 mW/mm²) per cell type. *PV interneurons*: no stimulation (“No stim”), percentiles, [-16.09, -7.91, -4.00, -1.36, 14.52], *n* = 44; low stimulation (“Low stim”), [1.94, 1.94, 2.63, 3.32, 3.32], *n* = 2; high stimulation (“High stim”), [-0.7751, 5.55, 14.36, 30.19, 67.69], *n* = 44. *NS interneurons*: no stimulation, [-16.09, -6.83, -3.07, -0.40, 14.52], *n* = 75; low stimulation, [-10.87, -6.60, -2.81, 1.94, 3.32], *n* = 6; high stimulation, [-20.15, -0.79, 3.59, 17.75, 67.69], *n* = 71. *WS interneurons*: no stimulation, [-8.37, -1.10, -0.22, 0.29, 10.08], *n* = 36; low stimulation, [-10.84, -0.30, 0.053, 0.52, 3.93], *n* = 11; high stimulation, [-10.87, -3.29, -0.82, -0.0094, 5.44], *n* = 36. *Pyramidal cells*: no stimulation, [-7.70, -0.17, 0.020, 0.20, 2.69], *n* = 36; low stimulation, [-1.94, -0.26, 0.00, 0.33, 2.47], *n* = 11; high stimulation, [-8.84, -0.47, -0.088, 0.016, 4.47], *n* = 36. Baseline was defined as the mean firing rate in the 1-second window before AZ or light onset; linear mixed-effects model. Data represent mean ± SEM. **P* < 0.05, ***P* < 0.01, ****P* < 0.005, *****P* < 0.001. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].






Figure 2-S7. Effects of optogenetic stimulation intensity on speed and licking activity. **A.** Top row, normalized change in lick rate from baseline levels as a function of time to optogenetic stimulation or AZ entry (dotted line at 0) for novel days 1-3, or “Nov Day 1-3” (green) and the last day of familiar navigation (black, “Fam Day 3”). Lick rate in the novel environment: day 1, no stimulation (“No stim”) percentiles = [-0.70, -0.0063, 0.077, 0.27, 0.68], $n = 118$ trials; low stimulation (“Low stim”) = [-0.45, 0.00, 0.079, 0.26, 0.64], $n = 59$; high stimulation (“High stim”) = [-0.69, 0.00, 0.16, 0.29, 0.68], $n = 62$; day 2, no stimulation = [-0.52, -0.0010, 0.11, 0.27, 0.56], $n = 115$; low stimulation = [-0.33, 0.00, 0.14, 0.25, 0.53], $n = 63$; high stimulation = [-0.47, 0.00, 0.082, 0.299, 0.59], $n = 67$; day 3, no stimulation = [-0.77, -0.027, 0.13, 0.28, 0.71], $n = 179$; low stimulation = [-0.39, -0.055, 0.072, 0.23, 0.73], $n = 96$; high stimulation = [-0.41, -0.032, 0.064, 0.24, 0.66], $n = 99$. Lick rate in the familiar environment: day 3, no stimulation = [-0.68, 0.0168, 0.19, 0.41, 0.78], $n = 559$; low stimulation = [-0.27, 0.056, 0.21, 0.35, 0.62], $n = 42$; high stimulation = [-0.27, 0.0082, 0.23, 0.40, 0.75], $n = 38$. Light to dark shades of color indicate low to high stimulation intensity. Data represent mean $\pm$ SEM. Bottom row, quantification of mean change in lick rate after stimulation or AZ onset. Each colored circle indicates a trial. Black line and error bars indicate mean $\pm$ SEM; linear mixed-effects model. **B.** Same as A, but for speed. Speed in the novel environment: day 1, no stimulation percentiles = [-0.48, -0.15, -0.034, 0.059, 0.50], $n = 179$; low stimulation = [-0.54, -0.15, -0.0381, 0.038, 0.45], $n = 99$; high stimulation = [-0.62, -0.16, -0.013, 0.10, 0.47], $n = 101$; day 2, no stimulation = [-0.62, -0.24, -0.081, 0.042, 0.64], $n = 179$; low stimulation = [-0.63, -0.24, -0.15, -0.048, 0.38], $n = 97$; high stimulation = [-0.52, -0.14, -0.047, 0.14, 0.47], $n = 103$; day 3, no stimulation = [-0.43, -0.10, -0.015, 0.054, 0.36], $n = 200$; low stimulation = [-0.52, -0.18, -0.090, 0.074, 0.44], $n = 104$; high stimulation = [-0.52, -0.18, -0.090, 0.074, 0.44], $n = 106$. Speed in the familiar environment: day 3, no stimulation = [-0.71, -0.33, -0.22, -0.13, 0.62], $n = 560$; low stimulation = [-0.66, -0.28, -0.20, -0.073, 0.14], $n = 43$; high stimulation = [-0.48, -0.32, -0.21, -0.10, 0.095], $n = 39$; linear mixed-effects model. **C.** Performance in the familiar environment based on speed (left), lick latency (middle), and lick rate (right) across different conditions and trial types; trials during baseline periods without visual cues (“No VR”),

trials with no stimulation (“No stim”), low stimulation (“L”), high stimulation (“H”), both low and high stimulation (“L+H”), and all trials (“All”). * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-S8. Sharp-wave ripple activity detected in WT mice. **A.** Example traces of local field potentials showing sharp-wave ripple events (marked with stars) detected in the familiar ("Fam," black) or novel ("Nov," green) environment. **B.** Distribution of ripple rate occurring in the familiar (black, percentiles, [0.00, 0.012, 0.036, 0.059, 0.17], $n = 41$ sessions) or novel (green, percentiles, [0.00, 0.018, 0.035, 0.061, 0.17], $n = 41$ sessions) environment.

0.16], $n = 42$ sessions) environment; linear mixed-effects model. Each data point indicates a recording session. **C.** Same as B, but for the probability of co-activation among pairs of goal-modulated pyramidal cells. Familiar percentiles, [0, 0, 0.063, 0.19, 1], $n = 4076$ pairs; novel = [0, 0, 0, 0.081, 0.89], $n = 3207$ pairs; linear mixed-effects model. Each data point indicates a pyramidal cell pair. **D.** Same as B, but for ripple power expressed as the number of standard deviations from the mean. Familiar percentiles, [3.00, 3.21, 3.54, 4.02, 7.63], $n = 598$ ripples; novel percentiles, [3.00, 3.21, 3.53, 4.09, 9.92], $n = 834$ ripples; linear mixed-effects model. Note our ripple detection requires the minimum ripple power of 3 standard deviations to be detected as a ripple event. Each data point indicates a ripple event. **E.** Same as B, but for the duration of ripples in seconds. Each data point indicates a ripple event. Familiar percentiles, [0.034, 0.061, 0.078, 0.1, 0.34], $n = 598$ ripples from 7 mice; novel percentiles, [0.034, 0.062, 0.081, 0.11, 0.31], $n = 834$ ripples from 7 mice; linear mixed-effects model. **F.** Change in normalized firing rate of NS interneurons ("NS Int.," dark blue), WS interneurons ("WS Int.," light blue), or pyramidal cells ("Pyr.," red) from baseline as a function of time to the mid-point of sharp-wave ripple (SWR), averaged across SWR events in the familiar (left) and novel (right) environment. Inset, zoomed-in view of normalized firing rate change, 1 second around SWR midpoint for each cell type. Data represent mean $\pm$ SEM. **G.** Comparison of normalized firing rate change between familiar (black) and novel (green) environments for NS interneurons (left), WS interneurons (middle), and pyramidal cells (right). Data represent mean $\pm$ SEM. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].




Figure 2-S9. Optogenetic stimulation does not significantly affect sharp-wave ripple generation, power, or duration. **A.** Change in normalized firing rate of NS interneurons (“NS Int.,” dark blue), WS interneurons (“WS Int.,” light blue), or pyramidal cells (“Pyr.,” red) from baseline as a function of time to the mid-point of sharp-wave ripple (SWR), averaged across SWR events in sessions with goal (“Novel Goal Stim,” left) or sham stimulation (“Novel Sham Stim,” right) in the novel environment. Inset, zoomed-in view of normalized firing rate change, 1 second around SWR midpoint for each cell type. Data represent

mean $\pm$ SEM. **B.** Comparison of normalized firing rate change between goal stimulation (“Goal stim,” blue) and sham stimulation (“Sham stim,” orange) environments for NS interneurons (left), WS interneurons (middle), and pyramidal cells (right). Data represent mean $\pm$ SEM. **C.** Distribution of ripple power expressed as the number of standard deviations from the mean with goal stimulation (blue) or no stimulation (black) in the familiar environment. Familiar goal stimulation percentiles, [3.0962, 3.3517, 3.7334, 3.9702, 4.4851], $n = 14$ ripples; familiar no stimulation percentiles, [3.0073, 3.171, 3.4287, 3.7714, 8.3701], $n = 347$ ripples; linear mixed-effects model. Note that part of the y-axis is removed to show the distributions; no data point was removed. **D.** Distribution of ripple duration with goal stimulation (blue) or no stimulation (black) in the familiar environment. Familiar goal stimulation percentiles, [0.037, 0.065, 0.080, 0.087, 0.26], $n = 14$ ripples; familiar no stimulation percentiles, [0.035, 0.059, 0.075, 0.095, 1.84], $n = 347$ ripples; linear mixed-effects model. Note that part of the y-axis is removed to show the distributions; no data point was removed. * $P < 0.05$, ** $P < 0.01$, *** $P < 0.005$, **** $P < 0.001$. Statistical details are available in the supplementary text. All percentiles represent [minimum, 25th, median, 75th, maximum].

Table 2-1. Genotypes, type of treadmill, recording sessions with novelty, single units, sharp-wave ripples, recording duration per animal

Animal	Genotype	Treadmill type	Single units	Narrow interneurons (PVs)	Wide interneurons	Pyramidal cells	Sharp-wave ripples	Stopped periods (min)
N11	WT	Linear	155	33 (N/A)	8	114	36	43.93
N18	WT	Linear	208	23 (N/A)	15	170	646	246.04
N21	WT	Linear	179	17 (N/A)	14	148	498	190.58
N24	WT	Linear	227	27 (N/A)	35	165	241	84.33
N45	WT	Spherical	201	34 (N/A)	8	159	8	14.66
N46	WT	Spherical	293	34 (N/A)	10	249	2	4.90
N47	WT	Spherical	223	35 (N/A)	12	176	1	15.87
N48	PVxAi32	Spherical	181	23 (13)	6	152	11	20.85
N50	PVxAi32	Spherical	210	34 (28)	16	160	32	21.46
N52	PVxAi32	Spherical	196	24 (12)	14	158	139	49.12
N53	PVxAi32	Spherical	197	21 (14)	14	162	3	7.09
N54	PVxAi32	Spherical	208	25 (10)	15	168	105	68.13
N57	PVxAi32	Spherical	192	15 (10)	27	150	201	78.82
N61	PVxAi32	Spherical	97	12 (4)	5	80	51	67.26
N62	PVxAi32	Spherical	154	20 (11)	9	125	108	106.13
N63	PVxAi32	Spherical	165	17 (7)	14	134	90	74.18
N65	PVxAi32	Spherical	109	5 (1)	15	89	108	59.87

Supplementary Text

Details on statistical analyses

Statistical significance abbreviations: ns (not significant) $P > 0.05$, $*P < 0.05$, $**P < 0.01$, $***P < 0.005$, $****P < 0.001$.

MAIN FIGURES

Fig. 2-1

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
Fig. 1E; fig. S1E	WT: speed	n=7	Linear mixed effects model; Tukey correction	Day	0.01	**	$F_{2,111} = 5.3245$
				Environment	$< 2.2e-16$	****	$F_{2,111} = 232.4114$
				Interaction	0.01	**	$F_{2,111} = 3.5107$
				No VR: Day 1 vs 2	0.81	ns	t-ratio = -0.621
				No VR: Day 1 vs 3	1.0	ns	t-ratio = 0.043
				No VR: Day 2 vs 3	0.78	ns	t-ratio = 0.665
				Familiar: Day 1 vs 2	0.76	ns	t-ratio = -0.699
				Familiar: Day 1 vs 3	0.98	ns	t-ratio = 0.207
				Familiar: Day 2 vs 3	0.64	ns	t-ratio = 0.906
				Novel: Day 1 vs 2	0.0002	****	t-ratio = -4.160
1I	WT: proportion of ramping units	NS: n=196 Wide: n=95 Pyr: n=1157	Linear mixed effects model; Tukey correction	Novel: Day 1 vs 3	0.0002	****	t-ratio = -4.189
				Novel: Day 2 vs 3	1.0	ns	t-ratio = -0.030
				Main position effect	$< 2.2e-16$	****	$F_{25,2351.9} = 12.763$
				Main position effect	0.00052	****	$F_{25,1050} = 2.2254$
				Main position effect	$< 2.2e-16$	****	$F_{25,10981} = 25.177$
				Day	0.17	ns	$F_{2,204.46} = 1.8055$
				Environment	$5.7e-11$	****	$F_{1,202.03} = 53.5878$
				Celltype	$2.9e-06$	****	$F_{2,203.37} = 13.6020$
				Day*Environment	0.15	ns	$F_{2,202.03} = 1.9426$
				Day*Celltype	0.01	*	$F_{4,203.18} = 2.8129$
				Environment*Celltype	0.0003	****	$F_{2,202.03} = 8.5743$
				Day*Environment*Celltype	0.26	ns	$F_{4,202.03} = 1.3280$
				NS: Fam 1 vs Nov 1	0.004	****	t-ratio = 4.287
				NS: Fam 2 vs Nov 2	0.13	ns	t-ratio = 3.194
				NS: Fam 3 vs Nov 3	1.0	ns	t-ratio = 0.353
				WS: Fam 1 vs Nov 1	0.006	**	t-ratio = 4.168
				WS: Fam 2 vs Nov 2	0.06	ns	t-ratio = 3.498
				WS: Fam 3 vs Nov 3	0.03	*	t-ratio = 3.685
				Pyr: Fam 1 vs Nov 1	1.0	ns	t-ratio = 0.645
				Pyr: Fam 2 vs Nov 2	1.0	ns	t-ratio = 0.704
				Pyr: Fam 3 vs Nov 3	1.0	ns	t-ratio = 0.791

Fig. 2-2

Panel	Data	Group size	Statistic method	Comparison	P-value	notation	F/t statistic
Fig. 2B	PVxAi32: speed, using <i>all</i> trials	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Day	0.01	**	$F_{2,17.091} = 6.4812$
				Stim condition	0.02	*	$F_{1,18.708} = 7.0143$
				Interaction	0.02	*	$F_{2,17.091} = 5.4060$
				Goal stim: Day 1 vs 2	0.99	ns	t-ratio = 0.150
				Goal stim: Day 1 vs 3	0.95	ns	t-ratio = -0.297
				Goal stim: Day 2 vs 3	0.90	ns	t-ratio = -0.447
				Sham stim: Day 1 vs 2	0.01	**	t-ratio = -3.523
				Sham stim: Day 1 vs 3	0.0003	****	t-ratio = -5.011
Fig. 2C, far left	PVxAi32: speed, using <i>no-stim</i> trials only	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Sham stim: Day 2 vs 3	0.32	ns	t-ratio = -1.487
				Day	0.05	*	$F_{1,17.134} = 3.6343$
				Stim condition	0.005	**	$F_{2,18.906} = 9.9883$
				Interaction	0.03	*	$F_{2,18.906} = 4.3143$
				Goal stim: Day 1 vs 2	0.76	ns	t-ratio = -0.348
				Goal stim: Day 1 vs 3	0.99	ns	t-ratio = -0.420
				Goal stim: Day 2 vs 3	0.67	ns	t-ratio = -0.072
				Sham stim: Day 1 vs 2	0.02	**	t-ratio = -2.973
Fig. 2C, center left	PVxAi32: speed, using <i>low-stim</i> trials only	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Sham stim: Day 1 vs 3	0.003	****	t-ratio = -4.913
				Sham stim: Day 2 vs 3	0.70	ns	t-ratio = -1.941
				Day	0.001	****	$F_{2,17.231} = 2.6652$
				Stim condition	0.010	*	$F_{1,17.073} = 10.3047$
				Interaction	0.04	*	$F_{2,17.073} = 3.9580$
				Goal stim: Day 1 vs 2	0.57	ns	t-ratio = -1.032
				Goal stim: Day 1 vs 3	0.52	ns	t-ratio = -1.097
				Goal stim: Day 2 vs 3	1.0	ns	t-ratio = -0.065
Fig. 2C, center	PVxAi32: speed, using <i>high-stim</i> trials only	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Sham stim: Day 1 vs 2	0.01	**	t-ratio = -3.367
				Sham stim: Day 1 vs 3	0.0001	****	t-ratio = -5.445
				Sham stim: Day 2 vs 3	0.12	ns	t-ratio = -2.078
				Day	0.22	ns	$F_{2,17.087} = 1.6424$
				Stim condition	0.10	ns	$F_{1,18.678} = 2.9299$
				Interaction	0.20	ns	$F_{2,17.087} = 1.7363$
				Goal stim: Day 1 vs 2	0.98	ns	t-ratio = 0.170
				Goal stim: Day 1 vs 3	1.0	ns	t-ratio = 0.045
Fig. 2C, center right	PVxAi32: speed, using <i>low+high stim</i> trials only	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Goal stim: Day 2 vs 3	0.99	ns	t-ratio = -0.126
				Sham stim: Day 1 vs 2	0.45	ns	t-ratio = -1.231
				Sham stim: Day 1 vs 3	0.03	*	t-ratio = -2.745
				Sham stim: Day 2 vs 3	0.30	ns	t-ratio = -1.514
				Day	0.006	**	$F_{2,15.367} = 7.2266$
				Stim condition	0.02	*	$F_{1,17.226} = 6.1607$
				Interaction	0.02	*	$F_{2,15.367} = 4.8395$
				Goal stim: Day 1 vs 2	0.93	ns	t-ratio = -0.367
Fig. 2C, far right	PVxAi32: speed, using <i>short stim</i> trials only	Goal stim: n=4; Sham stim: n=5	Linear mixed effects model; Tukey correction	Goal stim: Day 1 vs 3	0.90	ns	t-ratio = -0.443
				Goal stim: Day 2 vs 3	1.0	ns	t-ratio = -0.076
				Sham stim: Day 1 vs 2	0.006	**	t-ratio = -3.697
				Sham stim: Day 1 vs 3	0.0007	****	t-ratio = -4.682
				Sham stim: Day 2 vs 3	0.64	ns	t-ratio = -0.919
				Day	0.05	*	$F_{2,17.069} = 3.5468$
				Stim condition	0.0009	****	$F_{1,18.736} = 15.3374$
				Interaction	0.01	*	$F_{2,17.096} = 5.7282$
Fig. 2D	PvxAi32: normalized pyramidal cell	Goal stim: n=264; Sham stim: n=314	Linear mixed effects model; Tukey correction	Goal stim: Day 1 vs 2	0.61	ns	t-ratio = -0.367
				Goal stim: Day 1 vs 3	0.98	ns	t-ratio = -0.443
				Goal stim: Day 2 vs 3	0.74	ns	t-ratio = -0.076
				Sham stim: Day 1 vs 2	0.01	*	t-ratio = -3.267
				Sham stim: Day 1 vs 3	0.002	****	t-ratio = -4.221
				Sham stim: Day 2 vs 3	0.61	ns	t-ratio = -0.955
				Stim condition	0.41	ns	$F_{1,668.25} = 0.6926$
				Stim intensity	8.8e-14	****	$F_{2,1276.90} = 30.7748$
				Interaction	0.08	ns	$F_{2,1276.90} = 2.5956$

Fig. 2E	firing change from baseline			No stim: goal vs sham stim	0.79	ns	t-ratio = -1.285
				Low stim: goal vs sham stim	0.84	ns	t-ratio = 1.202
				High stim: goal vs sham stim	0.66	ns	t-ratio = -1.501
	PVxAi32: normalized PV firing change from baseline	Goal stim: n=12; Sham stim: n=15	Linear mixed effects model; Tukey correction	Stim condition	0.57	ns	$F_{1,20,767} = 0.3281$
				Stim intensity	< 2e-16	****	$F_{2,50,000} = 198.1915$
				Interaction	0.09	ns	$F_{2,50,000} = 2.4756$
				No stim: goal vs sham stim	0.71	ns	t-ratio = -1.434
				Low stim: goal vs sham stim	0.98	ns	t-ratio = -0.734
				High stim: goal vs sham stim	0.96	ns	t-ratio = 0.822

Fig. 2-3

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
Fig. 3B, top	WT: change in residual firing, novel day 1	NS: n=22	Linear mixed effects model	Main position effect	0.21	ns	$F_{25,525} = 1.2298$
	WT: change in residual firing, novel day 2	NS: n=34	Linear mixed effects model	Main position effect	1.5e-12	****	$F_{25,825} = 4.6196$
	WT: change in residual firing, novel day 3	NS: n=36	Linear mixed effects model	Main position effect	2.8e-09	****	$F_{25,875} = 3.7423$
	WT: change in residual firing, novel day 3	NS: n=36	Linear mixed effects model	Main position effect	2.8e-09	****	$F_{25,875} = 3.7423$
	WT: change in residual firing, all novel days combined	NS: n=90	Linear mixed effects model	Main position effect	2.2e-16	****	$F_{25,2276.9} = 6.4304$
Fig. 3B, middle	WT: change in residual firing, novel day 1	WS: n=17	Linear mixed effects model	Main position effect	5.5e-16	****	$F_{25,400} = 5.8982$
	WT: change in residual firing, novel day 2	WS: n=12	Linear mixed effects model	Main position effect	1.0	ns	$F_{25,275} = 0.336$
	WT: change in residual firing, novel day 3	WS: n=14	Linear mixed effects model	Main position effect	0.91	ns	$F_{25,325} = 0.6447$
	WT: change in residual firing, all	WS: n=43	Linear mixed effects model	Main position effect	1.3e-07	****	$F_{25,1050} = 3.2703$

	novel days combined						
Fig. 3B, bottom	WT: change in residual firing, novel day 1	Pyr.: n=108	Linear mixed effects model	Main position effect	0.014	*	$F_{25,2700.9} = 1.7242$
	WT: change in residual firing, novel day 2	Pyr.: n=144	Linear mixed effects model	Main position effect	0.0027	***	$F_{25,3652.9} = 1.973$
	WT: change in residual firing, novel day 3	Pyr.: n=181	Linear mixed effects model	Main position effect	1.3e-05	*****	$F_{25,4500} = 2.6685$
	WT: change in residual firing, all novel days combined	Pyr.: n=405	Linear mixed effects model	Main position effect	3.2e-08	****	$F_{25,10931} = 3.35$
Fig. 3C	WT: rate map correlation of goal cells	Familiar: n=179, 156, 202; Novel: n=133, 177, 192 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Novel Day	1.5e-05	****	$F_{2,1175.0} = 11.246$
				Environment	2.3e-07	****	$F_{1,1091.6} = 27.081$
				Familiar: Day 1 vs 2	0.90	ns	t-ratio = -1.061
				Familiar: Day 1 vs 3	<.0001	****	t-ratio = -4.514
				Familiar: Day 2 vs 3	0.01	**	t-ratio = -3.381
				Novel: Day 1 vs 2	0.90	ns	t-ratio = -1.061
				Novel: Day 1 vs 3	<.0001	****	t-ratio = -4.514
				Novel: Day 2 vs 3	0.01	**	t-ratio = -3.381
				Familiar 1 vs Novel 1	<.0001	****	t-ratio = -5.204
Fig. 3D	WT: spatial information of goal cells	Familiar: n=179, 156, 202; Novel: n=133, 177, 192 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Familiar 2 vs Novel 2	<.0001	****	t-ratio = -5.204
				Familiar 3 vs Novel 3	<.0001	****	t-ratio = -5.204
				Novel Day	0.0008	****	$F_{2,1182.7} = 7.2394$
				Environment	< 2e-16	****	$F_{1,1096.0} = 192.7330$
				Familiar: Day 1 vs 2	0.99	ns	t-ratio = 0.613
				Familiar: Day 1 vs 3	0.04	*	t-ratio = -2.911
				Familiar: Day 2 vs 3	0.01	**	t-ratio = -3.524
				Novel: Day 1 vs 2	0.99	ns	t-ratio = 0.613
				Novel: Day 1 vs 3	0.04	*	t-ratio = -2.911
Fig. 3E	WT: proportion of goal cells	Familiar: n=13, 15, 14; Novel: n=13, 14, 14 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Novel: Day 2 vs 3	0.006	**	t-ratio = -3.524
				Familiar 1 vs Novel 1	<.0001	****	t-ratio = 13.883
				Familiar 2 vs Novel 2	<.0001	****	t-ratio = 13.883
				Familiar 3 vs Novel 3	<.0001	****	t-ratio = 13.883
				Day	0.61	ns	$F_{2,70.237} = 0.4957$
Fig. 3E	WT: proportion of goal cells	Familiar: n=13, 15, 14; Novel: n=13, 14, 14 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Environment	0.15	ns	$F_{1,70.003} = 2.1038$
				Interaction	0.09	ns	$F_{2,70.003} = 2.5532$

Fig. 2-4

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
Fig. 4C	PvxAi32: rate map correlation of goal-representing cells	Goal stim: n=46, 36, 36; Sham stim: n= 51, 60, 54 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Day	0.09	ns	$F_{2,283.67} = 2.3965$
				Stim condition	0.58	ns	$F_{2,280.67} = 0.3068$
				Interaction	0.001	****	$F_{2,267.89} = 6.7213$
				Goal stim: Day 1 vs 2	0.97	ns	t-ratio = -0.787
				Goal stim: Day 1 vs 3	0.96	ns	t-ratio = -0.828
				Goal stim: Day 2 vs 3	1.0	ns	t-ratio = 0.004

				Sham stim: Day 1 vs 2	1.0	ns	t-ratio = -0.487
				Sham stim: Day 1 vs 3	0.006	**	t-ratio = -3.558
				Sham stim: Day 2 vs 3	0.04	*	t-ratio = -2.955
				Goal stim 1 vs Sham stim 1	1.0	ns	t-ratio = -0.529
				Goal stim 2 vs Sham stim 2	1.0	ns	t-ratio = -0.223
				Goal stim 3 vs Sham stim 3	0.02	*	t-ratio = -3.211
Fig. 4D	PvxAi32: spatial information of goal-representing cells	Goal stim: n=46, 36, 36; Sham stim: n= 51, 60, 54 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Day	0.09	ns	$F_{2,283.67} = 2.3965$
				Stim condition	0.58	ns	$F_{2,280.67} = 0.3068$
				Interaction	0.001	***	$F_{2,267.89} = 6.7213$
				Goal stim: Day 1 vs 2	0.99	ns	t-ratio = -0.590
				Goal stim: Day 1 vs 3	0.95	ns	t-ratio = 0.876
				Goal stim: Day 2 vs 3	0.73	ns	t-ratio = 1.399
				Sham stim: Day 1 vs 2	0.45	ns	t-ratio = -1.828
				Sham stim: Day 1 vs 3	0.0003	****	t-ratio = -4.331
				Sham stim: Day 2 vs 3	0.09	ns	t-ratio = -2.656
				Goal stim 1 vs Sham stim 1	0.18	ns	t-ratio = 2.333
				Goal stim 2 vs Sham stim 2	0.79	ns	t-ratio = 1.284
				Goal stim 3 vs Sham stim 3	0.12	ns	t-ratio = -2.523
Fig. 4E	PvxAi32: proportion of goal-representing cells	Goal stim: n=4; Sham stim: n=5 per day	Linear mixed-effects model	Day	0.44	ns	$F_{2,17.009} = 0.8580$
				Environment	0.24	ns	$F_{1,17.677} = 1.5097$
				Interaction	0.60	ns	$F_{2,17.009} = 0.5230$
Fig. 4F	PvxAi32: rate map correlation of non-goal-representing cells	Goal stim: n=54, 38, 45; Sham stim: n= 42, 41, 56 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Day	0.007	**	$F_{2,269.99} = 5.0851$
				Stim condition	0.04	*	$F_{1,253.74} = 4.2180$
				Interaction	0.07	ns	$F_{1,272.88} = 2.6775$
				Goal stim: Day 1 vs 2	0.97	ns	t-ratio = -0.787
				Goal stim: Day 1 vs 3	0.96	ns	t-ratio = -0.828
				Goal stim: Day 2 vs 3	1.0	ns	t-ratio = 0.004
				Sham stim: Day 1 vs 2	1.0	ns	t-ratio = -0.487
				Sham stim: Day 1 vs 3	0.006	**	t-ratio = -3.558
				Sham stim: Day 2 vs 3	0.04	*	t-ratio = -2.955
				Goal stim 1 vs Sham stim 1	1.0	ns	t-ratio = -0.529
				Goal stim 2 vs Sham stim 2	1.0	ns	t-ratio = -0.223
				Goal stim 3 vs Sham stim 3	0.02	*	t-ratio = -3.211
Fig. 4G	PvxAi32: spatial information of non-goal-representing cells	Goal stim: n=54, 38, 45; Sham stim: n= 42, 41, 56 for Day 1, 2, 3	Linear mixed-effects model; Tukey correction	Day	0.007	**	$F_{2,269.99} = 5.0851$
				Stim condition	0.04	*	$F_{1,253.74} = 4.2180$
				Interaction	0.07	ns	$F_{1,272.88} = 2.6775$
				Goal stim: Day 1 vs 2	0.97	ns	t-ratio = -0.787
				Goal stim: Day 1 vs 3	0.96	ns	t-ratio = -0.828
				Goal stim: Day 2 vs 3	1.0	ns	t-ratio = 0.004
				Sham stim: Day 1 vs 2	1.0	ns	t-ratio = -0.487
				Sham stim: Day 1 vs 3	0.006	**	t-ratio = -3.558
				Sham stim: Day 2 vs 3	0.04	*	t-ratio = -2.955
				Goal stim 1 vs Sham stim 1	1.0	ns	t-ratio = -0.529
				Goal stim 2 vs Sham stim 2	1.0	ns	t-ratio = -0.223
				Goal stim 3 vs Sham stim 3	0.02	*	t-ratio = -3.211
Fig. 4H	PvxAi32: proportion of non-goal-representing cells	Goal stim: n=4; Sham stim: n=5 per day	Linear mixed-effects model	Day	0.51	ns	$F_{1,19.008} = 0.4463$
				Environment	0.33	ns	$F_{1,19.116} = 0.9856$
				Interaction	0.52	ns	$F_{1,19.008} = 0.4269$

Fig. 2-5

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S5A, left	WT: AZ residual firing rate in Familiar	NS interneurons, n = 196, 75 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.05	*	$F_{1,124.84} = 3.9283$
fig. S5A, middle		WS interneurons, n = 95, 39 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.16	ns	$F_{1,65.966} = 1.9906$
fig. S5A, right		Pyramidal cells, n = 1168, 375 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.74	ns	$F_{1,691.2} = 0.109$
fig. S5B, left	WT: AZ residual firing rate in Novel	NS interneurons, n = 193, 146 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.11	ns	$F_{1,175.21} = 2.5855$
fig. S5B, middle		WS interneurons, n = 96, 84 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.03	*	$F_{1,93.442} = 4.9518$
fig. S5B, right		Pyramidal cells, n = 1144, 898 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.30	ns	$F_{1,1091.8} = 1.0851$

SUPPLEMENTARY FIGURES

fig. 2-S1

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S1E, left	WT: speed	n = 7	Linear mixed effects model; Tukey multiple correction	Day	0.006	**	$F_{2,111} = 5.3245$
				Environment	< 2.2e-16	****	$F_{2,111} = 232.4114$
				Interaction	0.01	**	$F_{2,111} = 3.5107$
				No VR: Day 1 vs 2	0.81	ns	t-ratio = -0.621
				No VR: Day 1 vs 3	1.0	ns	t-ratio = 0.043
				No VR: Day 2 vs 3	0.78	ns	t-ratio = 0.665
				Familiar: Day 1 vs 2	0.76	ns	t-ratio = -0.699
				Familiar: Day 1 vs 3	0.98	ns	t-ratio = 0.207
				Familiar: Day 2 vs 3	0.64	ns	t-ratio = 0.906
				Novel: Day 1 vs 2	0.0002	****	t-ratio = -4.160
fig. S1E, middle	WT: lick latency	n=7	Linear mixed effects model; Tukey multiple correction	Novel: Day 1 vs 3	0.0002	****	t-ratio = -4.189
				Novel: Day 2 vs 3	1.0	ns	t-ratio = -0.030
				Day	0.12	ns	$F_{2,110.03} = 2.1913$
				Environment	< 2.2e-16	†	$F_{2,110.03} = 53.2935$
				Interaction	0.006	**	$F_{2,110.03} = 3.7891$
				No VR: Day 1 vs 2	0.99	ns	t-ratio = -0.166
				No VR: Day 1 vs 3	0.92	ns	t-ratio = 0.397
				No VR: Day 2 vs 3	0.84	ns	t-ratio = 0.575
				Familiar: Day 1 vs 2	1.0	ns	t-ratio = 0.028
				Familiar: Day 1 vs 3	0.38	ns	t-ratio = 1.343
fig. S1E, right	WT: lick rate	n=7	Linear mixed effects model; Tukey multiple correction	Familiar: Day 2 vs 3	0.39	ns	t-ratio = 1.315
				Novel: Day 1 vs 2	0.002	***	t-ratio = -3.508
				Novel: Day 1 vs 3	0.001	***	t-ratio = -3.624
				Novel: Day 2 vs 3	0.99	ns	t-ratio = -0.116
				Day	0.003	****	$F_{2,111} = 6.2473$
				Environment	< 2.2e-16	****	$F_{2,111} = 58.5058$
				Interaction	0.11	ns	$F_{2,111} = 1.9910$
				No VR: Day 1 vs 2	1.0	ns	t-ratio = -0.050
				No VR: Day 1 vs 3	0.97	ns	t-ratio = -0.224
				No VR: Day 2 vs 3	0.98	ns	t-ratio = -0.175
				Familiar: Day 1 vs 2	0.03	*	t-ratio = -2.633
				Familiar: Day 1 vs 3	0.26	ns	t-ratio = -1.579
				Familiar: Day 2 vs 3	0.54	ns	t-ratio = 1.054
				Novel: Day 1 vs 2	0.03	*	t-ratio = -2.583
				Novel: Day 1 vs 3	0.002	****	t-ratio = -3.534
				Novel: Day 2 vs 3	0.61	ns	t-ratio = -0.951

fig. 2-S4

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S4A, speed quartile 1	WT: NS interneuron firing in Familiar vs No VR	n = 155	Wilcoxon sign-rank test; Bonferroni correction	Fam vs No-VR; quartile 1	3.0e-09	****	N/A
fig. S4A, speed quartile 2		n = 230		Fam vs No-VR; quartile 2	8.0e-07	****	N/A
fig. S4A, speed quartile 3		n = 241		Fam vs No-VR; quartile 3	0.0016	***	N/A
fig. S4A, speed quartile 4		n = 241		Fam vs No-VR; quartile 4	0.054	ns	N/A
fig. S4B	WT: firing rate at AZ during familiar reward approach	NS interneurons, n = 197; WS interneurons, n = 101; pyramidal cells, n = 1017	Linear mixed effects model; Tukey correction	Speed	0.16	ns	$F_{3,4708.4} = 1.7133$
				CellType	<2e-16	****	$F_{2,5862.6} = 2101.6979$
				Interaction	0.41	ns	$F_{6,4708.4} = 1.0190$
fig. S4C ₃	WT: lick rate	n = 43, 30 for normal and freeze trials	Two-sided t-test	Freeze vs Normal	0.97	ns	t=-0.0393
	WT: speed		Two-sided t-test	Freeze vs Normal	0.71	ns	t=0.3796
	WT: NS interneuron firing rates		Two-sided t-test	Freeze vs Normal	4.8e-05	****	t=7.9063
fig. S4C ₄	WT: lick rate	n = 34, 29 for normal and freeze trials	Two-sided t-test	Freeze vs Normal	0.91	ns	t=0.1146
	WT: speed		Two-sided t-test	Freeze vs Normal	0.64	ns	t=0.6385
	WT: NS interneuron firing rates		Two-sided t-test	Freeze vs Normal	0.004	*	t=4.0235

fig. 2-S5

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S5A, left	WT: AZ residual firing rate in Familiar	NS interneurons, n = 196, 75 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.05	*	$F_{1,124.84} = 3.9283$
fig. S5A, middle		WS interneurons, n = 95, 39 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.16	ns	$F_{1,65.966} = 1.9906$
fig. S5A, right		Pyramidal cells, n = 1168, 375 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.74	ns	$F_{1,691.2} = 0.109$
fig. S5B, left	WT: AZ residual firing rate in Novel	NS interneurons, n = 193, 146 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.11	ns	$F_{1,175.21} = 2.5855$
fig. S5B, middle		WS interneurons, n = 96, 84 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.03	*	$F_{1,93.442} = 4.9518$
fig. S5B, right		Pyramidal cells, n = 1144, 898 for correct and incorrect	Linear mixed effects model	TrialOutcome	0.30	ns	$F_{1,1091.8} = 1.0851$

fig. 2-S6

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S6C,	PVxAi32: change in		Linear mixed effects model;	Stim intensity	< 2.2e-16	****	$F_{2,118.33} = 148.32$

far left	normalized PV cell firing rates	n = 44, 2, 44 for no stim, low stim, high stim	Tukey correction	No stim vs low stim	0.002	***	t-ratio = -3.508
				No stim vs high stim	<.0001	ns	t-ratio = -16.068
				Low stim vs high stim	0.68	ns	t-ratio = -0.843
fig. S6C, center left	PVxAi32: change in normalized NS interneuron firing rates	n = 75, 6, 71 for no stim, low stim, high stim	Linear mixed effects model; Tukey correction	Stim intensity	3.8e-13	****	$F_{2,270} = 31.869$
				No stim vs low stim	0.65	ns	t-ratio = -0.883
				No stim vs high stim	<.0001	****	t-ratio = -7.343
fig. S6C, center right	PVxAi32: change in normalized WS interneuron firing rates	n = 36, 11, 35 for no stim, low stim, high stim	Linear mixed effects model; Tukey correction	Low stim vs high stim	0.22	ns	t-ratio = -1.663
				Stim intensity	0.008	**	$F_{2,175.68} = 4.9143$
				No stim vs low stim	1.0	ns	t-ratio = -0.048
fig. S6C, far right	PVxAi32: change in normalized pyramidal cell firing rates	n = 498, 90, 468 for no stim, low stim, high stim	Linear mixed effects model; Tukey correction	No stim vs high stim	0.007	**	t-ratio = 3.050
				Low stim vs high stim	0.18	ns	t-ratio = 1.777
				Stim intensity	2.153e-12	****	$F_{2,1794.8} = 27.27$
				No stim vs low stim	0.85	ns	t-ratio = 0.536
				No stim vs high stim	<.0001	***	t-ratio = 7.336
				Low stim vs high stim	0.008	**	t-ratio = 3.005

fig. 2-S7

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S7A, far left	PVxAi32: novel day1 lick rate	n = 179, 99, 101 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.08	ns	$F_{2,602} = 2.4829$
				No stim vs Low stim	0.18	ns	t-ratio = 1.768
				No stim vs High stim	0.76	ns	t-ratio = -0.699
				Low stim vs High stim	0.08	ns	t-ratio = -2.129
fig. S7A, center left	PVxAi32: novel day2 lick rate	n = 179, 97, 103 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.53	ns	$F_{2,670.07} = 0.6448$
				No stim vs Low stim	0.93	ns	t-ratio = 0.357
				No stim vs High stim	0.65	ns	t-ratio = -0.889
				Low stim vs High stim	0.52	ns	t-ratio = -1.085
fig. S7A, center right	PVxAi32: novel day3 lick rate	n = 200, 104, 106 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.17	ns	$F_{2,841.28} = 1.8018$
				No stim vs Low stim	0.14	ns	t-ratio = 1.896
				No stim vs High stim	0.84	ns	t-ratio = 0.565
				Low stim vs High stim	0.48	ns	t-ratio = -1.163
fig. S7A, far right	PVxAi32: familiar day3 lick rate	n = 560, 43, 39 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.35	ns	$F_{2,709.67} = 1.0504$
				No stim vs Low stim	0.60	ns	t-ratio = -0.956
				No stim vs High stim	0.48	ns	t-ratio = -1.154
				Low stim vs High stim	0.98	ns	t-ratio = -0.183
fig. S7B, far left	PVxAi32: novel day1 speed	n = 179, 99, 101 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.40	ns	$F_{2,602} = 2.4829$
				No stim vs Low stim	0.39	ns	t-ratio = -1.317
				No stim vs High stim	0.98	ns	t-ratio = -0.178
				Low stim vs High stim	0.58	ns	t-ratio = 1.002
fig. S7B, center left	PVxAi32: novel day2 speed	n = 179, 97, 103 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.0002	****	$F_{2,822.85} = 8.5027$
				No stim vs Low stim	0.89	ns	t-ratio = 0.464
				No stim vs High stim	0.0007	****	t-ratio = -3.683
				Low stim vs High stim	0.0009	****	t-ratio = -3.615
fig. S7B,	PVxAi32: novel	n = 200, 104, 106 trials	Linear mixed effects model; Tukey correction	Stim intensity	0.10	ns	$F_{2,880.13} = 2.3043$
				No stim vs Low stim	0.23	ns	t-ratio = -1.642
				No stim vs High stim	0.74	ns	t-ratio = 0.742

center right	day3 speed	across intensities		Low stim vs High stim	0.10	ns	t-ratio = 2.046
fig. S7B, far right	PVxAi32: familiar day3 speed	n = 560, 43, 39 trials across intensities	Linear mixed effects model; Tukey correction	Stim intensity	0.78	ns	$F_{2,712.83} = 0.2502$
				No stim vs Low stim	0.98	ns	t-ratio = -0.208
				No stim vs High stim	0.79	ns	t-ratio = 0.659
				Low stim vs High stim	0.80	ns	t-ratio = 0.645
fig. S7C, left	PVxAi32: AUC based on speed, familiar day 3	n = 5 mice per conditions	Linear mixed-effects model	Stimulation condition (Low to All trials)	0.62	ns	$F_{1,41.195} = 0.2462$
fig. S7C, middle	PVxAi32: AUC based on lick latency, familiar day 3	n = 5 mice per conditions	Linear mixed-effects model	Stimulation condition (Low to All trials)	0.15	ns	$F_{1,41.588} = 2.1977$
fig. S7C, bottom	PVxAi32: AUC based on lick rate, familiar day 3	n = 5 mice per conditions	Linear mixed-effects model;	Stimulation condition (Low to All trials)	0.18	ns	$F_{1,41.319} = 1.8338$

fig. 2-S8

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S8B	WT: ripple rate	n = 37, 38 for familiar and novel	Linear mixed effects model	Environment	0.59	ns	$F_{1,67.057} = 0.2897$
fig. S8C	WT: coactivation probability	n = 4076, 3207, for familiar and novel	Linear mixed effects model	Environment	< 2.2e-16	****	$F_{1,7280.6} = 435.4$
				Familiar vs Novel	<.0001	****	z-ratio = 20.866
fig. S8D	WT: ripple power	n = 527, 450 for familiar and novel	Linear mixed effects model	Environment	0.52	ns	$F_{1,1612} = 0.4207$
fig. S8E	WT: ripple duration	n = 527, 450 for familiar and novel	Linear mixed effects model	Environment	0.16	ns	$F_{1,1568.5} = 1.9816$

fig. 2-S9

Panel	Data	Group size	Statistic method	Comparison	P-value	Notation	F/t statistic
fig. S9C	PVxAi32: ripple power in familiar	n = 4, 39 for goal stim and no stim	Linear mixed effects model	Stimulation condition	0.96	ns	$F_{1,193.02} = 0.0034$
fig. S9D	PVxAi32: ripple duration in familiar	n = 14, 347 ripples for goal stim and no stim	Linear mixed effects model	Stimulation condition	0.90	ns	$F_{1,539.77} = 0.0174$

Chapter 3

Our memories of new and important things develop quickly and are under the tight regulation of specialized neurons in the brain. Previous studies primarily focused on excitatory populations in CA1. While useful, this prior work provided us with an incomplete mechanistic understanding about how we quickly prioritize and remember the most important things. In contrast to these prior studies, we focused on fast-spiking inhibitory interneurons in CA3 where the rapid formation of new memories occurs. We show for the first time physiological and learning-dependent reorganization in hippocampal inhibitory populations that exhibit both input selectivity and output response selectivity during novel spatial learning. Narrow spike-width and PV interneurons showed a spatially selective reduction in firing activity in response to goal-specific inputs. We found that this input-specific reduction in inhibitory drive onto excitatory cells was necessary for subsequent response bias in those excitatory cells to represent task-relevant goal information over less relevant information.

Our observations of goal-specific disinhibitory drive onto local excitatory pyramidal cells in CA3 directly oppose previous models of inhibition in spatial learning that assume spatial uniformity, non-feature selectivity, or prerequisite excitatory reorganization. The results of our optogenetic manipulations show a causal link between spatially selective disinhibition and learning-dependent activity changes at the single unit, population, and behavioral levels. First, we show that goal-specific inhibitory ramp downs are required for sharper and more stable global spatial representations by pyramidal cells. Second, goal-specific inhibitory ramp downs are required for the preferred representation of goals over other locations by groups of pyramidal cells during sharp-wave ripples known to be essential for spatial learning. Finally, goal-specific inhibitory ramp downs are necessary for animals' behavioral differentiation between goal and non-goal locations. Importantly, we compared similar optogenetic stimulation at goal and non-goal locations and found these changes specifically at goal locations. These results show that rapid learning-dependent

changes in inhibition, preceding changes in excitation, drive events critical for successful goal-directed navigational behavior. These results highlight a necessary inhibitory means of driving plasticity at precise locations relevant for new goal learning.

Caveats

We focused on PV interneurons for optogenetic targeting because we found the fast-acting changes in firing during learning in narrow spike-width interneurons that are most likely PV interneurons as determined by opto-tagging. Besides PV basket cells, PV expression is present in other classes of interneurons such as bistratified cells and axo-axonic cells that have cell bodies residing in or close to *stratum pyramidale* but with distinct morphological features and post-synaptic innervation patterns (Pelkey et al., 2017). While PV basket cells innervate principal cells primarily at the soma and proximal dendrites, bistratified cells target the basal and apical dendrites and axo-axonic cells exclusively target the axon initial segments. Thus, it is possible some of the effects of PV stimulation we observed in our study may be attributed to inhibition outside of the perisomatic domains. Prior work suggests that distinct classes of hippocampal PV interneurons have different phase preference and modulation strength during theta oscillations and different probabilities of action potential discharges during sharp-wave ripples (Varga et al., 2014). These class-specific characteristics within the PV inhibitory population likely have class-specific functional implications in disease. For example, PV basket cells, but not PV bistratified or axo-axonic cells in CA1, were found to have reduced spiking activity during sharp-wave ripples in a genetic mouse model of amyloid pathology (Caccavano et al., 2020). While we have not attempted to separate out different classes of PV interneurons in our study, it may be an interesting future avenue for examining relative contributions of these PV classes to new goal learning.

While our results support our hypothesis that disinhibition by PV interneurons is a necessary driver of new learning, we have not tested whether it is sufficient for such learning. It is challenging to strongly reduce inhibition during active behavior for two reasons. First, excessive reduction in inhibition can induce

unwanted seizures that could affect behavior and learning. In fact, disrupted balance in excitation and inhibition is thought to be the hallmark of many diseases including Alzheimer's disease (Maestú et al., 2021). Second, because reduced inhibition occurs naturally over goal learning, we may not observe a significant improvement in performance due to the ceiling effect in wild-type animals. Animal models of disease, specifically those with known deficits in both spatial learning and inhibitory firing, however, could be used in future studies to test whether spatial learning deficits could be rescued with optogenetic suppression of PV interneurons. In the genetic mouse model of amyloid deposition via overexpression of five familial Alzheimer's disease (5XFAD) mutations, we previously found significantly fewer and shorter sharp-wave ripple periods and weaker putative interneuron-to-pyramidal cell monosynaptic connections (Prince et al., 2021). In addition, we found deficits in inhibitory firing reduction around familiar goal locations in 5XFAD mice (**Fig. 3-1**). While this study used a relatively simple task without a novel learning component to enable learning in both 5XFAD mice and littermates, it suggests goal-specific inhibitory deficits exist even in the absence of major group differences in behavioral performance.

Future directions

Prior work in sensory cortices has suggested cholinergic inputs as a key neuromodulatory component of learning that is not examined in our study. In primary motor cortex, for example, higher cholinergic inputs from the basal forebrain were found during motor learning (Ren et al., 2022). This increased cholinergic input co-occurred with reduced activity in somatostatin-expressing (SST) interneurons and increased activity in vasoactive intestinal peptide (VIP) interneurons. Similarly, in primary visual cortex, running-induced increases in cholinergic input to VIP interneurons resulted in disinhibition of pyramidal cells from SST interneurons (Fu et al., 2014). This disinhibition improved visual responses, in line with sensory response modulation by behavioral state. The same group also showed that this improved sensory response is driven by VIP-SST circuit directly, independent of locomotion (Fu et al., 2015). In another brain region, primary barrel cortex, layer 2/3 VIP neurons were found to receive nicotinic input to depolarize, which in

turn disinhibits nearby excitatory pyramidal cells via reduced SST activity during active sensing. Given that PV interneurons target other interneurons including SST interneurons, the VIP-SST circuit extensively studied in the context of associative learning may extend to PV-SST interactions, though potentially cell-type-specific functional roles in learning must be investigated. In addition, recent findings of PV inhibition with sublamina- and subregion-specific differences suggest interesting future directions in examining how spatial and nonspatial operations in the hippocampus are delegated. For example, PV basket cells preferentially target deep pyramidal cells over more superficial pyramidal cells in CA1 (Danielson et al., 2016). Furthermore, CA3 PV interneurons have been shown to mediate the proximodistal gradient found in the inhibitory drive onto pyramidal cells (Sun et al., 2017). This gradient led to a lower inhibition-excitation ratio in CA3b, the region with more enhanced reactivation of fear memory traces, compared to CA3a or CA3c (Sun et al., 2017). These findings add to the complexity and heterogeneity associated with intrinsic synaptic properties and connectivity that affect hippocampal functions and propose an important avenue for future research.

Related to new learning, disinhibition is known to play an important role in experience-dependent synaptic plasticity, classically studied in sensory deprivation and subsequent adaptation (Chen et al., 2011; Kuhlman et al., 2013). For example, monocular deprivation leads to structural remodeling and adaptation, accompanied by a decrease in inhibitory inputs onto neighboring pyramidal cells (Chen et al., 2011). Interestingly, this study also found the antidepressant fluoxetine to be effective at reducing inhibitory tone and facilitating structural adaptation, suggesting a pharmacological avenue for reduced inhibition for learning-associated plasticity (Chen et al., 2011). In line with our interpretation of a role of PV disinhibition in plasticity events, another study proposed the reduction of PV firing and subsequent restoration of normal excitatory firing rates to permit competitive ocular dominance plasticity to proceed based on relative strengths of inputs to two eyes (Kuhlman et al., 2013). These studies highlight a therapeutic potential of reduced inhibition for deficits in learning-dependent plasticity and cognition, consistent with improved cognition following administration of GABA_A antagonists in a mouse model of Down's syndrome

(Fernandez et al., 2007). However, our study proposes important considerations for potential therapeutic effects of reduced inhibition in cognitive processing because spatially specific, not broad, inhibition is necessary for learning new specific information. Potential therapies of cognitive dysfunction should aim to match the timing of manipulated inhibitory activity with the specific periods of learning. Pharmacological compounds with half-lives of several hours to days may not allow for adequate temporal precision required for maximal therapeutic efficacy.

In conclusion, we found a novel inhibitory mechanism by which inhibition drives new learning of specific locations of importance. This work extends our understanding of how CA3 PV interneurons drive local circuit computations necessary for learning new important locations of new environments.

Figure





Figure 3-1. Deficits in inhibitory reduction at well-learned goal locations in the 5XFAD mouse model of amyloid pathology. Mean and SEM percent change from baseline firing of interneurons (left) and pyramidal cells (right) as a function of distance to reward zones (RZ, beige) for WT (black) and 5XFAD (light blue or pink) mice in a well-learned track. $N = 187$ 5XFAD, 142 WT interneurons, 708 5XFAD, 567 WT pyramidal cells.

Appendix A BrainWAVE: A Flexible Method for Noninvasive Stimulation of Brain Rhythms across Species

This appendix is a published article as it appears in eNeuro.

Abbreviated title

BrainWAVE: Noninvasive Stimulation of Brain Rhythms



Abstract

Rhythmic neural activity, which coordinates brain regions and neurons to achieve multiple brain functions, is impaired in many diseases. Despite the therapeutic potential of driving brain rhythms, methods to noninvasively target deep brain regions are limited. Accordingly, we recently introduced a noninvasive stimulation approach using flickering lights and sounds (“flicker”). Flicker drives rhythmic activity in deep

and superficial brain regions. Gamma flicker spurs immune function, clears pathogens, and rescues memory performance in mice with amyloid pathology. Here we present substantial improvements to this approach that is flexible, user-friendly, and generalizable across multiple experimental settings and species. We present novel open-source methods for flicker stimulation across rodents and humans. We demonstrate rapid, cross-species induction of rhythmic activity without behavioral confounds in multiple settings from electrophysiology to neuroimaging. This flicker approach provides an exceptional opportunity to discover the therapeutic effects of brain rhythms across scales and species.

Significance Statement

Despite many studies showing abnormal brain rhythms in multiple diseases, limited means to target deep brain regions noninvasively has restricted the therapeutic potential of driving brain rhythms. Accordingly, we developed a noninvasive millisecond precise sensory stimulation to drive brain rhythms. Here we introduce for the first time newly developed open-source software and instructions for building, testing, debugging, and using BrainWAVE (Brain Wide-spectrum Audio/Visual Exposure) stimulation. We demonstrate BrainWAVE stimulation across multiple species and different experimental settings. These methods constitute a customizable, open-source, accessible, and non-invasive technology that stimulates brain oscillations to causally test how rhythmic brain activity impacts brain function.

Introduction

Neural oscillations—rhythmic patterns of activity in the brain—have been observed across many species and extensively investigated in studies of sensory and cognitive processing. Studies of humans and animal

models of disease have uncovered neural deficits of different frequencies in multiple brain regions (Herrmann and Demiralp, 2005; Prince et al., 2021; Schneider et al., 2014; Shahriari et al., 2016; Solís-Vivanco et al., 2018; Tamura et al., 2017; Uhlhaas and Singer, 2006; Verret et al., 2012; Wang et al., 2013; Worrell et al., 2008). In particular, we and others have found reduced power of endogenous gamma-frequency oscillations in mouse models of Alzheimer's disease (Goutagny et al., 2013; Iaccarino et al., 2016; Mably et al., 2017; Martorell et al., 2019; Verret et al., 2012). We previously showed that enhancing gamma neural activity using noninvasive rhythmic sensory stimulation (hereinafter referred to as “flicker”), specifically lights and/or sounds turning on and off at 40 Hz, reduced Alzheimer's pathology, recruited immune cells, and improved memory performance in mice (Adaikkan et al., 2019; Iaccarino et al., 2016; Martorell et al., 2019). These studies highlight the potential therapeutic benefits of modulating neural oscillations using noninvasive sensory stimulation. Indeed, a growing number of studies have used flicker to ameliorate pathology beyond Alzheimer's disease. For example, 30-50 Hz light flicker was protective against neurons in a cerebral ischemia model (Zheng et al., 2020). Other studies have found flicker to be effective in correcting circadian rhythms (Chan et al., 2021; Yao et al., 2020). Furthermore, this sensory stimulation is useful for studying how rhythmic neural activity affects brain function including immune cells and signals (Garza et al., 2020; He et al., 2021; Iaccarino et al., 2016; Martorell et al., 2019; Venturino et al., 2021). These findings demonstrate the general applicability of flicker stimulation as promising means to treat multiple disorders and diseases (Huang et al., 2021; Jones et al., 2019; Park et al., 2020; Shi et al., 2021).

Flicker stimulation has significant advantages over existing methods of neuromodulation. First, flicker successfully produces reliable modulation in multiple species, including rodents and humans, in multiple brain regions, including difficult-to-target areas such as the hippocampus (Adaikkan et al., 2019; He et al., 2021; Iaccarino et al., 2016; Martorell et al., 2019; Quon et al., 2021). The ability to reach deeper brain regions beyond the superficial sensory areas is particularly exciting because many such brain regions are important for cognition and affected by disease but difficult to modulate noninvasively (Hescham et al.,

2020; Lozano et al., 2019; Qin et al., 2018; Spagnolo et al., 2019). Second, sensory flicker has been shown to alter neuroimmune signaling in mice and humans (Adaikkan et al., 2019; Garza et al., 2020; He et al., 2021; Iaccarino et al., 2016; Martorell et al., 2019; Venturino et al., 2021). Deficits in neuroimmune function are prevalent in many diseases, but traditional invasive techniques to manipulate neural or immune activity cause immune responses themselves thereby limiting their use or interpretation in the study of neuroimmune function (Chung et al., 2015; Hickman et al., 2018; Hong et al., 2016; Wohleb et al., 2016). Third, flicker is an attractive option for chronic at-home therapy in humans. This inexpensive flicker device is widely accessible to researchers and clinicians. Finally, flicker has limited risk and potential side effects in studies to date and offers an attractive model for individualized therapy programs.

Despite these advantages, earlier versions of flicker devices had limitations and room for improvement. Previous flicker stimulation devices used separate systems for audio and visual stimulation without the ability to synchronize the two signals. Furthermore, prior work did not describe methods to test and debug multimodal flicker. These early devices also required extensive knowledge of coding and circuitry to build and operate. As a result, these devices were not user-friendly or fully optimized for clinical and research use. Early users could not easily and quickly adjust brightness and volume based on participant comfort and study goals. Furthermore, previous visual flicker systems could not be used for MRI studies because they were either not MR-compatible or were not bright enough to sufficiently illuminate the field of view of a participant in the scanner bore from a safe distance. Furthermore, there was no established protocol for blinding, which is especially tricky when the intervention is visible to the experimenter and easily distinguished from control conditions. Earlier studies also did not describe the experience or potential side effects of flicker in healthy subjects, which is important to consider in how readily usable this stimulation is in a variety of participants.

For these reasons, we developed an easy-to-build, modular, and customizable BrainWAVE (Brain Wide-spectrum Audio/Visual Exposure) Stimulator to modulate neural activity across multiple species and experimental settings. Given the general utility of this device, here we introduce for the first time open-

source software and instructions for hardware assembly, testing, and debugging. We developed a user-friendly graphical user interface (GUI) to easily control and adjust flicker during experiments without programming. We demonstrate the feasibility, safety, and effectiveness of our newly optimized methods for use in both clinical and preclinical research. We extend our previously published work by detailing how to implement and troubleshoot flicker stimulation across multiple species and different experimental settings, including intracranial recordings from humans and mice, behavioral assays in mice, and assays of side effects in humans. We have also developed methods for minimizing signal interference during simultaneous audio and visual flicker with human EEG recordings or MRI scans while participants engage in behavior tasks. Additionally, we outline how to design and conduct a blinded flicker study and discuss considerations for human and animal experiments involving flicker that we hope will aid future research. These methods constitute a customizable, accessible, and noninvasive technology that stimulates brain oscillations to causally test how rhythmic brain activity impacts brain function.

Results

To deliver flicker stimulation to humans and mice, we developed custom BrainWAVE devices (Fig. A-1). Standard computer monitors and projectors do not have fast enough refresh rates to achieve 40 Hz flicker, thus custom-built LED BrainWAVE Stimulators are required. The BrainWAVE Stimulator interfaced with a variety of light sources and speakers as output components to accommodate different experimental needs (Fig. A-1, A,B,C). The output components included LED lights or a PC monitor with a high refresh rate (e.g. 165 Hz or greater) for visual signals, and speakers, headphones, or earbuds for auditory signals. To illustrate the flexibility and customizability of our device, we provide several output components we have used successfully. For mouse studies, we used LED light strips and speakers (Fig. A-1A) to expose animals to flicker stimulation while the mice were able to freely move within their cages. For human studies, we used LED goggles or LED monitor frames, and headphones or speakers (Fig. A-1, B and C). Flickering LED light strips attached to the edges of a monitor were used in studies where participants perform memory

and attention/reaction tasks on a PC while receiving flicker (Fig. A-1C). We also designed LED frames and panels that use extra bright LEDs to send light from a distance to the visual field of a participant laying down within an MRI scanner (Fig. A-1C). These different outputs interfaced with a common, compact, and portable circuit (Fig. A-1, D,E,F). A detailed parts list and instructions for BrainWAVE Stimulator circuit assembly are provided (see Extended Data, Table A-1). We developed a user-friendly application to run a variety of experiments involving visual and/or auditory stimulation (Fig. A-2). These examples demonstrate the feasibility and customizability of the BrainWAVE Stimulator in a variety of settings.

Modulation in Humans

To establish the effects of flicker stimulation in humans, we characterized neural activity with and without 40 Hz audio/visual stimulation during both scalp EEG recordings and intracranial recordings in human participants (Fig. A-4, A and D). Using scalp EEG, we found a significant increase in EEG power at 40 Hz during stimulation relative to no-stimulation baseline periods within the same subjects (Fig. A-4, B and C; $n = 10$ participants, $p=0.002$, two-sided Wilcoxon signed-rank test). This increase in EEG power during 40 Hz flicker was observed across multiple channels, with channels located over visual and auditory regions having higher modulation (Fig. A-4B). Importantly, elevated EEG power was specific to the frequency of stimulation; we found a significant difference in the power at 40 Hz compared to the mean power at neighboring frequencies (e.g. 40 Hz vs. 31-39 Hz and 41-49 Hz); (Fig. A-4C; $n = 10$ participants, $p=0.014$ two-sided Wilcoxon signed-rank test). Furthermore, elevated EEG power was due to sensory stimulation itself and not electrical artifacts since there was no significant increase in 40 Hz power during the occluded condition ($p=0.56$, paired t-test). While there was variability across subjects, modulation of at least three channels with at least one channel in each hemisphere was achieved within ten seconds after the onset of stimuli and lasted for the duration of flicker exposure (data not shown).

To determine the effects of flicker stimulation in humans with better spatial resolution, we recorded neural activity intracranially in treatment-resistant epileptic patients undergoing presurgical intracranial seizure monitoring with stereotactic EEG (Fig. A-4, D and E) and applied offline highly localizing Laplacian re-referencing to the LFP. We found that sensory flicker increased LFP oscillations, indicative of population dendritic activity, at the frequency of the flickering stimulus in auditory and visual cortices. As an example, in the early visual processing lingual gyrus, 40Hz visual stimulation induced an increase in power at the frequency of stimulation, which is not present in the random visual flicker condition, our control condition (Fig. A-4E; $n = 15$ trials in 1 subject, $p=0.0005$, paired t-test, power at 40Hz medians and quartiles: 4.05 dB, 3.47-5.07 dB during 40Hz flicker; 0.68 dB, 0.46-0.87 dB during random flicker; Cohen's $d = 2.05$). By applying Laplacian re-referencing to the LFP, where the average of adjacent contacts' signals is subtracted from the signal of the channel of interest we determined this modulation was local and was not due to distant volume conduction. Together, these findings demonstrate that the sensory flicker-induced modulation is reliable, frequency-specific, and efficiently induced.

Modulation in Mice

To assess the effects of flicker in sensory and memory circuits of mice with high temporal and spatial resolution, we recorded local field potentials and single neurons during flicker exposure. Using *in vivo* electrophysiology in awake, head-fixed mice, we quantified sensory flicker-induced changes in neural activity in mouse hippocampus (Fig. A-4, F-H). As previously reported in Martorell*, Paulson* et al. and Iaccarino*, Singer* et al., we found that LFP power was significantly elevated specifically at 40 Hz during 40 Hz stimulation, but not during no-stimulation baseline or random control conditions (Fig. A-4G; $n = 8$ trials in 1 animal, $p=0.010$ 40Hz vs baseline, $p=0.004$ 40Hz vs random, paired t-test, power at 40Hz medians and quartiles: 75.18 dB, 71.86-76.76 dB during 40 Hz flicker; 70.30 dB, 68.39-72.31 dB during baseline; 70.39 dB, 68.05-71.11 dB during random). Furthermore, exposing mice to 40 Hz auditory flicker led to

increased modulation of single-neuron spiking, meaning neurons were more likely to fire at a particular phase of the stimulus, in hippocampus (Fig. A-4H). These effects were not observed in no-stimulation baseline or random control conditions. Similar results were observed in auditory cortex and prefrontal cortex. These deeper regions are more difficult to target with other noninvasive stimuli, such as transcranial magnetic stimulation. These results demonstrate effective, frequency-specific, and noninvasive modulation of neural activity in multiple brain regions simultaneously using simple and customizable BrainWAVE Stimulator circuits.

Stimulation Side Effects and Behavioral Controls

One important consideration is whether flicker stimulation is aversive or has unintended effects on behavior. Accordingly, we asked our study participants to report any acute symptoms such as headache, dizziness, negative affect, and more in a survey immediately before and after a one-hour audio and visual flicker stimulation session. Out of eight participants, one reported that the light or sound was intolerable. After stimulation, some participants reported mild negative or positive effects including sleepiness/drowsiness (6), boredom (5), headache (2), increased (1) and decreased (2) ability to focus, and increased relaxation (1) (Fig. 5A-C). Some of these effects, such as drowsiness and boredom may be attributed to the experiment procedure which asked participants to remain still for over an hour, rather than being an effect caused by the stimulation itself.

In addition to the acute effects of flicker, we assessed the potential adverse effects of more chronic flicker exposure (He et al., 2021). In this study, older participants with prodromal Alzheimer's disease were exposed to 4-8 weeks of flicker stimulation. Overall, longer-term flicker exposure was well tolerated by most subjects. Out of 10 participants, 3 reported mild adverse events which may be attributed to flicker

exposure, including dizziness, tinnitus, headache, and worsened hearing loss. These mild adverse events were relatively rare, reported 5 times over the course of 4 to 8 weeks of daily 1hr flicker exposure.

One consideration is that different types of flicker stimulation indirectly affect neural or immune responses because that particular type of stimulus makes animals move more or less. Our results show that animals respond similarly in terms of the amount of activity and exploration of the environment with flicker at different frequencies and constant light. These results show that neural and immune responses to sensory flicker cannot solely be attributed to changes in activity levels during stimulation, although there may be more subtle behavioral differences that could not be quantified with our assays. To determine if the effects of sensory stimulation were confounded by changes in mouse behavior, we recorded the activity of mice during a one-hour session of visual stimulation (Garza et al., 2020 Fig. 5D, E). To do this, we took advantage of the fact that mice were allowed to move freely during stimulation. We quantified the amount of time spent in the center arena of the cage during each stimulation condition as a common measure of anxiety-like behavior in mice. We found no significant differences in both the time spent in the center of the arena and total activity across stimulation conditions (Fig. A-5, F and G; percent time in center: $F(3,15)=0.8754$, $p=0.4757$, RM-ANOVA, $n=5$ mice; percent time active: $F(3,15)=0.9306$, $p=0.4502$, RM-ANOVA). These results show that neural and immune responses to sensory flicker cannot solely be attributed to changes in behavior during stimulation, although there may be more subtle behavioral differences that could not be quantified with our assays.

Discussion

Here we introduce a newly optimized BrainWAVE Stimulator and user-friendly, open-source methods for assembling, testing, and implementing noninvasive sensory flicker across species and experimental designs. Building on our first example use of flicker in modifying disease pathology, here we developed several new methods to aid future research involving flicker. First, we integrated the audio and visual stimulation

systems to produce synchronized multimodal stimulation. Second, we created a user-friendly, intuitive GUI to easily control and adjust flicker during experiments and without requiring programming. Third, we designed an MR-compatible flicker device to remotely control flicker during MRI. Fourth, we established a new protocol for designing and administering a blinded flicker study with appropriate control conditions. Finally, we validated the safety of flicker by assessing potential adverse effects in healthy human subjects. Our flicker stimulation produced robust, rapid, and frequency-specific modulation of neural activity in both mice and humans with minimal side effects. We showed that side effects of stimulation were rare and mild. These results were not due to differences in overall activity levels or anxiety-like behavior during stimulation. Thus, this new and improved cross-species simulation tool provides a unique means to study and treat neural activity deficits noninvasively in a wide spectrum of brain regions and diseases.

This optimized BrainWAVE Stimulator described in the present work has several advantages to reach a wider scientific and medical community. We substantially improved our design to make our tools user-friendly, accessible, and therefore more impactful in research and medicine. With a simple and flexible design, sensory flicker is easy to integrate into a variety of experimental setups. These methodological improvements build on our previously demonstrated effects of sensory flicker. We previously found that audio and audio-visual sensory flicker noninvasively produces frequency-specific modulation in hippocampus and prefrontal cortex of mice (Martorell et al., 2019). The ability to target deep brain regions noninvasively is important because these regions are involved in memory and disease. Indeed, our prior study shows gamma (40 Hz) flicker rescues memory deficits in a mouse model of Alzheimer's disease (Martorell et al., 2019).

Our ability to flexibly alter the degree of synchronization between multiple sensory stimuli brings additional advantages to studying multimodal associative learning. For most of our previous studies, we have programmed our stimulator to generate auditory and visual signals simultaneously. However, for some experimental questions it is important to consider that the transmission times of auditory and visual information through their respective pathways differ (King et al., 1985) and the magnitude of these

differences can vary across species and individuals. Future experiments may examine how changing the onset and phase of these stimuli affect neural activity. Flexibly adjusting the delay between auditory and visual stimulation so that the sensory information reaches their respective cortices at the exact same time may indeed lead to more effective modulation in some brain structures. Furthermore, the BrainWAVE stimulator is capable of adjusting stimuli onset with millisecond precision in order to suit the researcher's experimental design or to personalize stimuli to suit an individual or specie's specific transmission speed. Prior work has shown that phase locking of visual and auditory stimuli enhances and predicts future long-term memory formation (Clouter et al., 2017; Wang et al., 2018). While we have not investigated this directly, enhancing phase locking may be part of the mechanism by which flicker improves memory in mice (Martorell, Paulson, et al., 2019).

Gamma flicker drastically reduces levels of amyloid beta levels, a peptide thought to initiate neurotoxic events in Alzheimer's disease, in sensory and memory circuits (Iaccarino et al., 2016; Martorell et al., 2019). Noninvasive sensory stimulation is currently being tested in several clinical trials in patients with neurodegenerative diseases (Cimenser et al., 2021; He et al., 2021; Murty et al., 2021). Human participants have successfully used flicker at home with minimal side effects, proving long-term studies are feasible and convenient (Cimenser et al., 2021; He et al., 2021). Future versions of BrainWAVE Stimulators may be integrated into existing wearable technology like smartwatches and virtual reality headsets.

When designing a human or animal flicker study, there are a few limitations and considerations to keep in mind. One of the recurring challenges is selecting a "control" condition or other stimuli for comparison. Possible control conditions include no sensory stimulation, alternate frequencies (such as 20 Hz or 80 Hz), constant (non-flickering) stimuli, and random (non-periodic) frequencies. Each type of stimulus condition controls for different aspects of the stimuli, such as periodicity, frequency, and total duration of stimulus exposure. Given these considerations, an ideal experimental design has multiple control groups with different types of stimulation parameters. When deciding on one or more control groups, group size and feasibility of an experiment may be limiting factors. As an additional limiting factor in human flicker use,

potential negative side effects must be considered and minimized. In this work and in a prior study, we excluded participants with a history of light-induced seizures or migraines in case the stimulus exacerbates these conditions (He et al., 2021). Though mild and rare, we noted some adverse side effects of acute flicker stimulation, including boredom, sensitivity to light, and headache which were similar to mild adverse events in our prior study on chronic stimulation (He et al., 2021). To help mitigate potential adverse effects, participants may be given the option to adjust the stimulus intensity to more comfortable levels. However, researchers should keep in mind that lowering the intensity may decrease the degree of modulation, and if lowered below a certain point, modulation may not be observed. Indeed, studies should establish the degree to which participants' neural activity modulates to the flicker stimulus before studying subsequent effects and establish baseline levels of acute modulation. In our studies, we first establish light and sound intensity levels at which the subject is comfortable as well as levels that show adequate modulation to the stimulus. We recommend testing multiple ranges of light and sound intensity prior to an experiment to include participants with robust modulation at tolerated intensities and exclude those reporting discomfort with the stimuli.

While most of our prior work has thus far focused on the effects of 40 Hz flicker in neurodegenerative disease, the effects of stimulation in other diseases and in the healthy brain are currently under investigation (Garza et al., 2020; He et al., 2021; Iaccarino et al., 2016; Martorell et al., 2019; Zheng et al., 2020). One study has reported that 30-50 Hz flicker protects hippocampal neurons in an animal model of ischemia (Zheng et al., 2020). Another study showed that 60Hz light flicker affects microglia remodeling of perineuronal nets, which play a key role in critical period plasticity, in healthy mice (Venturino et al., 2021). Our prior study shows that light flicker has frequency-specific effects on the expression of cytokines, an extracellular immune signaling protein, as well as intracellular immune signaling in healthy adult animals (Garza et al., 2020). These studies reveal that flicker could be used as a novel intervention in a variety of contexts. For example, this noninvasive means of driving brain rhythms is valuable for assessing immune effects of specific activity patterns without the confounding effects of invasive stimulation tools.

Inducing frequency-specific neural activity noninvasively using sensory flicker provides a novel approach to investigating the role of specific frequencies of neural activity in health and disease. Here we provide an easy-to-follow guide to build and implement such devices in experimental and clinical settings at low cost and with user-friendly software. These tools will be useful to future studies using our devices that will produce novel insights into the mechanisms of brain rhythms and immune function in health and disease.

Materials and Methods

Device Design

The hardware components of each type of BrainWAVE Stimulator were selected to suit many different types of subjects (animals/rodents, patients, healthy humans) and different types of studies (EEG, MRI, behavior, electrophysiological recordings, intracranial EEG, etc.). For mouse studies, we designed and built a BrainWAVE Stimulator using a strip of LEDs and a speaker to administer sensory stimulation to mice housed within a cage (Fig. A-1A) to assess behavioral and immunological effects. To deliver flicker stimulation to humans, we developed devices that consisted of headphones or earbuds and LED goggles or an LED frame (Fig. A-1, B and C). For audio stimulation during intracranial recordings or scalp EEG recordings small earbuds were advantageous compared to headphones with a headband since the headband interfered with electrodes. When incorporating flicker with computer-based behavioral tasks, we used an LED frame placed around a computer monitor (Fig. A-1C).

Each BrainWAVE Stimulator consisted of three types of hardware components: signal generators, signal modulators, and the audio and visual outputs (Fig. A-1, D and E; Table A-1). These three components were assembled to create either two independent audio and visual BrainWAVE Stimulator circuits or one combined circuit (Fig. A-1, D, E, F; Fig. A-S1). A signal generator produced the signal dictating the on/off timing of audio and visual outputs before sending it to a signal modulator. The signal modulator amplified

or attenuated the flicker control signal to adjust the intensity (e.g. brightness or volume) of the stimulation. The modulated signal was then sent to audio and visual output components, specifically the lights and speakers converting the electrical signals into visible and audible sensory stimuli.

For signal generation, we used a microcontroller (e.g., an Arduino) or digital acquisition (DAQ) hardware to create the audio and visual control signals (Fig. A-1, D and F; Fig. A-S1), though other signal generators with sub-millisecond precision may also be used. When selecting a signal generator, the advantages and disadvantages of an Arduino or DAQ were considered in regard to the temporal precision required for a particular experiment. A DAQ and PC system offers a higher sampling rate and better temporal precision (e.g., 400,000 samples per second for a NI-DAQ USB-6212 vs ~9600 Hz sampling rate for an Arduino Uno). Higher precision was required when aligning flicker stimulation to other signals with high temporal resolution, like electrophysiological signals in scalp EEG, intracranial EEG, or depth electrode recordings. If high sampling rate and precision were not necessary, then flicker signals were generated effectively and more inexpensively using an Arduino microcontroller (less than \$25). Arduino and DAQ signal generators were programmed with custom code (see Extended Data) in Arduino IDE software or MATLAB, respectively. For our studies, this code was either uploaded to an Arduino Uno (for an Arduino BrainWAVE Stimulator) from a PC or was run in MATLAB on a PC that sent signals to a NI-DAQ USB-6212 (for a DAQ BrainWAVE Stimulator).

Next, the generated output signal was typically amplified or attenuated via a signal modulator to produce the desired level of brightness or volume for a particular experiment. An audio amplifier (Fig. A-1D) was used to increase or decrease the volume of the audio signal before being sent to a speaker or earbuds. Visual signal modulation was achieved with a simple MOSFET circuit to power LEDs with a higher voltage source (Fig. A-1, E and F) because the output voltages of the signal generators were too low for some light sources. Arduino-generated visual signals were sent from the Arduino to a MOSFET circuit built on an Arduino shield and then the modified signals were sent to the visual outputs. We adjusted the brightness of the LEDs by installing a dimmer switch between the MOSFET circuit and the light source (see Extended Data Guide

1-1 for detailed instructions on how to assemble a BrainWAVE Stimulator). Auditory signals were similarly sent from the Arduino through an audio amplifier and then to the audio outputs. The ability to alter the level of audio and visual stimulus intensity was important especially in studies involving human subjects to ensure participant comfort and tolerance. For flexible adjustment of stimulus parameters, we developed a user-friendly GUI tool in MATLAB (software and associated code in the Extended Data).

A 3D-printed case was used to protect the Arduino and circuit (Fig. A-1F). Instead of the Arduino shield, circuits may be produced via a printed circuit board (PCB) or breadboard (Fig. A-S1). Using a PCB simplifies the assembly process, reduces the footprint of the circuit, and makes a more reliable circuit. To prevent electrical noise in electrophysiology or EEG recordings, the BrainWAVE Stimulator was shielded by placing it in a metal-lined box. Testing for electrical artifacts was performed prior to recording.

BrainWAVE Stimulator Code and Software

Custom software controlled the signals generated by the microcontroller or DAQ. For Arduinos, code was written in Arduino IDE and uploaded to an Arduino from a Windows 10 PC. Arduino code runs automatically whenever the Arduino is supplied with power (either from a PC via USB cable or from a wall power adapter) regardless of whether the Arduino is connected to the PC. For a DAQ BrainWAVE Stimulator, the DAQ was first connected to a PC with MATLAB using the data acquisition toolbox (see National Instruments for further instructions for NIDAQs). The DAQ BrainWAVE Stimulator was controlled using MATLAB software and unlike an Arduino, the DAQ typically must be connected to a PC while in use. While running the DAQ system, the signals generated in MATLAB were sent from the PC to the DAQ, which in turn sent the signals to the signal amplification/attenuation components and then the output components. We developed a user-friendly application (Fig. A-2) to run a variety of experiments involving visual and/or auditory stimulation. All chosen experiment details, and timing of trials, are saved in a MATLAB structure for offline data processing. Code used for Arduino and NI-DAQ BrainWAVE Stimulators is found on GitHub.

Code Accessibility

The code/software described in the paper is freely available online at <https://github.com/singerlabgt/BrainWAVE>. The code is available as Extended Data.

Measuring BrainWAVE Outputs

After constructing BrainWAVE Stimulators, testing was performed to determine if the devices generate appropriate stimulus intensity, timing, and other signal properties. Light illuminance and audio volume were measured with a light meter and decibel meter, respectively, with the distance between the sensor and meter approximating the distance from the sensory to the subjects' eyes and ears (Table A-3). For mouse studies, light intensity was set at ~150 lux and sound intensity at 60-65 dB (Garza et al., 2020; Martorell et al., 2019). For human studies, we adjusted stimulus intensity for each subject based on tolerance, with the levels ranging from 0-1400 lux for brightness and 0-80 dBA for sound (He et al., 2021). We measured the frequency and duty cycle of the audio and visual stimuli in real-time using an oscilloscope connected to the analog output ports of the light and decibel meters (Fig. A-3A). Alternatively, the timing of the light and sound stimulus may be measured with a photodiode and a microphone connected to an oscilloscope, or the stimulus may be recorded on a laptop and analyzed on a computer. Audio and visual signals were measured simultaneously to compare their duty cycle, frequency, and phase timing.

To modulate neural activity, we generated sensory signals at specific frequencies depending on the experimental design. Visual gamma flicker (40 Hz) was produced using a 5.17 V, 40 Hz square wave with a 50% duty cycle (Fig. A-3B). The voltage must be greater than 4 V to operate the MOSFET. Auditory gamma flicker was produced with a pure sinusoid tone signal that was modulated by a 40 Hz square wave with a 50% duty cycle for audio-visual stimulation, and a 4% duty cycle for audio-only stimulation (Fig. A-3C). The pure tone used was adjusted to fall within the center of the hearing range of the species tested: 10 kHz for mice and 7 or 8 kHz for humans (Heffner et al., 2007). We used a 4% duty cycle for audio-only stimulation to more closely match the timing of clicks in studies on auditory steady-state responses evoked

with 40-Hz click trains (Galambos et al., 1981; Ma et al., 2013; Osipova et al., 2006; Stapells et al., 1984; Thuné et al., 2016). Other frequencies of sensory were generated in a similar manner typically with a 50% duty cycle. Randomized stimulation was used to compare periodic to aperiodic flicker stimulation and had varying duty cycles (from 33% to 99%). Audio and visual signals were typically synchronized with similar duty cycles, but offset signals or different duty cycles may be desired in some cases (Fig. A-3, D, E, F, G, I).

Presenting BrainWAVE Stimuli to Humans and Mice

All human studies were approved and monitored by the Institutional Review Board. For EEG studies, male and female human subjects (ages 18-24) received the audio stimulus of a 7 or 8 kHz tone via headphones or earbuds. The visual stimulation was produced using an LED frame surrounding a PC monitor (Fig. A-4A) or via glasses lined with LEDs. Both the tone and LEDs were synchronized. Prior to experiments, subjects were presented with multiple levels of light and audio intensities to identify the optimal range the subject tolerated. Subjects were allowed to ask the researcher to change the stimulus intensity to a level they were comfortable with at any point during the study. After finding a comfortable stimulation level, we noted the new stimulation level and then checked the subject's EEG to see if they still met our modulation criteria. Any subjects that did not meet modulation criteria were excluded from the study. In neural recording studies, a "relative occluded" stimuli recording condition was performed to test for electrical noise from the BrainWAVE Stimulator. In human studies, during the relative occluded condition, the participant wore an opaque eye mask and earplugs that prevented exposure to the sensory stimuli. Neural activity of participants wearing the eye mask and ear plugs during this occluded condition session was not modulated when the sensory stimulation was turned on.

The MRI environment poses additional challenges because standard circuit components and stimulus devices are incompatible with the high magnetic field of the scanner. Thus, our lab developed novel methods for conducting flicker in the MRI scanner. First, we designed stimulus presentation methods that

were MR safe. During MRI scans, we delivered audio flicker stimulation using in-ear MR safe, headphones (brand: MRIAudio) to allow the participant to hear audio generated from a PC or BrainWAVE Stimulator over scanner noises. These headphones were designed to fit within the MRI head coil during scans and have a noise reduction rating (NRR) of 29 dB. For visual flicker, we created an MR safe LED frame that fits around a projection screen. This setup allowed participants to view pictures or videos or perform visual-based tasks while receiving visual flicker stimulation in their periphery. The translucent projection screen with the LED frame was constructed on an MR safe stand placed about two feet from the patient table of the MRI machine. We did not find MRI artifacts with this screen or the headphones.

Second, we designed a system such that the control circuit was far enough from the scanner (placed in the MRI control room) while the lights and sounds were near the MRI scanner so that the stimulation signals from the control circuit were not affected by the scanner's radio frequency pulses and so that the circuit's signals did not create noise in the MRI images. A shielded cable was used to connect the flicker stimulation device in the control room to the LED frame. We controlled and adjusted the volume and brightness of flicker from the control room using a computer or BrainWAVE Stimulator. A projector in the control room projected images through the control room window onto the back of the projection screen. A front-facing mirror placed above the head coil mirrored the screen and flicker frame to the subject lying down in the scanner. We also included simple attention tasks to determine if the subject is alert over the course of an experiment.

All animal work was approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology. For mouse studies, wild-type male 2- to 3-month-old mice were brought into the animal holding areas and left undisturbed for at least 30 minutes. Mice were then moved to a dark experimental room and individually placed in an empty enclosure with three opaque black sides (Fig. A-5D). One transparent side of the cage faced LED strips. An audio stimulus was presented through a speaker and synchronized to the onset and offset of the LEDs. The stimulus was presented for one hour or for one hour per day for multiple days. We used multiple different stimulation conditions as controls, namely,

random, 20 Hz, and constant light. The random group received sensory stimulation with randomized light-off intervals (duration ranging from 0-25 ms) while the total duration of the light-on phase was kept consistent with the 40 Hz group (12.5 ms). The 20 Hz group was exposed to 20 Hz light flicker with a 50% duty cycle. The random condition was used to assess the effects of periodic versus aperiodic stimuli and the 20 Hz condition was used to assess stimulus frequency-dependent effects. The constant light group was exposed to constant light for the entire duration of the session and therefore used to disambiguate the effects of constant versus flickering stimuli.

Blinding

A key consideration in flicker experiments is how to blind experimenters when different experimental stimuli are readily perceived during experiments. To address this, we developed a blinding system such that the experimenter could test and start flicker without being exposed to flicker. In this system, animals were monitored remotely using an infrared (IR) light source and a video camera with an IR-pass filter. The IR filter blocked visible light, thus preventing the experimenter from determining the light flicker frequency, while the IR light source provided illumination to monitor animal behavior. Animals were assigned to stimulation groups by a non-blinded third party and each animal identifier and the corresponding flicker group assignment were listed in a file that was read by the flicker device. To start flicker, a blinded experimenter entered the animal's unique identifier into the flicker device, which then loaded the correct stimulation condition corresponding to the animal's identifier. Prior to starting the flicker stimulation sequence or control condition sequence, the device first played a test sequence of light and sound so that the experimenter could confirm that the LEDs and speaker were working properly. After the test sequence, there was a pause in stimulation during which the experimenter exited the room to the remote monitoring computer. Flicker stimulation or the control stimulation condition then commenced while the experimenter monitored the animal via the IR video in another room.

Electrophysiology Recording, Preprocessing, and Analyses of Neural Data

We performed neural recording and analyses as described previously (He et al., 2021; Iaccarino et al., 2016; Martorell et al., 2019). In brief, for *in vivo* electrophysiology in mice, we made small craniotomies on the skull of the animal under isoflurane anesthesia using the coordinates for brain regions of interest. We recorded local field potential (LFP) and spiking activity using a 32-channel NeuroNexus probe with data acquired at a sampling rate of 20 kHz using an Intan RHD2000 Evaluation System with a ground pellet as reference (Fig. A-4). Additional recordings were performed with the electrode in saline above the craniotomy while animals were exposed to flicker to detect possible electrical artifacts. Recorded neural data were bandpass filtered from 300 to 6000 Hz for spikes to be clustered into single units using an automatic spike sorting algorithm (MountainSort). For LFP analyses, raw data were first downsampled to 2 kHz and bandpass filtered from 1 to 300 Hz. For power spectral density analyses, we used multitaper methods from the MATLAB Chronux toolbox and compared traces between 40 Hz and random frequency stimulation.

For the human scalp EEG recordings, we used a 32-channel BioSemi ActiveTwo system (Amsterdam, Netherlands) with data acquired at a sampling rate of 2,048 Hz (Fig. A-4). Signals were bandpass filtered from 1 to 100 Hz using Hamming windowed FIR filter (EEGLAB) for power spectral density analyses. We defined modulation of individual channels as elevated power at 40 Hz, at least 3 standard deviations above the mean power in neighboring frequencies of 31-39 Hz and 41-49 Hz. We defined modulation of an individual human subject as having at least three modulated channels with at least one modulated channel in either hemisphere over the course of a recording.

Human intracranial recordings were performed in treatment-resistant epileptic patients who underwent pre-surgical intracranial seizure monitoring to determine their seizure onset zones (Fig. A-4D). To record LFPs, these patients were usually implanted with a dozen (number and location determined based on clinical needs) depth electrodes, each containing up to 18 macro-contacts along their length. Some of these electrodes contained microwires that protruded at their tips and allowed recordings from single neurons. Typically, these patients were monitored for several days to weeks, providing a unique opportunity to

perform voluntary studies on intracranial human brain activity in between clinical care. We carried out all experiments in the patient's room. LFP recordings were acquired using the clinical system used by the hospital (XLTEK EMU 128FS; Natus Medical) at a rate ranging from 1024 to 2048 or higher Hz using subdermal contacts from an electrode array placed at the vertex (subgaleal) as ground and reference. Signals from microwires were recorded using the Blackrock NeuroPort system (Blackrock Microsystems, UTSW), at a rate of 30,000 samples/s, using a dedicated microwire as reference. Data were re-referenced using Laplacian reference and bandpass filtered between 2-300Hz, with a baseline correction over the duration of 12-second records segments. PSD was calculated over 2-100Hz, using the Chronux toolbox (Mitra and Bokil, 2009) (<http://chronux.org/>), with a time-bandwidth product of 3, and number of tapers of 5.

Behavioral Assessment

Human study participants were given a survey immediately before and after a one-hour audio and visual flicker stimulation to assess any acute symptoms such as headache, dizziness, and negative affect (Fig. A-5, A, B, C).

To monitor the effects of sensory flicker on animal behavior during stimulation, we recorded mouse behavior using an IR-sensitive camera (Fig. A-5D). The infrared lighting was required to avoid poor video tracking due to interference from the flickering visible light and to avoid experimenter bias. Behavior was analyzed via automated animal tracking (Ethovision XT v 14.0). We divided the arena into center and outer regions and quantified activity levels in both parts (Fig. A-5E). Animal activity was quantified by classifying periods as activity or inactivity/freezing. Inactivity was defined as having activity in 0.01% of the total arena for longer than 0.5 s (Fig. A-5, F and G). We performed a one-way ANOVA to assess group differences.

Figures




Figure A-1. BrainWAVE circuit design (A) Flicker presentation for mice in a clear enclosure with a strip of light-emitting diodes (LEDs) and a speaker. (B) Flicker presentation for humans with LED-lined goggles and headphones. (C) Flicker presentation for humans with an LED-frame surrounding a computer monitor and earbuds. (D) A circuit diagram of an audio flicker circuit with an audio amplifier to allow for volume adjustments. (E) A circuit diagram of an LED visual flicker circuit with a MOSFET to allow a 12-volt voltage source to power the LEDs. (F) Left: An Arduino Uno microcontroller (top) and Arduino shield

(bottom). A MOSFET and wires are soldered on the Arduino shield. Middle: The Arduino shield fits on top of the Arduino Uno and both sit in a custom 3D printed case. Right: The Arduino and BrainWAVE Stimulator circuit shield are enclosed in a 3D-printed case. Ports in the case allow a USB Type-B cable, an LED cable and a power cable to plug into the Arduino and shield. See Extended Data Table 1-1 for a list of BrainWAVE stimulator parts. See Extended Data Figure 1-1 for a physical diagram of an Arduino-Visual BrainWAVE circuit. See Extended Data Guide 1-1 for detailed instructions on how to assemble the device.




Figure A-2. A user-friendly BrainWAVE graphical user interface (GUI) to perform many types of experiments. We developed a user-friendly application to run a variety of experiments involving visual and/or auditory stimulation. Preprogramed tasks are found under the task dropdown menu and include four different tasks. First, a classical flicker task, with exposure to 5.5 Hz (theta-like), 40 Hz (gamma-like), 80 Hz, and random non-periodic flicker at visual, audio-visual and auditory modalities. Second, a flicker duration task, exposing subjects to a given modality and frequency of flicker for minutes at a time. Third, a flicker frequency task, which allows exposing subjects to up to 26 different frequencies of flicker of a given modality. Fourth, a single pulse evoked potential task, where subjects are exposed to single visual, audio-visual and auditory 12.5ms pulses. The stimuli parameters are set in entry boxes for stimulus duty cycle and tone (sound frequency). The comments box is used to write and save time-stamped experiment notes during the experiment. Developed for testing in human participants, each task includes tests for comfort to determine the optimal brightness and volume of the stimuli that are comfortable to the subject (adjusted on the device), tests for safety to determine if the intended flicker stimuli induce adverse events, experimental tasks, control occluded condition (where subjects wear a sleep mask and earplugs), and measures of brightness and volume used. See *Extended Data BrainWAVE Stimulator Guide* for instructions on how to set-up and run an experiment using the BrainWAVE GUI.




Figure A-3. Measuring BrainWAVE Stimulator output signals (A) Visual and auditory stimuli generated by signal generators are measured with a light meter and decibel meter, respectively. Connecting these devices to an oscilloscope allows visualization and quantification of frequency and duty cycle. (B) Ideal 40

Hz visual signal with a 50% duty cycle. (C) Ideal 40 Hz audio signal (8 kHz tone) with a 50% duty cycle. Inset, a zoomed-in view of the signal at the stimulus onset showing an 8 kHz sinusoid. (D) Ideal 40 Hz visual (blue) and audio (orange) signals in phase, both with a 50% duty cycle. (E) Example of an ideal 40 Hz visual signal (blue) with a 50% duty cycle in phase with a 40 Hz audio signal (orange) with a 4% duty cycle. (F) Example of a real 40 Hz visual signal (dark blue) and its light output (light blue) recorded with an oscilloscope. (G) Example of a real 40 Hz audio signal (dark orange) and its audio output (light orange) recorded with an oscilloscope. (H) Same as F, but of a BrainWAVE Stimulator with an inverted MOSFET, which turns the lights off with a high-voltage signal and turns LEDs on with a low-voltage signal. (I) Same as F, but with a BrainWAVE Stimulator that uses a light source with a slower onset time. A delay is observed in the onset timing between the visual signal (dark blue) and measured output (light blue). See Extended Data Table A-2 for recommended light and sound levels.




Figure A-4. Sensory flicker entrains brain regions in humans and rodents. (A) Top, schematic of a typical EEG recording setup. Participants with EEG caps sat in front of a computer monitor (about 2 feet away) framed by LEDs and received auditory stimulation through earbuds. Bottom, example EEG trace recorded

from the center electrode (Cz) during flicker stimulation. (B) Left, power spectral density averaged across all channels in an example EEG recording during baseline (top) and 40 Hz audio-visual flicker stimulation (bottom). Right, heatmaps of mean power at 40 Hz averaged across subjects before (top) and after (bottom) 40 Hz stimulation. (C) Mean power at 40 Hz during baseline (“Baseline”) and during 40 Hz audio-visual flicker stimulation (“40 Hz”), and mean power at neighboring frequencies (31-39 Hz and 41-49 Hz) during 40 Hz audio-visual stimulation (“Around 40 Hz”). Mean power at 40 Hz is significantly higher during 40 Hz stimulation than power at 40 Hz during baseline ($n = 10$ participants, $p=0.002$ two-sided Wilcoxon signed-rank test) or power at neighboring frequencies ($n = 10$ participants, $p=0.014$ two-sided Wilcoxon signed-rank test). Each colored line represents a single participant. (D) Example setup of a human intracranial recording during flicker stimulation. Local field potential recordings were obtained from treatment-resistant epileptic patients implanted with intracranial electrodes for seizure monitoring. A computer-controlled the delivery of sensory stimuli via a custom-made BrainWAVE Stimulator circuit, which sent its output to a set of LED-lined goggles and earbuds. (E) Example of modulation to visual flicker recorded in lingual gyrus. Left, axial slice of pre-implant T1 MRI overlaid with post-implant CT, showing the location of the recording depth electrode. Highlighted in red is the electrode for data on right. Example recording trace before and during stimulation where the start of stimulation is indicated with a dashed line (below). Right, power spectral density averaged across 15 trials of either 40 Hz visual (blue), or random visual stimulation (green). Shaded areas represent standard error of the mean. ($n = 15$ trials in 1 participant, $p=0.00005$, paired t-test). (F) Top, schematic of in vivo electrophysiology in head-fixed mice. Mice running on a spherical treadmill received sensory flicker stimulation through a strip of LEDs placed above the mouse and a speaker to the right. Bottom, example trace of local field potentials in mouse auditory cortex during 40 Hz audio-visual flicker. (G) Power spectral density comparison between 40 Hz flicker stimulation (blue) and baseline (grey), and between 40 Hz flicker and random (green) condition. (H) Firing rate modulation during 40 Hz audio-visual stimulation (left), baseline (middle), and random

(right) conditions in mouse hippocampus. Colors above indicate if light was on (yellow), off (black), or varied trial-to-trial (grey).




Figure A-5. Minimal adverse effects from flicker exposure. (A) Number of reported mild adverse events by participants ($n=22$) before or after one hour of audio-visual flicker stimulation during an EEG or MRI session. (B) Number of reported changes in mood, anxiety, energy, focus, or relaxation ($n=22$). (C) The distribution of responses from participants ($n=22$) rating the tolerability of flicker stimulation after one hour of flicker during an EEG or MRI session. (D) Schematic of flicker stimulation for freely moving mice in an enclosure. (E) Overhead view of mouse enclosure with a portion of the arena indicating the center zone which was used for analysis. (F) Percentage of time spent in the center zone during a one-hour session of constant light (yellow), 20 Hz (red), 40 Hz (blue), and random (green) conditions ($F(3,15)=0.8754$, $p=0.4757$, RM-ANOVA, $n=5$ mice). Error bars indicate the mean $\pm$ SEM. (G) Total percent of time mice were active during the one-hour session of flicker ($F(3,15)=0.9306$, $p=0.4502$, RM-ANOVA). See Movie 1 for mouse exposure to sensory flicker.

Extended Data

Table A-1. Parts List

Components for signal generation, modulation and sensory stimulation are needed to assemble and BrainWAVE stimulator. This table lists components that can be used and their estimated price.

Table A-2. BrainWAVE Light and Sound Levels

These volume for the illuminance of LEDs and the volume of a speaker producing flicker stimulation from have been successfully used in prior studies to comfortably produce gamma modulation in humans and mice.

Figure A-S1. Diagram of audio-visual BrainWAVE circuit

The audio and visual signal generators (green), an Arduino Uno, and a data acquisition device (DAQ) supply signals to the signal modulators (purple), an audio amplifier and a MOSFET, which then control the outputs (red) for this circuit, a speaker and LED strip lights. Eight 1.5-volt AA batteries supply power to a strip of LEDs.

Guide A-1. BrainWAVE Stimulator Guide

This guide provides more information on how to assemble a BrainWAVE device and set-up the software to produce the stimulation. The guide also provides information on how to measure the stimulation and troubleshooting tips.

Extended Data A-1 BrainWAVE Stimulator Code

These files contain code to generate and play flicker sensory stimulation with an Arduino Uno or NIDAQ BrainWAVE stimulator device.

Movie A-1. Mouse BrainWAVE exposure.

A mouse is exposed to 40Hz audio visual flicker.

FliCkER Circuit Parts List

Table A-1. *Parts List*

Part	Purpose	Part Number	Cost
<i>Signal Generation (select one)</i>			
Arduino Uno	Generates flicker signal	E.g., Arduino CC:7630049200074	\$21.90
Data Acquisition Device	Generates flicker signal	E.g., National Instruments: USB-6212	\$1,357.00
<i>Signal Modulation</i>			
n-channel MOSFET	Switches LED signal on and off	E.g., Digikey IRFZ44N	\$1.02
Eight AA Battery Holder	Supplies power to 12 V LED lights	E.g., Digikey: BH48AAW-ND	\$3.24
LED Dimmer	Allow adjustment of LED brightness	superbrightleds.com: LDK-8A	\$9.95
Audio Amplifier	Allows adjustment of audio volume	E.g., Amazon: B00ULRFQ1A	\$53.99
<i>Visual Output</i>			
LED Strip Lights	Delivers visual stimuli	E.g., superbrightleds.com: STN-A40K80-B3A- 08B5M-12V	\$28.95
<i>Audio Output</i>			
Speaker	Delivers audio stimuli	Amazon: B0007L8A7M	\$29.99

Circuit Components

12 V ACDC Power Adapter	Supplies power to the Arduino	E.g., Digikey: 364-1268-ND	\$7.96
Breadboard	Hold circuits components and wires	E.g., Digikey: 2183-4000-ND	\$2.55
Jumper Wires	Connects components	E.g., Digikey: BKWK-3-ND	\$3.90

FliCkER Light and Sounds Levels

Table A-2. Light and Sound Levels

<i>FliCkER Light Levels*</i>	
Mice – LED Strip	800 lux at side of enclosure near light source 100 lux at side of enclosure far from light source (~6 inches from front of cage)
Human – Goggles	100-1000 lux, measured from each eye of the LED goggles with light meter sensor directly on the goggle**
Human – LED Frame	500-800 lux, measured at participant’s head position**

<i>FliCkER Audio***</i>	
Mice – One Speaker	60-65 dB, measured at animal’s head position
Human – Earbuds or headphones	70-80 dB, earbuds, earbuds measured with both earbuds next to decibel meter, headphones measured with decibel meter between headphones**

*Light measurements were made using a Traceable® Light Meter set to “fast” *while lights flickering at 40 Hz.*

**For human studies, levels may be adjusted to comfort of participant.

***Measured with a dBA decibel meter.

Assembly Instructions

When using LEDs of 12 V or greater as visual outputs, the output signal typically required a MOSFET circuit to amplify the signal enough for the LEDs to reach a desired level of brightness. A FlickER MOSFET circuit for visual sensory stimulation includes a breadboard, an N-channel MOSFET, wires and a power source (see Supplementary Materials for a detailed parts list). We used either a DC power supply or another AC/DC wall power adapter as a power source (**Fig. A-S1**). For experiments sensitive to electrical noise, we recommend batteries as the voltage source for the LEDs.

The circuit can be assembled on a solderless or solderable breadboard.

Step 1. Place the MOSFET into the breadboard so that the Gate, Drain and Source pins are located in three different columns (**Fig. A-S1**).

Step 2. Connect the ground wire of the battery holder to the Source pin (right) of the MOSFET.

Step 3. Connect the positive end of the battery holder to the positive LED terminal.

Step 4. Connect the negative end of the LEDs to the breadboard in the Drain pin (middle) of the MOSFET.

Step 5. Use a wire to connect a ground pin of the Arduino to the breadboard in the Source (right) column of the MOSFET.

Step 6. Use a wire to connect the Signal pin (Pin 12) of your Arduino to the breadboard in the same column as the Gate (left) pin of the MOSFET.

Step 7. Plug the male USB type-B side of the cable into the Arduino. The other end of the USB cable is connected to either a PC or wall adapter.




Figure A-S1. Diagram of audio-visual FlickER circuit

The audio and visual signal generators (green), an Arduino Uno, and a data acquisition device (DAQ) supply signals to the signal modulators (purple), an audio amplifier and a MOSFET, which then control the outputs (red) for this circuit, a speaker and LED strip lights. Eight 1.5-volt AA batteries supply power to a strip of LEDs.



Video A-1. Mouse FlickER exposure. A mouse is exposed to 40Hz audio visual flicker.

Appendix B A rapid and generalizable goal-directed spatial learning paradigm

Abstract

Rapidly learning new spatial information to achieve goals is a key animal behavior directly relevant for everyday activities. Most electrophysiological studies measure behavior in well-learned tasks or outside a goal-directed task without clear learning outcomes, both of which make the investigation of rapid spatial learning difficult. To fill this need, we developed a novel goal-directed spatial learning paradigm that enables rapid learning of novel virtual environments (within 2-3 days) in mice. The use of a virtual reality system enables controlled presentation of visual cues and efficient switching between environments. Our paradigm provides clear quantitative metrics to monitor behavioral performance over learning. Importantly, mice performed our task on two different treadmills, indicating the robustness and generalizability of the task. The methods we describe here for training mice on the paradigm and analyzing behavior are easy to implement and may be useful to researchers studying memory formation in a variety of experimental settings.

Introduction

Animals rapidly learn to adapt to new contexts by identifying key relevant information and applying it to solve problems at hand. One way to study this essential ability to learn and use new information in a goal-directed manner is to examine how animals learn locations of importance in new environments. Studying the mechanisms of spatial learning may also help understand the mechanisms of well-documented deficits in spatial learning in Alzheimer's disease.

One challenge to studying spatial learning is that most electrophysiological studies use well-learned tasks and over-trained animals to standardize behavior and reduce potential uncontrolled variables in the

experiment. The use of virtual reality in experimental settings offers both precise control of the environment and the flexibility of switching between environments efficiently. Furthermore, most tasks are performed in one experimental setup for consistency, which does not provide information about the robustness or generalizability of the task. The ideal behavioral paradigms to test spatial learning would work well in multiple experimental setups, control the environment to assess the necessity of spatial cues to task performance, and produce sufficient numbers of trials to analyze learning-dependent changes in neural activity over time.

To address these needs, we developed a novel goal-directed virtual reality spatial learning paradigm that mice learn rapidly. In this task, mice first learned one familiar environment and then later learned a novel environment of distinct cues and reward locations. This learning of the new context was achieved rapidly (within 2-3 days) with clear and quantifiable learning-dependent changes in behavior. Importantly, mice could perform this task on two different treadmill systems, demonstrating the robustness of the task independent of experimental equipment. Here, we describe our behavioral training and analysis protocols with supplementary training videos. In addition, we provide methods for determining the performance criteria to include or exclude animals based on previous training history.

Results

Novel goal-directed spatial learning task in virtual reality

We developed a novel goal-directed virtual reality spatial navigation task in head-fixed mice (**Fig. B-1A**). In this task, mice were trained to run on a treadmill while licking for sweetened condensed milk rewards in familiar and unfamiliar environments with distinct visual cues (**Fig. B-1B**). Following habituation on the linear or spherical treadmill, mice underwent 3-5 days of passive training in which mice received automatically delivered rewards at three specific locations referred to as the reward zones (RZ) regardless of their licking activity (**Fig. B-1C**). Mice then transitioned to active training in which they were required to lick in the RZs to trigger reward delivery. To prevent nonspecific licking, mice received a 4-second

timeout if they licked more than 25-50 times outside the RZ and the anticipatory zone (AZ) that immediately precede the RZ. The flexibility of a virtual reality setup allows rapid switching between multiple environments and consistency of cues in the same environment across sessions. We typically trained mice on one track only (“Track A”) for at least 10 days, or until they showed clear signs of learned behavior in the familiar environment, before introducing them to a new environment (**Fig. B-1C**). To monitor position-specific speed over training days, we recorded movement speed in terms of forward advancement in the virtual environment per unit time (in degrees per second, or deg/s) and lick rate (in licks per second, or licks/s). We used receiver-operating characteristics to quantify behavioral performance by determining the likelihood of an animal being either in the AZ or in the non-reward control location (NRZ) based on animal-specific movement speed or lick-based metrics. The NRZ was defined as the zone that appeared 30 degrees after the RZ. The resulting receiver-operating characteristics curves were plotted against the true and false positive rates of binary classification, where larger areas under the curves indicate better performance (**Fig. B-1D**).

Rapid, position-specific learning in novel virtual environments

Mice performing in the familiar environment (Track A) slowed down approached the reward zone. This behavior showed that animals reliably differentiated between rewarded and unrewarded areas based on movement speed, with the average classification score close to 1 across three days (**Fig. B-1D**, shades of black for days 1-3). Mice performing with accuracy >75% for at least two consecutive sessions progressed to learning a novel environment (Track B or Track C). On the first day of exploration in the novel environment, however, mice typically did not slow down regardless of position (**Fig. B-1D**, lightest shade of green for day 1). After repeated exposure and trial and error, mice rapidly learned the novel reward locations within the next 2-3 days over 4-6 sessions of 20-40 minutes each (**Fig. B-1D**, darker shades of green for days 2-3). We quantified the areas under the receiver-operating characteristics curves based on movement speed, latency to lick in the RZ, and lick rate. Consistent with rapid learning of new goal

locations, we found a significant increase in the areas under the curves (**Fig. B-1E**). Although mice were exposed to two alternating environments (one novel and one familiar) multiple times on a single day, we found that behavior in the familiar environment was generally consistent across three days (**Fig. B-1E**). This consistency in familiar behavior indicated that the spatial memory previously developed had not been lost or interrupted by the newly developing memory of the second environment. Thus, our task represents a common form of adulthood learning during which new critical information is learned while maintaining previously stored memories. This novel spatial learning task can be relatively easily integrated into multiple existing experimental protocols and setups including optogenetics, *in vivo* imaging, and fiber photometry, improving the utility of the paradigm.

Robustness and generalizability of the task across treadmill types

Most behavioral tasks are performed in a single type of experimental setup. This provides stability and consistency of task experience but does not guarantee robustness or generalizability of the task across experimental settings. Therefore, we sought to test whether mice could learn the same task on two different treadmill systems. To examine generalizability of the behavioral paradigm across different setups, we trained mice on both the linear treadmill and spherical treadmills. We found that incremental learning of reward locations over sessions was reliably tracked on the both treadmill systems even though these systems had key distinctions. First, on the linear treadmill (**Fig. B-2A**), mice were placed on a flat surface and generally ran faster than on the spherical treadmill (**Fig. B-2F**). Stopping movement abruptly and completely was more likely to occur on the linear treadmill. Mice were allowed to turn their bodies more fully on the spherical treadmill. Our task does not require turning and therefore is comparable between treadmill systems. In the earlier phases of learning, mice slowed down only after entering the RZs and generally did not show a large difference in speed between the AZ and NRZ (**Fig. B-2, B-C**, session 1). By session 9, mice clearly slowed down before entering the RZ, indicating they had learned to anticipate the upcoming reward (**Fig. B-2, B-C**, sessions 9 and 17). The receiver-operating characteristics curves showed

larger deviation from chance-level performance in later sessions (**Fig. B-2D**). The magnitude of the reward zone-specific reduction in movement speed depended on the phase of learning, such that we observed a greater reduction in the later phases of learning (**Fig. B-2E**). The distance between two reward zones was never the same, thus, similar speed changes across three different reward zones could not be due to mice using nonspatial cues (e.g., distance-based strategy). The use of the spherical treadmill system produced similar results. Specifically, mice moved at similar speeds irrespective of position in the earlier training sessions and later slowed significantly in the AZs compared to the NRZs (**Fig. B-2, F-J**). Because we used unique thresholds for determining behavioral differentiation based on the behavior of individual animal in a given session, we could quantify and track learning over time, unaffected by variability across sessions or across animals. Thus, these results show that our task produces robust and quantifiable learning over time across multiple treadmill systems.

Next, we analyzed licking behavior as another measure of learning and tested whether licking behavior was similar on two different treadmill systems. Because mice were trained to lick in the reward zones to trigger a reward, a lick represents the chosen action of the mouse and determines the trial outcome (reward or no reward). Consistent with not knowing the specific reward locations, mice did not show a significant change in their lick rates regardless of position on the first day of training (**Fig. B-3, A and B**, session 1). After learning, however, mice showed a higher lick rate as they approached the RZ (**Fig. B-3B**, sessions 9 and 17). This increase in lick rate was rapidly abolished when the animal left the RZ. As with speed data, the distributions of lick rates in the AZ and NRZ had a larger overlap in the first session than later sessions (**Fig. B-3C**). Reflecting this learning-dependent change, the areas under the receiver-operating characteristics curves also increased significantly from early to late phases of learning (**Fig. B-3, D-E**).

Average lick rates varied per individual animal because some mice preferred to lick at all locations, and lick counts depended on movement speed and occupancy at different locations. Therefore, we used lick latency as another lick-based behavioral metric of interest, measuring the time between arriving at the RZ

and the first lick after the arrival. We hypothesized that well-trained mice would show a minimal to no latency to lick in the reward zone, indicating a higher level of confidence in locating the reward zone. In support of this hypothesis, we found that lick latencies in the reward zone had a negatively skewed distribution, where most trials had zero-second delay in licking. Since reward zone lick latencies do not take into consideration lick behavior outside the reward zone, it is possible that the animal licked in a spatially uniform manner. To eliminate that possibility, we also quantified lick latencies in the NRZs. We found that the distributions of lick latencies in the two zones were significantly different. While significantly different, individual variability in lick behavior resulted in large error bars in group averages (**Fig. B-3, E and J**). Despite using timeout periods to reduce the likelihood of animals using a distance-based learning strategy, some mice showed increased licking in an unrewarded area during earlier phases of learning (**Fig. B-3G**). This behavior was typically observed around the area away from a previous reward zone distanced similarly from another inter-reward distance. Mice eventually stopped unspecific licking activity with prolonged training and successfully progressed to novel spatial learning regardless of the treadmill system used. These data demonstrate that learning and performance in our tasks are not dependent on a particular treadmill system, demonstrating the robustness of the behavioral paradigm.

Discussion

Studying how we use new spatial information to our advantage requires an appropriate behavioral task that is robust and adaptable in many experimental settings. Prior studies have used tasks that require overtraining of animals, the use of aversive cues, or tasks that have neither clearly defined spatial information to be learned nor methods for animal-specific quantification of learning. Here, we introduce a novel goal-directed navigational task in which mice rapidly learn specific reward locations in new environments. We showed that incremental improvement in performance is clear and readily measurable across days. New learning occurred without forgetting the old memories. We demonstrated the generalizability of our task by comparing behavioral performance in two different treadmill systems and finding similar results. In

addition, we presented data-driven methods to establish laboratory-specific performance criteria for determining “learned status” prior to novel exploration.

Our behavioral paradigm presents several advantages. First, it tests and measures learning that is more nuanced, incremental, and of positive valence than fear conditioning. Second, the metrics of incremental learning produce more data points than a single trial and thus are suitable for studying both gradual and rapid learning. Third, the learning of novel environments is achieved rapidly within 2-3 days, which is crucial for experimental setups with limited time for investigation. Importantly, this new learning did not impair the animals’ performance in the familiar environment (**Fig. B-1E**). Thus, our task is well-suited for testing hypotheses regarding long-term memory formation and maintenance. Lastly, our task is robust and easy to implement in multiple experimental settings producing similar behavioral results. We have successfully used this task with acute electrophysiological recordings and optogenetic manipulation. In addition, this task can be used in several other experimental setups spanning spatial and temporal scales. Experimenters may wish to combine our behavioral task with chronic calcium imaging, neuroimaging, fiber photometry, or human intracranial recordings. Thus, our task may be useful for cross-scale or cross-species investigation.

There are important limitations of our task. First, it is important to consider differences between virtual reality and real-world environments when designing an experiment. For example, place cells recorded in the hippocampus are found in both types of environments but there are differences in number, field size, and firing activity. Second, some individual variability in strategy may be missed by our task or interpreted as poor learning. For example, since mice are not required to lick or slow in the AZ, some mice may not show anticipatory licking or slowing. Furthermore, we have mostly focused our task development with male mice at around 3 months of age. There is some evidence for sex- and age-dependent differences in strategy and performance in spatial tasks. Future work will investigate whether the task is uniformly performed by both sexes of wider age ranges. Finally, while common in the field, reward-based learning

tasks require close monitoring of food and/or water consumption as task performance is sensitive to day-to-day variations in motivation.

Given the flexibility and customizability of this novel behavioral paradigm to test rapid spatial learning and recall, we expect this task will be useful in various experimental contexts and designs.

Materials and Methods

Animals

We used C57BL/6J adult male mice (N = 7; 25-35 g) at 10-11 weeks of age (12-14 weeks at the time of behavioral training) in our study. All procedures involving animals were performed in accordance with the guidelines provided by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology. Animals were housed in a reverse dark-light cycle room (07:00 light on, 19:00 light off) with *ad libitum* access to food and water. We performed all behavioral training during the dark cycle.

Surgery

Mice were handled for at least three days prior to stereotaxic surgery. For head-plate implantations, mice were deeply anesthetized with isoflurane and head-plates were affixed to the skull with dental cement (Parkell C&B Metabond, Edgewood, NY). Animals were single-housed and monitored for at least three days post-surgery. Mice received a subcutaneous injection of ketoprofen (0.5 mg/kg) for the first two days and saline for the first three days post-surgery.

Handling and Habituation

For sensitive behavioral experiments, it was imperative that mice were comfortable with experimenters prior to behavioral training. To minimize neophobia to sweetened condensed milk rewards in mice, we introduced milk to mice in their home cages on the first day of habituation. We then placed our hand in the

cage with milk drops on the glove, taking care not to make abrupt movements. Only after a mouse started to approach and/or lick for rewards on the hand, we carefully scooped the animal up, never lifting it up by the tail. We watched for signs of mouse comfort on the experimenter's hand such as calm stillness, grooming, and/or consuming milk. We handled each mouse for ~5 minutes at a time, carefully returning it back into the home cage to rest for 5-10 minutes before handling again. We repeated the handling for 3 days immediately before surgery, and 3 days immediately before habituation on the treadmill.

We started food restriction on mice on the first day of habituation, gradually bringing down the body weight to 85-87% of the post-surgery and post-recovery body weight over ~5 days. For habituation on either a Styrofoam ball or a linear treadmill (PhenoSys SpeedBelt, Berlin, Germany), mice were head-fixed and introduced to the reward delivery needle, which was connected to either a piezoelectric sensor or a photointerruptor system for lick detection. On the first day of head-fixation, the main goal was to make sure the mouse does not vocalize discomfort and learn the position of the reward needle through manual reward delivery. No visuals were shown on the screen during this period of habituation and mice typically learned to lick and run consistently within a day or two on the ball and about 3-5 days on the belt. To encourage voluntary movement on the treadmill, we manually delivered rewards early and often, every time we saw an effort to move. Some mice tended to grab onto the reward dispenser during early sessions. To discourage this behavior, we delivered a higher frequency of rewards when the mouse walked on all four paws. Linear treadmill habituation additionally required ensuring that the body was centered and parallel to the belt, as some animals tended to anchor themselves using the edges of the stationary platform and walk diagonally as a result. We corrected the awkward body positioning by gently pushing the body into the right position and delivering rewards. Once mice could walk or run consistently, we delivered rewards less frequently, with longer time or distance traveled between rewards. It was critical for the animal to run sufficiently fast before introducing it to the virtual environment to ensure consistent delivery of rewards and motivation for

running. Following a successful habituation period, mice were introduced to a virtual reality environment for the first time.

Virtual Reality Apparatus

A cylindrical screen projected the virtual reality environment in front of the treadmill system. All virtual tracks were created with the freely available Virtual Reality MATLAB Engine (Aronov and Tank, 2014b). All annular tracks used in this study were of the same size but had distinct proximal wall cues and distal cues. There were 36 fixed-distance wall cues (10 degrees per cue) and 3 of them, termed “reward zones,” were associated with reward on each track. Reward zones were unevenly positioned so that inter-reward zone distance was unique. Speed gain was set such that the overall experience of virtual environments was similar on both treadmill systems.

Tracking of Position, Speed and Licking Behavior

Forward and sideways movement on the ball was detected by one or two optical mice placed in the front and/or side of the ball. Movement on the belt was detected as an analog signal sent to a NIDAQ and to MATLAB to automatically update the positional cues in virtual reality. Mice moved around the track unidirectionally and never skipped a cue. Licks were detected with either a photointerruptor system as previously described, or with a piezoelectric approach to convert mechanical or vibrational energy into changes in voltage. Either detection system was similarly accurate in detecting mouse licks.

Virtual Reality Behavioral Training and Familiar Track Training

Following treadmill habituation, each mouse received one or more 20-40-minute session(s) per day, 6-7 days a week. During behavioral training, mice were exposed to only one track, “Track A.” During recording, the same “Track A” was used as the familiar environment in comparisons between the familiar and two distinct novel tracks (“Track B” and “Track C”).

Behavioral training was divided into two stages, called passive and active training. Mice first entered passive training, in which visual feedback for movement was given through the virtual reality apparatus for the first time. This initial phase only required mice to travel to the reward zones to receive automatically delivered rewards. Over sessions, we looked for signs of anticipatory behavior, specifically slowing down and increasing lick rate when in the zone preceding the reward zone. We defined this zone, which was the same size as the reward zone, as the anticipatory zone to quantify learning over training sessions. Similarly, to analyze position-specific behavior, we defined the non-reward zone as the zone starting 30 degrees past the end of the reward zone and ending at the start of the following anticipatory zone. We wanted this zone to begin not too close to the reward zone, where the animal could still be consuming rewards, but also not too far away, where the behaviors might be confounded by the upcoming anticipatory zone of a different reward zone. In the non-reward zone, we monitored for consistently high speeds as well as low lick rates.

Once consistent discrimination between reward zones and non-reward zones was evident (typically taking anywhere from 3-5 days), mice were moved to active training, in which reward dispense was behavior dependent. In these sessions, animals were required to lick in the reward zones to trigger a reward delivery. Once again, we looked for consistent discriminatory behavior between reward and non-reward zones during sessions. To prevent generalized licking irrespective of position on training days, licking more than 25-50 cumulative times in zones outside the anticipatory and reward zones triggered a 4-second time-out period. While we used both 25 and 50 licks as the timeout threshold, we found the latter to be more effective and generalizable across mice. During a time-out period, the animal was teleported to a grey box with no reward delivery and teleported back to the original location and resumed where it had left off. Occasionally, a mouse would miss multiple rewards in a row during these active sessions; to disrupt this pattern of behavior, we would temporarily move the mouse back to passive training until discriminatory behavior ensued. Once a mouse consistently demonstrated expert behavior, qualitatively defined as high

levels of lick and speed discrimination between reward zones and non-reward zones in the familiar environment, the animal was exposed to a novel environment.

Novel Track Training

After learning the familiar track, mice were exposed to two different tracks, familiar Track A and a novel track (either Track B or Track C), alternating between the novel and familiar environments throughout the day for three consecutive days. We started each day with the same familiar track before the novel track. Mice initially spent 15-20 minutes in the familiar environment, followed by a baseline period of ~10 minutes in the dark. During the baseline periods, we tracked movement around the familiar track and licking behavior, but the projector was covered so no cues were visible to the animal. After the baseline period, mice spent 20-30 minutes on the novel track, followed by another baseline period. This sequence was repeated at least twice per day. On the first day of novel exposure, the lick threshold for a timeout was set very high (over 100 licks) as not to trigger timeouts and not to discourage animals from sampling different positions as they sought novel reward zones. By novel day 3, this threshold was brought down to the level consistent with normal training days.

Quantification of Running and Licking Behaviors

Raw behavioral data were divided based on an animal's position (in degrees) into either 360-degree "laps" spanning the entire environment or 70-degree "trials" that spanned the area 30 degrees before and after each reward zone. Each reward zone was 10 degrees. These laps or trials were then divided into 2-degree position bins (~3 cm on the belt). Mean speed (in degrees/s) was calculated by dividing the total distance traveled within a position bin by the total time spent in that bin. Lick rate (in licks/s) was calculated by dividing the number of total licks per position by the total time spent in that bin over all position bins. Lick probability was calculated by taking the ratio of the number of licks per position bin to the total number of licks within the lap or trial so that the numbers would sum up to 1. Speed and lick behavior for each lap or trial was

smoothed using the robust lowess method with a moving average window of 11 bins. Once learned, mice tended to show stereotyped behavior in all three reward zones. Therefore, speed and lick behavioral data during all trials were concatenated and averaged for each session regardless of the absolute position on the track.

Receiver Operating Characteristic Curves Based on Speed and Licking Behavior

We used receiver operating characteristic (ROC) curves to quantify and visualize learning performance over sessions or days. We used position-based binned speed or lick rate data in the anticipatory and non-reward zones for binary classification between these two zones at varying discrimination threshold levels of mean speed or lick rate. This ROC metric is useful because it uses animal-specific criteria for each session independently as opposed to using a fixed criterion applied to all sessions. This approach is especially useful because mice usually increase their maximum speed with training making a fixed criterion unsuitable. The areas under the curves (AUC) range between 0.5 and 1, with 0.5 being chance-level classification, or differentiation between curves, and 1 being perfect classification. We averaged AUC values across animals for each environment. Because mice differed in their learning speeds and total number of sessions received, we divided total learning sessions into three blocks as early, mid, and late phases of learning. To quantify performance over these phases of learning across animals, we used the average of the first and last two sessions as the early and late phases, respectively, and two sessions in the mid-point as the mid-phase of learning.

Statistical Analyses

Analyses of behavioral performance were performed with custom scripts in MATLAB (MathWorks, MA). Unless noted otherwise, we used the Wilcoxon signed-rank test and Bonferroni correction for multiple comparison where appropriate.

Figures




Figure B-1. Novel rapid virtual reality spatial learning paradigm. **A.** Illustration of head-fixed mouse surrounded by a virtual reality screen (left) and close-up of the head-fixed setup. **B.** Top, example virtual reality environments with distinct cues. Bottom, example timeline of track exposure. **C.** Experimental timeline used in the study. **D.** Speed-based receiver-operating characteristic curves across three days in the familiar ("Fam D1-3," shades of black) or novel ("Nov D1-3," shades of green) environments in an example animal. **E.** Areas under the receiver-operating characteristic curves based on speed (left), lick latency (middle), or lick rate (right) in familiar (black) and novel (green) environments across three days.




Figure B-2. Generalizable movement speed-based quantification of learning across two different treadmill systems. **A.** Illustration of linear treadmill setup. **B.** Session average speed as a function of track position across three example sessions where darker shades of color indicate later phases of training of an example animal. **C.** Distributions of speed in the AZ (blue bars) and the NRZ (white bars) in the three sessions indicated in B. **D.** Receiver-operating characteristic curves for each of the three sessions indicated

in B and C. **E.** Average speed-based behavioral performance across all trained animals ($N = 7$). Early, middle, and late stages of training indicate the first, second, and last third of sessions, respectively. **F.** Illustration of spherical treadmill setup. **G.** Same as B, but for an example session with the spherical treadmill system. **H.** Same as C, but for an example session with the spherical treadmill system. **I.** Same as D, but for an example session with the spherical treadmill system. **J.** Same as E, but for the spherical treadmill system.




Figure B-3. Generalizable lick rate-based quantification of learning across two different treadmill systems. **A.** Illustration of linear treadmill setup. **B.** Session average lick rate as a function of track position across three example sessions where darker shades of color indicate later phases of training of an example animal. **C.** Distributions of lick rates in the AZ (pink bars) and the NRZ (white bars) in the three sessions indicated in **B**. **D.** Receiver-operating characteristic curves for each of the three sessions indicated in **B** and **C**. **E.** Average lick-based performance across all trained animals ($N = 7$). Early, middle, and late

*stages of training indicate the first, second, and last third of sessions, respectively. **F.** Illustration of spherical treadmill setup. **G.** Same as B, but for an example session with the spherical treadmill system. **H.** Same as C, but for an example session with the spherical treadmill system. **I.** Same as D, but for an example session with the spherical treadmill system. **J.** Same as E, but for the spherical treadmill system.*

References

- Adaikkan, C., Middleton, S.J., Marco, A., Pao, P.-C., Mathys, H., Kim, D.N.-W., Gao, F., Young, J.Z., Suk, H.-J., Boyden, E.S., et al. (2019). Gamma Entrainment Binds Higher-Order Brain Regions and Offers Neuroprotection. *Neuron* 102, 929-943.e8.
- Allen, K., and Monyer, H. (2015). Interneuron control of hippocampal oscillations. *Curr. Opin. Neurobiol.* 31, 81–87.
- Allen, B.D., Singer, A.C., and Boyden, E.S. (2015). Principles of designing interpretable optogenetic behavior experiments. *Learn. Mem.* 22, 232–238.
- Allen, K., Fuchs, E.C., Jaschonek, H., Bannerman, D.M., and Monyer, H. (2011). Gap Junctions between Interneurons Are Required for Normal Spatial Coding in the Hippocampus and Short-Term Spatial Memory. *J. Neurosci.* 31, 6542.
- Anderson, M.C., and Floresco, S.B. (2021). Prefrontal-hippocampal interactions supporting the extinction of emotional memories: the retrieval stopping model. *Neuropsychopharmacol.* 2021 471 47, 180–195.
- Andrews-Zwilling, Y., Gillespie, A.K., Kravitz, A. V., Nelson, A.B., Devidze, N., Lo, I., Yoon, S.Y., Bien-Ly, N., Ring, K., Zwilling, D., et al. (2012). Hilar GABAergic Interneuron Activity Controls Spatial Learning and Memory Retrieval. *PLoS One* 7, e40555.
- Antonoudiou, P., Tan, Y.L., Kontou, G., Louise Upton, A., and Mann, E.O. (2020). Parvalbumin and Somatostatin Interneurons Contribute to the Generation of Hippocampal Gamma Oscillations. *J. Neurosci.* 40, 7668–7687.
- Aronov, D., and Tank, D.W. (2014a). Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron* 84, 442–456.

- Aronov, D., and Tank, D.W. (2014b). Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. *Neuron* 84, 442–456.
- Artinian, J., Jordan, A., Khlaifia, A., Honore, E., Fontaine, A. La, Racine, A.S., Laplante, I., and Lacaille, J.C. (2019). Regulation of Hippocampal Memory by mTORC1 in Somatostatin Interneurons. *J. Neurosci.* 39, 8439–8456.
- Bannon, N.M., Chistiakova, M., and Volgushev, M. (2020). Synaptic Plasticity in Cortical Inhibitory Neurons: What Mechanisms May Help to Balance Synaptic Weight Changes? *Front. Cell. Neurosci.* 14, 204.
- Bates, D., Mächler, M., Bolker, B.M., and Walker, S.C. (2015). Fitting linear mixed-effects models using lme4. *J. Stat. Softw.* 67, 1–48.
- Best, P.J., and White, A.M. (1998). Hippocampal cellular activity: A brief history of space. *Proc. Natl. Acad. Sci. U. S. A.* 95, 2717.
- Bhatia, A., Moza, S., and Bhalla, U.S. (2019). Precise excitation-inhibition balance controls gain and timing in the hippocampus. *Elife* 8.
- Bissonette, G.B., Schoenbaum, G., Roesch, M.R., and Powell, E.M. (2015). Interneurons are necessary for coordinated activity during reversal learning in orbitofrontal cortex. *Biol. Psychiatry* 77, 454.
- Bittner, K.C., Grienberger, C., Vaidya, S.P., Milstein, A.D., Macklin, J.J., Suh, J., Tonegawa, S., and Magee, J.C. (2015). Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat. Neurosci.* 2015 18, 1133–1142.
- Bittner, K.C., Milstein, A.D., Grienberger, C., Romani, S., and Magee, J.C. (2017). Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* (80-.). 357, 1033–1036.
- Busche, M.A., Eichhoff, G., Adelsberger, H., Abramowski, D., Wiederhold, K.H., Haass, C., Staufenbiel,

- M., Konnerth, A., and Garaschuk, O. (2008). Clusters of hyperactive neurons near amyloid plaques in a mouse model of Alzheimer's disease. *Science* (80-.).
- Butler, J.L., and Paulsen, O. (2015). Hippocampal network oscillations — recent insights from in vitro experiments. *Curr. Opin. Neurobiol.* 31, 40–44.
- Buzsáki, G. (2010). Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron* 68, 362–385.
- Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus* 25, 1073.
- Caccavano, A., Bozzelli, P.L., Forcelli, P.A., Pak, D.T.S., Wu, J.-Y., Conant, K., and Vicini, S. (2020). Inhibitory parvalbumin basket cell activity is selectively reduced during hippocampal sharp wave ripples in a mouse model of familial Alzheimer's disease. *J. Neurosci.*
- Carey, A.A., Tanaka, Y., and van der Meer, M.A.A. (2019). Reward revaluation biases hippocampal replay content away from the preferred outcome. *Nat. Neurosci.* 2019 229 22, 1450–1459.
- Chan, D., Suk, H.-J., Jackson, B., Milman, N.P., Stark, D., Klerman, E.B., Kitchener, E., Avalos, V.S.F., Banerjee, A., Beach, S.D., et al. (2021). Gamma Frequency Sensory Stimulation in Probable Mild Alzheimer's Dementia Patients: Results of a Preliminary Clinical Trial. *MedRxiv* 15, 2021.03.01.21252717.
- Chen, J.L., Lin, W.C., Cha, J.W., So, P.T., Kubota, Y., and Nedivi, E. (2011). Structural basis for the role of inhibition in facilitating adult brain plasticity. *Nat. Neurosci.* 2011 145 14, 587–594.
- Chung, H., Park, K., Jang, H.J., Kohl, M.M., and Kwag, J. (2020). Dissociation of somatostatin and parvalbumin interneurons circuit dysfunctions underlying hippocampal theta and gamma oscillations impaired by amyloid β oligomers in vivo. *Brain Struct. Funct.* 225, 935–954.
- Chung, W.-S., Welsh, C.A., Barres, B.A., and Stevens, B. (2015). Do glia drive synaptic and cognitive

- impairment in disease? *Nat. Neurosci.* *18*.
- Cimenser, A., Hempel, E., Travers, T., Strozewski, N., Martin, K., Malchano, Z., and Hajós, M. (2021). Sensory-Evoked 40-Hz Gamma Oscillation Improves Sleep and Daily Living Activities in Alzheimer's Disease Patients. *Front. Syst. Neurosci.* *15*, 103.
- Clouter, A., Shapiro, K.L., and Hanslmayr, S. (2017). Theta Phase Synchronization Is the Glue that Binds Human Associative Memory. *Curr. Biol.* *27*, 3143-3148.e6.
- Colgin, L.L. (2016). Rhythms of the hippocampal network. *Nat. Rev. Neurosci.* *17*, 239–249.
- Contestabile, A., Magara, S., and Cancedda, L. (2017). The GABAergic hypothesis for cognitive disabilities in down syndrome. *Front. Cell. Neurosci.* *11*, 54.
- Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving Rat. *J. Neurosci.*
- Danielson, N.B., Zaremba, J.D., Kaifosh, P., Bowler, J., Ladow, M., and Losonczy, A. (2016). Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1. *Neuron* *91*, 652–665.
- Delorme, J., Wang, L., Kuhn, F.R., Kodoth, V., Ma, J., Martinez, J.D., Raven, F., Toth, B.A., Balendran, V., Medina, A.V., et al. (2021). Sleep loss drives acetylcholine- And somatostatin interneuron-mediated gating of hippocampal activity to inhibit memory consolidation. *Proc. Natl. Acad. Sci. U. S. A.* *118*.
- Deng, X., Gu, L., Sui, N., Guo, J., and Liang, J. (2019). Parvalbumin interneuron in the ventral hippocampus functions as a discriminator in social memory. *Proc. Natl. Acad. Sci.* *116*, 16583–16592.
- Doron, A., Rubin, A., Benmelech-Chovav, A., Benaim, N., Carmi, T., Refaeli, R., Novick, N., Kreisel, T.,

- Ziv, Y., and Goshen, I. (2022). Hippocampal astrocytes encode reward location. *Nat.* 2022 6097928 609, 772–778.
- Dudok, B., Szoboszlay, M., Paul, A., Klein, P.M., Liao, Z., Hwaun, E., Szabo, G.G., Geiller, T., Vancura, B., Wang, B.S., et al. (2021). Recruitment and inhibitory action of hippocampal axo-axonic cells during behavior. *Neuron* 109, 3838-3850.e8.
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* 13, 995–1002.
- Dupret, D., O'Neill, J., Csicsvari, J., O'Neill, J., and Csicsvari, J. (2013). Dynamic reconfiguration of hippocampal interneuron circuits during spatial learning. *Neuron* 78, 166–180.
- Dvorak, D., Radwan, B., Sparks, F.T., Talbot, Z.N., and Fenton, A.A. (2018). Control of recollection by slow gamma dominating mid-frequency gamma in hippocampus CA1. *PLOS Biol.* 16.
- Ego-Stengel, V., and Wilson, M.A. (2007). Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus* 17, 161–174.
- English, D.F., McKenzie, S., Evans, T., Kim, K., Yoon, E., and Buzsáki, G. (2017). Pyramidal Cell-Interneuron Circuit Architecture and Dynamics in Hippocampal Networks. *Neuron* 96, 505-520.e7.
- Evangelista, R., Cano, G., Cooper, C., Schmitz, D., Maier, N., and Kempter, R. (2020). Generation of Sharp Wave-Ripple Events by Disinhibition. *J. Neurosci.* 40, 7811–7836.
- Fernández-Ruiz, A., Oliva, A., Oliveira, E.F. de, Rocha-Almeida, F., Tingley, D., and Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. *Science* (80-.). 364, 1082–1086.

- Fernandez, F., Morishita, W., Zuniga, E., Nguyen, J., Blank, M., Malenka, R.C., and Garner, C.C. (2007). Pharmacotherapy for cognitive impairment in a mouse model of Down syndrome. *Nat. Neurosci.* 2007 104 10, 411–413.
- Filice, F., Janickova, L., Henzi, T., Bilella, A., and Schwaller, B. (2020). The Parvalbumin Hypothesis of Autism Spectrum Disorder. *Front. Cell. Neurosci.* 14, 415.
- Francavilla, R., Villette, V., Luo, X., Chamberland, S., Muñoz-Pino, E., Camiré, O., Wagner, K., Kis, V., Somogyi, P., and Topolnik, L. (2018). Connectivity and network state-dependent recruitment of long-range VIP-GABAergic neurons in the mouse hippocampus. *Nat. Commun.* 2018 91 9, 1–17.
- Francavilla, R., Villette, V., Martel, O., and Topolnik, L. (2019). Calcium dynamics in dendrites of hippocampal CA1 interneurons in awake mice. *Front. Cell. Neurosci.* 13, 98.
- Fu, Y., Tucciarone, J.M., Espinosa, J.S., Sheng, N., Darcy, D.P., Nicoll, R.A., Huang, Z.J., and Stryker, M.P. (2014). A Cortical Circuit for Gain Control by Behavioral State. *Cell* 156, 1139–1152.
- Fu, Y., Kaneko, M., Tang, Y., Alvarez-Buylla, A., and Stryker, M.P. (2015). A cortical disinhibitory circuit for enhancing adult plasticity. *Elife* 2015.
- Fukuda, T., and Kosaka, T. (2000). Gap Junctions Linking the Dendritic Network of GABAergic Interneurons in the Hippocampus. *J. Neurosci.* 20, 1519–1528.
- Galambos, R., Makeig, S., and Talmachoff, P.J. (1981). A 40-Hz auditory potential recorded from the human scalp. 78.
- Garza, K.M., Zhang, L., Borron, B., Wood, L.B., and Singer, A.C. (2020). Gamma Visual Stimulation Induces a Neuroimmune Signaling Profile Distinct from Acute Neuroinflammation. *J. Neurosci.* 40, 1211–1225.
- Gauthier, J.L., and Tank, D.W. (2018a). A Dedicated Population for Reward Coding in the Hippocampus.

Neuron 99, 179-193.e7.

Gauthier, J.L., and Tank, D.W. (2018b). A Dedicated Population for Reward Coding in the Hippocampus.

Neuron 99, 179-193.e7.

Geiller, T., Vancura, B., Terada, S., Troullinou, E., Chavlis, S., Tsagkatakis, G., Tsakalides, P., Ócsai, K.,

Poirazi, P., Rózsa, B.J., et al. (2020). Large-Scale 3D Two-Photon Imaging of Molecularly

Identified CA1 Interneuron Dynamics in Behaving Mice. *Neuron* 0.

Geiller, T., Sadeh, S., Rolotti, S. V., Blockus, H., Vancura, B., Negrean, A., Murray, A.J., Rózsa, B.,

Polleux, F., Clopath, C., et al. (2021). Local circuit amplification of spatial selectivity in the

hippocampus. *Nature* 1–5.

Góis, Z.H.T.D., and Tort, A.B.L. (2018). Characterizing Speed Cells in the Rat Hippocampus. *Cell Rep.*

25, 1872-1884.e4.

Gordon, J.A., Lacefield, C.O., Kentros, C.G., and Hen, R. (2005). State-Dependent Alterations in

Hippocampal Oscillations in Serotonin 1A Receptor-Deficient Mice. *J. Neurosci.* 25, 6509–6519.

Goutagny, R., Gu, N., Cavanagh, C., Jackson, J., Chabot, J.-G., Emi Quirion, R., Krantic, S., and

Williams, S. (2013). Alterations in hippocampal network oscillations and theta–gamma coupling

arise before Ab overproduction in a mouse model of Alzheimer’s disease. *Eur. J. Neurosci.* 37,

1896–1902.

Grienberger, C., Milstein, A.D., Bittner, K.C., Romani, S., and Magee, J.C. (2017). Inhibitory suppression

of heterogeneously tuned excitation enhances spatial coding in CA1 place cells. *Nat. Neurosci.*

20, 417.

Hangya, B., Li, Y., Muller, R.U., and Czurkó, A. (2010). Complementary spatial firing in place cell-

interneuron pairs. *J. Physiol.*

- Harris, K.D., Hochgerner, H., Skene, N.G., Magno, L., Katona, L., Bengtsson Gonzales, C., Somogyi, P., Kessaris, N., Linnarsson, S., and Hjerling-Leffler, J. (2018). Classes and continua of hippocampal CA1 inhibitory neurons revealed by single-cell transcriptomics. *PLOS Biol.* 16, e2006387.
- He, Q., Colon-Motas, K.M., Pybus, A.F., Piendel, L., Seppa, J.K., Walker, M.L., Manzanares, C.M., Qiu, D., Miocinovic, S., Wood, L.B., et al. (2021). A feasibility trial of gamma sensory flicker for patients with prodromal Alzheimer's disease. *Alzheimer's Dement. Transl. Res. Clin. Interv.* 7.
- Hennequin, G., Agnes, E.J., and Vogels, T.P. (2017). Inhibitory Plasticity: Balance, Control, and Codependence. <https://doi.org/10.1146/annurev-neuro-072116-031005> 40, 557–579.
- Herrmann, C.S., and Demiralp, T. (2005). Human EEG gamma oscillations in neuropsychiatric disorders.
- Hescham, S., Liu, H., Jahanshahi, A., and Temel, Y. (2020). Deep brain stimulation and cognition: Translational aspects. *Neurobiol. Learn. Mem.* 174, 107283.
- Hickman, S., Izzy, S., Sen, P., Morsett, L., and El Khoury, J. (2018). Microglia in neurodegeneration. *Nat. Neurosci.* 21, 1359–1369.
- Higley, M.J., and Contreras, D. (2006). Balanced Excitation and Inhibition Determine Spike Timing during Frequency Adaptation. *J. Neurosci.* 26, 448–457.
- Hollup, S.A., Molden, S., Donnett, J.G., Moser, M.B., and Moser, E.I. (2001). Accumulation of Hippocampal Place Fields at the Goal Location in an Annular Watermaze Task. *J. Neurosci.* 21, 1635–1644.
- Hong, S., Dissing-Olesen, L., and Stevens, B. (2016). New insights on the role of microglia in synaptic pruning in health and disease. *Curr. Opin. Neurobiol.* 36, 128–134.
- Huang, X., Huang, P., Huang, L., Hu, Z., Liu, X., Shen, J., Xi, Y., Yang, Y., Fu, Y., Tao, Q., et al. (2021). A Visual Circuit Related to the Nucleus Reunions for the Spatial-Memory-Promoting Effects of

- Light Treatment. *Neuron* 109, 347-362.e7.
- Iaccarino, H.F., Singer, A.C., Martorell, A.J., Rudenko, A., Gao, F., Gillingham, T.Z., Mathys, H., Seo, J., Kritskiy, O., Abdurrob, F., et al. (2016). Gamma frequency entrainment attenuates amyloid load and modifies microglia. *Nature* 540.
- Jackson, J.C., Johnson, A., and Redish, A.D. (2006). Hippocampal Sharp Waves and Reactivation during Awake States Depend on Repeated Sequential Experience. *J. Neurosci.*
- Jadhav, S.P., Kemere, C., German, P.W., and Frank, L.M. (2012). Awake Hippocampal Sharp-Wave Ripples Support Spatial Memory. *Science* (80-.). 336, 1454–1458.
- Jones, M., Mcdermott, B., Oliveira, B.L., O'brien, A., Coogan, D., Lang, M., Moriarty, N., Dowd, E., Quinlan, L., McGinley, B., et al. (2019). Gamma Band Light Stimulation in Human Case Studies: Groundwork for Potential Alzheimer's Disease Treatment. *J. Alzheimer's Dis.* 70, 171–185.
- Karlsson, M.P., and Frank, L.M. (2009). Awake replay of remote experiences in the hippocampus. *Nat. Neurosci.* 12, 913–918.
- Kim, S.M., and Frank, L.M. (2009). Hippocampal Lesions Impair Rapid Learning of a Continuous Spatial Alternation Task. *PLoS One* 4, e5494.
- Klausberger, T., and Somogyi, P. (2008). Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations. *Science* (80-.). 321, 53–57.
- Knierim, J.J. (2002). Dynamic interactions between local surface cues, distal landmarks, and intrinsic circuitry in hippocampal place cells. *J. Neurosci.* 22, 6254–6264.
- Krishnan, S., Heer, C., Cherian, C., and Sheffield, M.E.J. (2022). Reward expectation extinction restructures and degrades CA1 spatial maps through loss of a dopaminergic reward proximity signal. *Nat. Commun.* 2022 131 13, 1–19.

- Kuhlman, S.J., Olivas, N.D., Tring, E., Ikrar, T., Xu, X., and Trachtenberg, J.T. (2013). A disinhibitory microcircuit initiates critical-period plasticity in the visual cortex. *Nat.* 2013 5017468 501, 543–546.
- Kuznetsova, A., Brockhoff, P.B., and Christensen, R.H.B. (2017). lmerTest Package: Tests in Linear Mixed Effects Models. *J. Stat. Softw.* 82, 1–26.
- Kvitsiani, D., Ranade, S., Hangya, B., Taniguchi, H., Huang, J.Z., and Kepecs, A. (2013). Distinct behavioural and network correlates of two interneuron types in prefrontal cortex. *Nature* 498, 363–366.
- Lee, I., and Kesner, R.P. (2002). Differential contribution of NMDA receptors in hippocampal subregions to spatial working memory. *Nat. Neurosci.* 5, 162–168.
- Lee, H., Ghim, J.W., Kim, H., Lee, D., and Jung, M.W. (2012). Hippocampal neural correlates for values of experienced events. *J. Neurosci.* 32, 15053–15065.
- Lee, J.S., Briguglio, J.J., Cohen, J.D., Romani, S., Lee Correspondence, A.K., and Lee, A.K. (2020). The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value. *Cell* 183, 1–16.
- Leung, L., Andrews-Zwilling, Y., Yoon, S.Y., Jain, S., Ring, K., Dai, J., Wang, M.M., Tong, L., Walker, D., and Huang, Y. (2012). Apolipoprotein E4 Causes Age- and Sex-Dependent Impairments of Hilar GABAergic Interneurons and Learning and Memory Deficits in Mice. *PLoS One*.
- Li, G., Bien-Ly, N., Andrews-Zwilling, Y., Xu, Q., Bernardo, A., Ring, K., Halabisky, B., Deng, C., Mahley, R.W., and Huang, Y. (2009). GABAergic Interneuron Dysfunction Impairs Hippocampal Neurogenesis in Adult Apolipoprotein E4 Knockin Mice. *Cell Stem Cell*.
- Li, H., Zhao, J., Lai, L., Xia, Y., Wan, C., Wei, S., Liang, J., Chen, Y., and Xu, N. (2022). Loss of SST and PV positive interneurons in the ventral hippocampus results in anxiety-like behavior in

- 5xFAD mice. *Neurobiol. Aging* 117, 165–178.
- Liu, G. (2004). Local structural balance and functional interaction of excitatory and inhibitory synapses in hippocampal dendrites. *Nat. Neurosci.* 2004 7 4, 373–379.
- Liu, B. hua, Li, Y. tang, Ma, W. pei, Pan, C. jie, Zhang, L.I., and Tao, H.W. (2011). Broad inhibition sharpens orientation selectivity by expanding input dynamic range in mouse simple cells. *Neuron* 71, 542–554.
- Lovett-Barron, M., Kaifosh, P., Kheirbek, M.A., Danielson, N., Zaremba, J.D., Reardon, T.R., Turi, G.F., Hen, R., Zemelman, B. V., and Losonczy, A. (2014). Dendritic inhibition in the hippocampus supports fear learning. *Science* (80-.). 343, 857–863.
- Lozano, A.M., Lipsman, N., Bergman, H., Brown, P., Chabardes, S., Chang, J.W., Matthews, K., McIntyre, C.C., Schlaepfer, T.E., Schulder, M., et al. (2019). Deep brain stimulation: current challenges and future directions. *Nat. Rev. Neurol.* 15, 148.
- Ma, L., Tai, X., Su, L., Shi, L., Wang, E., and Qin, L. (2013). The Neuronal Responses to Repetitive Acoustic Pulses in Different Fields of the Auditory Cortex of Awake Rats. *PLoS One* 8, e64288.
- Mably, A.J., Gereke, B.J., Jones, D.T., and Colgin, L.L. (2017). Impairments in spatial representations and rhythmic coordination of place cells in the 3xTg mouse model of Alzheimer’s disease. *Hippocampus*.
- MacDonald, C.J., Lepage, K.Q., Eden, U.T., and Eichenbaum, H. (2011). Hippocampal “time cells” bridge the gap in memory for discontinuous events. *Neuron* 71, 737–749.
- Madisen, L., Mao, T., Koch, H., Zhuo, J.M., Berenyi, A., Fujisawa, S., Hsu, Y.W.A., Garcia, A.J., Gu, X., Zanella, S., et al. (2012). A toolbox of Cre-dependent optogenetic transgenic mice for light-induced activation and silencing. *Nat. Neurosci.* 15, 793.

- Maestú, F., de Haan, W., Busche, M.A., and DeFelipe, J. (2021). Neuronal excitation/inhibition imbalance: core element of a translational perspective on Alzheimer pathophysiology. *Ageing Res. Rev.* 69, 101372.
- Mann, E.O., Radcliffe, C.A., and Paulsen, O. (2005). Hippocampal gamma-frequency oscillations: from interneurons to pyramidal cells, and back. *J. Physiol.* 562, 55.
- Marín, O. (2012). Interneuron dysfunction in psychiatric disorders. *Nat. Rev. Neurosci.* 2012 132 13, 107–120.
- Martinez-Losa, M., Tracy, T.E., Ma, K., Verret, L., Clemente-Perez, A., Khan, A.S., Cobos, I., Ho, K., Gan, L., Mucke, L., et al. (2018). Nav1.1-Overexpressing Interneuron Transplants Restore Brain Rhythms and Cognition in a Mouse Model of Alzheimer's Disease. *Neuron*.
- Martorell, A.J., Paulson, A.L., Suk, H.-J., Abdurrob, F., Drummond, G.T., Guan, W., Young, J.Z., Kim, D.N.-W., Kritskiy, O., Barker, S.J., et al. (2019). Multi-sensory Gamma Stimulation Ameliorates Alzheimer's-Associated Pathology and Improves Cognition. *Cell* 0.
- McKenzie, S., Huszár, R., English, D.F., Kim, K., Christensen, F., Yoon, E., and Buzsáki, G. (2021). Preexisting hippocampal network dynamics constrain optogenetically induced place fields. *Neuron* 109, 1040-1054.e7.
- McNaughton, B.L., and Morris, R.G.M. (1987). Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends Neurosci.* 10, 408–415.
- Milstein, A.D., Li, Y., Bittner, K.C., Grienberger, C., Soltesz, I., Magee, J.C., and Romani, S. (2021). Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife* 10.
- Minderer, M., Harvey, C.D., Donato, F., and Moser, E.I. (2016). Virtual reality explored. *Nat.* 2016 5337603 533, 324–325.

- Mitra, P., and Bokil, H. (2009). *Observed Brain Dynamics* (Oxford University Press).
- Moser, E., Moser, M.B., and Andersen, P. (1993). Spatial learning impairment parallels the magnitude of dorsal hippocampal lesions, but is hardly present following ventral lesions. *J. Neurosci.* *13*, 3916–3925.
- Murty, D.V.P.S., Manikandan, K., Kumar, W.S., Ramesh, R.G., Purokayastha, S., Nagendra, B., Abhishek, M.L., Balakrishnan, A., Javali, M., Rao, N.P., et al. (2021). Stimulus-induced gamma rhythms are weaker in human elderly with mild cognitive impairment and alzheimer’s disease. *Elife* *10*.
- Nakazawa, K., Sun, L.D., Quirk, M.C., Rondi-Reig, L., Wilson, M.A., and Tonegawa, S. (2003). Hippocampal CA3 NMDA receptors are crucial for memory acquisition of one-time experience. *Neuron* *38*, 305–315.
- Nissen, W., Szabo, A., Somogyi, J., Somogyi, P., and Lamsa, K.P. (2010). Cell Type-Specific Long-Term Plasticity at Glutamatergic Synapses onto Hippocampal Interneurons Expressing either Parvalbumin or CB1 Cannabinoid Receptor. *J. Neurosci.* *30*, 1337–1347.
- Nitz, D., and McNaughton, B. (2004). Differential Modulation of CA1 and Dentate Gyrus Interneurons During Exploration of Novel Environments. *J. Neurophysiol.* *91*, 863–872.
- Norimoto, H., Makino, K., Gao, M., Shikano, Y., Okamoto, K., Ishikawa, T., Sasaki, T., Hioki, H., Fujisawa, S., and Ikegaya, Y. (2018). Hippocampal ripples down-regulate synapses. *Science* (80-.). *359*, 1524–1527.
- O’keefe, J., and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*.
- O’Keefe, J., and Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* *34*, 171–175.

- Oliveira da Cruz, J.F., Busquets-Garcia, A., Zhao, Z., Varilh, M., Lavanco, G., Bellocchio, L., Robin, L., Cannich, A., Julio-Kalajzić, F., Lesté-Lasserre, T., et al. (2020). Specific Hippocampal Interneurons Shape Consolidation of Recognition Memory. *Cell Rep.* 32, 108046.
- Osipova, D., Pekkonen, E., and Ahveninen, J. (2006). Enhanced magnetic auditory steady-state response in early Alzheimer's disease. *Clin. Neurophysiol.* 117, 1990–1995.
- Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Harris, K.D. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels (Cold Spring Harbor Labs Journals).
- Palop, J.J., Chin, J., Roberson, E.D., Wang, J., Thwin, M.T., Bien-Ly, N., Yoo, J., Ho, K.O., Yu, G.-Q.Q., Kreitzer, A., et al. (2007). Aberrant Excitatory Neuronal Activity and Compensatory Remodeling of Inhibitory Hippocampal Circuits in Mouse Models of Alzheimer's Disease. *Neuron*.
- Park, S.S., Park, H.S., Kim, C.J., Kang, H.S., Kim, D.H., Baek, S.S., and Kim, T.W. (2020). Physical exercise during exposure to 40-Hz light flicker improves cognitive functions in the 3xTg mouse model of Alzheimer's disease. *Alzheimer's Res. Ther.* 12, 1–15.
- Paterno, R., Casalia, M., and Baraban, S.C. (2020). Interneuron deficits in neurodevelopmental disorders: Implications for disease pathology and interneuron-based therapies. *Eur. J. Paediatr. Neurol.* 24, 81.
- Paterno, R., Marafija, J.R., Ramsay, H., Li, T., Salvati, K.A., and Baraban, S.C. (2021). Hippocampal gamma and sharp-wave ripple oscillations are altered in a *Cntnap2* mouse model of autism spectrum disorder. *Cell Rep.* 37.
- Pelkey, K.A., Chittajallu, R., Craig, M.T., Tricoire, L., Wester, J.C., and McBain, C.J. (2017). Hippocampal gabaergic inhibitory interneurons. *Physiol. Rev.* 97, 1619–1747.
- Petersen, P.C., Siegle, J.H., Steinmetz, N.A., Mahallati, S., and Buzsáki, G. (2021). CellExplorer: A

- framework for visualizing and characterizing single neurons. *Neuron* 109, 3594-3608.e2.
- Pfeiffer, B.E., and Foster, D.J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* 497, 74–79.
- Poo, C., and Isaacson, J.S. (2009). Odor Representations in Olfactory Cortex: “Sparse” Coding, Global Inhibition, and Oscillations. *Neuron* 62, 850–861.
- Prince, S.M., Paulson, A.L., Jeong, N., Zhang, L., Amigues, S., and Singer, A.C. (2021). Alzheimer’s pathology causes impaired inhibitory connections and reactivation of spatial codes during spatial navigation. *Cell Rep.* 35, 109008.
- Qin, H., Fu, L., Hu, B., Liao, X., Lu, J., He, W., Liang, S., Zhang, K., Li, R., Yao, J., et al. (2018). A Visual-Cue-Dependent Memory Circuit for Place Navigation. *Neuron*.
- Quon, R.J., Leslie, G.A., Camp, E.J., Meisenhelter, S., Steimel, S.A., Song, Y., Ettinger, A.B., Bujarski, K.A., Casey, M.A., and Jobst, B.C. (2021). 40-Hz auditory stimulation for intracranial interictal activity: A pilot study. *Acta Neurol. Scand.* 144, 192–201.
- Rao-Ruiz, P., Yu, J., Kushner, S.A., and Josselyn, S.A. (2019). Neuronal competition: microcircuit mechanisms define the sparsity of the engram. *Curr. Opin. Neurobiol.* 54, 163–170.
- Reid, H.M.O., Chen-Mack, N., Snowden, T., and Christie, B.R. (2021). Understanding Changes in Hippocampal Interneurons Subtypes in the Pathogenesis of Alzheimer’s Disease: A Systematic Review. *Brain Connect.* 11, 159–179.
- Ren, C., Peng, K., Yang, R., Liu, W., Liu, C., and Komiyama, T. (2022). Global and subtype-specific modulation of cortical inhibitory neurons regulated by acetylcholine during motor learning. *Neuron* 110, 2334-2350.e8.
- Robinson, N.T.M., Descamps, L.A.L., Russell, L.E., Nutbrown, R., Schmidt-Hieber, C., Häusser

- Correspondence, M., Buchholz, M.O., Bicknell, B.A., Antonov, G.K., Lau, J.Y.N., et al. (2020). Targeted Activation of Hippocampal Place Cells Drives Memory-Guided Spatial Behavior. *Cell* 183, 1586-1599.e10.
- Rolls, E.T. (2018). The storage and recall of memories in the hippocampo-cortical system. *Cell Tissue Res.* 373, 577–604.
- Rolotti, S. V., Ahmed, M.S., Szoboszlay, M., Geiller, T., Negrean, A., Blockus, H., Gonzalez, K.C., Sparks, F.T., Solis Canales, A.S., Tuttman, A.L., et al. (2022). Local feedback inhibition tightly controls rapid formation of hippocampal place fields. *Neuron* 110, 783-794.e6.
- Roux, L., and Buzsaki, G. (2015). Tasks for inhibitory interneurons in intact brain circuits. *Neuropharmacology* 88, 10–23.
- Royer, S., Zemelman, B. V, Losonczy, A., Kim, J., Chance, F., Magee, J.C., and Buzsaki, G. (2012). Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* 15, 769–775.
- Ruden, J.B., Dugan, L.L., and Konradi, C. (2020). Parvalbumin interneuron vulnerability and brain disorders. *Neuropsychopharmacol.* 2020 462 46, 279–287.
- de Salas-Quiroga, A., García-Rincón, D., Gómez-Domínguez, D., Valero, M., Simón-Sánchez, S., Paraíso-Luna, J., Agüeroles, J., Pujadas, M., Muguruza, C., Callado, L.F., et al. (2020). Long-term hippocampal interneuronopathy drives sex-dimorphic spatial memory impairment induced by prenatal THC exposure. *Neuropsychopharmacology* 45, 877–886.
- Schlingloff, D., Kali, S., Freund, T.F., Hajos, N., and Gulyas, A.I. (2014). Mechanisms of Sharp Wave Initiation and Ripple Generation. *J. Neurosci.*
- Schneider, F., Baldauf, K., Wetzell, W., and Reymann, K.G.G. (2014). Behavioral and EEG changes in male 5xFAD mice. *Physiol. Behav.* 135, 25–33.

- Shahriari, Y., Krusienski, D., Dadi, Y.S., Seo, M., Shin, H.-S., and Choi, J.H. (2016). Impaired auditory evoked potentials and oscillations in frontal and auditory cortex of a schizophrenia mouse model. *World J. Biol. Psychiatry* 1–10.
- Shi, J., Phensy, A.J., and Sohal, V.S. (2021). Rhythmic auditory stimulation rescues cognitive flexibility in mutant mice with impaired gamma synchrony. *BioRxiv* 2021.11.15.468681.
- Sigurdsson, T., and Duvarci, S. (2015). Hippocampal-Prefrontal Interactions in Cognition, Behavior and Psychiatric Disease. *Front. Syst. Neurosci.* 9.
- Sik, A., Penttonen, M., Ylinen, A., and Buzsáki, G. (1995). Hippocampal CA1 interneurons: an in vivo intracellular labeling study. *J. Neurosci.* 15, 6651.
- Singer, A.C., and Frank, L.M. (2009). Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. *Neuron* 64.
- Singer, A.C., Carr, M.F., Karlsson, M.P., and Frank, L.M. (2013). Hippocampal SWR Activity Predicts Correct Decisions during the Initial Learning of an Alternation Task. *Neuron* 77.
- Skaggs, W.E., McNaughton, B.L., Wilson, M.A., and Barnes, C.A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*.
- Solís-Vivanco, R., Rodríguez-Violante, M., Cervantes-Arriaga, A., Justo-Guillén, E., and Ricardo-Garcell, J. (2018). Brain oscillations reveal impaired novelty detection from early stages of Parkinson's disease. *NeuroImage Clin.*
- Spagnolo, P.A., Wang, H., Srivannitchapoom, P., Schwandt, M., Heilig, M., and Hallett, M. (2019). Lack of target engagement following low-frequency deep Transcranial Magnetic Stimulation of the anterior insula. *Neuromodulation* 22, 877.
- Stapells, D.R., Suffield, J.B., Picton, T.W., Linden, D., Suffield, J.B., Hamel, G., and Picton, T.W.

- (1984). Human Auditory Steady State Potentials. *Ear Hear.* 5.
- Stark, E., Koos, T., and Buzsáki, G. (2012). Diode probes for spatiotemporal optical control of multiple neurons in freely moving animals. *J. Neurophysiol.* 108, 349–363.
- Stark, E., Roux, L., Eichler, R., Senzai, Y., Royer, S., and Buzsáki, G. (2014). Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron*.
- Stefanelli, T., Bertollini, C., Lüscher, C., Muller, D., and Mendez, P. (2016). Hippocampal Somatostatin Interneurons Control the Size of Neuronal Memory Ensembles. *Neuron* 89, 1074–1085.
- Strüber, M., Sauer, J.-F., and Bartos, M. (2022). Parvalbumin expressing interneurons control spike-phase coupling of hippocampal cells to theta oscillations. *Sci. Rep.* 12, 1–10.
- Sun, Q., Sotayo, A., Cazzulino, A.S., Snyder, A.M., Denny, C.A., and Siegelbaum, S.A. (2017). Proximodistal Heterogeneity of Hippocampal CA3 Pyramidal Neuron Intrinsic Properties, Connectivity, and Reactivation during Memory Recall. *Neuron* 95, 656-672.e3.
- Szabo, A., Somogyi, J., Cauli, B., Lambolez, B., Somogyi, P., and Lamsa, K.P. (2012). Calcium-Permeable AMPA Receptors Provide a Common Mechanism for LTP in Glutamatergic Synapses of Distinct Hippocampal Interneuron Types. *J. Neurosci.* 32, 6511–6516.
- Tamura, M., Spellman, T.J., Rosen, A.M., Gogos, J.A., and Gordon, J.A. (2017). Hippocampal-prefrontal theta-gamma coupling during performance of a spatial working memory task. *Nat. Commun.* 8, 1–9.
- Tanaka, K.Z., He, H., Tomar, A., Niisato, K., Huang, A.J.Y., and McHugh, T.J. (2018). The hippocampal engram maps experience but not place. *Science* (80-.). 361, 392–397.
- Tennant, S.A., Clark, H., Hawes, I., Tam, W.K., Hua, J., Yang, W., Gerlei, K.Z., Wood, E.R., and Nolan, M.F. (2022). Spatial representation by ramping activity of neurons in the retrohippocampal

cortex. *Curr. Biol.*

Thuné, H., Recasens, M., and Uhlhaas, P.J. (2016). The 40-Hz Auditory Steady-State Response in Patients With Schizophrenia: A Meta-analysis. *JAMA Psychiatry* 73, 1145–1153.

Tolman, E.C. (1948). Cognitive maps in rats and men. *Psychol. Rev.* 55.

Turi, G.F., Li, W.-K., Chavlis, S., Pandi, I., O'Hare, J., Priestley, J.B., Grosmark, A.D., Liao, Z., Ladow, M., Zhang, J.F., et al. (2019). Vasoactive Intestinal Polypeptide-Expressing Interneurons in the Hippocampus Support Goal-Oriented Spatial Learning. *Neuron* 101, 1150-1165.e8.

Tyan, L., Chamberland, S., Magnin, E., Camiré, O., Francavilla, R., Suzanne David, L., Deisseroth, K., and Topolnik, L. (2014). Dendritic Inhibition Provided by Interneuron-Specific Cells Controls the Firing Rate and Timing of the Hippocampal Feedback Inhibitory Circuitry. *J. Neurosci.* 34, 4534–4547.

Uchida, K., Taguchi, Y., Sato, C., Miyazaki, H., Kobayashi, K., Kobayashi, T., and Itoi, K. (2014). Amelioration of improper differentiation of somatostatin-positive interneurons by triiodothyronine in a growth-retarded hypothyroid mouse strain. *Neurosci. Lett.* 559, 111–116.

Udakis, M., Pedrosa, V., Chamberlain, S.E.L., Clopath, C., and Mellor, J.R. (2020). Interneuron-specific plasticity at parvalbumin and somatostatin inhibitory synapses onto CA1 pyramidal neurons shapes hippocampal output. *11*, 1–17.

Uhlhaas, P.J., and Singer, W. (2006). Review Neural Synchrony in Brain Disorders: Relevance for Cognitive Dysfunctions and Pathophysiology. *Neuron* 52, 155–168.

Umschweif, G., Medrihan, L., McCabe, K.A., Sagi, Y., and Greengard, P. (2021). Activation of the p11/SMARCA3/Neurensin-2 pathway in parvalbumin interneurons mediates the response to chronic antidepressants. *Mol. Psychiatry* 2021 267 26, 3350–3362.

- Varga, C., Oijala, M., Lish, J., Szabo, G.G., Bezaire, M., Marchionni, I., Golshani, P., and Soltesz, I. (2014). Functional fission of parvalbumin interneuron classes during fast network events. *Elife* 3, 1–23.
- Venturino, A., Schulz, R., De Jesús-Cortés, H., Maes, M.E., Nagy, B., Reilly-Andújar, F., Colombo, G., Cubero, R.J.A., Schoot Uiterkamp, F.E., Bear, M.F., et al. (2021). Microglia enable mature perineuronal nets disassembly upon anesthetic ketamine exposure or 60-Hz light entrainment in the healthy brain. *Cell Rep.* 36, 109313.
- Verret, L., Mann, E.O., Hang, G.B., Barth, A.M.I., Cobos, I., Ho, K., Devidze, N., Masliah, E., Kreitzer, A.C., Mody, I., et al. (2012). Inhibitory Interneuron Deficit Links Altered Network Activity and Cognitive Dysfunction in Alzheimer Model. *Cell* 149, 708–721.
- Wallace, D.J., and Kerr, J.N.D. (2010). Chasing the cell assembly. *Curr. Opin. Neurobiol.* 20, 296–305.
- Wang, C.H., Monaco, J.D., and Knierim, J.J. (2020a). Hippocampal Place Cells Encode Local Surface-Texture Boundaries. *Curr. Biol.* 30, 1397-1409.e7.
- Wang, D., Clouter, A., Chen, Q., Shapiro, K.L., and Hanslmayr, S. (2018). Single-Trial Phase Entrainment of Theta Oscillations in Sensory Regions Predicts Human Associative Memory Performance. *J. Neurosci.* 38, 6299–6309.
- Wang, F., Schoenbaum, G., and Kahnt, T. (2020b). Interactions between human orbitofrontal cortex and hippocampus support model-based inference. *PLOS Biol.* 18, e3000578.
- Wang, J., Zuo, X., Dai, Z., Xia, M., Zhao, Z., Zhao, X., Jia, J., Han, Y., and He, Y. (2013). Disrupted Functional Brain Connectome in Individuals at Risk for Alzheimer’s Disease. *Biol. Psychiatry* 73, 472–481.
- Wehr, M., and Zador, A.M. (2003). Balanced inhibition underlies tuning and sharpens spike timing in auditory cortex. *Nat.* 2003 4266965 426, 442–446.

- Wiener, S.I., Paul, C.A., and Eichenbaum, H. (1989). Spatial and behavioral correlates of hippocampal neuronal activity. *J. Neurosci.* 9, 2737–2763.
- Wilent, W.B., and Nitz, D.A. (2007). Discrete Place Fields of Hippocampal Formation Interneurons. *J. Neurophysiol.* 97, 4152–4161.
- Wilson, M.A., McNaughton, B.L., Filipkowski, R.K., Silva, A.J., and Eichenbaum, H. (1993). Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055–1058.
- Wohleb, E.S., Franklin, T., Iwata, M., and Duman, R.S. (2016). Integrating neuroimmune systems in the neurobiology of depression.
- Worrell, G.A., Gardner, A.B., Stead, S.M., Hu, S., Goerss, S., Cascino, G.J., Meyer, F.B., Marsh, R., and Litt, B. (2008). High-frequency oscillations in human temporal lobe: Simultaneous microwire and clinical macroelectrode recordings. *Brain*.
- Wu, F., Stark, E., Ku, P.C., Wise, K.D., Buzsáki, G., and Yoon, E. (2015). Monolithically Integrated μ LEDs on Silicon Neural Probes for High-Resolution Optogenetic Studies in Behaving Animals. *Neuron* 88, 1136–1148.
- Wu, G.K., Arbuckle, R., Liu, B. hua, Tao, H.W., and Zhang, L.I. (2008). Lateral Sharpening of Cortical Frequency Tuning by Approximately Balanced Inhibition. *Neuron* 58, 132–143.
- Xu, H., BaracsKay, P., O'Neill, J., and Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze. *Neuron* 101, 119-132.e4.
- Yao, Y., Ying, Y., Deng, Q., Zhang, W., Zhu, H., Lin, Z., Zhang, S., Ma, J., and Zhao, Y. (2020). Non-invasive 40-Hz Light Flicker Ameliorates Alzheimer's-Associated Rhythm Disorder via Regulating Central Circadian Clock in Mice. *Front. Physiol.* 11, 294.

Ylinen, A., Bragin, A., Nadasdy, Z., Jando, G., Szabo, I., Sik, A., Buzsáki, G., Buzsaki, G., and Buzsáki, G. (1995). Sharp Wave-Associated High-Frequency Oscillation (200 Hz) in the Intact Hippocampus: Network and Intracellular Mechanisms. *J. Neurosci.* *15*.

Zaremba, J.D., Diamantopoulou, A., Danielson, N.B., Grosmark, A.D., Kaifosh, P.W., Bowler, J.C., Liao, Z., Sparks, F.T., Gogos, J.A., and Losonczy, A. (2017). Impaired hippocampal place cell dynamics in a mouse model of the 22q11.2 deletion. *Nat. Neurosci.* *20*, 1612–1623.

Zheng, L., Yu, M., Lin, R., Wang, Y., Zhuo, Z., Cheng, N., Wang, M., Tang, Y., Wang, L., and Hou, S.-T. (2020). Rhythmic light flicker rescues hippocampal low gamma and protects ischemic neurons by enhancing presynaptic plasticity. *Nat. Commun.* 2020 111 *11*, 1–16.