Title

Non-spatial hippocampal behavioral timescale synaptic plasticity during working memory is gated by entorhinal inputs

Author

Dorian, Conor C

Publication Date

2025

UNIVERSITY OF CALIFORNIA
Los Angeles

Non-spatial hippocampal behavioral timescale synaptic plasticity during working memory is gated by entorhinal inputs

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Neuroscience

by

Conor Cunningham Dorian

2025

ABSTRACT OF THE DISSERTATION

Non-spatial hippocampal behavioral timescale synaptic plasticity during working memory is gated by entorhinal inputs

by

Conor Cunningham Dorian
Doctor of Philosophy in Neuroscience
University of California, Los Angeles, 2025
Professor Peyman Golshani, Chair

Working memory is the ability to temporarily store and manipulate information. Many brain regions have been implicated in sensory representations necessary for working memory, but the hippocampus's sequential activity and well-known role in many types of memory provides a unique model to study the formation of neural ensemble representations. Most hippocampal research has focused on its representations of visuospatial sensory information, which leaves many unanswered questions about how representations of non-spatial and internally generated temporal representations may differ. Therefore, this dissertation focuses on non-spatial sensory and temporal encoding within hippocampal CA1 region, and the hippocampus's main input areas - lateral and medial entorhinal cortex (LEC and MEC) - during an olfactory working memory task.

Chapter 1 will outline key evidence for the roles of entorhinal cortex and hippocampus for working memory and sequential activity that encodes sensory and temporal information. Chapter 2 will describe our recent finding that behavioral timescale synaptic plasticity (BTSP) is a rapid form of plasticity that generalizes to non-spatial tasks with the rapid formation of odor-selective responses in individual CA1 neurons during our olfactory working memory task. Chapter 3 will describe our findings showing how sequential activity in CA1 simultaneously encodes sensory working memory and temporal information in a novel rodent working memory implicit timing task. Chapter 4 will describe our findings that LEC or MEC optogenetic inhibition does not impair learning or performance on our olfactory working memory task even when we make the task dramatically more difficult. Chapter 5 will describe our findings that the inhibitory opsin stGtACR2 can lead to rapid kindling via light pulses that likely are excitatory on axon terminals in the hippocampus when opsin expression was expected to be limited to somas in EC. Altogether, this dissertation contributes to the field's understanding of how LEC and MEC drive non-spatial and temporal representations in the hippocampus during working memory.

The dissertation of Conor Cunningham Dorian is approved.

Hugh T. Blair

Dean Buonomano

Anne K. Churchland

Peyman Golshani, Committee Chair

University of California, Los Angeles

2025

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

First, I would like to acknowledge and thank my advisor and committee chair, Peyman Golshani. His steady positivity and curiosity led me to follow the science wherever it took us. This guidance led to a diverse set of experiences: fruitful funding opportunities, collaborations, and the eventual central finding connecting my non-spatial research to the spatial field. His vast knowledge and personal connections through various subfields contributed to my well-rounded growth and learning in the interdisciplinary field of neuroscience. Thank you Peyman for all the support, guidance, and reminders to always keep pushing.

Second, I would like to thank my committee members: Dean Buonomano, Anne Churchland, and Tad Blair. I give an extra thank you to Dean Buonomano for the mentorship and critical eye beyond just our collaboration. The statistical knowledge that I gained during graduate school is almost all thanks to him.

Next, I would like to thank the entire Golshani Lab family. Particular thanks are necessary for many colleagues: Jiannis Taxidis, for being my secondary mentor throughout my entire PhD and providing valuable guidance all along the way after getting me on my feet in the lab. Blake Madruga, for letting me try to absorb a small portion of your optics and engineering knowledge, and always being willing to help and listen to my technical or scientific challenges. Ahmet Arac, for the fruitful collaborations and scientific career discussions. Douglas Vormstein-Schneider, for all the training room and hallway discussions. Samara Hecht, Zoë Day, Hesper Chen, and Celina Yang, for helping with all aspects of the lab and being the glue that holds it together. Liron Sheintuch, Arash Bellafard, and Diego Espino, for helpful lab meeting questions that led to many more in-depth discussions in the office. Pingping Zhao, Long Yang, Ahmet Arac, Tamiris Prizon, Chrysa Iordanidou, Stephanie Cheung, Daniel Okobi, Mina Shahi, Ronen Reshef, and everyone else in the lab, for the interactions that fostered an environment for my scientific growth.

Outside of the Golshani lab, I must thank everyone on the 7th floor, for the random hallway interactions that fostered a lovely scientific environment. I also thank Dr. Carolyn Houser, Dr. Dean Buonomano, and Benjamin Liu for their help on the collaboration in chapter 5. I thank UC Irvine, Dr. Tallie Baram, Dr. Robert Hunt, and the epilepsy T32 program.

I acknowledge the NSIDP graduate program, the UCLA Brain Research Institute, and the funding sources that made this work possible: T32DGE1829071, 5T32NS045540-20, R01NS116589, and 1P50HD103557-01.

Finally, I would like to acknowledge the mice who sacrificed their lives for this work. Since I spent the vast majority of my time working alone watching them learn, they were a continuous inspiration to try to unlock the secrets hidden in their brains.

VITA

2014–2018	B.S. in Neuroscience, Rhodes College
2018–2025	Ph.D. Candidate, Neuroscience, University of California Los Angeles

PUBLICATIONS

Dorian, C.C., Taxidis, J., Arac, A., & Golshani, P. Non-spatial hippocampal behavioral time-scale synaptic plasticity during working memory is gated by entorhinal inputs. *Manuscript in Preparation*, **2025**

Dorian, C.C., Taxidis, J., Buonomano, D.V., & Golshani, P. Sequential CA1 activity represents working memory and time during a novel differntial-delay nonmatch-to-sample (dDNMS) task. *Manuscript in Preparation*, **2025**

Madruga, B.A., **Dorian, C.C**., Sehgal, M., Silva, A.J., Shtrahman, M., Aharoni, D., & Golshani, P. Open-source, high performance miniature multiphoton microscopy systems for freely behaving animals. *In Review, Nature Communications*, **2025**

Taxidis, J., Madruga, B.A., Safaryan, K., **Dorian, C.C**., Melin, M.D., Day, Z., Lin, M.Z., & Golshani, P. Voltage imaging reveals hippocampal inhibitory dynamics during memory-encoding pyramidal sequences. *Accepted for Publication, Nature Neuroscience*, **2025**

Evans, S.W., Shi, D.Q., Chavarha, M., Plitt, M.H., Taxidis, J., Madruga, B., Fan, J.L., Hwang, F.J., van Keulen, S.C., Suomivuori, C.M., Pang, M.M., Su, S., Lee, S., Hao, Y.A., Zhang, G., Jiang, D., Pradhan, L., Roth, R.H., Liu, Y., **Dorian, C.C.**, Reese, A.L., Negrean, A., Losonczy, A., Makinson, C.D., Wang, S., Clandinin, T.R., Dror, R.O., Ding, J.B., Ji, N., Golshani, P., Giocomo, L.M., Bi, G.Q., & Lin, M.Z. A positively tuned voltage indicator for extended electrical recordings in the brain. *Nature Methods*, **2023**

Taxidis, J., Pnevmatikakis, E.A., **Dorian, C.C.**, Mylavarapu, A.L., Arora, J.S., Samadian, K.D., Hoffberg, E.A., & Golshani, P. Differential emergence and stability of sensory and temporal representations in context-specific hippocampal sequences. *Neuron*, **2020**

Honig, M.G., **Dorian, C.C.**, Worthen, J.D., Micetich, A.C., Mulder, I.A., Sanchez, K.B., Pierce, W.F., Del Mar, N.A., & Reiner, A. Progressive long-term spatial memory loss following repeat concussive and subconcussive brain injury in mice, associated with dorsal hippocampal neuron loss, microglial phenotype shift, and vascular abnormalities. *European Journal of Neuroscience*, **2020**

CHAPTER 1

Introduction

1.1 Memory

The brain is a complex organ, but its primary function is as an input/output system. It must take in sensory inputs from the external world and guide motor outputs. However, the timescale between an input and an output can vary from less than a second to many years. Memory is the brain's ability to bridge the gap between input and output by storing information.

1.1.1 Patient H.M and Different Memory Systems

However, there are many different types of memory depending on the type of information being stored and the length of time it must be stored. This was first most clearly demonstrated by Dr. Brenda Milner and her patient H.M. (Henry Moliason), one of the most famous neuroscience patients in history (Squire, 2009). Following a bike accident at age 7, patient H.M. began having seizures in his childhood that worsened to the point of being resistant to anticonvulsant medications. Patient H.M. was unable to lead a normal life, so at the age of 27 in 1953 he chose to undergo an experimental procedure by Dr. William Scoville to bilaterally resect large portions of the medial temporal lobe believed to be driving the seizures. Although the surgery was successful in that patient H.M. never suffered from seizures again, Dr. Brenda Milner noted an extreme memory impairment (Milner et al., 1968). His intellect and cognition were largely unaltered, but he was unable to form new long-term memories - anterograde amnesia. He could remember details of his childhood and events that preceded approximately three years before the surgery, yet no new facts or events could be remembered. He no longer had the ability to encode semantic or episodic memories in the long-term. He was able to hold conversations and appear entirely normal for several minutes, but as soon his attention was redirected, he would forget the entire previous interaction.

It was Dr. Milner's careful studies and behavioral tests that first gave strong evidence to the different memory systems based on the type of information being stored (Squire, 2009). For example, patient H.M could retain strings of numbers for 15 minutes but nonverbal stimuli like faces and designs for less than a minute. Additionally, he could not learn new facts, but he could learn new visuomotor skills. Despite claiming that he had never seen the task before, patient H.M. could steadily improve at a difficult drawing over days and weeks of practice. These findings were the first to suggest that some types of memory are not reliant on the medial temporal lobe.

The other critical distinction between memory systems that Dr. Milner made was between 'immediate memory' and long-term memory (Squire, 2009). Patient H.M.'s memory impairment was characterized as an inability to convert immediate memory into long-term memory. It was proposed that the medial temporal lobe is critical for this conversion. Given patient H.M.'s intact immediate memory, it was believed to rely on brain regions outside of the medial temporal lobe. Also given his ability to remember events from before the surgery, it was believed that long-term memories are stored outside the area too. These findings established a foundation that continues to guide memory research today.

1.1.2 Working Memory

Some of the language has evolved since Dr. Milner's first descriptions, and additional types of memory have been classified. One such important type of short-term memory is called working memory, which is the ability to temporarily store and manipulate information

(Baddeley, 2012). Patient H.M.'s seemingly spared working memory allowed him to hold a short conversation, store numbers for a short period, and think critically. Because the ability to transiently store and use information is the basis of most cognitive functions, working memory is critical for our intellect.

Although there are discrepancies on the exact timeframe of working memory, it is generally considered on the timescale of several seconds (Baddeley, 2012). Human working memory abilities vary based on the individual's age, attentional state, and the type of information. Visuospatial working memory is thought to be one of the longer lasting types of up to 30 seconds (Posner and Konick, 1966), while auditory working memory is thought to only last about two seconds (Baddeley et al., 1975). While there was been much work characterizing the limits of human working memory and what factors contribute to its difference from other types of memory, there is still a large gap of knowledge in the neuroscience field as to how working memory is maintained within the brain. The first question is what brain regions are involved?

1.2 The Hippocampal Formation

The human brain is a 3-pound organ with 100 billion neurons and 1 trillion other supporting cells (Herculano-Houzel, 2009). Individual neurons communicate to each other through electrical and chemical signals in the axon, and they receive inputs from other neurons in their dendrites. These connections are made throughout and across the entire brain, but scientists use terminology to divide the brain into different regions and subregions based on physical architecture and function. When patient H.M. underwent his surgery, the medial temporal lobe was understood to likely be important for causing his seizures, but the precise functional divisions have since been further characterized. Starting in 1992, he began having magnetic resonance imaging (MRI) tests to better understand the damage that led to his memory deficits (Squire, 2009). These tests and post-mortem analysis confirmed that the original surgery lesioned the medial temporal lobe and specifically the hippocampus, amygdala, entorhinal cortex (EC), perirhinal cortex, and parahippocampal cortex (Squire and Zola-Morgan, 1991). While many of the original hypotheses focused on the idea of the hippocampus's role for memory, the discussion became focused on the interaction of these many regions.

1.2.1 Hippocampal Structure

Firstly, at the architecture level, the hippocampus is named because its shape resembles a seahorse, which is genus *Hippocampus* (Knierim, 2015). The hippocampus is further divided into four subregions: dentate gyrus (DG), cornu ammonis 1 (CA1), CA2, and CA3. While each of these four subregions have inputs from outside the hippocampus, the DG is the one with the strongest inputs outside of the hippocampus. The DG sends projections to CA3 via the mossy fibers (Knierim, 2015). CA3 projects to CA1 and also many collaterals that connect to other CA3 neurons, which forms a strong recurrent network (Li et al., 1994; Le Duigou et al., 2014). CA1 is the output region of the hippocampus as it projects to the entorhinal cortex (EC) and many other cortical and subcortical regions across the brain (Knierim, 2015).

1.2.2 Entorhinal Cortex Structure

The entorhinal cortex (EC) is the primary output of CA1 and the primary input to DG (Canto et al., 2008). Like most cortical areas in the human brain, it consists of 6 layers with distinct anatomical structure and function. Layer I is the most superficial, while layer VI is the deepest and borders white matter tracts. Most neurons in layer I are inhibitory and project to layer II and III to inhibit the excitatory cells there. Layers II and III have the densest organization of excitatory neurons and these project to other layers of EC and outside of EC. Layer II sends projections to DG and CA3, while layer III sends projections directly to CA1. The pathway from EC to DG forms the perforant path, and CA1's projections back to EC complete the canonical trisynaptic loop. Meanwhile the direct projection from EC layer III to CA1 is called the temporammonic pathway. Layer IV is very thin and only contains small amounts of inhibitory neurons. Layer V is another dense layer of excitatory pyramidal neurons, and these project to superficial layers and outside of the EC. Finally, layer VI contains excitatory neurons with projections to layer V and VI and the subiculum.

This oversimplification of the hippocampal and EC circuit does not address the rich diversity of many different types of excitatory and inhibitory neurons at each region, subregion, and layer (Varga et al., 2010; Chamberland and Topolnik, 2012). While work is still being conducted to classify all the different subtypes of neurons based on their structure and morphology, most modern work is focused on understanding the functional properties that define each area and cell type. Both the hippocampus and EC still hold many mysteries to their function in various contexts.

1.2.3 Hippocampal and EC Function

In order to follow up on Dr. Brenda Milner's conclusions about hippocampus from patient H.M., the neuroscience community had to causally test its necessity for types of memory. Fortunately, the structure and function of the hippocampus and EC are largely evolutionarily conserved across mammalian species, which has allowed scientists to use rodent models (Zemla and Basu, 2017). Early lesion studies supported Dr. Brenda Milner's conclusions that the hippocampus is necessary for encoding of long-term memories, but this storage happens in cortical areas outside of the hippocampus (McKenzie and Eichenbaum, 2011). However, the story is more complicated with consideration of different contexts and information types. Overall, the hippocampus has been shown to be important for learning and memory of sensory, spatial, and temporal information over the timescale of seconds to years (Zemla and Basu, 2017). The dynamics of neural activity and contributions of EC heavily depend on the parameters of experience, modality, and timescale.

1.2.3.1 Spatial Memory

One of the most well-known hippocampal and EC findings was the 1971 discovery of ‘place cells’ in the hippocampus (O’Keefe and Dostrovsky, 1971) and the 2005 discovery of ‘grid cells’ in medial EC (Hafting et al., 2005). In 2014, the Nobel Prize in Physiology or Medicine was awarded to John O’Keefe, May-Britt Moser, and Edvard I. Moser for these discoveries that helped us understand the role of the hippocampus and EC in spatial memory and navigation. In 1971, John O’Keefe and Jonathan Dostrovsky published their finding that individual neurons in the CA1 region of hippocampus would fire at specific positions within an environment as rats were recorded *in-vivo* (O’Keefe and Dostrovsky, 1971). These cells were called ‘place cells’ as they each had different ‘place fields’ that filled the environment; together it is possible to determine the location of the animal based on the activity of these groups of place cells. In 2005, research in the labs of May-Britt Moser and Edvard I. Moser was published describing ‘grid cells’ in medial EC, which are individual neurons that exhibit a grid-like pattern of many place fields instead of the singular place field that O’Keefe’s place cells exhibited (Hafting et al., 2005). For the last 50 years, much research has been dedicated to understanding the interaction between the hippocampus and EC in spatial memory and navigation. The discoveries of ‘head-direction cells’ (Taube et al., 1990), ‘border cells’ (Lever et al., 2009), ‘speed cells’ (Kropff et al., 2015), and others have clarified how individual neurons encode spatial information of the environment. Some of these representations are necessary for behavior as there has been extensive work evaluating the necessity of the hippocampus or EC for performance of various spatial memory tasks (Eichenbaum, 2017).

Spatial tasks and contexts have proven convenient models for many other significant discoveries regarding the hippocampus and learning. For example, at the molecular and cellular level, discoveries of the roles of N-methyl-D-aspartate receptors (NMDAR) and calcium-calmodulin-dependent kinase II (CaMKII) in synaptic plasticity and learning were first demonstrated in spatial contexts (Morris et al., 1986; Silva et al., 1992; Nakazawa et al., 2004). However, not all memories are spatial. While focusing on one modality has led to a wealth of knowledge for the field, it leaves questions for how these findings can be generalized to other sensory modalities.

1.2.3.2 Non-spatial Memory

Early on the hippocampus's role in memory was understood to be very multimodal, but the focus on medial EC for its 'grid cells' drew attention to the functional division between medial and lateral entorhinal cortices (MEC and LEC). The most common understanding of the functional division classifies MEC as 'where' and LEC as 'what'; MEC sends spatial information to the hippocampus, while LEC sends non-spatial information. However, this distinction is oversimplified and nuanced understandings are more appropriate (Knierim et al., 2014; Save and Sargolini, 2017).

Within spatial tasks and environments, the simplified distinction seems appropriate. For example, in rodent foraging tasks, MEC neurons will fire in grid like patterns while LEC neurons will be contextually sensitive to wall colors (Lu et al., 2013) or the amount of time that has elapsed (Tsao et al., 2018). However, in less space-centric tasks or ones with less cognitive demand make the distinction less clear (Save and Sargolini, 2017). For example, in object recognition tasks, LEC only becomes necessary when the cognitive demand is high (Kuruvilla and Ainge, 2017; Ku et al., 2017). When cognitive demands are low, both LEC and MEC encode similar features and may not be necessary for task performance.

Taken together, this suggests that is most helpful to consider the distinction between MEC and LEC in its relationship to the hippocampus. It is often difficult to disentangle their relative contributions because of the substantial projections between the two regions and back from the CA1 region of the hippocampus. Also given that functional differences vary between tasks and environments, newer proposals frame LEC and MEC as 'content' and 'context' as opposed to 'what' and 'where' (Knierim et al., 2014). Together, MEC and LEC provide the content and contextual information for the hippocampus to form a holistic representation of the external world. However, the content and context will vary highly depending on what representation must be maintained. Chapter 4 will investigate this question by evaluating working memory performance when LEC or MEC activity is silenced during a olfactory working memory task.

1.3 Population Dynamics of Working Memory

With the advancement of technology, the field of neuroscience continues to evolve. Many early studies relied on lesions to test causality of a brain region's role in a particular task. Early electrophysiological recordings could only record from a small population of cells, so more analyses focused on understanding how individual cells represent information of the external world. However, these approaches make it difficult to understand the interaction between multiple regions. Additionally, what are the activity patterns at a large population level that lead to behavior, and how does this change with experience?

1.3.1 Persistent Activity

In addition to the hippocampus and EC, early studies also focused on the prefrontal cortex as an area important for working memory. Single cell electrophysiological recordings in primate prefrontal cortex were the first to demonstrate neural activity capable of storing information in the form of individual cell activity patterns (Fuster and Alexander, 1971; Kubota and Niki, 1971). Compared to low firing rates in pretrial periods, individual cells would maintain persistently high firing rates during the delay period. Later research would also show that cells could have cue-specific firing; a cell would fire persistently during the delay only following its preferred cue (Funahashi et al., 1989). This effectively stores the identity of the first cue so that it can then be compared to another stimuli arriving after the delay. Therefore, persistent activity became the first demonstrated 'neural code' for how the brain can maintain information across delays of several seconds. While much research has continued to expand on our knowledge of persistent activity in the prefrontal cortex and connect it to rodent and human research (Curtis and D'Esposito, 2003; Curtis and Sprague, 2021), other brain regions do not use this same neural code.

1.3.2 Sequential Activity

The discovery of 'time cells' in hippocampal CA1 (Pastalkova et al., 2008; MacDonald et al., 2011) provides an alternative neural code for working memory by also temporally organizing the memories. In tasks that require information to be held across a delay period, CA1 neurons have been shown to have peaks of activity at specific times in the delay. Unlike prefrontal cortical cells that persistently fire throughout the entire delay, individual CA1 'time cells' each have a 'time field' - which is the timepoint at which they prefer to be the most active. Together, many time cells can tile the entire delay period with each one having a unique time field. This population dynamic is referred to as sequential activity. Additionally, these sequences were demonstrated to be cue-specific, such that distinct populations of cells would respond at specific timepoints in the delay only following their preferred cue (MacDonald et al., 2013; Taxidis et al., 2020). Therefore, sequential activity is an additional neural code for working memory. Other models have also been shown in other areas such as ramping activity, which is when cells steadily increase of decrease their firing across a delay period (Chafee and Goldman-Rakic, 1998; Barak et al., 2010). Alternative 'activity silent' models have also been proposed that rely on temporary synaptic plasticity to maintain information in the absence of neural activity (Stokes, 2015; Masse et al., 2019). While both persistent and sequential activity exist in areas of the brain, it is valuable to understand the environmental factors and circuit connections that lead to preferring one or another. Persistent activity is likely robust, while sparser activity patterns like ramping or sequential activity may be more energetically efficient.

1.3.3 Multiplexing with Time

Critically, another feature of ramping and sequential activity is that these patterns not only encode the previous stimuli, but also how much time has elapsed. In sequences, since each neuron has a different preferred time field, it is possible to determine the timepoint within the delay by evaluating which cells are currently active. With persistent activity, this temporal information is not available because the population dynamics are stable throughout the delay period. This has led scientists to ask if sequential activity can be a shared neural mechanism for the encoding of working memory and timing information (Zhou et al., 2023). It is hypothesized that persistent activity is more present when it is not valuable to track time; for example, when delays lengths are random it is not possible to predict when the delay may end. On the other hand, sequential activity may be more present when it is valuable to track time. Even during a working memory task with no explicit timing component, humans and animals show evidence of implicitly learning the temporal structure of the task (Nobre and Van Ede, 2018; Zhou et al., 2023). Temporal predictions help to build anticipation at the appropriate time so that the subject is maximally prepared to respond behaviorally. However, this theory that sequential activity multiplexes working memory and time has yet to be clearly demonstrated in neural data. Chapter 3 will address this by examining our findings of CA1 sequential activity in a novel rodent working memory task with an implicit timing component.

1.3.4 Sequential Activity in CA1 during olfactory working memory

The first step to understanding the role of sequential activity is to observe how sequences vary between different contexts. Given the bias of hippocampal research to focus on spatial information, most ‘time cells’ have been demonstrated in spatial tasks (Pastalkova et al., 2008; MacDonald et al., 2011; Eichenbaum, 2014; Buzsáki and Tingley, 2018) Often enclosed spaces with linear treadmills are used to disentangle space and movement from elapsed time. However, relatively less research has evaluated temporal encoding during completely non-spatial tasks. Previous work within the Golshani lab addressed this gap of knowledge (Taxidis et al., 2020). Additionally, taking advantage of technological advancements, they recorded larger numbers of CA1 neurons by performing two-photon calcium imaging. This provided unique advantages over the electrophysiological tetrode recordings most of the previous spatial research relied on.

1.3.4.1 Calcium Imaging

With a sacrifice of temporal resolution of neural activity, calcium imaging has the main advantages of ease of recording larger number of neurons in a cell-type specific manner and the ability to record the exact same neurons for weeks or months (Grienberger and Konnerth, 2012; Grienberger et al., 2022). When a neuron fires an action potential, voltage-gated calcium ion channels open to allow the influx of calcium ions. Calcium binds calmodulin as the start of mechanism to trigger synaptic vesicle release of neurotransmitters into the synapse. Protein engineers have designed modified calcium sensors with calmodulin components that fluoresce a specific color when calcium binds. With viral delivery of these genetically engineered calcium indicators (GECI), researchers can visualize the fluorescence intensity as a proxy for neural activity. This technique has revolutionized systems neuroscience as we are able to record calcium signals from thousands of cells simultaneously and track their activity as animals learn across days, weeks, or months. Additionally, calcium imaging of dendrites or axons allows us to visualize the activity being sent from one brain region to another.

In behavioral tasks with delay periods of several seconds, the lack of temporal resolution to identify single action potentials is outweighed by the tremendous benefits of using calcium imaging. This tool is ideal for asking questions about how neural dynamics of specific neuron populations change across weeks of learning a working memory task.

1.3.4.2 Major Previous Findings

In 2020, the Golshani lab published research showing how ‘odor cells’ and ‘time cells’ comprise odor-specific sequences in dorsal CA1 during an olfactory delayed non-match-to-sample (DNMS) task (Taxidis et al., 2020). Each trial of the task consisted of two 1-second odor presentations separated by a 5-second delay period. Either odor presentation could be odor A or odor B. Mice were trained to refrain from licking after match trials and to lick for a water reward after non-match trials. Using two-photon calcium imaging, the activity of thousands of dorsal CA1 neurons was recorded during learning of the task. Together the odor cells and time cells created sequences that tiled the entire 5-second delay, and there was a distinct sequence following both odor A and odor B. One main finding was that odor cells were more stable and robust as compared to time cells.

Additionally, the number of time cells increased with experience and better performance, while the number of odor cells remained the same (Taxidis et al., 2020). This suggests that sequential activity is dependent on familiarity of the context. Perhaps the ability to keep track of time during the delay period becomes better as the sequence fills in additional time cells. While chapter 3 will address how an increase in the number of time cells may help to better encode temporal information at specific points in the delay period, another major question is how the representations are formed. What brain regions are involved in the formation of odor cells and time cells in this task? Given that LEC is known for its role in olfaction encoding (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024), is LEC necessary for the formation of these odor-specific sequence cells? Does MEC play a complementary role?

1.4 Representations Change with Learning and Experience

These large questions about how neural representations change across learning and experience extend beyond memory tasks, as they are major questions for the entire field of systems neuroscience. Therefore, it is important to consider how neurons change their activity in response to their inputs. This ability to modify strengths of connections is termed plasticity, and there are several different types and subtypes.

1.4.1 Hebbian Plasticity

The most famous and widely accepted of plasticity is Hebbian plasticity, named after Donald Hebb published his book "The Organization of Behavior" in 1949 (Hebb, 2005). Often summarized by the phrase "neurons that fire together, wire together", he postulated a theory that requires causality and repetition in order to increase synaptic strength between two connected neurons. Presynaptic cells must causally drive the postsynaptic cell to fire, and if this is done repeatedly, the strength of the connection will grow stronger. This was first demonstrated in the hippocampus (Bliss and Lømo, 1973; Levy and Steward, 1983) and termed long-term potentiation (LTP). Specifically, it was demonstrated that the presynaptic cell must fire just before the postsynaptic cell, and this must be repeat (Bliss and Lømo, 1973; Levy and Steward, 1983). This idea of temporal precedence became more clearly defined with the classification of spike-timing-dependent plasticity (STDP) (Markram et al., 1997). Here, the amplitude of the strengthening is dependent on how close in time the presynaptic spike and postsynaptic spike occurred. If the presynaptic spike occurs after the postsynaptic spike, the connection is weakened as the presynaptic cell could not have caused the postsynaptic spike. This is termed long-term depression (LTD). Together LTP and LTD provide a mechanism for Hebb's original theory. This strengthening of efficient connections and weakening of inefficient connections is thought to be the basis of learning and memory.

Lots of work has demonstrated how the strengthening and weakening of synapses occurs at the molecular level with AMPA and NMDA receptors (Watt et al., 2000). Yet, making the connection to behavior is difficult often because of the timescale required for STDP and Hebbian forms of plasticity. If many repetitions are required for synapse strengthening, how is it that the brain can learn after single experiences? If the postsynaptic spike must happen only milliseconds after the presynaptic spike, how is it that we can learn associations between events that are separated by seconds?

1.4.2 Behavioral Timescale Synaptic Plasticity (BTSP)

The discovery of a new form of plasticity by the Magee lab provides a unique answer to how the brain can learn on behavioral timescales. Behavioral timescale synaptic plasticity (BTSP) does not require causality or repetition. It has been demonstrated as a robust mechanism for the formation of 'place cells' in the hippocampus (Bittner et al., 2015, 2017; Grienberger et al., 2017; Zhao et al., 2020; Magee and Grienberger, 2020; Milstein et al., 2021; Grienberger and Magee, 2022; Xiao et al., 2023). During spatial learning tasks, a single calcium plateau potential can serve as an induction event. When a CA1 neuron has a plateau potential, it increases the synaptic strengths of connections that were active in the seconds around the plateau. Inputs that occur up to 4 seconds before the plateau potential are potentiated. Although it is asymmetric as only inputs that came within one second after the plateau are potentiated. In linear track tasks, driving a single plateau potential on one lap is enough to induce strong place cell firing at the place field on the subsequent laps.

Work has shown that MEC inputs are critical for driving these plateaus in CA1 (Grienberger and Magee, 2022), yet it is unknown if this type of plasticity generalizes to other contexts. Is BTSP a mechanism for the formation of non-spatial representations like odor cells and time cells? Is the role of MEC and LEC in driving BTSP different in non-spatial tasks? Chapter 2 addresses these questions.

1.4.3 Representational Drift

The existence of 'place cells', 'grid cells', 'odor cells', and 'time cells' tell the story of individual neurons in the hippocampus and how these individuals contribute to the population dynamics. However, these representations are rarely stable. Throughout the brain, it has been demonstrated that representations drift even when the environment and behavior are entirely stable (Rule et al., 2019; Micou and O'Leary, 2023). Since calcium imaging allows the tracking of neurons across days, place cells have been shown to emerge and fade across multiple days (Ziv et al., 2013; Hainmueller and Bartos, 2018). Similar has been shown with odor cells and time cells in the hippocampus as well (Taxidis et al., 2020). It is believed that this drift allows the brain to continuously learn by being flexible and plastic. However, it is still unknown how downstream regions continue to extract stable information from highly dynamic neural patterns. It is also unknown what types of plasticity contribute to the continuous turnover of cell representations. The final part of chapter 2 will address this question.

1.5 Summary and Overview of Following Chapters

This chapter has summarized some of the relevant research on the hippocampus and EC, working memory, and plasticity mechanisms that drive the formation of neural representations. This chapter posed many questions and highlighted areas of future research for the field to address. The following chapters will address a few of the questions and gaps of knowledge mentioned here.

In chapter 2, we demonstrated that BTSP is a plasticity mechanism that generalizes to non-spatial tasks with the formation of 'odor cells' in dorsal CA1 following large single calcium plateaus. These BTSP events shared many characteristics with the ones previously described in the spatial literature. They had asymmetric formation with MEC activity playing a role in driving the plateaus. However, we also expanded on previous knowledge by demonstrating LEC's role in the formation of the odor cells. Additionally, we made a connection to representational drift and proposed that BTSP may be underlying the turnover of odor cells during expert performance of our olfactory DNMS task.

In chapter 3, we developed a novel rodent working memory implicit timing task that validated a behavioral effect seen in humans. We then showed that sequential activity in CA1 during this task was modulated by temporal expectation. Importantly, the increase in cells that fire during the delay demonstrated anticipation and learning of the implicit timing component of the task.

In chapter 4, we followed up on a surprising result from chapter 2. Despite strong effects of LEC or MEC inhibition on CA1 neural representations, we did not observe a behavioral effect. Neither chemogenetic nor optogenetic inhibition led to behavior impairments. We validated this finding by tuning many different parameters of the DNMS task and training protocol, so that we are confident to conclude that LEC or MEC are not necessary for learning or performance of the task.

Finally, in chapter 5, we tell a cautionary tale about the use of optogenetic tools. We discuss the importance of understanding the complexities of the tools being used. We demonstrate that the inhibitory optogenetic opsin stGtACR2 can be inhibitory at the soma and excitatory at the axon terminals. While this had been shown previously in slices, we found that this led to seizure-like activity *in-vivo* in a rapid kindling protocol.

CHAPTER 2

Non-Spatial Hippocampal Behavioral Timescale Synaptic Plasticity during Working Memory is Gated by Entorhinal Inputs

2.1 Abstract

Behavioral timescale synaptic plasticity (BTSP) is a form of synaptic potentiation where the occurrence of a single large plateau potential in CA1 hippocampal neurons leads to the formation of reliable place fields during spatial learning tasks. We asked whether BTSP could also be a plasticity mechanism for generation of non-spatial responses in the hippocampus and what roles the medial and lateral entorhinal cortex (MEC and LEC) play in driving non-spatial BTSP. By performing simultaneous calcium imaging of dorsal CA1 neurons and chemogenetic inhibition of LEC or MEC while mice performed an olfactory working memory task, we discovered BTSP-like events which formed stable odor-specific fields. Critically, the success rate of calcium events generating a significant odor-field increased with event amplitude, and large events exhibited asymmetrical formation with the newly formed odor-fields preceding the time at which the large induction event occurred. We found that MEC and LEC play distinct roles in modulating BTSP: MEC inhibition reduced the frequency of large calcium events, while LEC inhibition reduced the success rate of odor-field generation. Using two-photon calcium imaging of LEC and MEC temporammonic axons projecting to CA1, we found that LEC projections to CA1 were strongly odor selective even early in task learning, while MEC projection odor-selectivity

increased with task learning but remained weaker than LEC. Finally, we found that LEC and MEC inhibition both slowed representational drift of odor representations in CA1 across 48 hours. Altogether, odor-specific information from LEC and strong odor-timed activity from MEC are crucial for driving BTSP in CA1, which is a synaptic plasticity mechanism for generation of both spatial and non-spatial responses in the hippocampus that may play a role in explaining representational drift and one-shot learning of non-spatial information.

2.2 Introduction

In many situations, learning is not a smooth or gradual process. In fact, our ability to make associations after a single experience is critical for survival. While there have been dramatic improvements in artificial intelligence and machine learning algorithms that implement ‘one-shot learning’ (Lake et al., 2015; Song et al., 2023), the neural underpinnings of this abrupt form of learning have remained elusive. In the hippocampus, a region recognized for its importance in learning and memory, behavioral timescale synaptic plasticity (BTSP) has emerged as a robust mechanism for the rapid generation of spatial representations (place fields) following the occurrence of plateau potentials (Bittner et al., 2015, 2017; Grienberger et al., 2017; Zhao et al., 2020; Magee and Grienberger, 2020; Milstein et al., 2021; Grienberger and Magee, 2022; Xiao et al., 2023). However, the hippocampus not only represents the location of animals in space (O’Keefe and Dostrovsky, 1971; Wood et al., 1999; Buzsáki and Tingley, 2018; Moser et al., 2017), but also non-spatial sensory information (Hampson et al., 1999; Squire, 1992; Aronov et al., 2017; Taxidis et al., 2020). The hippocampus dynamically links these sensory experiences across time through sequential firing that tracks the passage of time after specific events (Buzsáki and Tingley, 2018; Taxidis et al., 2020; Pastalkova et al., 2008; MacDonald et al., 2011; Eichenbaum, 2014). It is unclear whether BTSP also drives the formation of non-spatial sensory-driven or internally generated hippocampal representations, which can form the basis for ‘one-shot learning'.

Hebbian plasticity mechanisms such as spike-timing dependent plasticity require causality and repetition to potentiate synapses when presynaptic spikes precede postsynaptic action potentials by a few milliseconds (Bliss and Lømo, 1973; Caporale and Dan, 2008; Mayford et al., 2012; Feldman, 2012). While this mechanism may play a role in modulating hippocampal responses, BTSP on the other hand, has many features which could make it a more robust and rapid mechanism for the generation of non-spatial receptive fields. During spatial learning tasks, a single calcium plateau potential can serve as the induction event, asymmetrically boosting synaptic inputs that occur several seconds before the induction event, leading to a membrane potential ramp and reliable spatial firing on subsequent trials (Bittner et al., 2017; Magee and Grienberger, 2020; Milstein et al., 2021). Plateau potentials occurring during non-spatial tasks could also boost synaptic inputs at specific time points in the task leading to the rapid formation of stable representations of sensory stimuli, time, and reward/outcome. Yet, whether these events occur during non-spatial tasks, and whether they can rapidly induce stable task-related responses is not known.

The role of the entorhinal cortex (EC) in inducing BTSP events (Grienberger and Magee, 2022) and relaying sensory information during non-spatial tasks (Aronov et al., 2017; Deshmukh and Knierim, 2011; Knierim et al., 2014) is poorly understood. CA1 receives direct layer III EC input via the temporammonic (TA) pathway and indirect input via the perforant path from layer II EC to dentate gyrus, which then projects to CA3, which in turn projects to CA1 (Andersen et al., 2007; Van Strien et al., 2009). Lateral and medial EC (LEC and MEC) have distinct inputs and behaviorally relevant response properties: LEC robustly represents olfactory information (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024), while MEC is more recognized for its encoding of visuo-spatial information (Fyhn et al., 2004; Hafting et al., 2005; Sargolini et al., 2006; Cholvin et al., 2021). Furthermore, the MEC plays a major role in the induction of plateau potential 'teaching signals' during BTSP induced during spatial learning tasks

(Grienberger and Magee, 2022). Yet, whether MEC and LEC play a differential role in the generation of BTSP during non-spatial tasks remains to be determined.

To address these questions, we investigated multimodal representations within CA1 and EC during a non-spatial olfactory delayed non-match-to-sample (DNMS) working memory task (Taxidis et al., 2020). We have previously shown that CA1 pyramidal neurons fire sequentially in response to specific odors and across the 5-second delay period during DNMS performance (Taxidis et al., 2020). We hypothesized that non-spatial BTSP can generate odor representations in CA1 and that this process would be modulated by MEC and LEC inputs. Using two-photon calcium imaging of dorsal CA1, we recorded non-spatial BSTP-like events that formed odor-specific fields in CA1 during expert performance of the DNMS task. Through simultaneous chemogenetic inhibition of LEC or MEC and calcium imaging of CA1, we investigated the role of each EC region in driving odor-specific BTSP-like events. We found that MEC inhibition decreased the frequency of large calcium induction events, while LEC inhibition reduced the success rate of odor-field generation. Additionally, we performed two-photon calcium imaging of LEC and MEC axons in the TA pathway projecting onto dorsal CA1 during learning of the task to investigate how EC inputs to CA1 change with experience and improved performance. Altogether, we demonstrate that MEC's strong firing to odor presentations drives large plateau-like calcium events in CA1, and LEC's odor-selectivity mediates plasticity in the formation of odor-fields after the large calcium induction event.

2.3 Results

We used *in-vivo* two-photon calcium imaging to record the activity of neurons in the pyramidal layer of dorsal CA1 while animals performed an olfactory delayed non-match-to-sample (DNMS) working memory task (Figure 2.1 A-E). Adult male and female mice (n=17) were injected with AAV1-Syn-jGCaMP8f into the right dorsal CA1 and implanted with a 3mm diameter glass-bottomed titanium cannula above the intact alveus after aspiration of the overlying cortex and corpus callosum (Taxidis et al., 2020) (Figure 2.1 E). After one week of recovery, mice were water-restricted and trained on the olfactory DNMS working memory task (Taxidis et al., 2020; Liu et al., 2014), while head-fixed on a spherical treadmill (Figure 2.1 A-B). Each trial consisted of two 1-second odor presentations separated by a 5-second delay period. One second after the offset of the 2nd odor there was a 3-second reward period during which the choice of the animal was determined. Mice were trained to lick the lickport to release water during this reward period if the two odors did not match (correct ‘hit’). Mice learned to refrain from licking the lickport if the odors matched (correct ‘rejection’), and overall performance was quantified as the percentage of correct ‘hits’ and correct ‘rejections’ out of all trials (Figure 2.1 C). We considered performance above 85% to be expert level. Each session of the DNMS task consisted of 5 blocks of 20 trials, with pseudorandomly distributed odor combinations (Figure 2.1 D). Mice were recorded for 8 days during expert performance for a total of 136 recording sessions yielding an average of 312 ± 125 (mean ± standard deviation) active neurons per day. We successfully imaged the same field of view (FOV) for each of the 8 days for all animals. Calcium signals were extracted and deconvolved using Suite2p (Pachitariu et al., 2017) (see methods).

2.3.1 Non-spatial BTSP-like events in CA1 formed stable odor-specific fields

In our previous work, we found that a population of hippocampal neurons fired during specific epochs of the DNMS task (Taxidis et al., 2020). CA1 pyramidal neurons fired during the presentation of specific odors or at time points during the delay period after presentation of specific odors (Taxidis et al., 2020). Here, we observed CA1 neurons with activity patterns consistent with BTSP during expert DNMS performance (Figure 2.1 G-H and Figures 2.5-2.7). Namely, cells without a clear odor or time-field developed a stable field after a single spontaneous large calcium event as the induction event (putative plateau potential). To quantify these rare events, we developed strict criteria for a single calcium event to be considered an ‘induction event’ that could potentially generate an odor-field (see methods).

With increasing ‘induction event’ amplitude, success rate for induction of an odor-field increased (Figure 2.1 I), strongly suggesting a causal role for these induction events in driving the formation of odor-fields. We also found that only events larger than 10 STD exhibited a significantly asymmetrical formation with the newly formed odor-fields preceding the trial time at which the large induction event occurred, (Figure 2.1 K) suggesting that these subset of induction events were true BTSP events. Based on these findings, we set the criteria for an ‘induction event’ to be considered a ‘plateau-like event’ to be that the large calcium induction event must have an amplitude greater than 10 STD. For these plateau-like events, nearly all of the successful induction events peaked during the odor presentation or immediately after the offset, with the success rate reaching 15% during the second half of the odor presentation period (Figure 2.1 J). We observed only 26 events (8% of the 323 successful events) yielding time-fields beyond 0.5 seconds after odor offset, and only 24 events yielding reward-related fields in separate analysis (Figure 2.7). The newly formed odor-specific fields peaked at 0.42 ± 0.14 seconds prior to the onset of the putative plateau potential (n=323 successful ‘plateau-like’ events) (Figure 2.1 L). The small events that represented the random chance of an event passing our strict criteria only had a success rate that peaked at 0.6% during odor presentation and they did not have a significant backward drift (Figure 2.8 A-B).

To determine if motor movements of the animal influenced non-spatial BTSP events, we recorded the movements of the spherical Styrofoam treadmill during performance of the task (Figure 2.8 C). Mice exhibited a range of movement patterns while performing the task with some mice rarely moving on the treadmill and many mice primarily flinching or twitching at the onset of odor presentations (Figure 2.8 C). However, nearly all mice had bouts of running that we defined as periods of continuous locomotion for greater than 2 seconds. As expected, the frequency of low amplitude calcium events was greater during periods of running (1.83 ± 0.21 small events per cell per minute during non-running periods compared to 3.31 ± 0.68 small events per cell per minute during running bouts) (Figure 2.8 F). Additionally, running increased the rate of ‘plateau-like’ events (0.012 ± 0.005 ‘plateau-like’ events per cell per minute during non-running periods compared to 0.028 ± 0.010 ‘plateau-like’ events per cell per minute during running bouts) (Figure 2.8 G). However, running during the ‘plateau-like’ event did not impact the success rate of formation of an odor-field (2.07 ± 1.25 % success rate following ‘plateau-like event’ during non-during periods compared to 2.00 ± 2.63 % success rate during running bouts) (Figure 2.8 H). Therefore, running increased the rate of ‘plateau-like’ events but not the formation of odor-fields, which suggests that the ‘BTSP-like’ events we observed are non-spatial in nature. Together these findings suggest that BTSP can generate non-spatial hippocampal representations.

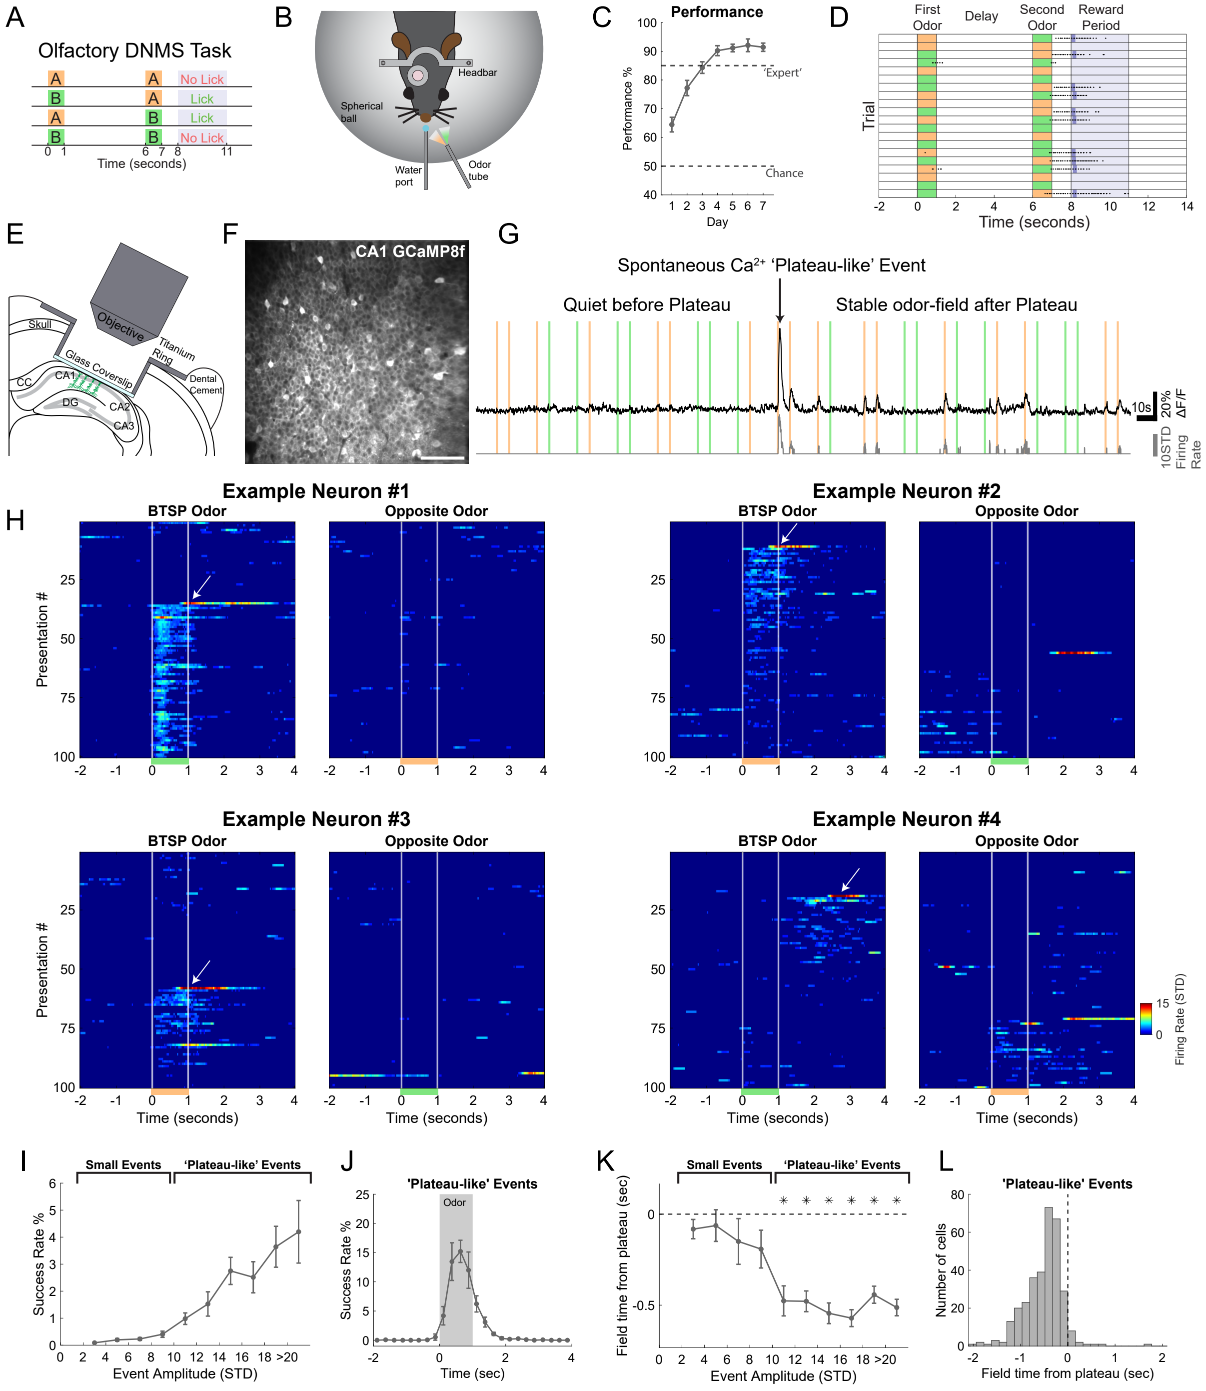
A
Olfactory DNMS Task
A A No Lick
B A Lick
A B Lick
B B No Lick
0 1 6 7 8 11
Time (seconds)
B
Headbar
Spherical ball
Water port
Odor tube
C
Performance
Performance %
'Expert'
Chance
Day
D
First Odor
Delay
Second Odor
Reward Period
Trial
Time (seconds)
E
Skull
Objective
Titanium Ring
Glass Coverslip
CC
CA1
DG
CA2
CA3
Dental Cement
F
CA1 GCaMP8f
G
Spontaneous Ca2+ 'Plateau-like' Event
Quiet before Plateau
Stable odor-field after Plateau
10s
20% ΔF/F
10STD Firing Rate
H
Example Neuron #1
Example Neuron #2
Example Neuron #3
Example Neuron #4
BTSP Odor
Opposite Odor
Presentation #
Time (seconds)
Firing Rate (STD)
15
0
I
Small Events
'Plateau-like' Events
Success Rate %
Event Amplitude (STD)
J
'Plateau-like' Events
Odor
Success Rate %
Time (sec)
K
Small Events
'Plateau-like' Events
Field time from plateau (sec)
Event Amplitude (STD)
L
'Plateau-like' Events
Number of cells
Field time from plateau (sec)

Figure 2.1: **Behavioral timescale synaptic plasticity (BTSP) events in a non-spatial working memory task**

A) Schematic of the olfactory delayed non-match-to-sample (DNMS) task. Water delivery and licking behavior was assessed during the 3 second reward period. **B)** Mice were head-fixed above a Styrofoam spherical ball to allow running. **C)** Behavioral performance across 7 days of learning (n=33). Chance level performance was 50%, and we considered 85% to be 'expert' performance. **D)** Example block of perfect performance for 20 trials. Dots indicate licks and dark blue bars indicate water delivery. **E)** Schematic of two-photon imaging of dorsal CA1 pyramidal neurons. **F)** Example field of view of CA1 neurons expressing GCaMP8f. Scale bar is 100µm. **G)** Example trace of one neuron with a BTSP event and odor-field formed. Colored bars indicate odor presentations. Black trace is $\Delta F/F$, and gray is z-scored deconvolved signal. **H)** Four examples of BTSP-like events. White vertical lines indicate odor onset and offset, and white arrows point to spontaneous induction 'plateau-like' events. **I)** Success rate of a calcium event generating an odor-field increases with induction-event amplitude. Success rate is defined as percentage of events that generate a significant odor-field. Standard error bars represent the standard error of the mean across the 17 animals. Events above 10 STD are considered 'plateau-like' events. **J)** Success rate is highest during odor presentation (for only 'plateau-like' events). **K)** Asymmetrical field formation with trial time difference between 'plateau-like' event peak and formed odor-field peak. This difference is only significant for 'plateau-like' events. Thus, the 'small events' represent chance events that passed our criteria and were likely not BTSP. **L)** Histogram showing asymmetrical distribution for all 323 successful 'plateau-like' events.

2.3.2 Chemogenetic inhibition of entorhinal cortex disrupted non-spatial BTSP

Entorhinal inputs can drive BTSP induction events during spatial navigation (Grienberger and Magee, 2022; Takahashi and Magee, 2009). To determine whether entorhinal inputs may also play a role in the generation of 'plateau-like' events during non-spatial BTSP, we used a chemogenetic strategy to inhibit lateral entorhinal cortex (LEC) or medial entorhinal cortex (MEC), while imaging CA1 calcium activity during the working memory task. Mice were injected with AAV1-Syn-jGCaMP8f in the dorsal CA1 and were subsequently implanted with an optical canula over CA1 as in the previous section. These mice also underwent injection of AAV5-CaMKII-PSAM4 into either LEC (n=6 mice) or MEC (n=5 mice) to express the potent chemogenetic inhibitor PSAM4 (Magnus et al., 2019) in excitatory neurons of either structure. Control mice underwent injections of AAV5-CaMKII-mCherry into either LEC (n=3 mice) or MEC (n=3 mice). Animals were water-restricted, trained on the task, and imaged 3 weeks after viral expression. Each animal was recorded for 8 days after reaching expert level performance. Between 10-20 minutes before two-photon calcium imaging began each day, mice received an intraperitoneal (IP) injection of saline or uPSEM (the effector molecule for PSAM4). Saline and uPSEM injections were alternated daily and animals were counter-balanced such that half of the mice received injections of uPSEM on the first day and the other half of mice received injections of saline. We compared the activity of matched neurons over 4 pairs of days, where animals received saline on one day and uPSEM on the other day.

Despite a lack of a behavioral effect with LEC or MEC inhibition (Figure 2.9), both strongly affected non-spatial BTSP-like events. MEC inhibition significantly reduced the number of 'plateau-like' events from 1.91 ± 0.95 per cell per day to 1.60 ± 1.09 per cell per day (Figure 2.2 B). In contrast, LEC did not affect the number of 'plateau-like' events, but dramatically reduced the success rate of 'plateau-like' events in inducing a new odor-field from 2.13 ± 2.44% to 0.89 ± 1.32%, while MEC inhibition did not significantly alter the success-rate (Figure 2.2 C). Importantly, neither LEC nor MEC inhibition affected locomotion or the percentage of time spent running, so these effects could not be explained by differences in animal movement (Figure 2.9 B-C). Together, these findings suggest that MEC affects the generation of large ‘plateau-like’ events in CA1, while LEC activity increases the likelihood that these events result in successful field generation.

2.3.3 LEC inhibition reduced strength of odor representations in dorsal CA1, while MEC inhibition increased strength

Given that LEC inputs have been previously shown to encode odor-related information (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024), we hypothesized that they could convey odor-related information to CA1 in our DNMS task. If so, we would expect inhibition of LEC but not MEC to decrease odor selectivity in CA1, potentially driving the decrease in success rate of ‘plateau-like’ events in generating odor-fields. Indeed, LEC chemogenetic inhibition significantly decreased odor selectivity values and the percentage of odor selective cells (Figure 2.2 D and Figure 2.10 A). None of the 6 mCherry controls animals showed a shift in odor selectivity (Wilcoxon signed-rank test pairing all cells, $p > 0.05$ for each animal). First and second odor selectivity were similarly modulated (Figure 2.10 B). The proportion of significantly odor-selective neurons (based on comparisons with shuffled controls; see methods) was $24.5 \pm 10.7\%$ on saline control days and only $16.2 \pm 10.9\%$ on uPSEM inhibition days in LEC experimental PSAM4 animals. In contrast, MEC inhibition showed a slight trend in the opposite direction with $20.7 \pm 5.8\%$ of neurons being odor selective on saline control days and $21.8 \pm 5.5\%$ on uPSEM inhibition days in MEC experimental PSAM4 animals. Therefore, LEC inhibition weakened CA1 neuron odor selectivity.

To further confirm this effect, we performed binary support vector machine (SVM) decoding training and testing on the same day to evaluate the relative strength of odor encoding on saline days compared to uPSEM days. Overall, LEC inhibition significantly decreased odor decoding accuracy, while MEC inhibition significantly improved odor decoding (Figure 2.2 E-F). During the odor presentation period (subsampling only 100 neurons), the decoding accuracy in LEC experimental PSAM4 animals was 76.4 ± 11.0% on saline control days and decreased to 66.7 ± 13.9% on uPSEM inhibition days. In contrast, MEC experimental PSAM4 mice had decoding accuracy of 79.0 ± 5.8% on saline control days, which increased to 82.5 ± 5.5% on uPSEM inhibition days. Odor decoding of control animals expressing mCherry was unaffected by uPSEM administration (Figure 2.10 A-B). MEC inhibition only increased decoding accuracy during odor presentation, while LEC inhibition decreased odor decoding accuracy during the earlier part of the delay period as well (Figure 2.2 E). Increasing the number of subsampled cells for decoding led to improvements in odor decoding, but in general differences between MEC and LEC inhibition were observed for a large range of neuron numbers subsampled for decoding analysis (Figure 2.10 D).

Collectively, LEC inhibition strongly decreased whereas MEC inhibition modestly increased odor selectivity and decodability in CA1. LEC's effect may have driven the reduction in the success rate of 'plateau-like' events in generating odor-fields.

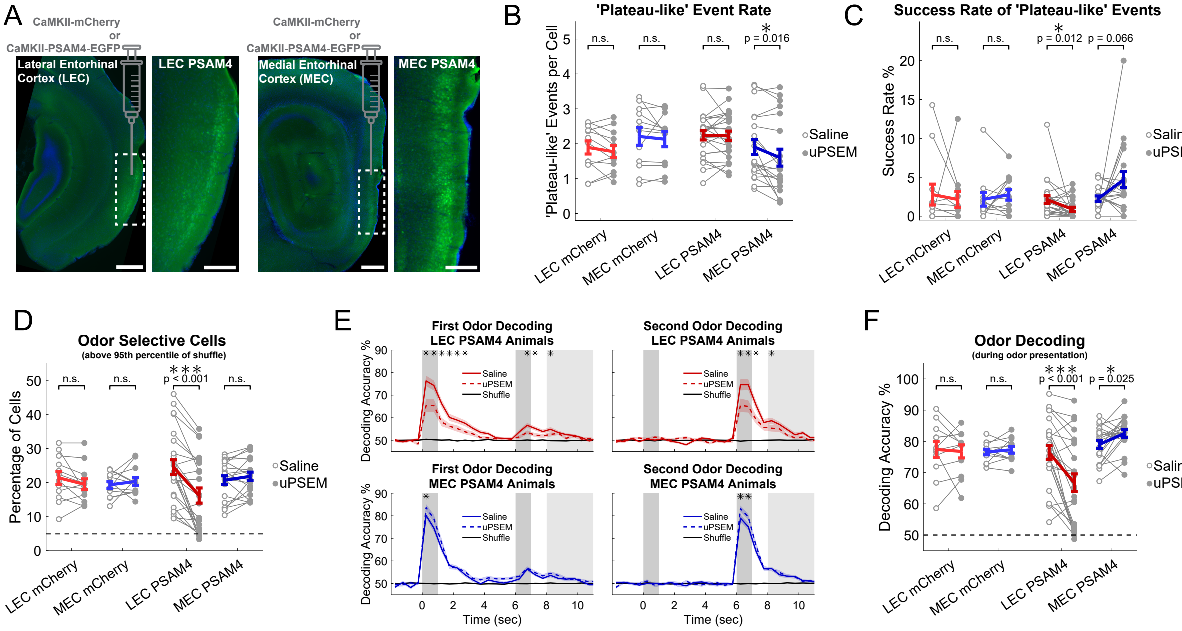

Figure 2.2: **LEC and MEC inhibition differentially modulated BTSP, and LEC inhibition weakened odor selectivity in CA1**

A) Injections of virus to drive the expression of mCherry or PSAM4 were delivered to either LEC or MEC. Images showing LEC are from coronal sections, while MEC are from sagittal sections. For both LEC and MEC, the larger image on the left has a 500µm scale bar and the right image is a zoom of the white outline with a 200µm scale bar. **B)** Number of events greater than 10 STD per cell per day. Paired dots represent the pairs of imaging days (4 per animal). Statistics are two-way ANOVA (animal and pair) with repeated measures on the saline/uPSEM condition. **C)** Success rate of 'plateau-like' events generating an odor-field. **D)** Percentage of cells that had a selectivity value above 95th percentile of shuffle. Statistics are also two-way ANOVA (animal and pair) with repeated measures on the saline/uPSEM condition. **E)** Binary support vector machine (SVM) decoding of first and second odor across the trial structure with 0.5 second bins for experimental animal groups (repetitive subsampling of 100 neurons for each recording session). Thinner gray bars indicate odor presentation and wider bar from seconds 8-11 is the reward period. Statistics are the same, and asterisks indicate bins with $p < 0.05$ (corrected for multiple comparisons using the Benjamini-Hochberg procedure). **F)** Odor decoding performance only during the odor presentation period.

2.3.4 Two-photon calcium imaging of entorhinal cortical axons in dorsal CA1 revealed differential sequential activity in LEC and MEC inputs

The EC is the primary cortical input to the hippocampus; CA1 receives direct layer III EC input via the temporammonic (TA) pathway and indirect input via the perforant path from layer II EC to dentate gyrus, which then projects to CA3, which in turn projects to CA1 (Andersen et al., 2007; Van Strien et al., 2009). Given the contribution of MEC in driving 'plateau-like' events, we asked if there are differences in timing of LEC and MEC TA inputs. Also, given the strong differences in odor decodability observed in dorsal CA1 with LEC versus MEC inhibition, we asked whether TA inputs from LEC and MEC differ in the sensory and task-related information they convey to CA1. Do LEC and MEC TA inputs change as mice learn the task?

To address these questions, we performed two-photon calcium imaging of LEC or MEC TA axons in dorsal CA1 as mice learned the DNMS task. Adult male and female mice were injected with AAV1-CaMKII-Cre and AAV1-CAG-FLEX-jGCaMP7s in either LEC (n=8) or MEC (n=8) (Figure 2.3 A and C). Mice were implanted with hippocampal windows as in the previous experiments. After 3 weeks of expression, confocal imaging demonstrated extensive GCaMP7s axonal expression of TA inputs within the stratum lacunosum-moleculare (SLM) layer, as well as layer II EC perforant path axons ramifying deeper within the stratum moleculare (MOL) layer of the dentate gyrus. *In-vivo*, we could selectively image TA EC axons 300 and 400µm beneath the alveus. Post-hoc histology after two-photon imaging experiments also confirmed that all mice had extensive expression of GCaMP7s in axons within the SLM layer of hippocampus and somatic expression restricted to either LEC or MEC. In these experiments, imaging experiments began on the first day of training when mice are presented with matched pairs and began learning to refrain from licking on these trials (see methods). Recordings were processed with Suite2p (Pachitariu et al., 2017) using parameters optimized for axonal imaging, followed by post hoc fusion of axon segments with highly correlated activity which were branches of the same axon (see methods) (Figure 2.3 B and D).

In trained animals, a proportion of LEC and MEC TA axons responded reliably to different task variables. Some axons had responses which had significant peaks during the odor presentation, some during odor offset, and others during the delay period (Figure 2.3 E and F). Altogether, the LEC or MEC axonal populations had sequential activity that tiled the entire first odor presentation and delay period. However, these sequences differed drastically between LEC and MEC; a much higher proportion of MEC axons had significant peaks during odor presentations compared to LEC axons (Figure 2.3 G-J).

To investigate how LEC and MEC inputs to CA1 change with learning, we first visualized the sequential firing of significantly modulated neurons (see methods) during expert performance and below expert performance (Figure 2.3 G). We noticed stark differences in the proportion of axons with peak activity during the odor compared to during the delay between LEC and MEC. There were also clear differences in these proportions when comparing poor performance to expert performance. During days of expert performance, MEC had more axons with peak firing during odor presentation compared to LEC ($26.9 \pm 8.7\%$ compared to $10.1 \pm 7.2\%$, ANOVA $p < 0.001$), while LEC had more axons with peak firing during the delay period compared to MEC ($5.9 \pm 1.9\%$ compared to $2.5 \pm 1.1\%$, ANOVA $p < 0.001$) (Figure 2.3 I). As a result, MEC axons also showed greater trial reliability as compared to LEC (Figure 2.11 C). The percentage of MEC axons with peak firing during odor presentation increased across learning (Pearson's $r = 0.311$), while those with peak firing during the delay period decreased (Pearson's $r = -0.433$). Meanwhile, the proportion of LEC axons with peak firing during odor and delay periods remained stable with learning (Pearson's $r = 0.082$ and 0.074).

In summary, timing of LEC inputs to CA1 were stable with learning, while MEC inputs became more tuned to the odor presentation period. Given that MEC inhibition reduced the rate of 'plateau-like' events in CA1, we hypothesize that this strong MEC input timed to the odor presentation is likely key for driving 'plateau-like' events.

2.3.5 LEC odor representations were stable during learning, while MEC tuned firing to odor presentation and odor selectivity emerged

Although a higher proportion of MEC axons were tuned to firing during the odor presentation, examination of sequential firing patterns suggested that these axons were firing with less odor-specificity (Figure 2.3 G-H and Figure 2.11 A-B). To quantify odor information carried by EC axons, we calculated odor selectivity and odor decoding accuracy. Despite the increased number of MEC axons with peak firing during the odor presentation, LEC had a greater proportion of significantly odor-selective axons (14.3 ± 6.8% in LEC and 10.3 ± 4.6% in MEC, ANOVA $p < 0.001$) (Figure 2.3 K). This effect was strongest early in learning, as MEC odor selectivity increased with DNMS performance (Pearson's $r = 0.373$). The larger number of odor selective axons in LEC resulted in better odor decoding during odor presentation when repetitively subsampling only 100 axons from each recording session (for all recordings LEC decoding accuracy is 57.0 ± 4.2% and MEC accuracy is 53.5 ± 2.2%, ANOVA $p < 0.001$; for only expert sessions LEC accuracy is 57.0 ± 4.3% and MEC accuracy is 54.4 ± 2.1%, ANOVA $p = 0.032$) (Figure 2.3 L-M). Again, decoding accuracy for LEC was stable across days and performance levels (Pearson's $r = 0.093$), but dramatically improved for MEC (Pearson's $r = 0.539$) (Figure 2.3 N). While odor decoding was worse during the delay period than during the odor presentation, decoding accuracy during the delay period remained significantly greater for LEC compared to MEC (for all recordings LEC accuracy was 54.8 ± 3.8% and MEC accuracy was 51.7 ± 1.5%, ANOVA $p < 0.001$; for only expert sessions LEC accuracy was 54.5 ± 4.1% and MEC accuracy was 52.0 ± 1.7%, ANOVA $p = 0.030$) (Figure 2.3 L-M and Figure 2.11 H). Increasing the number of subsampled axons for decoding led to improvements in odor decoding, but in general differences between LEC and MEC were similar across a large range of axon numbers subsampled for decoding analysis (Figure 2.11 I-J).

To understand if LEC and MEC encode other task relevant representations in our working memory task, we asked whether EC axons can encode whether the two odors matched or did not match. Both MEC and LEC axons showed an increase in SVM decoding accuracy of match versus non-match trials with increasing performance (LEC Pearson's r = 0.382 and MEC Pearson's r = 0.746); however, MEC accuracy was dramatically higher even during the 2nd odor which is one second before the start of the reward period (67.0 ± 7.4% for MEC, and 57.6 ± 6.2% for LEC, ANOVA $p < 0.001$). During poor performance sessions (<85% performance), match versus non-match trial decoding accuracy peaked during the middle of the reward period as the outcome was encoded, and MEC decoding accuracy was significantly higher than LEC (Figure 2.12 D). Interestingly, once mice reached expert performance, decoding accuracy of match versus non-match trials peaked during the 2nd odor for MEC, but still peaked during the reward period for LEC (Figure 2.12 E).

Altogether, these findings suggest that LEC temporammonic axonal odor representations were strong in novice animals and did not improve with performance, while MEC axonal firing became strongly tuned to firing at odor presentations during learning but had relatively weaker odor selectivity. Meanwhile MEC but not LEC axons showed emergence of robust working memory representations related to reward, choice, or trial types.

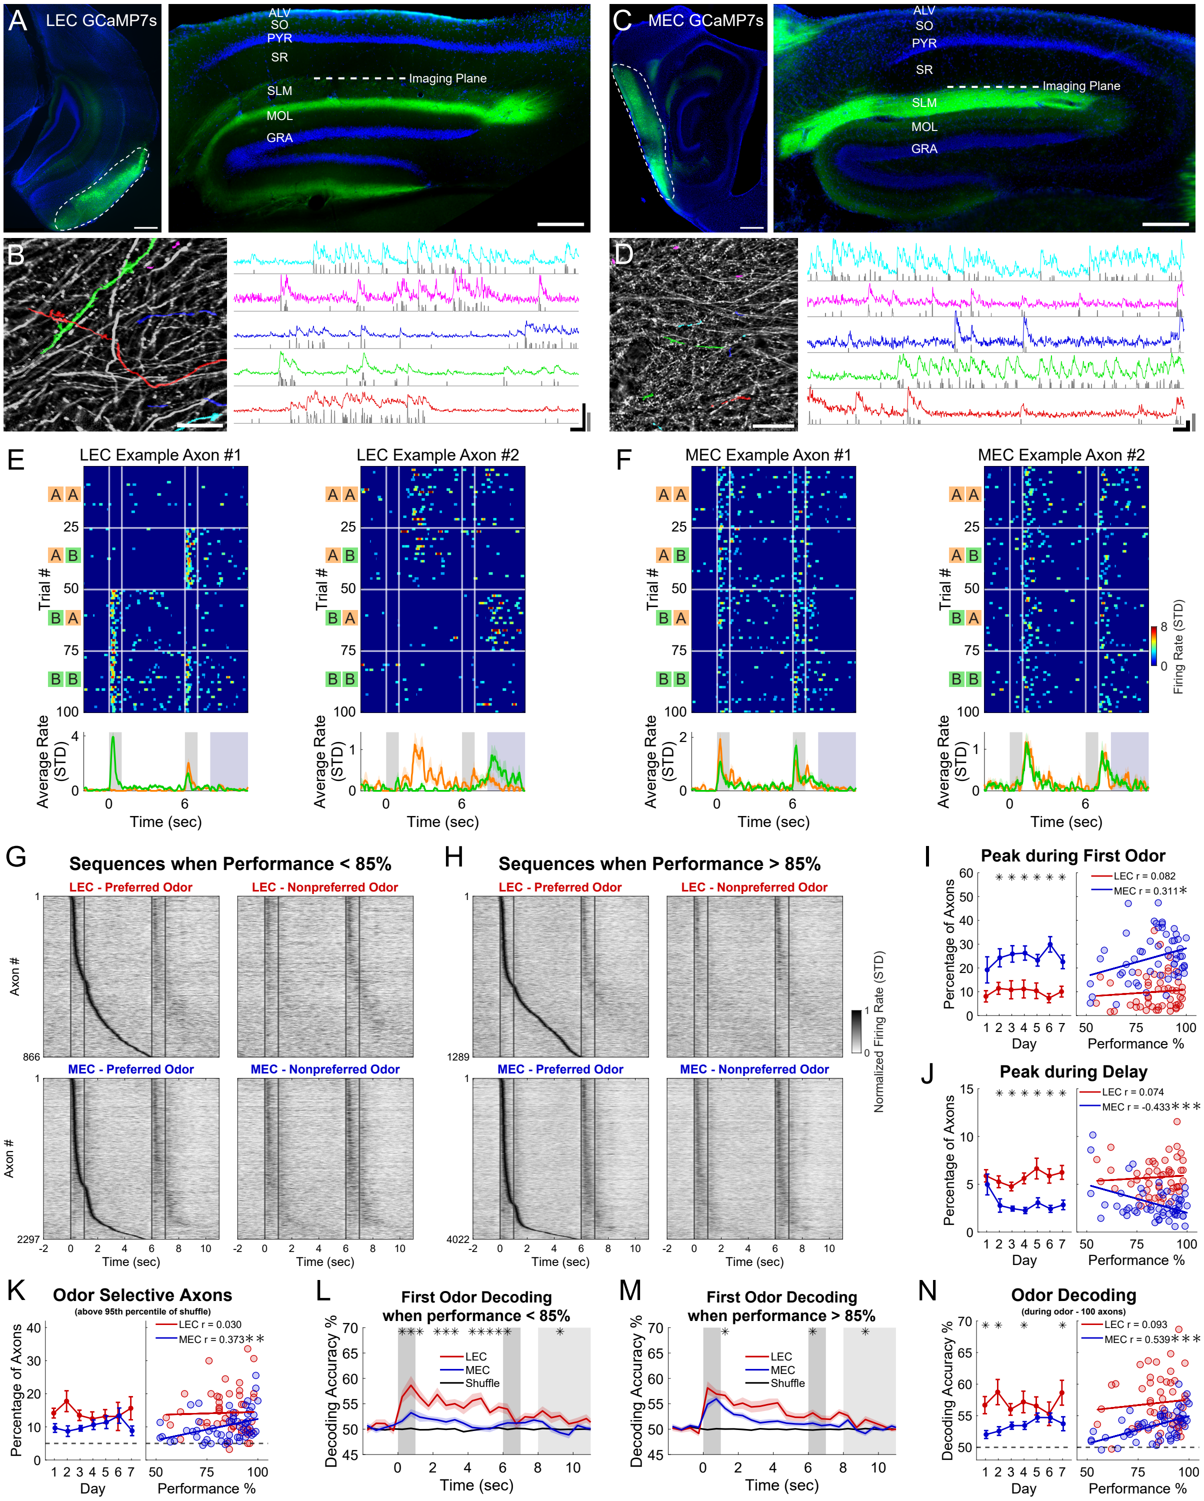
A
LEC GCaMP7s
ALV
SO
PYR
SR
Imaging Plane
SLM
MOL
GRA
C
MEC GCaMP7s
ALV
SO
PYR
SR
Imaging Plane
SLM
MOL
GRA
B
D
E
LEC Example Axon #1
LEC Example Axon #2
A A
25
A B
Trial #
50
B A
75
B B
100
Average Rate (STD)
4
0
1
Time (sec)
0
6
F
MEC Example Axon #1
MEC Example Axon #2
2
8
0
Firing Rate (STD)
G
Sequences when Performance < 85%
LEC - Preferred Odor
LEC - Nonpreferred Odor
MEC - Preferred Odor
MEC - Nonpreferred Odor
Axon #
1
866
2297
-2
0
2
4
6
8
10
Time (sec)
H
Sequences when Performance > 85%
1289
4022
Normalized Firing Rate (STD)
1
0
I
Peak during First Odor
Percentage of Axons
60
50
40
30
20
10
0
1 2 3 4 5 6 7
Day
LEC r = 0.082
MEC r = 0.311*
50
75
100
Performance %
J
Peak during Delay
15
10
5
0
LEC r = 0.074
MEC r = -0.433***
K
Odor Selective Axons
(above 95th percentile of shuffle)
40
30
20
10
0
LEC r = 0.030
MEC r = 0.373**
L
First Odor Decoding
when performance < 85%
Decoding Accuracy %
70
65
60
55
50
45
LEC
MEC
Shuffle
M
First Odor Decoding
when performance > 85%
N
Odor Decoding
(during odor - 100 axons)
LEC r = 0.093
MEC r = 0.539***

Figure 2.3: **Two-photon calcium imaging of entorhinal cortical axons in dorsal CA1 revealed differential sequential activity in LEC and MEC inputs**

A) Coronal sections showing GCaMP7s expression LEC (left panel, scale bar = 500µm) and in dorsal hippocampus (right panel, scale bar = 250µm). Blue is DAPI. Imaging plane is at the superficial part of the SLM layer, which is the first layer of axons visible when lowering into the tissue roughly 300-400µm beneath the coverglass. ALV = alveus, SO = stratum oriens, PYR = stratum pyramidale, SR = stratum radiatum, SLM = stratum lacunosum-moleculare, MOL = stratum moleculare, GRA = stratum granulare. **B)** Field of view from the same animal (scale bar = 50µm), with 5 example masks and their corresponding fluorescence traces. Gray is z-scored deconvolved signal. Black horizontal scale bar = 10 seconds. Black vertical scale bar = 5% $\Delta F/F$. Gray vertical scale bar = 10 STD normalized deconvolved signal. **C-D)** Same as (A-B) but for MEC and showing sagittal sections. All scale bars are the same. **E)** Two example axon segments showing odor-specific firing. The left axon had its peak during the odor presentation, while the right one had its peak during the delay period. Heatmaps show deconvolved signal on each trial that was grouped into trial type. Average traces at bottom show difference in average firing rate split by trials that started with Odor A and those that started with Odor B. **F)** Same as (E) but for two representative MEC axon segments with less odor-selectivity. The right axon had its peak following the offset of the odor presentation. **G)** Sequential activity of only axon segments that had a significant peak during the first odor presentation or delay period from recording sessions with performance less than 85%. Each row is the average trace of trials with the preferred or nonpreferred first odor (normalized to peak). Blue lines indicate odor onset and offset. **H)** Same as (G) but when performance was at least 85%. **I)** Percentage of axons with a significant peak during the first odor presentation period. Statistics for left panel are two-sample t-tests and p-values were corrected for multiple comparisons using the Benjamini-Hochberg procedure. Statistics for right panel are Pearsons' R correlation with performance. **J)** Same as (I) but percentage of axons with a significant peak during the delay period. **K)** Percentage of axons that had a selectivity value above 95th percentile of shuffle. **L)** Binary SVM decoding of first odor across the trial structure with 0.5 second bins (repetitive subsampling 100 axons for each recording session), only on recordings sessions with behavior performance less than 85%. Statistics are two-way ANOVA (animal and day), and asterisks indicate bins with $p < 0.05$ (corrected for multiple comparisons using the Benjamini-Hochberg procedure). **M)** Same as (L) but when performance was at least 85%. **N)** Binary SVM odor decoding only during the odor presentation period (repetitive subsampling 100 axons for each recording session).

2.3.6 LEC and MEC inhibition slow representational drift of odor representations in dorsal CA1

Despite the similarity of LEC TA axon population dynamics across days, our previous work revealed that CA1 odor representations drift over days (Taxidis et al., 2020) with new cells forming sensory relevant fields and other cells losing their responsiveness or selectivity. Given that non-spatial BTSP can result in rapid generation of odor-selective responses, we hypothesized that it could play a role in representational drift. Since MEC inhibition reduces the frequency of 'plateau-like' events and LEC inhibition reduces the success rate of odor-field formation, we hypothesized that the reduction of BTSP events through EC inhibition may result in increased stability of representations. To address this hypothesis and compare representations over days, we matched dorsal CA1 neurons across 8 days of alternating saline and uPSEM administrations (Figure 2.4 A). While some BTSP events formed odor-fields that fade within the recording session (Figure 2.1 H and Figures 2.5-2.6), others formed fields that lasted for several days (Figure 2.4 A-B). We used binary SVM decoders trained on the activity of 100 randomly chosen neurons on the day before EC inhibition (Saline Day X) and tested on the same neurons the day after EC inhibition (48 hours later on Saline Day X+2). We compared these results to same-day decoding on Saline Day X. The higher the success rate of the decoder for across-day decoding, the more stable the representation.

In control mice expressing mCherry, decoder accuracy declined quickly (77.0 ± 6.4% for same-day decoding (Saline Day X) to 58.2 ± 9.7% for across-day decoding two days later (Saline Day X+2) (Figure 2.4 B-C)), suggesting substantial representational drift. In experimental mice expressing PSAM4, decoder accuracy dropped substantially less for LEC (76.4 ± 11.0% for same-day decoding to 65.1 ± 10.1% for across-day decoding two days later) and MEC (79.0 ± 5.8% for same-day decoding to 65.0 ± 4.6% for across-day decoding two days later). In addition, the percentage of Saline Day X neurons that remain significantly odor-selective on Saline Day X+2 was higher in LEC experimental PSAM4 animals (45.4 ± 10.9%) than in control mCherry animals (36.3 ± 9.9%) (Figure 2.4 E). This indicates that PSAM4 inhibition of both LEC or MEC decreased representational drift across a 48-hour period.

In summary, MEC inhibition reduced large calcium events in CA1, LEC inhibition decreased success rate of these ‘plateau-like’ events, and inhibition of either LEC or MEC slowed representational drift of odor in CA1. These findings suggest that drift of CA1 olfactory representations is modulated by EC inputs, potentially from decreased non-spatial BTSP.

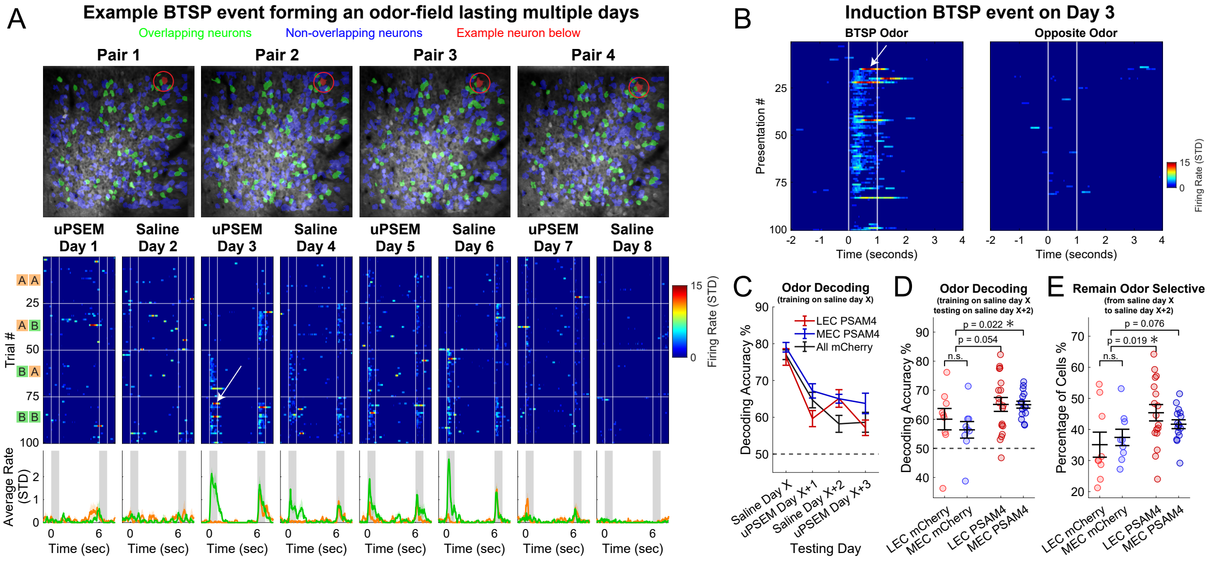

Figure 2.4: **LEC and MEC inhibition both slow representational drift of odor representations in dorsal CA1**

A) Example CA1 neuron across 8 days of expert performance. The 4 FOV images show the masks used for each 'pair' (see methods), and neurons were aligned across pairs with CellReg (Sheintuch et al., 2017). Green masks are cells that overlap in all 4 pairs; blue masks do not overlap in all 4 pairs, and the single red mask is the example neuron with activity below. Heatmaps show deconvolved signal on each trial with trials grouped according to DNMS odor combinations. Average traces at bottom show average firing rates for trials that started with Odor A (orange) or with Odor B (green). **B)** Visualization of BTSP event in Day 3 from (A) that is likely reinforced by several strong 'plateau-like' events. The white arrow points to the same induction 'plateau-like' event as the white arrow in (A). **C)** Binary SVM decoding of odor (only during odor presentation periods) after training on saline day X (repetitive random subsampling of 100 neurons for each recording session). **D)** Highlighting the effect in (C) on saline day X+2 with each circle representing a recording session. Black bars represent mean and standard error of the mean (SEM). Statistics are two-way ANOVA (animal and pair). **E)** Percentage of cells that have odor selectivity values that remained above 90th percentile of shuffle for their preferred odor.

2.4 Discussion

Using two-photon calcium imaging of dorsal CA1 pyramidal neurons during an olfactory working memory task, we find that non-spatial sensory representations can form on single trials following large calcium events. These events have characteristics of BTSP reported previously during spatial tasks (Bittner et al., 2015, 2017; Grienberger et al., 2017; Zhao et al., 2020; Magee and Grienberger, 2020; Milstein et al., 2021; Grienberger and Magee, 2022; Xiao et al., 2023), suggesting that BTSP may be a general plasticity mechanism for formation of hippocampal representations during both spatial and non-spatial cognition. Additionally, MEC and LEC inhibition differentially modulate non-spatial BTSP during working memory performance. MEC inhibition decreases the frequency of large 'plateau-like' calcium events, while LEC inhibition reduces the success rate of these 'plateau-like' events generating an odor-field. LEC inputs are critical for generation of odor representations in CA1, with LEC inhibition dramatically weakening CA1 odor selectivity and odor encoding. This may contribute to their modulation of BTSP success rate in generating odor-selective responses. By performing two-photon calcium imaging of LEC or MEC temporammonic pathway axons to CA1, we show that LEC relays stronger odor-specific information to CA1 that is invariant in learning and expert stages, while MEC axonal activity shows greater plasticity with learning, increasing odor and match/non-match selectivity and tuning to more reliably fire during the odor presentations. Finally, inhibition of both LEC or MEC leads to reduced representational drift of CA1 odor representations, suggesting that BTSP (or another EC-dependent plasticity process) can modulate representational drift.

This is to our knowledge the first description of behavioral timescale synaptic plasticity (BTSP) occurring in a non-spatial context. Non-spatial BTSP described in this paper and spatial BTSP described in spatial contexts (Bittner et al., 2015, 2017; Grienberger et al., 2017; Zhao et al., 2020; Magee and Grienberger, 2020; Milstein et al., 2021; Grienberger and Magee, 2022; Xiao et al., 2023) share many attributes. First, they are both induced by large calcium events. Second, like spatial BTSP, odor-responsive fields typically form around 0.5 seconds before the time of onset of the ‘plateau-like’ event. This temporally asymmetric induction of fields is characteristic of BTSP in CA1. Membrane potential (Vm) recordings in CA1 during spatial BTSP demonstrate potentiation causing the induction of an asymmetric Vm ramp extending back nearly 4 seconds from the timepoint of induction. Voltage recordings would be required to determine whether a Vm ramp extending several seconds is also induced by non-spatial BTSP. There are notable differences between spatial and non-spatial BTSP, however. While spatial BTSP can induce place fields anywhere in the virtual track, during our non-spatial BTSP, 86% of successful fields were formed during or immediately after the odor presentations, with few fields formed during the delay and reward periods. It is possible that this occurs because subthreshold inputs potentiated by BTSP in the delay period fail to reach action potential threshold. This could be explained by the fewer EC inputs activated during the delay period as LEC has nearly twice as many axons and MEC nearly 10 times as many axons with peak firing during the odor period compared to the delay period. Recordings of Vm during the task would be necessary to find whether the magnitude of synaptic potentiation is similar during the different phases of the task. It also remains to be determined whether TA inputs, CA3 inputs, or both are potentiated during non-spatial BTSP. Finally, while inhibitory interneuron subtypes have been characterized by their roles in gating spatial EC and CA3 inputs to CA1 (Klausberger and Somogyi, 2008; Kepecs and Fishell, 2014; Milstein et al., 2015; Basu et al., 2016), it remains unclear how the different interneuron subtypes within the different layers of CA1 contribute to BTSP and gate non-spatial sensory inputs. Future recordings and manipulations of the activity of these neurons will further elucidate the complex mechanisms underlying non-spatial BTSP in CA1.

We find that inhibition of LEC and MEC have distinct effects on non-spatial BTSP. While MEC inhibition reduces the frequency of large calcium events, LEC inhibition has no impact on the frequency or amplitude of these events but reduces their success rate in generating odor-fields. Therefore, while it is clear that MEC plays a major role in generating the plateau potential teaching signal with most of its activity timed to stimulus presentations, the exact mechanism through which LEC regulates the success of BTSP events is less clear. There are several possibilities. It is possible that BTSP potentiates the LEC inputs on the distal dendrites of CA1 pyramidal neuron which aids in generating odor-selective responses. Alternatively, it is possible that LEC inhibition reduces odor-selectivity and the amplitude of odor responses in dentate gyrus granule cells or in CA3, which in turn reduces the potentiation of CA3 inputs to CA1. Our results are in line with studies which have shown the importance of MEC inputs for generation of teaching signals to drive BTSP during spatial navigation (Magee and Grienberger, 2020; Grienberger and Magee, 2022), but our results describe the further complexity given the distinct roles of LEC and MEC.

We found that CA1 population odor representations were more stable the day after MEC or LEC inhibition, suggesting that EC inhibition slows representational drift. This reduction can potentially occur through a reduction in frequency or success rate of BTSP events, as shown in our work, or may occur through a different plasticity mechanism governed by EC activation (Rule et al., 2019; Driscoll et al., 2022; Micou and O'Leary, 2023). While BTSP can clearly explain the appearance of a new field, the mechanism for erasure of existing fields remains less clear. One possibility is that decreases in synaptic weights can occur following mistimed plateau potentials, given that spatial BTSP has been shown to increase synaptic weights of inputs within 2 seconds of a plateau and decrease synaptic weights of inputs between 2 and 5 seconds of the plateau (Milstein et al., 2021). However, a continuing challenge for the field will be to understand the complex interplay of other plasticity mechanisms implicated in representational drift that operate on different and longer timescales, such as Hebbian spike timing-dependent plasticity that includes long-term potentiation (LTP) and long-term depression (LTD) (Kappel et al., 2015; Aitchison et al., 2021; Qin et al., 2023) or dendritic spine turnover (Yasumatsu et al., 2008; Minerbi et al., 2009; Attardo et al., 2015). There is some evidence that there are distinct pools of CA1 neurons with short or long place field lifetimes, which may be related to BTSP success rate (Vaidya et al., 2023). Whether similar pools exist for non-spatial representation remains to be determined.

Our findings support the structural and functional connectivity of LEC and the hippocampus in olfactory based tasks (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024), but further experiments with other modalities would be valuable in establishing LEC and MEC's unique roles in driving plateau potentials and forming non-spatial representational fields. CA1 is also well known for its internal representations (Pastalkova et al., 2008; MacDonald et al., 2011; Eichenbaum, 2014). Although we observed some BTSP events that form odor-specific fields during the delay period, future recordings should investigate if LEC and MEC inputs coincide with the output from recurrent CA3 networks capable of generated temporal codes (Liu and Buonomano, 2009; Holtmaat and Caroni, 2016; Salz et al., 2016) to drive BTSP for internally generated representations.

2.5 Method Details

Animals

All of the experiments were conducted according to the National Institute of Health (NIH) guidelines and with the approval of the Chancellor's Animal Research Committee of the University of California, Los Angeles. A total of 9 adult male and 8 female mice (8-16 weeks old) were used for *in-vivo* calcium CA1 neuron imaging experiments, and a total of 7 adult male and 9 female mice (8-16 weeks old) were used for *in-vivo* calcium EC axon imaging experiments. CA1 imaging mice are divided into 4 groups: LEC mCherry n=3, MEC mCherry n=3, LEC PSAM4 n=6, MEC PSAM4 n=5. Axon imaging mice are divided into 2 groups: LEC n=8, MEC n=8. All were C57BL/6J (Jackson Laboratory, 000664), experimentally naïve, and housed in the vivarium under a 12-hour light/dark cycle. All mice were group housed (2-4 per cage) with the exception of 2 that had to be separated following surgery because of fighting.

Surgical Procedures

Mice (8-12 weeks old) were subcutaneously administered pre-operative drugs (carprofen 5 mg/kg, dexamethasone 0.2 mg/kg, lidocaine 5 mg/kg) 30 minutes before surgery. Mice were anaesthetized with isoflurane (5% induction, 1-2% for maintenance), and anesthesia was continuously monitored and adjusted as necessary. The scalp was shaved, and mice were placed into a stereotactic frame (David Kopf Instruments, Tujunga, CA) on a feedback-controlled heating pad (Harvard Apparatus) set to maintain body temperature at 37°C. Eyes were protected from desiccation using artificial tear ointment. The surgical incision site was cleaned three times with 10% povidone-iodine and 70% ethanol. Fascia was removed by applying hydrogen peroxide, connective tissue was cleared from the skull, and the skull was scored to facilitate effective bonding with adhesives at the end of surgery. After stereotactically aligning the skull, a single or several burr holes were made depending on the experiment performed and virus was injected.

CA1 calcium imaging experiments: Control virus (500 nL of 1:5 saline dilution of pAAV1-CaMKIIa-mCherry into all 4 sites) or experimental virus (500 nL of 1:5 saline dilution of AAV5-CaMKII-PSAM4-GlyR-IRES-EGFP into all 4 sites) was injected into LEC (bilaterally 3.4 and 3.9 mm posterior, 4.35 mm lateral, and 4.3 ventral from bregma) or MEC (bilaterally 4.7 mm posterior, 3.35 mm lateral, and 3.8 and 3.0 mm ventral from bregma). Additionally, pGP-AAV1-syn-jGCaMP8f-WPRE (1000nL of 1:5 saline dilution) was injected into the right dorsal CA1 (2.0 mm posterior from bregma, 1.8 lateral from bregma, and 1.3 ventral from dura).

EC axon calcium imaging experiments: pENN.AAV1.CaMKII.0.4.Cre.SV40 and pGP-AAV1-CAG-FLEX-jGCaMP7f-WPRE were mixed immediately before the injection (500 nL of 1:1 mix) into right LEC (3.5 mm posterior, 4.35 mm lateral, and 4.3 ventral from bregma) or right MEC (4.7 mm posterior, 3.35 mm lateral, and 3.5 mm ventral from bregma). All viruses were injected using a Nanoject II microinjector (Drummond Scientific) at 60nL per minute.

For mice in all experiments, following virus injection, a circular craniotomy (3 mm diameter) was made centered around a point made 2.0 mm posterior and 1.8 lateral to bregma. Dura beneath the craniotomy was removed and cortical tissue above dorsal CA1 was carefully aspirated using a 27-gauge blunt needle. Corpus callosum was spread to the sides of the craniotomy to expose the alveus. Cortex buffer (NaCl = 7.88g/L, KCl = 0.372g/L, HEPES = 1.192g/L, $CaCl_2$ = 0.264g/L, $MgCl_2$ = 0.204g/L, at a pH of 7.4) was continuously flushed during aspiration and until bleeding stopped. A titanium ring with a 3 mm diameter circular thin #0 coverglass attached to its bottom was implanted into the aspirated craniotomy and the overhanging flange was secured to the skull with vetbond (3M). A custom-made lightweight stainless-steel headbar was attached to posterior skull and secured with cyanoacrylate glue. Dental cement (Ortho-Jet, Lang Dental) was applied to seal and cover any remaining skull, and to form a small well around the titanium ring for holding immersion water for the objective during imaging. Following surgery, all animals were given post-operative care (carprofen 5 mg/kg and dexamethasone 0.2 mg/kg for 48 hours after surgery) and provided amoxicillin-treated water at 0.5 mg/mL for 7 days. All mice recovered for 7-14 days before experiments began.

Experimental setup

The entire behavioral setup is as described in Taxidis et al. (Taxidis et al., 2020). Mice were head-fixed above an 8-inch spherical Styrofoam ball (Graham Sweet) which can rotate about one axis for 1D locomotion that was recorded with a sensor (Avago ADNS-9500). A continuous stream of clean air (∼1 L/min) was delivered toward the animal's nose via Tygon PVC clear tubing and a custom-made port that held the air tube and water port. At the onset of the odor presentation period, a dual synchronous 3-way valve (NResearch) switched to the odorized one for 1 second. Odorized air was created by using a 4-ports olfactometer (Rev. 7c; Biology Electronics, Caltech) supplying air to either of two glass vials containing odor A (70% isoamyl acetate basis, FCC; Sigma Aldrich) or odor B

((-)-a-Pinene $\geq$ 97%, FCC; Sigma Aldrich), which were both diluted in mineral oil at 5% concentration. Water droplets (~10µl) were released by a 3-way solenoid valve (Lee Company), and licks were detected by using a custom battery-operated circuit board with one end of the circuit connected to the headbar and the other to the lickport. The behavioral rig was controlled with custom written software (MATLAB) and through a data acquisition board (USB-6341: National Instruments).

Behavioral training

After 7-14 days recovering from surgery, mice were handled and began water-restriction to 85% of their original weight before water-restriction. After one day of handling, mice were habituated to being head-fixed above the spherical treadmill (can rotate about one metal axis for 1D locomotion that is recorded) for two days. On the 4th day of training, mice began learning to lick from the lickport as water was automatically delivered at the beginning of the reward period following only non-matched odor trials (AB or BA, with water delivery at time point of 8 seconds). Trials were delivered in blocks of 20 trials. This phase was always 2 days except for the rare mouse that needed one extra day to reach motivation level and lick water from port for at least 50 trials. In the next phase, water was only delivered if the mouse licked during the response period, and mice learned to reliably lick in anticipation of the reward following the 2nd odor. This phase was also 2 or 3 days, dependent on the mouse licking during the response period of at least 50 trials. The final phase was the full delayed non-match-to-sample (DNMS) task in which matched odor trials (AA and BB) were introduced and mice learned to refrain from licking the port following these trials. There was no punishment or timeout following an incorrect lick; the water was simply not delivered. The first day of this final full DNMS task was considered ‘Day 1’ in the axon imaging experiments (6-8 days from the start of water-restriction). A total of 100 trials delivered in five blocks of 20 trials were given each day, and we considered ‘expert performance’ to be any day with performance greater than or equal to 85%. In the CA1 imaging experiments, two-photon calcium imaging only began after the mouse had 2 consecutive days of 'expert performance'. Mice underwent 5-7 days of learning the full DNMS task before recording began.

In-vivo **two-photon imaging**

All two-photon calcium imaging was conducted using a resonant scanning two-photon microscope (Scientifica) fitted with a 16x 0.80 NA objective (Nikon) to record 512x512 pixel frames at 30.9 Hz. CA1 imaging fields of view were 500x500 µm and axonal imaging fields were 250x250 µm. Excitation light was delivered with a Ti:sapphire excitation laser (Chameleon Ultra II, Coherent), operated at 920 nm. GCaMP8f and GCaMP7s fluorescence was recorded with a green channel gallium arsenide photomultiplier tube (GaAsP PMT; Hamamatsu). Microscope control and image acquisition were performed using LabView-based software (SciScan). Imaging and behavioral data were synchronized by recording TTL pulses generated at the onset of each imaging frame and olfactory stimulation digital signals at 1 kHz, using WinEDR software (Strathclyde Electrophysiology Software).

For CA1 imaging experiments, a single field of view (FOV) was imaged for 8 consecutive days of expert performance. Careful attention was given to aligning the FOV to the previous day's as perfectly as possible, and animals were not included in analysis if successful alignment was not possible. We used rotating stages, a motor for adjusting mouse head angle, and a tiltable objective attachment with two degrees of freedom to fine-tune the alignment. For axonal imaging experiments, the same alignment was always attempted for 7 consecutive days of learning, but the extra difficulty of alignment made it not always possible. Therefore, axon segments were not registered between days; however, FOVs were typically very similar. Laser power and PMT settings were kept consistent between days, except for rare occasions when it was necessary to keep similar signal-to-noise. Out of the 16 axonal imaging animals included in analysis (each recorded for 7 days), 7 recording sessions were not included because of poor signal-to-noise.

For each day of recording, imaging was halted between each of the 5 blocks of 20 trials. This allowed fine-tuning of alignment, and it also prevented brain heating or photo-toxicity. Laser power was kept as minimal as possible (60-80mW for CA1, and 100-200mW for EC axons) without sacrificing signal-to-noise ratio, and only mild photo-bleaching was observed in some axonal imaging animals.

Chemogenetic inhibition

All CA1 imaging animals received subcutaneous injections of saline for at least 5 days prior to imaging to habituate them to the injection prior to being head-fixed. For the 8 days of imaging, mice received alternating injections of saline and uPSEM (ultrapotent PSEM 792 hydrochloride binds to PSAM4 to cause strong inhibition). Half of the mice started with saline, and the other half started with uPSEM on the first day of imaging. The uPSEM powder was dissolved into saline at a concentration of 0.3 mg/mL, and injections were administered to achieve 3 mg/kg. After weighing the mouse to calculate the appropriate volume of saline or uPSEM, the mouse was injected intraperitoneally and head-fixed under the microscope. 10-20 minutes elapsed between the injection and the start of behavior.

Histology

Following all experiments, mice were deeply anaesthetized under isoflurane and transcardially perfused with 30 mL 1x PBS followed by 30 mL 4% paraformaldehyde in 1x PBS at a rate of approximately 4 mL/min. After perfusion, the brains were extracted and post-fixed in 4% paraformaldehyde. Sections of 80 µm were collected using a vibratome, 24-48 hours after perfusion. For animals with LEC viral expression, coronal sections were taken, while sagittal sections were taken from animals with MEC viral expression. The sections were mounted onto glass slides and cover-slipped with DAPI mounting medium. Images were acquired on an Apotome2 microscope (Zeiss; 5x, 10x, 20x objectives) to confirm proper expression and location of viral expression. For CA1 imaging experiments, GCaMP8f was confirmed to be in dorsal CA1, and sufficient PSAM4 or mCherry expression was found restricted to either LEC or MEC. In axonal imaging experiments, somatic GCaMP7s was confirmed to be restricted to only LEC or MEC, and axonal expression was found in the SLM layer of dorsal hippocampus. Mice with insufficient PSAM4/mCherry expression or PSAM4/mCherry/GCaMP7s that spread to outside of their desired target were excluded from analysis.

2.6 Quantification and Statistical Analysis

Calcium imaging data pre-processing

For CA1 imaging experiments, the 8 days of recordings were divided into 4 pairs of days, so that each pair consisted of one saline day and one uPSEM day. Both recordings from a single pair were concatenated before processing so that the same neurons could be detected within the pair of imaging days. Concatenated movies were processed using the Python implementation of Suite2P 0.9.2 (Pachitariu et al., 2017) to perform non-rigid motion registration, neuron segmentation, extraction of fluorescence signals, and deconvolution with parameters optimized to our GCaMP8f CA1 recordings. We used the default classifier and an 'iscell' threshold of 0.1 to only include masks that were likely neurons. Neuron masks were then aligned across the 4 pairs of days using CellReg (Sheintuch et al., 2017). Because FOVs themselves were more helpful than the cell masks alone, we modified the CellReg code to do alignment based on the Suite2P registered mean image of the FOV. This yielded excellent registration for all animals with the maximal centroid distance set to 5 µm.

For axonal imaging experiments, the 7 days of recordings were all processed separately. Movies were also processed using Suite2P but with parameters optimized to our GCaMP7s axonal recordings. An additional step of axon merging was taken to decrease the number of duplicates (as an axon could appear as multiple segments within the FOV); this also increased signal-to-noise by increasing the number of pixels for a single mask. By visualizing axon correlation values and their fluorescence traces within the Suite2P, we chose axon segments to merge based on correlation values and footprint distributions. Using custom Python code with functions from Suite2P's source code, we 'merged' axons by generating new ROIs with these new pixels. The old axon segments were then eliminated from analysis and deconvolution was run on the new axon masks.

For all experiments, deconvolved signals were taken as the selected output from Suite2P and taken to MATLAB 2021a for further processing. Deconvolved signals were smoothed by a rolling mean of 10 frames (0.32 seconds), then z-scored, and finally values below 2 were set to zero. The resulting signal was what was used for all analysis and referred to as 'firing rate (STD)' as a proxy for spiking activity. Signals were aligned to the trial structure (odor presentations, reward period, lick timing) and the recorded locomotion as mice ran on the spherical ball.

BTSP event detection and analysis

First, 6-second periods were extracted for each odor presentation period (2 seconds before odor and 3 seconds after) and divided for Odor A and Odor B regardless of whether it was the first or second odor presented in the trial. Since each recording had 5 blocks of 20 trials, we have 100 odor presentations of each odor per cell per day. Next, we identified each 'event'; which we define as a group of consecutive timepoints with a non-zero deconvolved signal. The size and timing of that event is counted as the peak value within the event and that timepoint's time relative to the odor, respectively.

Next, we identified which events satisfied criteria to be considered as a possible induction event. This detection was performed separately for Odor A and Odor B presentations. Events in the first 10 or last 10 odor presentations were not considered for analysis because we needed enough odor presentations before and after the event to detect BTSP events. There were two criteria for an event to be considered a possible induction event: during the previous 3 odor presentations the cell must show no activity within 2 seconds before or after the event in question and there must not be a significant peak firing field. To determine the significance of a firing field, we took 6-second periods of all previous odor presentations and found the peak of the average activity. We then circularly shuffled each odor presentation and found the peak of the average activity from this shuffle data. This was repeated 2000 times to generate 2000 peak values from shuffle data. For a possible induction event, the real peak of average activity must not have been greater than the 90th percentile of the shuffle.

If an event passed criteria to be considered as a possible induction event, we analyzed if it is successful in forming a field. There were four criteria for a successful field formation: 1. The resulting field must have been significant above the 95th percentile of the shuffle; 2. The resulting field occurred within 2 seconds of the peak of the induction event; 3. The cell must have fired (have value above 2 STD) within 0.5 seconds of the resulting field for the next 3 odor presentations; 4. The cell must have fired within 0.5 seconds of the resulting field for at least 7 out of the next 10 odor presentations. All these criteria were decided by visually inspecting all successful events and improving based on our expectations of how BTSP events would look in our task. The strict criteria for activity in the previous 3, following 3, and following 10 odor presentations improves the likelihood that the event in question does induce the resulting field. The ± 2 second window for the difference between the event peak and field peak allowed us to look for backward drift without any bias. The lack of any criteria regarding the amplitude of the induction event allowed us to probe the relationship of amplitude to success rate and backward drift. Success rate increases continuously with amplitude (Figure 2.1 I), but only events with amplitude above 10 STD had statistically significant backward drift. Therefore, we considered any event above 10 STD to be ‘plateau-like’, and successful ‘plateau-like’ events are what we considered to be BTSP events. We considered any event between 2 and 10 STD to be a ‘small event’.

Locomotion analysis

1D locomotion that was recorded with a sensor (Avago ADNS-9500) at 1kHz was binned to match the frame rate of calcium imaging. Binned signals were smoothed by a rolling mean of 10 frames (0.32 seconds), then z-scored, and finally values below 1 were set to zero. These binned signals are displayed as 'locomotion (a.u.)'. Since most of the locomotion was small movements around the onset and offset of odors, in other analysis we binarized locomotion into 'not running' and 'running' bouts. A bout of running must have been at least 2 seconds of locomotion values above 1; and all other periods were considered to be 'not running'.

Selectivity analysis

We calculated the odor selectivity index value for each ROI as: $SI = (R_a - R_b) / (R_a + R_b)$; where R_a is the firing rate at a given bin for Odor A trials and R_b is the same for Odor B trials. The same approach was taken for selectivity of match trials versus non-match trials. Bin sizes were always 0.5 seconds, and performance was never considered so all trials are included. For each ROI, a distribution of 2000 shuffled index values were also calculated by randomly shuffling the trial type assignment 2000 times for each bin. The maximal absolute value index is chosen from all the bins (for the real ROI and all 2000 shuffles), and the bin is noted. ROI's with an absolute value index value above the 95^{th} percentile of absolute value shuffled index values are considered to be 'significantly selective'.

Support vector machine decoding

Binary support vector machine (SVM) decoding was performed in MATLAB 2021a (default parameters) using bin sizes of 0.5 seconds (averaging the deconvolved signal for those frames within the bin). Unless otherwise noted, the number of ROIs was controlled by randomly subsampling 100 ROIs out of all possible ROIs. This 100 was chosen as it is the largest multiple of 50 that is smaller than the number of ROIs in each recording (CA1 and EC axons). In all cases, the result of 20 subsamples of ROIs were averaged for each data point. For each bin and subsample, 80% of trials were used for training the decoder, and the remaining 20% were used for testing. This was repeated 4 more times so that each block of 20 trials was used as the 20% for testing. For each training of the decoder, another training was done with a shuffled assignment of trial type to confirm a shuffle comparison of data yields decoder accuracy of ∼50%. For odor decoding, the trials were broken down into odor presentations (same as in BTSP detection analysis) to evaluate odor decoder accuracy regardless of the order of the odors. When specific timepoints were mentioned, such as ‘during odor presentation’ or ‘during delay period’, the average accuracies of the 0.5 second bins were averaged and not trained/tested with larger bins.

To evaluate the relationship of the number of subsampled ROIs and decoder accuracy, all the previous steps were repeated using different numbers of subsampled ROIs. Again 20 subsamples for each were used. If a recording session had fewer than the chosen number of ROIs, all ROIs were used. For axonal decoding in Figure 2.11 I-J we pooled axons from the different days of the same animal only in panels showing ‘number of subsampled axons’ on the x-axis. This was done simply to illustrate improved decoder performance with many more ROIs, but all other decoding figure panels were done by subsampling 100 axons and treating each recording session separately. For CA1 decoding in Figure 2.11 D, most recording sessions had more than 300 neurons, so no pooling of days was necessary.

Sequence-axon detection and analysis

To evaluate peak firing timing in EC axons, we performed sequence-axon detection similar to the previously described approach in CA1 neurons in our DNMS task, Taxidis et al. (Taxidis et al., 2020). First, trials that begin with Odor A and those that begin with Odor B are separated, and the one with a larger peak of the average activity was considered further. Additionally, only the 6-second period including first odor presentation and the delay period was considered. In the same way as described in BTSP-event detection, the peak of average activity within this period and a given trial type was determined to be significant if the peak was greater than the 95th percentile of 2000 circular shuffles. The cell must also have had a trial reliability of at least 20% and have fired above 2 STD for 20% of the preferred trials within 0.5 seconds of the peak frame found in the previous step. If an ROI passed both criteria, it was considered to be a 'sequence-axon' regardless of its odor selectivity, as that was a separate analysis. An ROI was considered to have a peak during the odor presentation if the peak was within the odor presentation period. An ROI was considered to have a peak during the odor offset (sometimes referred to as immediately after the odor) if the peak was in the first second of the delay period. An ROI was considered to have a peak during the delay period only if the peak was during the last 4 seconds of the seconds of the delay. This was done to not include the large population of ROIs that fired to the offset of the odor (likely the auditory cue of the clicking of the valve).

Analysis across days

For CA1 imaging, CellReg output registration maps were used to align cells across pairs. Within each pair saline and uPSEM days had the exact same cell indices. For decoding across days (Figure 2.4 C-D), the same binary SVM decoding was performed on subsamples of 100 neurons that overlap between the 2 days in question. Binning and all parameters were identical, with the exception that 100% of trials from the training day were used for training the decoding and 100% of trials from testing day were used for testing the decoder. To evaluate the percentage of overlapping neurons that remained odor selective (Figure 2.4 E), we used a threshold of the 90th percentile for odor selectivity. If a neuron was found to have had a selectivity value above the 90th percentile and preferred the same odor in both days in question, then it was considered to 'remain odor selective'. All axonal analysis was performed separately for each recording session as alignment across days was difficult to achieve for each animal.

Statistical analysis

For CA1 imaging figures that show paired points, a single line connects the two days within a pair, so there are 4 times as many lines as animals. However, all statistics were performed as a Two-Way ANOVA (animal x pair) with repeated measures on the saline-uPSEM condition (using 'fitrm' and 'ranova' functions in MATLAB 2021), so as not to treat each pair as entirely independent. For non-paired points (Figure 2.4 D-E), Two-Way ANOVA (animal x pair) were performed. For when CA1 imaging groups were not compared (Figure 2.1 I-L), all 17 animals were treated independently, and statistics were one-sample t-tests (Figure 2.1 K). The Wilcoxon signed rank test was conducted using pairings of all cells in CA1 recordings to evaluate the change of the distribution of selectivity values.

For axonal imaging figures, significance was determined on each day by two-sample t-tests for each day. ANOVA p-values are reported in the text as the overall significance using a Two-Way ANOVA (animal and day). For correlations with performance, Pearson's R was calculated with its corresponding p-value.

For all figures, no asterisks were shown if $p \geq 0.05$, 1 asterisk if $p < 0.05$, 2 asterisks if $p < 0.01$, 3 asterisks if $p < 0.001$. If the $p \geq 0.1$, 'n.s.' is displayed, but if $p < 0.05$ the p-value was typically displayed in the figure. On occasions when single asterisks were displayed above a curve or trace, p-values were corrected for multiple comparisons using the false discovery rate Benjamini-Hochberg procedure. In all cases in the text, values were written in the format 'mean ± standard deviation' (STD), while error bars in all figures show the mean and standard error of the mean (SEM). No statistical methods were used to determine appropriate sample sizes but were chosen as being comparable to sizes used in similar publications.

2.7 Supplemental Figures

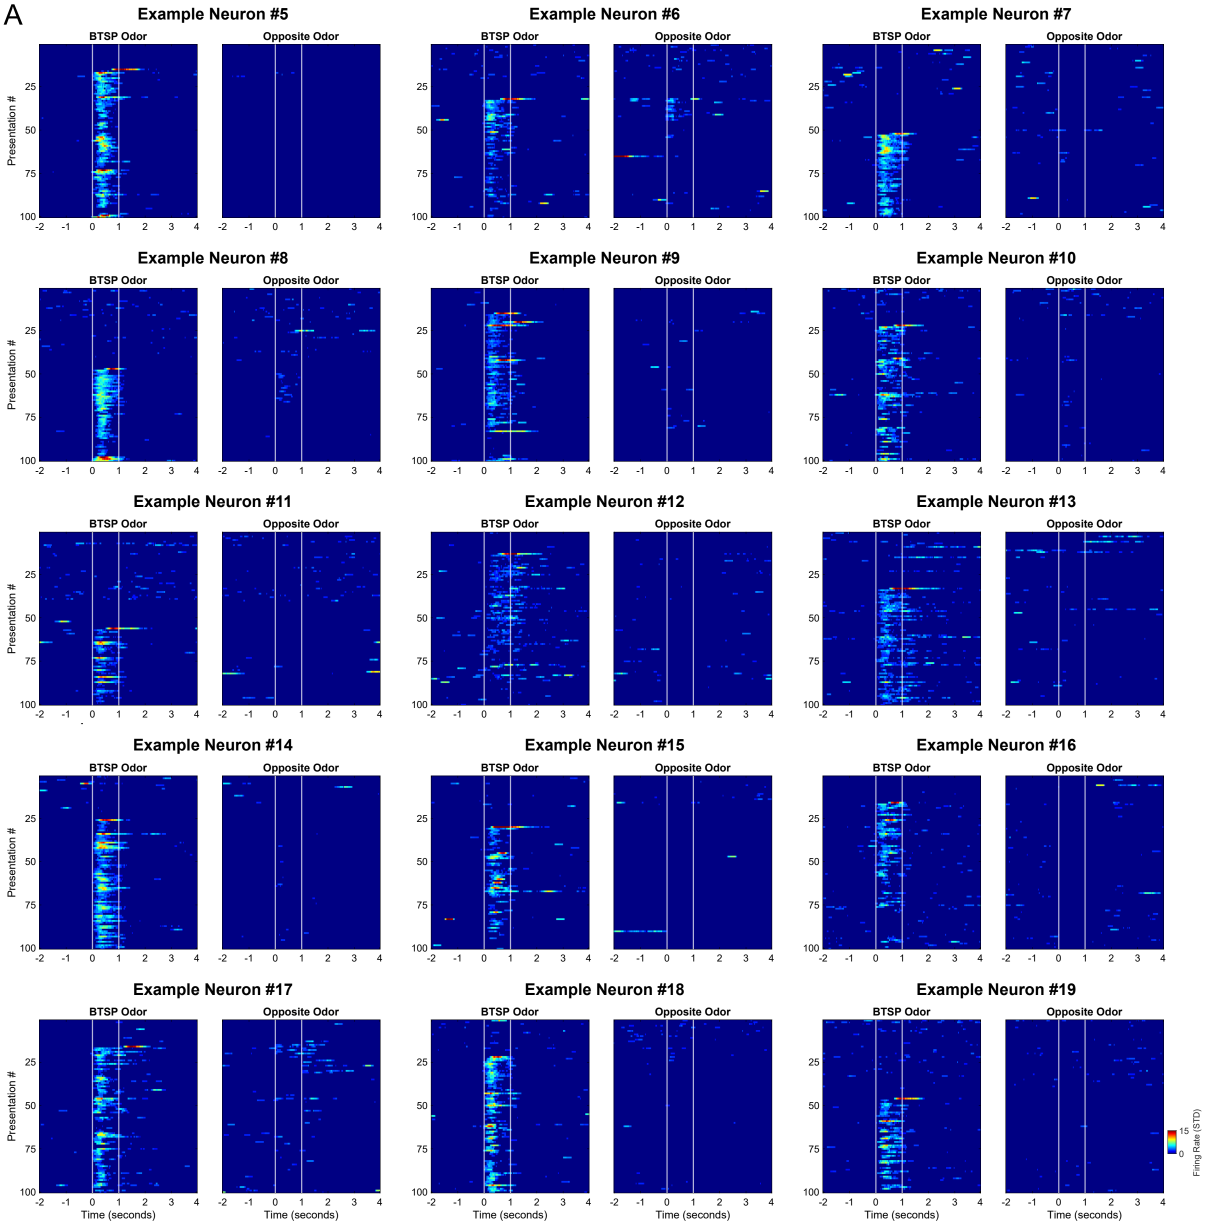

Figure 2.5: **Additional examples of BTSP-like events**
Supplementing Figure 2.1 H. **A)** 15 more example neurons showing BTSP-like events.

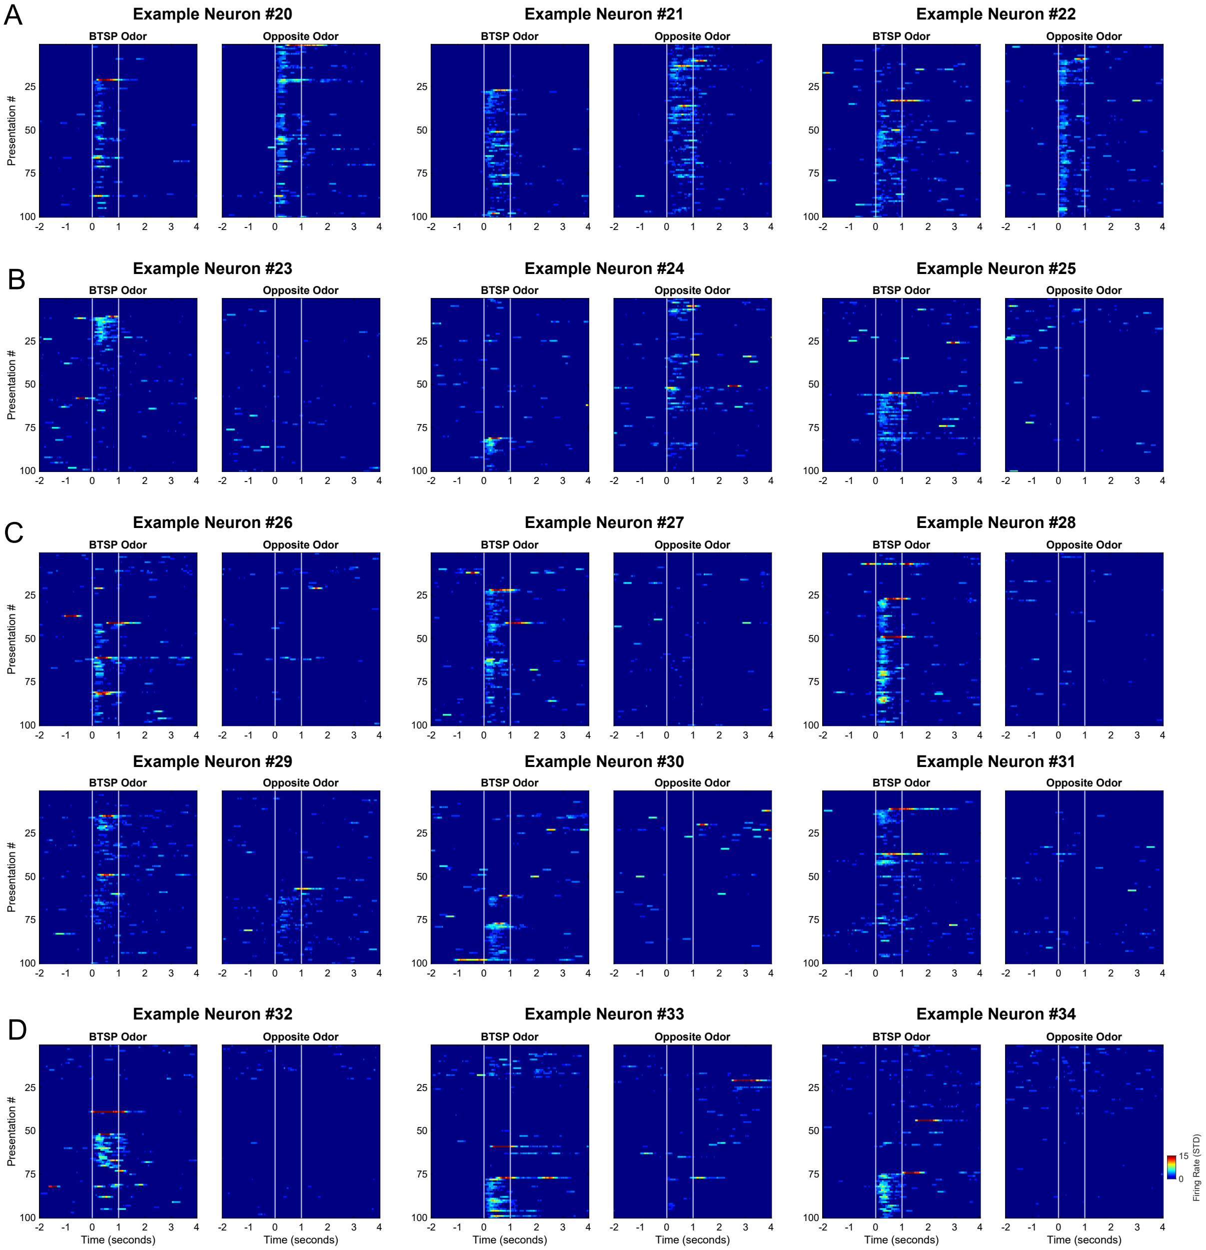

Figure 2.6: **More unique examples of BTSP-like events**

Supplementing Figure 2.1 H. **A)** 3 examples of an odor-selective cell becoming non-selective because the new BTSP induced field. **B)** 3 examples of cells that formed an odor-field after a 'plateau-like' event, but the field faded quickly. **C)** 6 examples of cells with multiple 'plateau-like' events that seem to reinforce the odor-field. **D)** 3 examples of failed 'plateau-like' events followed by successful ones.

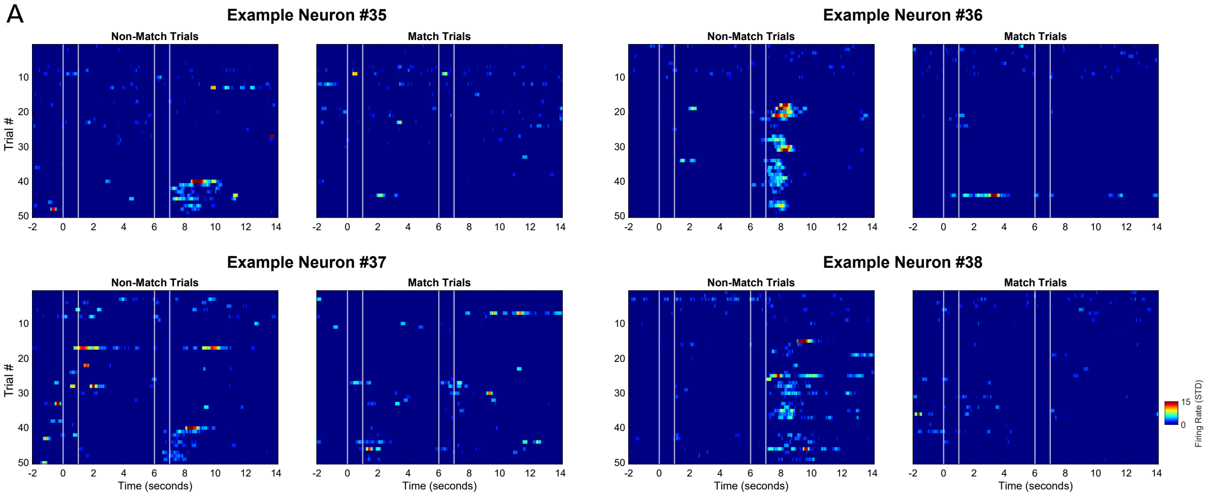

Figure 2.7: **Examples of BTSP-like events that formed reward encoding fields** Supplementing Figure 2.1 H. **A)** 4 examples of ‘plateau-like’ events forming representations following the 2nd odor and selective to non-match trials.

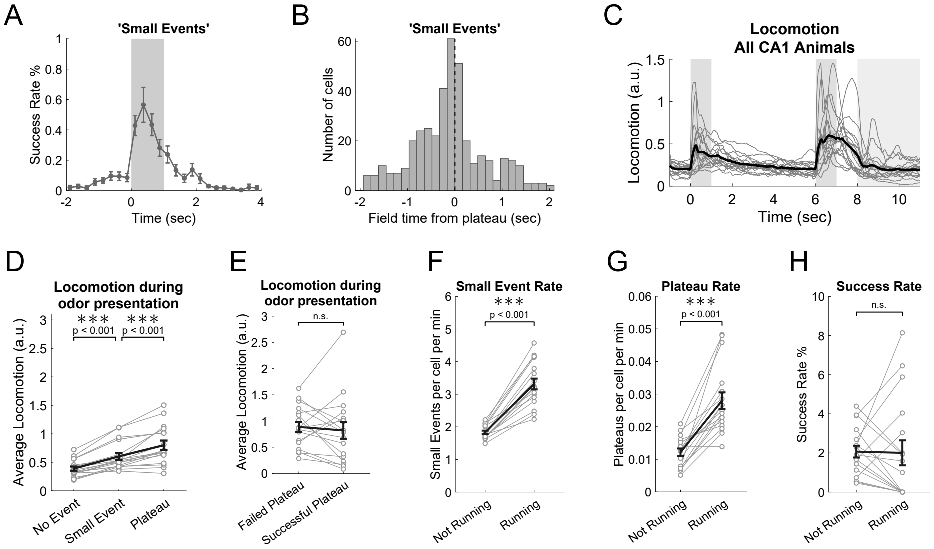

Figure 2.8: **Small events were not BTSP, and locomotion drove 'plateau-like' events but not successful field generation**

Supplementing Figure 2.1 I-L. **A)** Supplementing Figure 2.1 J showing success rate was very low for small events. These represent randomness of non-BTSP events passing criteria for BTSP. **B)** Supplementing Figure 2.1 L showing that small events did not have backward drift that was seen for 'plateau-like' events. **C)** Average locomotion for all 17 CA1 animals. Locomotion was normalized voltage signal from sensor (arbitrary units). **D)** Locomotion values during odor presentations split by when there is no calcium event, a small event, and 'plateau-like' event (data points are each animal, averaging all cells). **E)** Same as (D) but split by when a 'plateau-like' event failed or succeeded in generating a field (paired t-tests). **F)** Comparing small event rate, after binarizing locomotion values to find 2 second bouts of 'running' (paired t-tests). **G)** Same as (F) but for 'plateau-like' event rate. **H)** Similar to (F) and (G) but for success rate of 'plateau-like' events forming an odor-field.

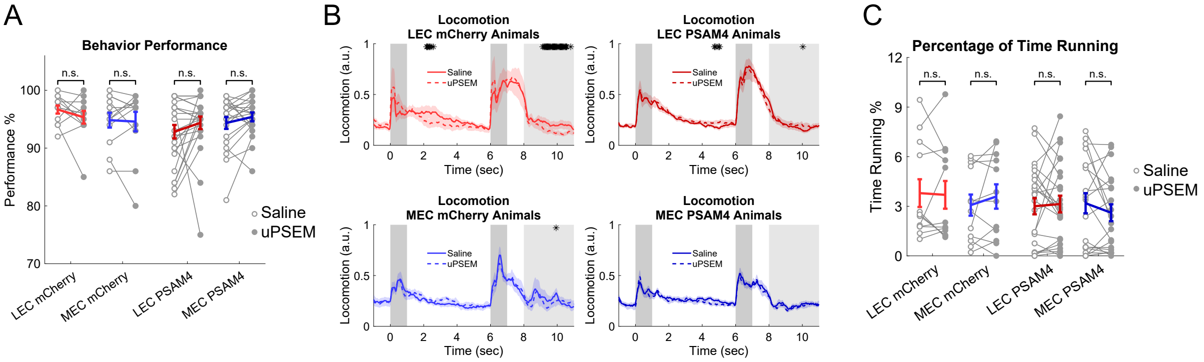

Figure 2.9: **LEC or MEC inhibition had no effect on behavioral performance, locomotion, or running**

Supplementing Figure 2.2 B-C. **A)** Average behavioral performance was unaffected by uPSEM. **B)** Locomotion was also largely unaffected by uPSEM for all groups of animals. There were some minor significant bins, but these were random and cannot explain any other findings.

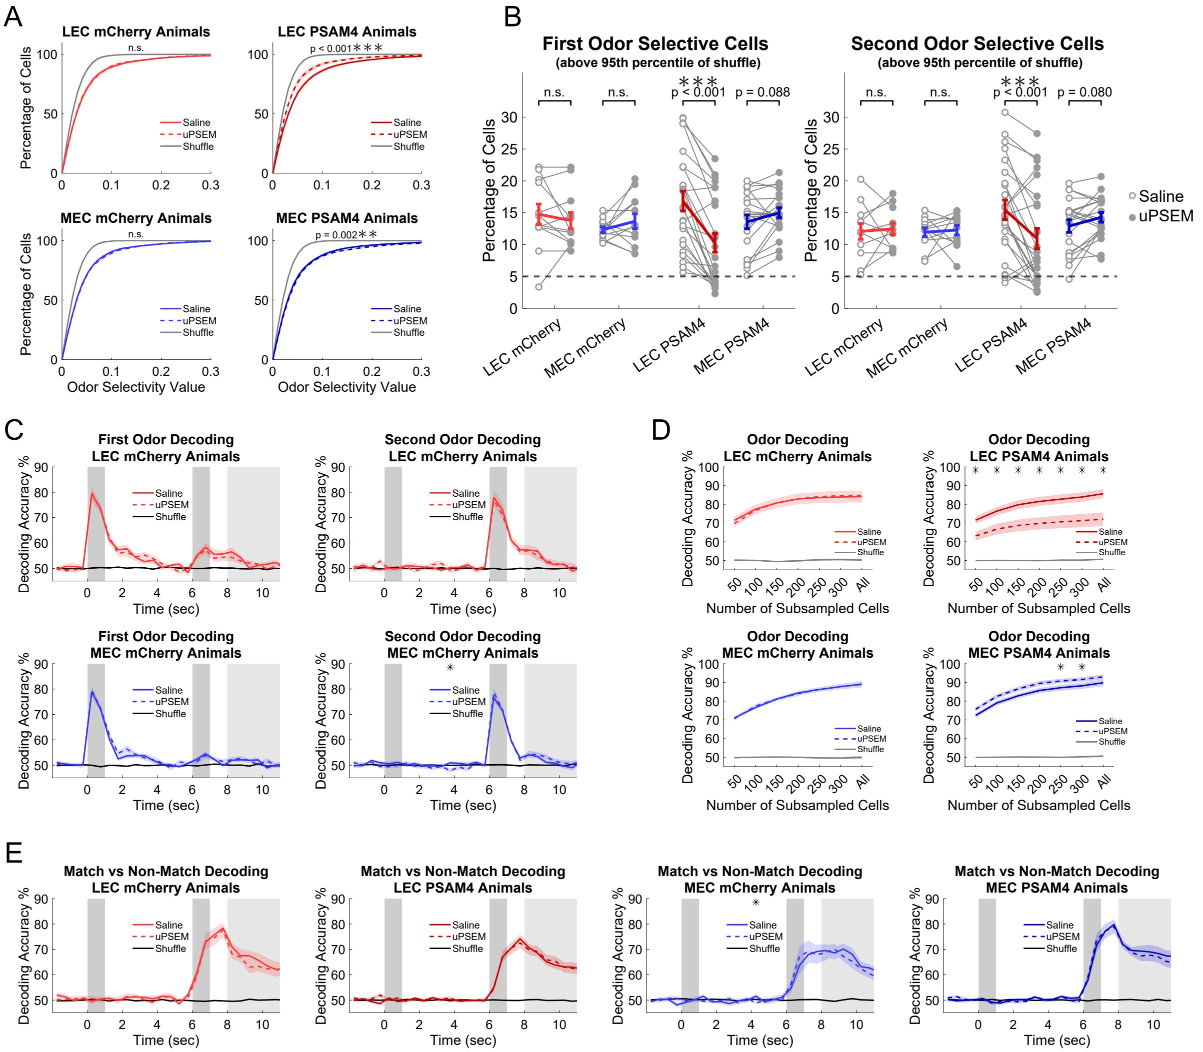

Figure 2.10: **uPSEM inhibition does not affect odor decoding of mCherry controls or match versus nonmatch trial decoding for any groups** Supplementing Figure 2.2. **A)** Cumulative distributions of odor selectivity values for all 4 groups of animals compared to shuffle distributions. Statistics are Wilcoxon signed-rank test for all neurons. **B)** Same as Figure 2.2 D but split between first odor and second odor (cells could belong in both groups of significant neurons). **C)** Panels complementing those in Figure 2.2 E, but for mCherry control animals. **D)** Odor decoding performance increased with the number of neurons subsampled. The number of 100 neurons was chosen for all main figures because all recordings have at least 100 neurons or axons. Statistics are the same as (C). **E)** SVM decoding of match trials versus non-match trials was unaffected by uPSEM inhibition.

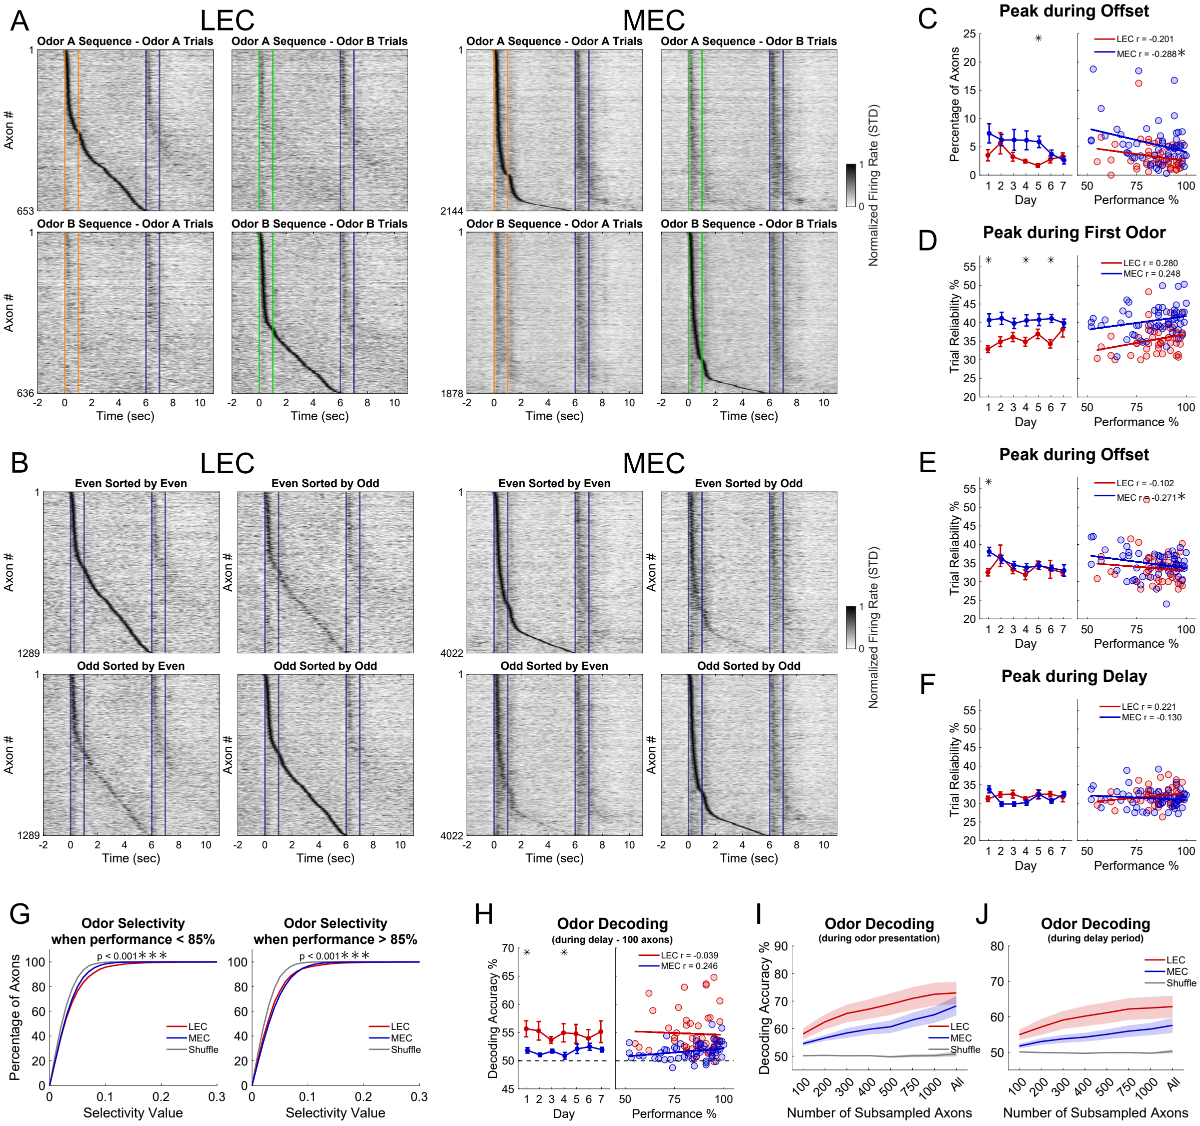

Figure 2.11: **LEC and MEC sequences; odor axon reliability greater for MEC** Supplementing Figure 2.3. **A)** LEC and MEC sequences for expert performance split by Odor A and Odor B to show the similarities and highlight odor specificity differences. **B)** Even/odd trial validation for LEC and MEC preferred odor sequences. **C)** Percentage of axons with a significant peak during the first odor ‘offset’ (first second of delay period). **D-F)** Trial reliability (percentage of trials with a calcium event at its field on the preferred odor trials) for axons that peak during the first odor, first odor offset, or delay period. **G)** Cumulative distributions of odor selectivity values during sub-expert and expert performance. Statistics are two-sample Kolmogorov-Smirnov test for all axons. **H)** Binary SVM odor decoding only during the delay period. **I)** Odor decoding (during odor presentation period) accuracy increased with more axons subsampled (similar to Figure 2.10 D. To achieve larger number of axons, axons were pooled for each animal. **J)** Same as I, but only decoding during the delay period.

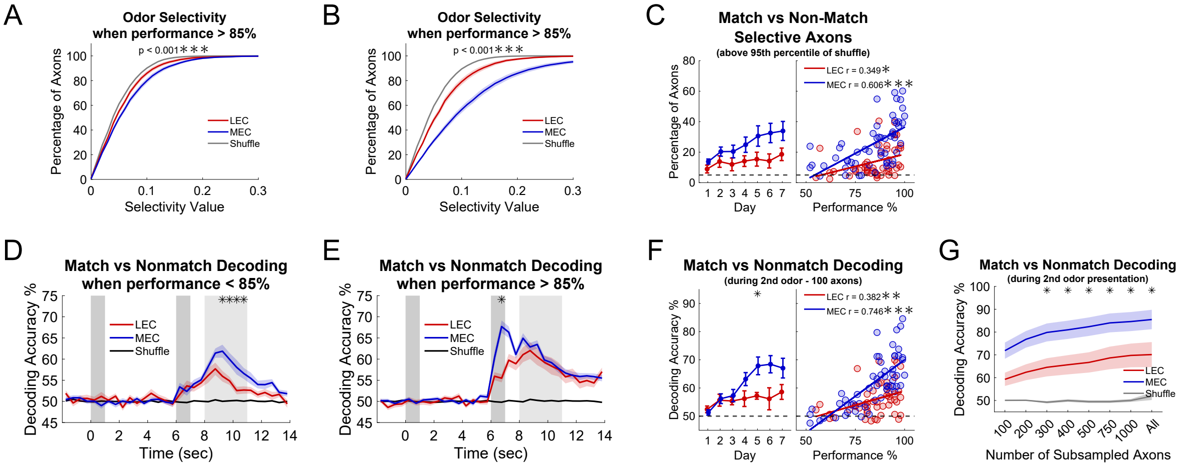

Figure 2.12: **MEC encodes match vs nonmatch trial more strongly even before the reward period**

Supplementing Figure 2.3. **A)** Cumulative distributions of match/non-match selectivity during sub-expert performance. **B)** Same as (A) but during expert performance. **C)** Percentage of axons with a match/non-match selectivity value above 95th percentile of shuffle, split by day and correlation with performance. **D)** SVM decoding of match trials versus non-match trials was greater in MEC later in the reward period during sub-expert performance (likely encoding outcome of receiving reward or not). **E)** Decoding of match trials versus non-match trials increased and peaked before reward period during expert performance. **F)** Decoding of match trials versus non-match trials during the 2nd odor period increased more rapidly for MEC. **G)** Same decoding but during expert performance as number of subsampled axons increased.

CHAPTER 3

Sequential CA1 Activity Represents Working Memory and Time during a Novel Differential-Delay Non-Match-to-Sample (dDNMS) Task

3.1 Abstract

Working memory (WM) and timing are considered distinct cognitive functions, yet the neural signatures underlying both are similar. Recent proposals have suggested that neural sequential activity is a possible shared mechanism of WM and timing. To answer if WM and timing rely on a shared neural signature, we developed a novel rodent differential-delayed nonmatch-to-sample (dDNMS) task, in which the identity of the first odor stimulus predicts the delay duration. The cued-differential delays are irrelevant to the WM task, yet we found that WM performance decreased if delay expectations were violated with reversed delay lengths. By performing calcium imaging of dorsal CA1 neurons, we found odor-specific sequential activity tiling the short and long delays. While encoding the identity of the first odor, the trajectories and shapes of the sequences also reflected an expectation or anticipation of the timing of the 2$^{\text{nd}}$ odor. Together, our findings suggest that working memory and implicit timing representations are multiplexed in dorsal CA1.

3.2 Introduction

Working memory (WM) is the ability to temporarily store and manipulate information, and it is essential for cognition (Baddeley, 1992, 2012). In many cases, reading out the WM trace requires appropriate timing such that the additional stimuli are appropriately anticipated, and the behavioral response is made at the right time. WM and timing are often studied separately because they are thought to have distinct neural mechanisms in different brain regions. However, many similar features of WM and timing suggest that they could have shared neural mechanisms. Both working memory and timing are functions that operate on the scale of seconds and require information to be transiently stored. In traditional delay WM tasks, WM during the delay period is necessary to maintain retrospective information and prepare for potential prospective information. Timing also keeps track of the time elapsed since the first stimulus was delivered (retrospective information) and how much time needs to elapse until a second stimulus arrives (prospective information). Recent computational research also has suggested that WM and timing could be implemented by the same circuit mechanism (Rajan et al., 2016; Cueva et al., 2020; Zhou et al., 2023). However, *in-vivo* evidence of a shared mechanism for both working memory and timing has yet to be demonstrated in a behavioral task with both WM and timing components.

Early pioneering studies of WM suggested that steady-state persistent activity is the dominant neural mechanism for how the brain maintains information across a delay period (Fuster and Alexander, 1971; Kubota and Niki, 1971; Funahashi et al., 1989). Since then, however, other WM mechanisms have been reported: sequential activity (Pastalkova et al., 2008; Rajan et al., 2016; Taxidis et al., 2020) and ramping activity (Inagaki et al., 2018, 2019; Cueva et al., 2020). Critically, these same time-varying patterns of neural activity have also been reported for the encoding of time (Narayanan, 2016; Zhou et al., 2020; Cueva et al., 2020; Ma et al., 2024).

However, it has been difficult to link the mechanisms of WM and timing because neural dynamics have been recorded in distinct tasks designed to elucidate mechanisms of either

WM or timing. To make this link, a WM task is needed that allows the researcher to probe if subjects are learning the temporal structure of the task itself (Cueva et al., 2020). While most timing research has focused on explicit timing tasks where subjects must keep track of time correctly to perform the task; subjects can also perform implicit timing where they learn the task-irrelevant temporal structure of the task. This implicit timing can be learned during WM tasks with fixed delay lengths, during which subjects can learn to predict when the delay is ending and another stimulus is arriving. Recent human research has shown that WM performance can be impaired when temporal expectations are violated (Cravo et al., 2011; Vangkilde et al., 2013; Zhou et al., 2023). This suggests an interaction between WM and implicit timing. Importantly, computational studies have demonstrated how time-varying neural trajectories can encode both WM and timing information (Liu and Buonomano, 2009; Murray and Escola, 2017; Zhou et al., 2023); yet direct experimental neural evidence for a shared mechanism for WM and timing has been lacking.

To find a shared neural mechanism of WM and timing, we developed a novel olfactory WM and implicit timing task for rodents. This modification of a rodent olfactory working memory task (Taxidis et al., 2020) allowed us to observe behavioral deficits in WM when implicit timing expectations were violated, replicating a previously found human behavioral result (Zhou et al., 2023). We have previously found that during performance of an olfactory delayed non-match-to-sample (DNMS) WM task, CA1 pyramidal neurons fire sequentially across a fixed 5-second delay following specific odors (Taxidis et al., 2020). We hypothesized that these sequences would be modulated by temporal expectation in our novel differential delayed non-match-to-sample (dDNMS) WM task where each odor is followed by a delay period of distinct duration (2.5 or 5.0 seconds). We performed two-photon calcium imaging of CA1 pyramidal neurons as mice performed the task. We found that CA1 odor-specific sequential activity is shaped differently when animals are expecting different delay lengths with an over-representation of CA1 neuronal activity prior to the expected arrival of the 2^{nd} odor, suggesting that implicit timing is encoded within these sequences. Neural trajectories sped up in anticipation of the 2nd odor following a short delay, and this corresponded with better time encoding on short delay trials. Altogether, we demonstrate that sequential activity in CA1 is a shared mechanism of working memory and implicit timing within a novel task that shows how implicit timing violations can impair working memory performance.

3.3 Results

We developed a novel working memory and implicit timing task for rodents, by modifying a standard olfactory DNMS working memory task (Taxidis et al., 2020; Liu et al., 2014). In this dDNMS task, either odor A or odor B is presented for 1 second followed by a delay period whose duration depends on the identity of the first odor. If the first odor presented is odor A, the reward period is 2.5 seconds. If the first odor presented is odor B, the reward period is 5 seconds. After the delay period, either odor A or B is presented as the second odor. If the odors delivered before and after the delay are the same, the animal learns to refrain from licking. If the odors do not match, the animal learns to lick the reward port for a water reward. Adult male and female mice (n = 23) were water-restricted and trained to perform the dDNMS task (Figure 3.1 A). After the animals learned the task to criterion (see Methods), we introduced 'reverse' trials in which the cue-delay contingency was reversed for 20% of all trials (Figure 3.1 A-B) to investigate whether mice implicitly learned the temporal structure of the task and whether unexpected delay lengths would alter WM performance. Animals performed the task with reverse trials for 2 days, and each day consisted of 7 blocks of 20 trials with pseudorandomly distributed odor combinations.

3.3.1 Novel rodent dDNMS working memory task showed behavioral evidence of implicit timing

To investigate whether mice showed behavioral evidence of learning the implicit temporal structure of the dDNMS task, we compared performance accuracy and lick latency between standard and reverse trials on the first 2 days that reverse trials were introduced. Mice showed decreased performance accuracy on reverse trials (87.00 ± 10.42 % on standard trials and 83.77 ± 12.10 % on reverse trials, $p = 0.0039$, Figure 3.1 C). However, this effect was driven entirely by reverse long trials (86.49 ± 9.47 % on standard long trials and 80.75 ± 14.19 % on reverse long cued-short trials, $p = 0.0009$), as the reverse short trials showed no effect on performance accuracy (87.50 ± 11.88 % on standard short trials and 86.80 ± 12.25 % on reverse short cued-long trials, $p = 0.70$). Therefore, performance accuracy was only impaired when mice expected a short delay and had to maintain the memory of the first odor for longer than anticipated. Additionally, by measuring the lick latency on each trial when mice made a choice to lick, we found that mice took a longer time to respond on reverse trials (563 ± 243 milliseconds on standard trials and 584 ± 252 milliseconds on reverse trials, $p = 0.037$, Figure 3.1 D). However, this effect was driven equally strongly by both the reverse short and reverse long trials. Together, the results suggest that mice learned the temporal structure of the differential delays. Violating their temporal expectations with reverse trials demonstrated behavioral results similar to those obtained in humans performing a comparable task (Zhou et al., 2023).

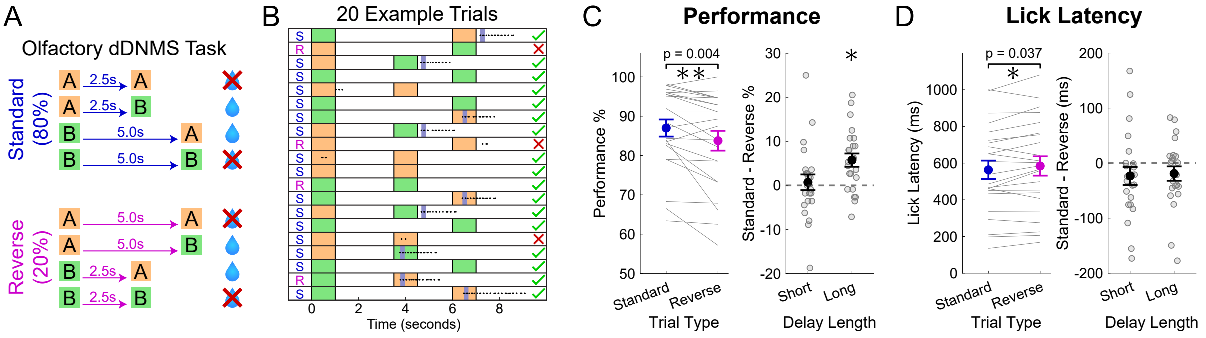

Figure 3.1: **Novel rodent dDNMS working memory task showed behavioral evidence of implicit timing**
A) Schematic of the olfactory differentially delayed non-match-to-sample (DNMS) task. Water delivery and licking behavior was assessed during a 3 second reward period starting at 2nd odor onset. Mice were head-fixed above a Styrofoam spherical ball to allow running. **B)** Example block of perfect performance for 20 trials. S and R letters on the left indicate standard or reverse trials. Dots indicate licks and dark blue bars indicate water delivery. **C)** Behavioral performance accuracy calculated as percentage of trials correct. Large dots represent the mean, and error bars indicate standard error of mean (SEM). All statistics are paired t-tests. The left panel data was split by delay length for the right panel and difference between standard and reverse trials is displayed. The asterisk above long delay in right panel indicates significance ($p = 0.0009$). **D)** Same as (C), but for median lick latency after onset of the 2nd odor.

3.3.2 Differential sequential activity in CA1 reflects implicit timing during delay period

To determine if there were shared neural mechanisms of WM and implicit timing representations, we performed *in-vivo* two-photon calcium imaging of the dorsal CA1 region of the hippocampus as mice performed the dDNMS task (Figure 3.2 A-B). In our previous work using a standard olfactory DNMS task with fixed 5 second delays, we found that a population of dorsal CA1 neurons fired during specific epochs of the task (Taxidis et al., 2020). While some cells fired during the presentation of specific odors, others fired during specific points during the delay period following a specific odor yielding odor-specific sequential activity tiling the entire delay period (Taxidis et al., 2020). Here, we recorded dorsal CA1 neural activity to investigate how differential temporal expectations influence hippocampal sequential activity.

Adult male and female mice (n=11) were injected with AAV1-Syn-jGCaMP8f or jGCaMP7f into the right dorsal CA1 and implanted with a 3mm diameter glass-bottomed titanium cannula above the intact alveus after aspiration of the overlying cortex and corpus callosum (Taxidis et al., 2020). Neural activity was recorded for each mouse on its last two days of receiving only standard trials and the first two days with the inclusion of reverse trials, yielding 44 recording sessions with an average of 405 ± 187 active neurons per day (Figure 3.2 A-B). A proportion of CA1 neurons responded reliably during different epochs of the trial. Some had responses with significant activity peaks during the odor presentation, while others had peaks during the delay period (Figure 3.2 C). Altogether, the population of CA1 neurons formed sequential activity that tiled the first odor presentation and delay period. One group of neurons with significant activity peaks during or after Odor A formed the short-trial sequence, while another group with peaks during or after Odor B formed the long-trial sequence. Out of the total population of active neurons, 13.9 ± 12.1 % had significant peaks and were assigned short trial (odor A) or long trial (odor B) sequences.

To visualize the sequential activity, we pooled all neurons with significant peaks during the first two days with reverse trials. After grouping short trial neurons and long trial neurons separately and then sorting them based on the peak firing, we observed that cells tiled the entire delay period (left panels of Figure 3.2 D). These sequences were strongly odor-selective, as most cells were quiet during or after the presentation of their non-preferred odors (middle panels of Figure 3.2 D). On reverse trials, the long trial sequence was cut short, while the short-trial sequence failed to reach the second odor (right panels of Figure 3.2 D).

Visually, we noticed an over-representation of neurons with peak firing at the end of the short delay, and a modest over-representation at the end of the long delay. To quantify if differences in expected delay duration influenced the shape of sequential activity, we compared the shapes of the standard short trial sequence and standard long trial sequence only during the first 3.5 seconds of the trials (first odor and short delay length). This allowed a fair comparison of equivalent time lengths to determine if the hippocampus allocates additional resources to the anticipation of the second odor arrival in the short trial sequence. The distribution of these pooled significant neurons was different between the short and long trial sequences (Kolmogorov-Smirnov test $p = 4\text{x}10^{-6}$). This effect was driven by the large over-representation in the last 0.5 seconds of the short delay period (3.0 to 3.5 seconds trial time) for the short trial sequence neurons, with 13.9 ± 12.1 % of short trial neurons with a peak during the last 0.5 seconds compared to only 13.9 ± 12.1 % for long trial sequence neurons (out of only neurons with a peak during the first 3.5 second period, $n = 11$, paired t-test $p = 0.0315$). Together, these findings suggest that odor-specific sequential activity in CA1 reflected temporal expectations as sequence shapes were different between short and long delay lengths.

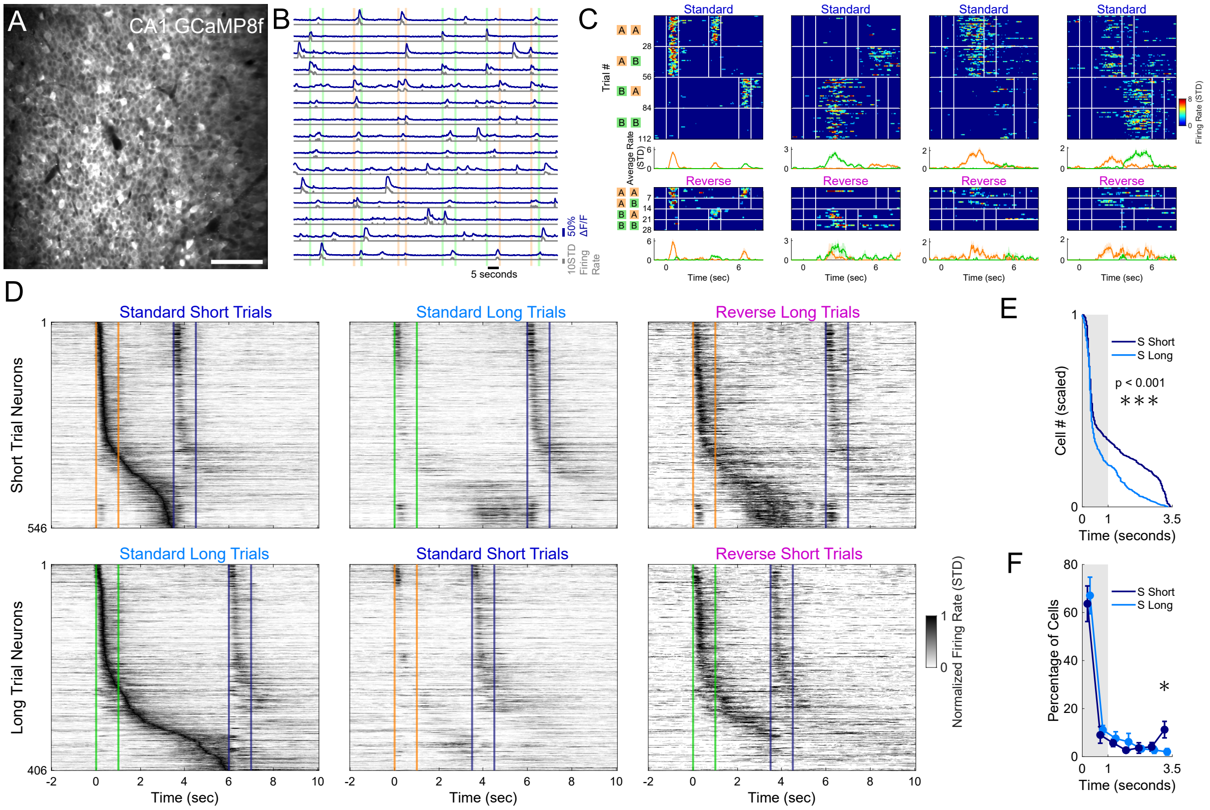

Figure 3.2: **Differential sequential activity in CA1 reflects implicit timing during delay period**

A) Example field of view of dorsal CA1 neurons expressing GCaMP8f. Scale bar is 100µm. **B)** Example activity traces of 15 neurons across 6 standard trials. Dark blue traces indicate the ΔF/F, while the gray trace represents the normalized and floored deconvolved signal used to estimate firing rate. Vertical colored bars represent the odor delivery periods. **C)** Heatmaps of 4 example neurons from the first day including reverse trials. Heatmaps show deconvolved signal on each trial that was grouped into trial type. Vertical white lines indicate onset and offset of each odor delivery. Average traces at bottom show difference in average firing rate split by trials that started with Odor A and those that started with Odor B. **D)** Sequential activity of only neurons that had a significant peak of firing during the first odor presentation or delay period. Each row is the average trace of selected trials normalized and sorted based on its peak firing rate on preferred trials (left most panels). The top three panels show neurons that have a significant peak during odor A or the delay period following odor A (short trial neurons). The bottom three panels are the same but for neurons with a significant peak during or after odor B (long trial neurons). Orange and green lines indicate the first odor onset and offset, while blue lines indicate second odor onset and offset (which could be either odor). **E)** Distribution of significant peaks were different between short trial and long trial neurons (pooling all cells, Kolmogorov-Smirnov test, $p = 4x10^{-6}$). **F)** Location of significant peaks into 0.5 second bins ($n = 11$, repeated paired t-tests corrected for multiple comparisons with Benjamini-Hochberg procedure, asterisk indicates $p = 0.0315$).

3.3.3 Neural trajectories increase in speed in anticipation of second odor

To further evaluate the differences between CA1 population activity between short and long delays, we evaluated neural trajectory dynamics of all significant short-trial or long-trial sequence neurons. First, we performed principal component analysis (PCA) and visualized the first two principal components (PCs) of population activity during the first odor and delay period. By splitting the neural trajectories based on the trial type, we revealed differences in speed of change and distance from a baseline resting state before the presentation of the first odor (Figure 3.3 A). Following a dramatic change of population activity during the odor presentation, the neural trajectory slowed and returned toward the baseline starting point. However, on standard short trials, the trajectory turned away from the baseline and increased its speed as if in anticipation of the second odor (Figure 3.3 B and D). When comparing the trajectories between standard short and standard long trials, the difference was significant over 1 second before the arrival of the 2nd odor on the short trials. On reverse-long trials when the second odor comes later than expected, the trajectory slowed and returned toward baseline (Figure 3.3 A, C, and E). When comparing standard-long and reverse-long trials (both with 5 second delays), the distance from baseline and trajectory speed was significantly higher for reverse-long trials in the middle of the delay period when the second odor was expected. This is likely driven by the over-representation of significant sequence cells that have their peak activity at the end of the short delay following odor A. All together, these findings further suggest that CA1 neural trajectories show differences in temporal expectations with the speed at which the population neural activity changes across the delay period.

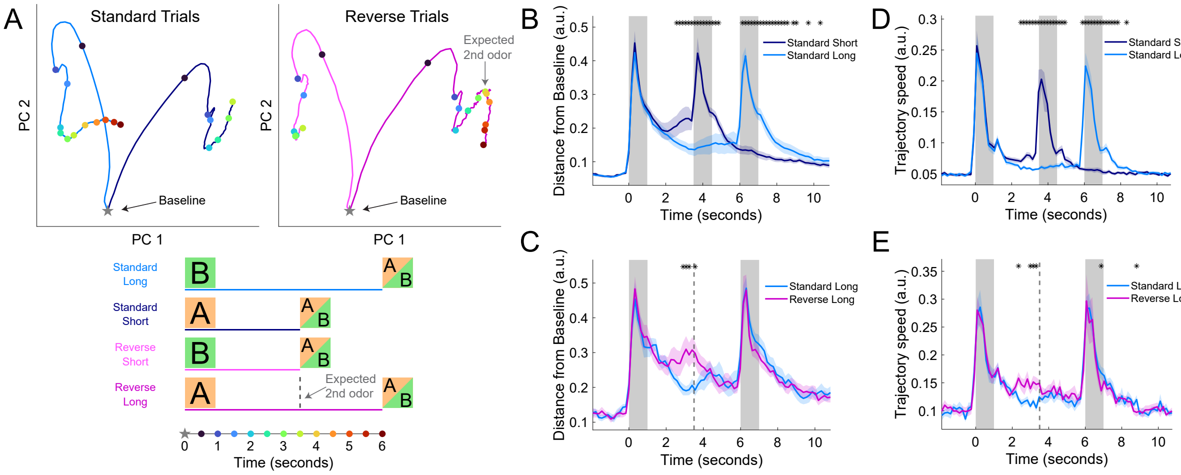

Figure 3.3: **Neural trajectories increase in speed in anticipation of second odor** **A)** Principal component analysis of neural trajectories during first odor and delay period. Top left panel are standard long and standard short trials averaged, while top right panel being reverse short and reverse long trials averaged. Bottom shows timepoints of colored dots that correspond to half-second increments in top panels. Baseline is calculated as the activity that precedes the first odor (from -2 to 0 seconds). Dotted line for 'expected 2nd odor' refers to the expected arrival of 2nd odor of reversal long trials with the cued-short odor A. **B)** Distance from baseline is calculated as the Euclidian distance between neural activities at a point in time compared to baseline 2 second period before the first odor. Comparing traces of significant short trial and long trial sequence neurons (from Figure 3.2 D-F) by averaging standard short and standard long trials (normalizing number of significant neurons (n) by dividing by $\sqrt{n}$). Thick line represents mean of 22 recording sessions (first two days with inclusion of reverse trials), and shaded area represents standard error of the mean. Asterisks represent bins of 1/6 second that are significantly different (Two-Way ANOVA animal and day, corrected for multiple comparisons with Benjamini-Hochberg procedure, $p < 0.01$). **C)** Same as (B) but comparing standard long and reverse long trials. The dotted line is same as (A) for expected 2nd odor arrival on reverse long trials. To control for the same number of trials, the nearest neighbor standard long trials were chosen. **D-E)** Same as (B-C) but for trajectory speed, which is the Euclidian distance between neighboring bins of size 1/6 second.

3.3.4 CA1 encoding of elapsed time is stronger with short delay expectation

Given that the neural trajectories of CA1 sequential activity were different between expected short and long delay lengths, we asked if this yields a difference in the strength of time encoding at the population level. Does an increase of trajectory speed and over-representation of cells firing at the end of the short delay lead to a better ability to decode the current time within the delay period? To answer this question, we trained a support vector machine (SVM) error-correcting output codes (ECOC) multiclass decoding on the activity of all neurons. To allow fair comparison between short and long trials, we only performed decoding during the short delay period from the 1.1 second to 3.4 second timepoints to exclude any directly cue-related activity (Figure 3.4 A). Using 7 bins of 1/3 seconds, we calculated the correlation of real and predicted bins as the quantitative measure to compare between standard short and standard long trials. Across all four days of recording (last 2 days of only standard trials and first 2 days with reverse trials), time decoding correlations were higher in the standard short trials suggesting that the neural activity is more predictive during these trials (Figure 3.4 B). Together, these findings suggest that in the first 2.5 seconds of the delay periods, CA1 population activity more strongly encodes time when expecting a shorter delay. This suggests that the neural sequences and differences in trajectory speed led to significant differences in the ability to decode time within the delay period.

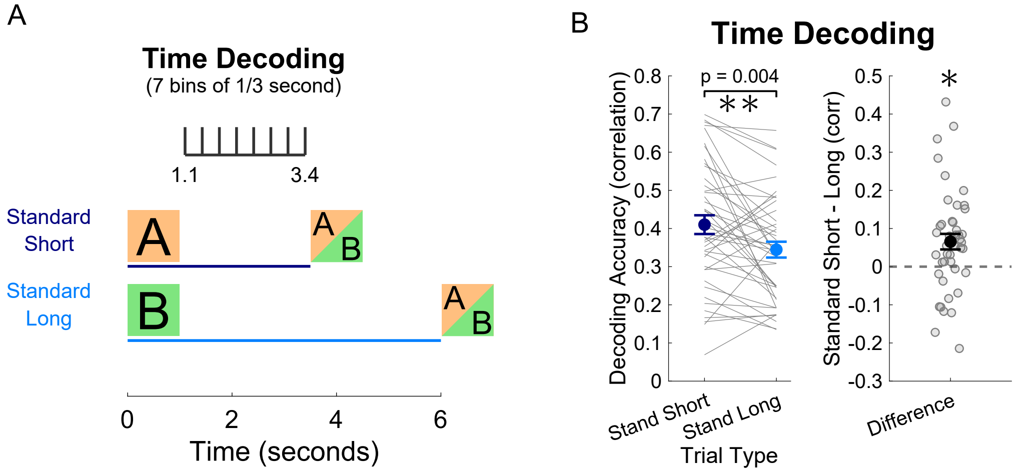

Figure 3.4: **CA1 encoding of elapsed time is stronger with short delay expectation**

A) Support vector machine (SVM) error-correcting output codes (ECOC) multiclass decoding of time bin within short delay period between 1.1 and 3.4 second time points with 7 bins of size 1/3 seconds. **B)** Decoding accuracy calculated as the correlation of real bins to the decoder predicted bins (0 being chance, and 1 being perfect). Two-Way ANOVA animal and day for the 44 recording sessions, p = 0.004).

3.4 Discussion

By developing a novel rodent olfactory-based differential-delayed nonmatch-to-sample (dDNMS) task, we demonstrate that mice learn an implicit cue-delay association. Similar to previous human experiments (Zhou et al., 2023), mice performed worse and responded later when delay length expectations were violated in reverse trials. These behavioral results suggest an interaction between the required WM demands of the task and the implicit timing that mice learn. Using two-photon calcium imaging of dorsal CA1 pyramidal neurons during performance of the dDNMS task, we found odor-specific sequential activity encoding WM and timing representations during the delay period. First, odor-specific sequential activity being odor-specific and tiling the entire delay lengths allowed the first odor representation to be maintained until the arrival of the 2nd odor. Second, sequential activity kept track of elapsed time as different neurons become active at different time points within the delay period. Third, the different slopes and speed of change of these sequences encoded the distinct expected delay lengths. When expecting the arrival of the 2nd odor following a short delay, there was an over-representation of timepoints prior to the arrival of the 2nd odor and neural trajectories increased their speed. Finally, this anticipatory increase in trajectory speed improved decodability within the short delay, suggesting a role in the animal's ability to predict the arrival of the 2nd odor.

While sequential activity has been proposed as possibly multiplexing working memory and timing representations (Liu and Buonomano, 2009; Rajan et al., 2016; Murray and Escola, 2017; Zhou et al., 2023), our findings demonstrate this at both the behavioral and neural level. When implicit timing expectations were violated, WM performance suffered despite no change to the WM component of the task. Future manipulation studies will be necessary to probe the role of increasing trajectory speed in CA1 for anticipation or maintenance of working memory olfactory representations.

Different population level neural dynamics of working memory have also been observed in many areas of the brain: steady-state persistent activity (Fuster and Alexander, 1971;

Kubota and Niki, 1971; Funahashi et al., 1989; Riley and Constantinidis, 2016) and ramping activity (Inagaki et al., 2018, 2019; Cueva et al., 2020). Although persistent activity does not allow for the encoding of temporal information, both ramping activity and complex dynamics could be alternative or additional ways the brain may multiplex working memory and timing. Despite our robust findings including only neurons that passed our sequence-cell detection analysis, it is possible that other dynamics may also play a role in CA1. Voltage or electrophysiological recording would be better suited to distinguish sequential activity from persistent, ramping, or complex activity in tasks like ours with relatively short delay periods of 2.5 seconds.

While dorsal CA1 was chosen for this study because of its known representations of time through sequential activity (Pastalkova et al., 2008; Itskov et al., 2011; MacDonald et al., 2011, 2013; Taxidis et al., 2020), its relatively weak recurrent connectivity (Knowles and Schwartzkroin, 1981; Deuchars and Thomson, 1996; Yang et al., 2014) makes it unlikely to be generating the temporal representations. The odor-specific working memory information observed in CA1 is likely integrated from multiple areas such as lateral entorhinal cortex (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024) and piriform cortex (Howard et al., 2009; Bekkers and Suzuki, 2013; Strauch and Manahan-Vaughan, 2020), but it is unclear which brain regions play a role in maintaining this information during the delay period. The CA3 region of the hippocampus - given its stronger recurrent connectivity and inputs to CA1 (Li et al., 1994; Le Duigou et al., 2014) - is likely playing a critical role in the formation or the sequential activity we observed in CA1. Alternatively, both the medial and lateral entorhinal cortices have been implicated in timing representations as well (Kraus et al., 2015; Tsao et al., 2018; Umbach et al., 2020). Future manipulation studies will be necessary to understand which regions play a role in generating the timing-dependent dynamics we have recorded in this study.

Finally, this study focused on excitatory pyramidal neurons, but the role of inhibition within the CA1 region during our task is likely critical for the formation of the sequential activity we observed. The balance of excitation and inhibition is suggested to be necessary for sequential activity, but recent computation theories suggest that inhibitory plasticity and connection may be more important than excitatory ones (Zhou et al., 2023). Although some research has focused on the differential role of interneuron subtypes on CA1 sequential activity (Taxidis et al., 2023), it will be valuable to understand their causal role in driving sequential formation across learning.

3.4.1 Method Details

Animals

All of the experiments were conducted according to the National Institute of Health (NIH) guidelines and with the approval of the Chancellor's Animal Research Committee of the University of California, Los Angeles. A total of 9 adult male and 3 female mice (8-34 weeks old) were used for *in-vivo* behavioral experiments, and a total of 6 adult male and 5 female mice (8-16 weeks old) were used *in-vivo* calcium CA1 neuron imaging experiments. All were C57BL/6J (Jackson Laboratory, 000664), experimentally naïve, and housed in the vivarium under a 12-hour light/dark cycle. All mice were group housed (2-4 per cage) with the exception of 3 that had to be separated following surgery because of fighting.

Surgical Procedures

Mice (8-30 weeks old) were subcutaneously administered pre-operative drugs (carprofen 5 mg/kg, dexamethasone 0.2 mg/kg, lidocaine 5 mg/kg) 30 minutes before surgery. Mice were anaesthetized with isoflurane (5% induction, 1-2% for maintenance), and anesthesia was continuously monitored and adjusted as necessary. The scalp was shaved, and mice were placed into a stereotactic frame (David Kopf Instruments, Tujunga, CA) on a feedback-controlled heating pad (Harvard Apparatus) set to maintain body temperature at 37°C. Eyes were protected from desiccation using artificial tear ointment. The surgical incision site was cleaned three times with 10% povidone-iodine and 70% ethanol. Fascia was removed by applying hydrogen peroxide, connective tissue was cleared from the skull, and the skull was scored to facilitate effective bonding with adhesives at the end of surgery.

Behavioral experiments: After removing fascia and connective tissue, a custom-made lightweight stainless-steel headbar was attached to posterior skull and secured with cyanoacrylate glue. Dental cement (Ortho-Jet, Lang Dental) was applied to seal and cover any remaining skull.

CA1 calcium imaging experiments: After stereotactically aligning the skull, a single burr hole was made above right dorsal CA1. Either pGP-AAV1-syn-jGCaMP8f-WPRE (1000nL of 1:5 saline dilution) or pGP-AAV1-syn-jGCaMP8f-WPRE (1000nL of 1:5 saline dilution) was injected into the right dorsal CA1 (2.0 mm posterior from bregma, 1.8 lateral from bregma, and 1.3 ventral from dura). Viruses were injected using a Nanoject II microinjector (Drummond Scientific) at 60nL per minute. Following virus injection and waiting for 45 minutes, a circular craniotomy (3 mm diameter) was made centered around a point made 2.0 mm posterior and 1.8 lateral to bregma. Dura beneath the craniotomy was removed and cortical tissue above dorsal CA1 was carefully aspirated using a 27-gauge blunt needle. Corpus callosum was spread to the sides of the craniotomy to expose the alveus. Cortex buffer (NaCl = 7.88g/L, KCl = 0.372g/L, HEPES = 1.192g/L, $CaCl_2$ = 0.264g/L, $MgCl_2$ = 0.204g/L, at a pH of 7.4) was continuously flushed during aspiration and until bleeding stopped. A titanium ring with a 3 mm diameter circular thin #0 coverglass attached to its bottom was implanted into the aspirated craniotomy and the overhanging flange was secured to the skull with vetbond (3M). A custom-made lightweight stainless-steel headbar was attached to posterior skull and secured with cyanoacrylate glue. Dental cement (Ortho-Jet, Lang Dental) was applied to seal and cover any remaining skull, and to form a small well around the titanium ring for holding immersion water for the objective during imaging.

Following surgery, all animals were given post-operative care (carprofen 5 mg/kg and dexamethasone 0.2 mg/kg for 48 hours after surgery) and provided amoxicillin-treated water at 0.5 mg/mL for 7 days. All mice recovered for 7-14 days before experiments began.

Experimental setup

The entire behavioral setup is as described in Taxidis et al. (Taxidis et al., 2020). Mice were head-fixed above an 8-inch spherical Styrofoam ball (Graham Sweet) which can rotate about one axis for 1D locomotion that was recorded with a sensor (Avago ADNS-9500). A continuous stream of clean air (∼1 L/min) was delivered toward the animal's nose via Tygon PVC clear tubing and a custom-made port that held the air tube and water port. At the onset of the odor presentation period, a dual synchronous 3-way valve (NResearch) switched to the odorized one for 1 second. Odorized air was created by using a 4-ports olfactometer (Rev. 7c; Biology Electronics, Caltech) supplying air to either of two glass vials containing odor A (70% Citral, FCC; Sigma Aldrich) or odor B ((-)-a-Pinene $\geq$ 97%, FCC; Sigma Aldrich), which were both diluted in mineral oil at 5% concentration. Water droplets (∼10µl) were released by a 3-way solenoid valve (Lee Company), and licks were detected by using a custom battery-operated circuit board with one end of the circuit connected to the headbar and the other to the lickport. The behavioral rig was controlled with custom written software (MATLAB) and through a data acquisition board (USB-6341: National Instruments).

Behavioral training

After 7-14 days recovering from surgery, mice were handled and began water-restriction to 85% of their original weight before water-restriction. After one day of handling, mice were habituated to being head-fixed above the spherical treadmill (can rotate about one metal axis for 1D locomotion that is recorded) for two days. On the 4th day of training, mice began learning to lick from the lickport as water was automatically delivered at the beginning of the reward period following only standard non-matched odor trials (AB or

BA, with water delivery at the offset of the 2nd odor). Trials were delivered in blocks of 20 trials. This phase was always 2 days except for the rare mouse that needed one extra day to reach motivation level and lick water from port for at least 50 trials. In the next phase, water was only delivered if the mouse licked during the response period, and mice learned to reliably lick in anticipation of the reward following the 2nd odor. This phase was also 2 or 3 days, dependent on the mouse licking during the response period of at least 50 trials. The next phase was the full differential delayed non-match-to-sample (DNMS) task in which standard matched odor trials (AA and BB) were introduced and mice learned to refrain from licking the port following these trials. There was no punishment or timeout following an incorrect lick; the water was simply not delivered. For this phase and previous phases, the response window was 3 seconds starting at the offset of the 2nd odor. After 6 days of this phase or until behavior reached 85 %, we began sliding the response window forward by 1/3 of second. For example, the first day with the new response window started 2/3 of a second after the onset of the second odor. On the following day, the response window would start 1/3 of a second after the onset of the second odor. Finally, mice performed the dDNMS task with the response window starting at the onset of the second odor. At this stage, always 7 blocks of 20 trials were delivered. Mice performed 3 days of this phase with only standard trials delivered, except for 3 mice that were given an extra 4th day because of performance below 85%. This phase is referred to as the standard trial days. Following this phase, reverse trials were introduced at a 20% rate for 3 more days. Again, 7 blocks of 20 trials were delivered. For simplicity throughout the figures and text, odor A (citral) always predicts short delay and odor B (pinene) always predicts long delay. However, we counterbalanced mice with half trained the opposite way with odor A (citral) predicting long delay and odor B (pinene) predicting short delay.

In-vivo **two-photon imaging**

All two-photon calcium imaging was conducted using a resonant scanning two-photon microscope (Scientifica) fitted with a 16x 0.80 NA objective (Nikon) to record 512x512 pixel frames at 30.9 Hz. CA1 imaging fields of view were 500x500 µm and axonal imaging fields were 250x250 µm. Excitation light was delivered with a Ti:sapphire excitation laser (Chameleon Ultra II, Coherent), operated at 920 nm. GCaMP8f and GCaMP7f fluorescence was recorded with a green channel gallium arsenide photomultiplier tube (GaAsP PMT; Hamamatsu). Microscope control and image acquisition were performed using LabView-based software (SciScan). Imaging and behavioral data were synchronized by recording TTL pulses generated at the onset of each imaging frame and olfactory stimulation digital signals at 1 kHz, using WinEDR software (Strathclyde Electrophysiology Software).

For imaging experiments, dorsal CA1 was imaged for at least 5 consecutive days of task performance. This includes 2 days of standard trials only and 3 days with reversal trials (however, the last reversal trial day was not included in any analysis presented here). While careful attention was given to aligning the FOV to the previous day's when possible, sometimes FOV's needed to be changed to optimize the number of active neurons. We used rotating stages, a motor for adjusting mouse head angle, and a tiltable objective attachment with two degrees of freedom to fine-tune the alignment. Laser power and PMT settings were kept consistent between days, except for rare occasions when it was necessary to keep similar signal-to-noise.

For each day of recording, imaging was halted between each of the 7 blocks of 20 trials. This allowed fine-tuning of alignment to keep the same FOV within the day, and it also prevented brain heating or photo-toxicity. Laser power was kept as minimal as possible (60-90mW) without sacrificing signal-to-noise ratio, and no significant photo-bleaching was observed.

Histology

Following all experiments, mice were deeply anaesthetized under isoflurane and transcardially perfused with 30 mL 1x PBS followed by 30 mL 4% paraformaldehyde in 1x PBS at a rate of approximately 4 mL/min. After perfusion, the brains were extracted and post-fixed in 4% paraformaldehyde. Coronal sections of 80 µm were collected using a vibratome, 24-48 hours after perfusion. The sections were mounted onto glass slides and cover-slipped with DAPI mounting medium. Images were acquired on an Apotome2 microscope (Zeiss; 5x, 10x, 20x objectives) to confirm proper expression and location of viral expression. For CA1 imaging experiments, GCaMP8f or GCaMP7f was confirmed to be in dorsal CA1 with no damage to the hippocampal formation.

3.4.2 Quantification and Statistical Analysis

Performance and lick-latency

For all 23 mice (behavior and CA1 imaging cohorts), performance was measured as the percentage of trials with the correct response. Responses were only considered for performance during the response window (3 second period starting at the onset of 2nd odor). Within the response period, the time elapsed between the onset of the 2nd odor and the first lick is considered the ‘lick latency’. To minimize the effect of extremely late licking outlier events, the median lick latency was taken for each recording session.

Calcium imaging data pre-processing

For CA1 imaging experiments, each day was analyzed separately as neurons were not aligned across days. Movies were processed using the Python implementation of Suite2P 0.9.2 (Pachitariu et al., 2017) to perform non-rigid motion registration, neuron segmentation, extraction of fluorescence signals, and deconvolution with parameters optimized to our GCaMP8f and GCaMP7f CA1 recordings (with only the decay time-constant being different between the two). We used the default classifier and an ‘iscell’ threshold of 0.3 to only include masks that were likely neurons. Deconvolved signals were taken as the selected output from Suite2P and taken to MATLAB 2021a for further processing. Deconvolved signals were smoothed by a rolling mean of 10 frames (0.32 seconds), then z-scored, and finally values below 2 were set to zero. The resulting signal was what was used for all analysis and referred to as ‘firing rate (STD)’ as a proxy for spiking activity. Signals were aligned to the trial structure (odor presentations and lick timing) and the recorded locomotion as mice ran on the spherical ball.

Sequence neuron detection and analysis

To evaluate peak firing timing, we performed sequence detection similar to the previously described approach in our DNMS task, Taxidis et al. (Taxidis et al., 2020). First, standard trials that begin with Odor A and those that begin with Odor B are separated, and the one with a larger peak of the average activity was considered further. Additionally, only the 3.5 or 6-second period including first odor presentation and the delay period was considered. The peak of average activity within this period and a given trial type was determined to be significant if the peak was greater than the 95th percentile of 2000 circular shuffles. The cell must also have had a trial reliability of at least 20% and have fired above 2 STD for 20% of the preferred trials within 0.5 seconds of the peak frame found in the previous step. If a cell passed both criteria, it was considered to be a ‘short trial neuron’ or ‘long trial neuron’ regardless of its odor selectivity.

PCA analysis and Trajectory Speed

For figure 3.3 A, all active neurons were taken, and principal component analysis (PCA) was conducted on the average deconvolved signal of each of the 4 different trial types. The baseline point was calculated as the average position of the trajectories in the 2 seconds preceding the onset of the first odor. While figure 3.3 A only shows the first 2 principal components (PCs), figure 3.3 B-E calculates the Euclidian distances without performance any dimensionality reduction. These figure panels only use the significant 'short trial neurons' and 'long trial neurons' detected in the previous figure. Using bins of 5 frames (0.16 seconds), distance from baseline is defined as the distance in full dimensional space between the baseline point and the bin in question. Trajectory speed is the distance in full dimensional space from one bin to the next. Arbitrary units (a.u.) are displayed because distance and speed are normalized for number of significant neurons (n) by dividing by $\sqrt{\text{n}}$.

Time decoding

As described in the text, support vector machine (SVM) error-correcting output codes (ECOC) multiclass decoding was used to decode the time bin within the short delay period between 1.1 and 3.4 second time points with 7 bins of size 1/3 seconds. The MATLAB function 'fitcecoc' was used to perform the decoding using the default one vs one method. Correlation was chosen as the measure to achieve one value for accuracy for each recording session.

Statistical analysis

All statistical tests are described in the legend text. For all figures, individual p-values are listed in the figure or legend text. On occasions when single asterisks were displayed above a curve or trace, p-values were corrected for multiple comparisons using the false discovery rate Benjamini-Hochberg procedure, and asterisks indicate bins in which the adjusted p-value < 0.01. In all cases in the text, values were written in the format 'mean ± standard deviation' (STD), while error bars in all figures show the mean and standard error of the mean (SEM). No statistical methods were used to determine appropriate sample sizes but were chosen as being comparable to sizes used in similar publications.

CHAPTER 4

LEC or MEC inhibition does not impair learning or performance on an olfactory working memory task

4.1 Introduction

This hippocampus is known to integrate inputs from many regions during the formation of memories (Canto et al., 2008; Knierim, 2015; Zemla and Basu, 2017). The primary inputs to the hippocampus are from the entorhinal cortex (EC), which is further divided into the lateral and medial parts (LEC and MEC). Within spatial tasks, MEC is known for its importance in spatial encoding (Hafting et al., 2005; Knierim et al., 2014; Save and Sargolini, 2017), while LEC is recognized for its encoding of non-spatial elements in the environment such as wall color (Lu et al., 2013) and time elapsed (Tsao et al., 2018). While older models characterize the functional difference between MEC and LEC as being 'where' and 'what', newer models have a more appropriate division being 'context' and 'content' (Knierim et al., 2014). This distinction is more easily applied to non-spatial contexts.

Additionally, LEC is recognized for its strong olfactory encoding (Igarashi et al., 2014; Li et al., 2017; Woods et al., 2020; Zhang et al., 2024). LEC receives direct inputs from the olfactory bulb and extensive inputs from the piriform cortex, and LEC neurons have been shown to encode odor-specific information in various tasks and contexts (Kerr et al., 2007; Chapuis et al., 2013; Igarashi et al., 2014). Furthermore, the direct pathway of LEC to CA1 has been shown to be necessary for learning and performance on a simple odor association task (Li et al., 2017).

Therefore, we predicted that LEC would be necessary for learning and performance of an olfactory-based delayed non-match-to-sample (DNMS) working memory task. The result in Figure 2.9 A that LEC or MEC chemogenetic inhibition did not impair behavioral performance of expert mice was unexpected. To address this unexpected result, we performed optogenetic inhibition of LEC and MEC to investigate if either region is necessary for learning. However, we consistently observed that LEC or MEC inhibition had no effect on learning or performance of the olfactory DNMS working memory task. Neither inhibition of excitatory neurons with stGtACR2 nor excitation of inhibitory neurons with channelrhodopsin-2 yielded a behavioral effect. Making the task more difficult by increasing delay length to 15-seconds, diluting odors to 0.05%, and changing the pre-training shaping protocol also did not yield an effect. Finally, no effect was observed with stGtACR2 inhibition of excitatory neurons during an olfactory Go-No-Go association task. Together these findings support the chemogenetic inhibition finding in Figure 2.9 A, but ask many further questions about the causal role of LEC in olfactory tasks.

4.2 Results

We trained a total of 64 adult male and female mice on various versions of olfactory working memory and olfactory association tasks. For all DNMS tasks, mice were water-restricted and trained while head-fixed on a spherical treadmill. Each trial consisted of two 1-second odor presentations separated by a delay period. One second after the offset of the 2nd odor there was a 3-second reward period during which the choice of the animal was determined. Mice were trained to lick the lickport to release water during this reward period if the two odors did not match (correct 'hit'). Mice learned to refrain from licking the lickport if the odors matched (correct 'rejection'), and overall performance was quantified as the percentage of correct 'hits' and correct 'rejections' out of all trials. Each session of the

DNMS task consisted of 5 blocks of 20 trials, with pseudorandomly distributed odor combinations. For all variants of the DNMS task, 470nm light was delivered through fiber cannulas to the brain region of interest continuously during both odor periods and the delay period (Figure 4.1 A).

4.2.1 stGtACR2 inhibition of LEC or MEC does not impair learning or performance on the standard olfactory DNMS task

In the first cohort of 18 animals, animals performed the standard DNMS task with a 5 second delay (Figure 4.1). Mice were injected with viruses (pAAV-CaMKIIa-EGFP or pAAV-CaMKIIa-stGtACR2-FusionRed) in LEC or MEC (Figure 4.1 B). A standard shaping protocol that consisted of 5 days of only non-match trials was used to train mice to lick after the 2nd odor. Following these 5 days of shaping, 'Day 1' was considered the first day of learning as match trials were introduced and mice had to learn to refrain from licking after these trials. On every trial starting on this Day 1, blue light was delivered for the 7-second period including both odors and the delay period. We observed no significant behavioral effect for 6 days of learning (individual t-tests: day 1 $p = 0.15$, day 2 $p = 0.29$, day 3 $p = 0.93$, day 4 $= 0.10$) (Figure 4.1 C). The same comparison for MEC mice also showed no significant difference, although we were not sufficiently powered to observe a statistically significant difference (Figure 4.1).

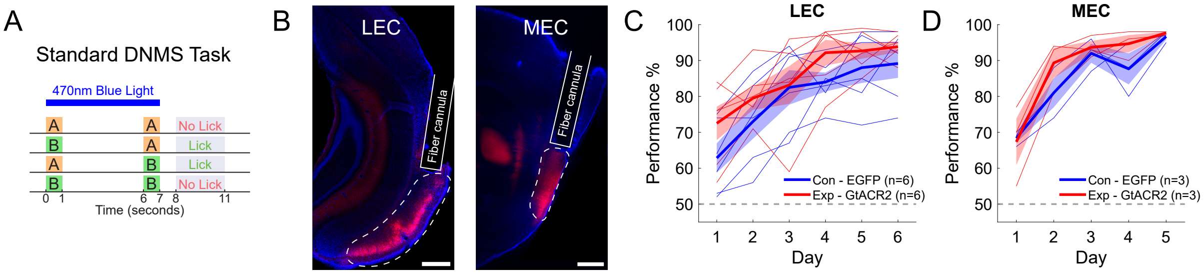

Figure 4.1: **stGtACR2 inhibition of LEC or MEC does not impair learning or performance on the standard olfactory DNMS task**

A) Blue light delivered continuously for 7 seconds during odor and delay periods of standard DNMS task. **B)** Injections of virus to drive the expression of EGFP or stGtACR2 were delivered to either LEC or MEC. Blue is Dapi, and red is FusionRed in stGtACR2. Images showing LEC are from coronal sections, while MEC are from sagittal sections. For both LEC and MEC, the scale bar is 500µm. **C)** LEC inhibition yields no significant behavioral effect. Thin lines represent individual animals, while thick lines are group average. Shaded areas represent SEM (standard error of mean). **D)** Same as (C), but MEC also not significant.

4.2.2 Longer delays and odor-free shaping protocol make task more difficult but no effect of LEC inhibition

The fast learning curves observed in the standard DNMS raised questions regarding whether low-level of task difficulty was impairing our ability to detect behavioral differences with LEC or MEC inhibition. As a first measure in making the task more difficult, we increased the delay period from 5 to 15 seconds. It is known that working memory demands are higher with longer delays and impairments may only be observed on longer delay tasks (Olson et al., 2006). With all other parameters identical to Figure OPTO1 C, the 15-second delay LEC cohort of mice also showed no behavioral difference between control EGFP and experimental stGtACR2 mice (Figure 4.2 A). Additionally, the mice still learned the task remarkably quickly, which forced us to evaluate the pre-training shaping period of days that precede 'Day 1' of match trials and blue light delivery.

In the next cohort of mice, all parameters remained the same to the standard DNMS task, except the olfactometer remained turned off during shaping. Termed 'odor-free' shaping, mice were forced to learn the temporal structure (to lick after 2nd odor delivery) from only the auditory clicks of the valve for odor period onset and offset. Because the auditory cues are salient, mice learned the temporal structure in 'odor-free' shaping similarly to the standard shaping protocol. With 5-second delays and all other parameters the same, this 'odor-free' shaping variant of the task was more difficult since the entire cohort of animals learned much more slowly than cohorts trained on the standard DNMS task (Figure 4.2 B). All mice remained at chance levels on day 1 because this was the first day of receiving any odors. However, again we observed no significant differences between experimental stGtACR2 mice and control EGFP mice. While we were likely not powered to observed significance, the experimental group trended on learning the task more quickly.

Next, in another cohort of mice, we combined the 'odor-free' shaping protocol with 15-second delays. Mice learned more slowly, but we still observed no behavioral difference between groups. These results suggest that making the task more difficult with longer delays was not sufficient to observe behavioral effects. Furthermore, our concern about odor exposure for days before blue light delivery resulted in a new shaping protocol that slowed learning but did not yield behavioral impairments with LEC inhibition.

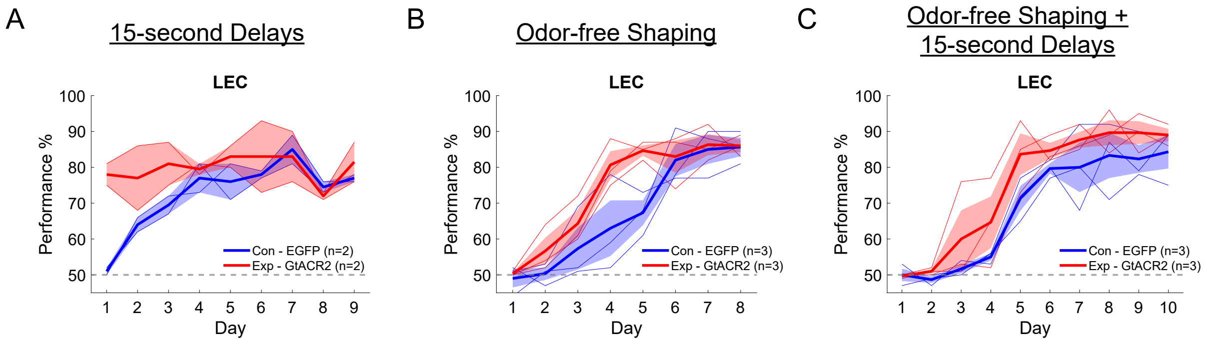

Figure 4.2: **Longer delays and odor-free shaping protocol make task more difficult but no effect of LEC inhibition**

A) Blue light delivered continuously for 17 seconds during odor and delay periods of DNMS task with 15 second delays. Visualization same as Figure 4.1. **B)** Odor-free shaping in days that preceded 7-second blue light delivery during standard DNMS task. **C)** Odor-free shaping and 15 second delays with same 17-second blue light delivery.

4.2.3 Excitation of inhibitory neurons also does not impair learning or performance

Given that task difficulty did not affect the impact of LEC inhibition, we chose to test a different form of inhibition. Instead of inhibiting excitatory neurons with stGtACR2, we chose to excite inhibitory neurons. Gad2Cre::Ai14 (expressing tdTomato in all inhibitory neurons) and Gad2Cre::Ai27 (expressing channelrhodopsin-2 in all inhibitory neurons) were used in the next experiment. Fiber cannulas were lowered above LEC just as previously done. First, with the standard DNMS task and shaping protocol, we observed a small trend of experimental Gad2Cre::Ai27 animals learning more slowly, but we were not powered to see an effect.

Next, we used Gad2Cre::Ai27 negative and positive cage mates as control and experimental mice following 'odor-free' shaping. Again, we observed slower learning compared to the standard shaping protocol, but we observed no difference between the control and experimental groups. Together, these experiments exciting inhibitory neuron further support our previous findings that LEC may not be necessary for learning of the olfactory DNMS task.

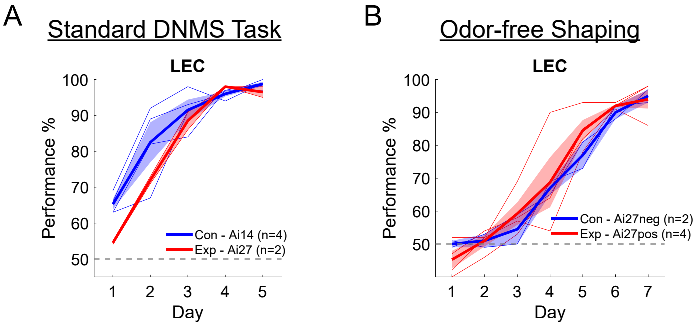

Figure 4.3: **Excitation of inhibitory neurons also does not impair learning or performance**
A) Control mice express Ai14 tdTomato in inhibitory neurons, while experimental mice express Ai27 channelrhodopsin-2 in inhibitory neurons. Blue light delivered continuously for 7 seconds during odor and delay periods of standard DNMS task. **B)** Control mice were cage mate negatives to the experimental positive Ai27 channelrhodopsin-2 mice. Odor-free shaping in days that preceded 7-second blue light delivery during standard DNMS task.

4.2.4 Diluting odors in DNMS or Go-No-Go task also does not impair learning or performance

Although LEC may not be necessary for the working memory component of the task, we hypothesized that it still may be necessary for olfaction perception. Therefore, we chose to make the olfaction perception component more difficult by dramatically reducing the concentration of the odors. All previous experiments were performed with 5% odorant concentration diluted in mineral oil, as was done previously in the lab (Taxidis et al., 2020). Previous experiments from that publication demonstrated 0.05% as being the lowest concentration that mice could still perform the DNMS task above chance level (Taxidis et al., 2020). Therefore, we repeated the 'odor-free' shaping protocol from Figure 4.2 B and conducted the standard DNMS task with 5-second delays with odorant concentrations at 0.05%. Mice had the slowest learning curves of all cohorts, suggesting that we successfully made the task much more difficult to learn; however, we observed no behavioral difference between control EGFP and experimental stGtACR2 mice.

Finally, we attempted to validate the behavioral finding from Li et al., 2017 showing behavioral impairments in a Go-No-Go olfactory association task (Li et al., 2017). Instead of only inhibiting the LEC axon terminals in CA1 with halorhodopsin, we continued to inhibit LEC directly with stGtACR2. However, other behavioral parameters were identical, and we continued to use the diluted odors at 0.05%. Mice were first pre-trained to lick following a LED response cue. The task consisted of odor A or B being presented for 1 second with a 2 second delay period until the response cue, which initiated a 1 second response window. If mice licked during the response window following odor A, they received a water reward. If they licked during following odor B, no reward was dispensed. Assignment was randomized between mice as to which odor was rewarded. We trained the mice on 5 blocks of 20 trials per day, and mice rapidly learned to refrain from licking after odor B. We observed similar learning curves as Li et al., 2017 (Li et al., 2017), however, we observed no difference between control EGFP and experimental stGtACR2 mice.

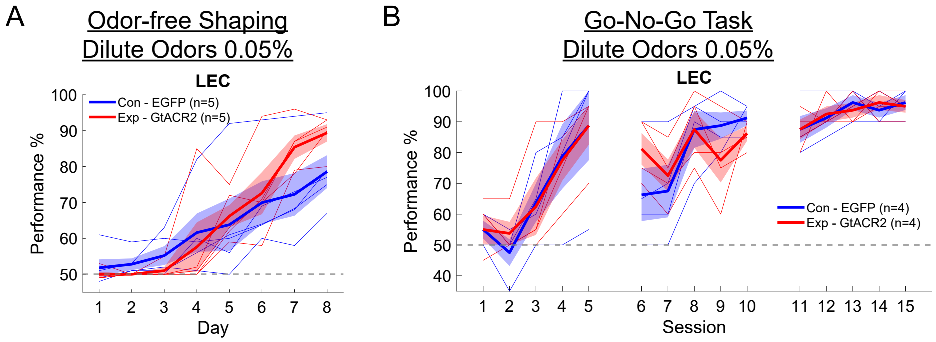

Figure 4.4: **Diluting odors in DNMS or Go-No-Go task also does not impair learning or performance**
A) Odors diluted to 0.05% instead of 5%. Blue light delivered continuously for 7 seconds during odor and delay periods of standard DNMS task. **B)** 3 days of Go-No-Go task with the same 0.05% odors. Each day was divided into 5 sessions of 20 trials.

4.3 Discussion

Using optogenetic silencing approaches, we asked if LEC or MEC is necessary for learning of an olfactory DNMS working memory task. Throughout many different iterations and versions of experiments, we consistently observed no behavioral difference between control animals and inhibited animals. Neither inhibiting excitatory neurons with stGtACR2 nor exciting inhibitory neurons with channelrhodopsin-2 yielded a behavioral effect. We also made the task more difficult with longer delays, more dilute odors, or an odor-free shaping protocol. We even attempted to replicate a previous finding of LEC to CA1 projections being necessary for learning of an olfactory association task. In the same task, our stGtACR2 silencing of LEC itself did not yield a learning or performance effect.

Many questions remain regarding the necessity of LEC or MEC in learning or performance of our olfactory DNMS task. Although some of the individual experiments were not sufficiently powered to observe statistical significance, overall, we did not see a trend toward LEC inhibition impairing performance. We performed two inhibition methods, but it remains possible that other opsins or inhibition approaches may yield behavioral impairments. Another possibility is simply that inhibiting all of LEC or MEC does not impair performance on our task. Finer manipulations to inhibit specific pathways such as LEC to CA1 may be required to observe behavioral effects.

Future experiments testing various optogenetic opsins, chemogenetic tools, or other inhibition methods may be required to confirm the lack of behavioral necessity we observed. Additionally, alternate versions of pretraining protocol and task itself may be required.

4.4 Method Details

Animals

All of the experiments were conducted according to the National Institute of Health (NIH) guidelines and with the approval of the Chancellor's Animal Research Committee of the University of California, Los Angeles. A total of 52 mice (8-18 weeks old) were used for EGFP and stGtACR2 inhibition experiments, and all were C57BL/6J (Jackson Laboratory, 000664). A total of 12 mice (8-18 weeks old) were used for channelrhodopsin-2 excitation of inhibitory neurons. These were all Gad2-IRES-Cre (Jackson Laboratory, 010802) and were additionally bred with either Ai14 (Jackson Laboratory, 007914) or Ai27 (Jackson Laboratory, 012567). All mice were experimentally naïve and housed in the vivarium under a 12-hour light/dark cycle. All mice were group housed (2-4 per cage) with the exception of few that had to be separated following surgery because of fighting.

Surgical Procedures

Mice (8-18 weeks old) were subcutaneously administered pre-operative drugs (carprofen 5 mg/kg, dexamethasone 0.2 mg/kg, lidocaine 5 mg/kg) 30 minutes before surgery. Mice were anaesthetized with isoflurane (5% induction, 1-2% for maintenance), and anesthesia was continuously monitored and adjusted as necessary. The scalp was shaved, and mice were placed into a stereotactic frame (David Kopf Instruments, Tujunga, CA) on a feedback-controlled heating pad (Harvard Apparatus) set to maintain body temperature at 37°C. Eyes were protected from desiccation using artificial tear ointment. The surgical incision site was cleaned three times with 10% povidone-iodine and 70% ethanol. Fascia was removed by applying hydrogen peroxide, connective tissue was cleared from the skull, and the skull was scored to facilitate effective bonding with adhesives at the end of surgery. After stereotactically aligning the skull, a single or several burr holes were made depending on the experiment performed and virus was injected.

EGFP and stGtACR2 experiments: Control virus (800 nL of 1:5 saline dilution of pAAV1-CaMKIIa-EGFP into all 4 sites) or experimental virus (800 nL of 1:5 saline dilution of pAAV1-CaMKIIa-EGFP into all 4 sites) was injected into LEC (bilaterally 3.4 and 3.9 mm posterior, 4.35 mm lateral, and 4.3 ventral from bregma) or MEC (bilaterally 4.7 mm posterior, 3.35 mm lateral, and 3.8 and 3.0 mm ventral from bregma). After waiting at least 20 minutes following the last injection, 400µm diameter 0.39NA fiber cannulas were slowed lowered to 0.5 mm above the injection sites of LEC (bilaterally 3.5 mm posterior, 4.35 mm lateral, and 3.8 ventral from bregma) or MEC (bilaterally 4.7 mm posterior, 3.35 mm lateral, and 2.5 mm ventral from bregma).

Gad2Cre channelrhodopsin-2 experiments: Cannula placement was identical to EGFP and stGtACR2 experiments.

All mice recovered for 7-14 days before experiments began.

Experimental setup

The entire behavioral setup is as described in Taxidis et al. (Taxidis et al., 2020). Mice were head-fixed above an 8-inch spherical Styrofoam ball (Graham Sweet) which can rotate about one axis for 1D locomotion that was recorded with a sensor (Avago ADNS-9500). A continuous stream of clean air (∼1 L/min) was delivered toward the animal's nose via Tygon PVC clear tubing and a custom-made port that held the air tube and water port. At the onset of the odor presentation period, a dual synchronous 3-way valve (NResearch) switched to the odorized one for 1 second. Odorized air was created by using a 4-ports olfactometer (Rev. 7c; Biology Electronics, Caltech) supplying air to either of two glass vials containing odor A (70% isoamyl acetate basis, FCC; Sigma Aldrich) or odor B ((-)-a-Pinene $\geq$ 97%, FCC; Sigma Aldrich), which were both diluted in mineral oil at 5% concentration. Water droplets (∼10µl) were released by a 3-way solenoid valve (Lee Company), and licks were detected by using a custom battery-operated circuit board with one end of the circuit connected to the headbar and the other to the lickport. The behavioral rig was controlled with custom written software (MATLAB) and through a data acquisition board (USB-6341: National Instruments). Blue light 470nm from fiber-coupled LED (Thorlabs, M470F3) was delivered via 400µm diameter fiber which was coupled to the 400µm 0.39NA fiber cannula surgically attached to the skull. Light power was measured to be 4mW coming out of each fiber cannula.

Behavioral training

After 7-14 days recovering from surgery, mice were handled and began water-restriction to 85% of their original weight before water-restriction. After one day of handling, mice were habituated to being head-fixed above the spherical treadmill (can rotate about one metal axis for 1D locomotion that is recorded) for two days. These first 3 days were consistent for all experiments.

For standard pretraining / shaping: On the 4th day, mice began learning to lick from the lickport as water was automatically delivered at the beginning of the reward period following only non-matched odor trials (AB or BA, with water delivery at time point of 8 seconds). Trials were delivered in blocks of 20 trials. This phase was always 2 days except for the rare mouse that needed one extra day to reach motivation level and lick water from port for at least 50 trials. In the next phase, water was only delivered if the mouse licked during the response period, and mice learned to reliably lick in anticipation of the reward following the 2nd odor. This phase was also 2 or 3 days, dependent on the mouse licking during the response period of at least 50 trials. The final phase was the full delayed non-match-to-sample (DNMS) task in which matched odor trials (AA and BB) were introduced and mice learned to refrain from licking the port following these trials. There was no punishment or timeout following an incorrect lick; the water was simply not delivered. The first day of this final full DNMS task was considered ‘Day 1’ and the first day of blue light delivery (6-8 days from the start of water-restriction). A total of 100 trials delivered in five blocks of 20 trials were given each day.

For ‘odor-free’ shaping: All steps are identical except that the olfactometer remains off for the days before ‘Day 1’ of full DNMS task.

Go-No-Go Association Task: Different from the standard DNMS task, an orange LED was installed just above the lickport. Similar to the standard pretraining / shaping, there were 2-3 days of automatic water delivery. Odor valves clicks were present but the olfactometer remained off, and 2 seconds after the offset the LED turned on for 0.5 seconds. The one second response started at the same time as the LED turned on. In the first stage of automatic water delivery, the water was delivered at the offset of the LED. However, on the following stage where mice needed to trigger water delivery, the entire response window was available to trigger. After 2-3 days on this stage, mice transitioned to ‘Day 1’, which is when the olfactometer was turned on and blue light was delivered. Again a total of 100 trials delivered in five blocks of 20 trials were given for 3 days total.

Quantification and Statistical Analysis

All data was analyzed and visualized in MATLAB 2021a. Two-sample t-tests were performed when noted in the text.

CHAPTER 5

Paradoxical optogenetic kindling with ‘inhibitory’ opsin stGtACR2 in excitatory neurons

5.1 Introduction

Optogenetics is a powerful tool for evaluating causal roles of circuits by controlling neural activity with light and genetic engineering (Emiliani et al., 2022). However, as is true with any technique, there are certain limitations that must be considered. This can sometimes be challenging with a tool like optogenetics when so many different opsins and protocols are used in labs across the world.

Here we tell a short cautionary tale to warn others about potential problems with use of an inhibitory opsin. In early pilot optogenetics experiments to silence excitatory neurons in lateral entorhinal cortex or medial entorhinal cortex (LEC or MEC), we chose stGtACR2 as the opsin of choice because of its reported potency of inhibition (Mahn et al., 2018). This soma-targeted version of the anion-conducting channelrhodopsin is also widely used. However, we mistakenly chose to pulse the 470nm blue light at 20Hz in the same previous experiments in the lab had done with halorhodopsin. We noticed that mice rapidly developed seizure-like activity within minutes of the first blue-light pulse trains, even though we only expected to be inhibiting excitatory neurons. While the creators of stGtACR2 acknowledge the possibility of rebound spiking following the offset of light, the soma targeted version was designed to limit the likelihood this occurring.

We followed up this discovery with a systematic experiment to confirm that pulsed light delivery to inhibit excitatory neurons in LEC or MEC with stGtACR2 could elicit seizure-like activity. To accomplish this, we modified an existing protocol for rapid optogenetic kindling (Cela et al., 2019). Kindling is a model for epileptogenesis, which is a gradual process for a brain to develop synchronous activity following an electrical stimulation (Goddard, 1967). While at first the electrical stimulation has no behavioral effect, if repeated epileptiform behavior will gradually arise as synchronous activity builds to the electrical stimulation. Recent work has shown kindling protocols are also possible using optogenetics to target specific cells (Cela et al., 2019). We delivered at 20Hz for 5 second trains 4 times per day separated by 20 minutes, and all mice expressing stGtACR2 in LEC or MEC fully kindled to criteria within 4 to 11 days. This originally unexpected result suggests that stGtACR2 and potentially other similar opsins should be used with caution, especially in circuits known to be involved in epilepsy like the EC or hippocampus (Thom, 2014). Additionally, it behaviorally supports previous observations that differences in chloride reversal potentials between the soma and axon terminals can lead the opening of chloride channels with opsins like stGtACR2 to be hyperpolarizing and inhibitory at the soma while depolarizing and excitatory at axon terminals (Messier et al., 2018). Although the soma targeted version of GtACR2 increases the ratio of expression to the soma, some remaining expression in the axons could drive the effects we observed.

5.2 Results

5.2.1 Rapid optogenetic kindling by 'inhibition' of LEC or MEC with stGtACR2 expressing in excitatory neurons

We performed a rapid optogenetic kindling protocol on a total of 11 adult mice. Mice were injected with viruses (pAAV-CaMKIIa-EGFP or pAAV-CaMKIIa-stGtACR2-FusionRed) in LEC or MEC and fiber cannulas were inserted just above the region (Figure 5.1 A). Following 3-5 weeks of viral expression, mice were habituated to an open arena chamber with fibers connected bilaterally to the fiber cannulas attached to their skull. After one day of habituation, the experiment began with 470nm blue light being delivered for 5 seconds at 20Hz with a 20% duty cycle (Figure 5.1 B). These 5 second pulse trains were delivered 4 times per day separated by 20 minutes (Figure 5.1 B). Using a modified Racine scale (Racine, 1972), stages of epileptic seizures were categorized from 0 to 5 following each of the pulse trains (Figure 5.1 B). Mice continued until a Racine scale 5 was reached.

On the first day, no LEC or MEC stGtACR2 expressing mice exhibited any behavioral changes to the blue light pulse trains. However, by day 2 most of the mice began exhibiting freezing behavior. Seizure severity gradually increased until all 7 mice eventually reached Racine scale 5 seizure with wild running and jumping. Notably, 3 control EGFP expressing mice for both LEC and MEC followed the same protocol for 7 days. None of these 6 control mice ever showed any behavioral changes; they scored a 0 on each pulse train. Additionally, no mice showing any behavioral changes to continuous light delivery (see chapter 4).

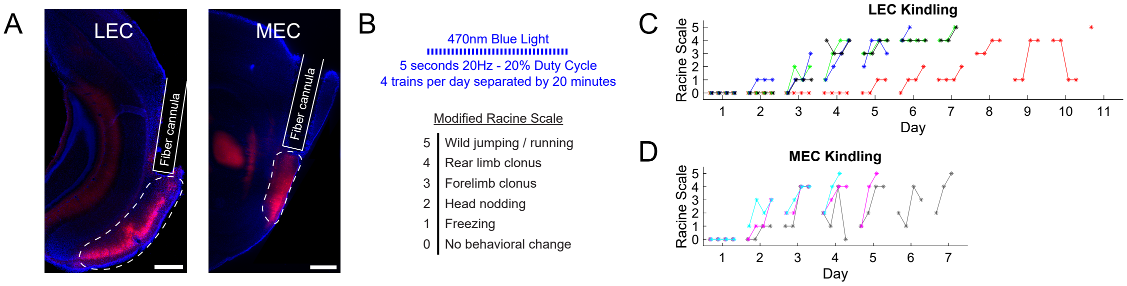

Figure 5.1: **Rapid optogentic kindling by 'inhibition' of LEC or MEC with stGtACR2 expressing in excitatory neurons**
A) Injections of virus to drive the expression of EGFP or stGtACR2 were delivered to either LEC or MEC. Blue is Dapi, and red is FusionRed in stGtACR2. Images showing LEC are from coronal sections, while MEC are from sagittal sections. For both LEC and MEC, the scale bar is 500µm. **B)** Details for rapid optogenetic kindling protocol. **C)** 4 LEC mice all fully kindle to a Racine scale of 5. Each color is a different mouse. Each day consisted of 4 light trains with each having a Racine scale measure. **D)** Same as (C), but MEC also had 3 mice fully kindle.

5.3 Discussion

We demonstrated that stGtACR2 expression in LEC or MEC can drive seizure-like activity if activation blue-light is pulsed. Using a rapid optogenetic kindling protocol, we fully kindled all mice in 4 to 11 days. The results are paradoxical given that stGtACR2 is an inhibitory anion-conducting channelrhodopsin, and we restricted expression to excitatory neurons with the CaMKIIa promoter. However, slice electrophysiological studies of GtACR2 point to the likely explanation. Messier et al., observed that GtACR2 activation drove hyperpolarization and inhibition at the soma, but drove depolarization and excitation at axon terminals (Messier et al., 2018). This is due to differences in resting chloride membrane potentials. While the soma targeted version does enhance expression at the soma, we still observe significant expression in the axon terminals (Figure 5.1 A).

Additional experiments would be needed to confirm the pathway that is driving the seizure activity we observe as it could be axon terminals in EC that are being depolarized or terminals in other areas like the ventral hippocampus that may also be receiving enough activation blue light. However, we demonstrate sufficient evidence to tell the cautionary tale for using pulsed light with stGtACR2 and we remind all scientists of the value of researching and fully understanding the limitations of the tools being used.

5.4 Method Details

Mice were not experimentally naïve. Kindling experiments took place after the optogenetic experiments described in chapter 4 (see Methods for that chapter). Following the conclusion of the previous experiments, mice were given 2 days to recover and hydrate following water restriction. The kindling protocol is as described in the main text, and all surgical and light delivery methods can be found in the Methods section of Chapter 4.

Aitchison, L., Jegminat, J., Menendez, J. A., Pfister, J.-P., Pouget, A., and Latham, P. E. (2021). Synaptic plasticity as bayesian inference. *Nature neuroscience*, 24(4):565–571. 41

Andersen, P., Morris, R., Amaral, D., Bliss, T., and O'Keefe, J. (2007). *The hippocampus book*. Oxford university press. 19, 30

Aronov, D., Nevers, R., and Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543(7647):719–722. 18, 19

Attardo, A., Fitzgerald, J. E., and Schnitzer, M. J. (2015). Impermanence of dendritic spines in live adult ca1 hippocampus. *Nature*, 523(7562):592–596. 41

Baddeley, A. (1992). Working memory. *Science*, 255(5044):556–559. 65

Baddeley, A. (2012). Working memory: Theories, models, and controversies. *Annual review of psychology*, 63(1):1–29. 3, 65

Baddeley, A. D., Thomson, N., and Buchanan, M. (1975). Word length and the structure of short-term memory. *Journal of verbal learning and verbal behavior*, 14(6):575–589. 3

Barak, O., Tsodyks, M., and Romo, R. (2010). Neuronal population coding of parametric working memory. *Journal of Neuroscience*, 30(28):9424–9430. 9

Basu, J., Zaremba, J. D., Cheung, S. K., Hitti, F. L., Zemelman, B. V., Losonczy, A., and Siegelbaum, S. A. (2016). Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science*, 351(6269):aaa5694. 40

Bekkers, J. M. and Suzuki, N. (2013). Neurons and circuits for odor processing in the piriform cortex. *Trends in neurosciences*, 36(7):429–438. 78

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. *Nature neuroscience*, 18(8):1133–1142. 14, 18, 39

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*, 357(6355):1033–1036. 14, 18, 19, 39

Bliss, T. V. and Lømo, T. (1973). Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of physiology*, 232(2):331–356. 13, 19

Buzsáki, G. and Tingley, D. (2018). Space and time: the hippocampus as a sequence generator. *Trends in cognitive sciences*, 22(10):853–869. 10, 18

Canto, C. B., Wouterlood, F. G., and Witter, M. P. (2008). What does the anatomical organization of the entorhinal cortex tell us? *Neural plasticity*, 2008(1):381243. 4, 87

Caporale, N. and Dan, Y. (2008). Spike timing–dependent plasticity: a hebbian learning rule. *Annu. Rev. Neurosci.*, 31(1):25–46. 19

Cela, E., McFarlan, A. R., Chung, A. J., Wang, T., Chierzi, S., Murai, K. K., and Sjöström, P. J. (2019). An optogenetic kindling model of neocortical epilepsy. *Scientific Reports*, 9(1):5236. 104

Chafee, M. V. and Goldman-Rakic, P. S. (1998). Matching patterns of activity in primate prefrontal area 8a and parietal area 7ip neurons during a spatial working memorytask. *Journal of neurophysiology*, 79(6):2919–2940. 9

Chamberland, S. and Topolnik, L. (2012). Inhibitory control of hippocampal inhibitory neurons. *Frontiers in neuroscience*, 6:165. 5

Chapuis, J., Cohen, Y., He, X., Zhang, Z., Jin, S., Xu, F., and Wilson, D. A. (2013). Lateral entorhinal modulation of piriform cortical activity and fine odor discrimination. *Journal of Neuroscience*, 33(33):13449–13459. 87

Cholvin, T., Hainmueller, T., and Bartos, M. (2021). The hippocampus converts dynamic entorhinal inputs into stable spatial maps. *Neuron*, 109(19):3135–3148. 19

Cravo, A. M., Rohenkohl, G., Wyart, V., and Nobre, A. C. (2011). Endogenous modulation of low frequency oscillations by temporal expectations. *Journal of neurophysiology*, 106(6):2964–2972. 66

Cueva, C. J., Saez, A., Marcos, E., Genovesio, A., Jazayeri, M., Romo, R., Salzman, C. D., Shadlen, M. N., and Fusi, S. (2020). Low-dimensional dynamics for working memory and time encoding. *Proceedings of the National Academy of Sciences*, 117(37):23021–23032. 65, 66, 78

Curtis, C. E. and D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in cognitive sciences*, 7(9):415–423. 9

Curtis, C. E. and Sprague, T. C. (2021). Persistent activity during working memory from front to back. *Frontiers in Neural Circuits*, 15:696060. 9

Deshmukh, S. S. and Knierim, J. J. (2011). Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Frontiers in behavioral neuroscience*, 5:69. 19

Deuchars, J. and Thomson, A. (1996). Ca1 pyramid-pyramid connections in rat hippocampus in vitro: dual intracellular recordings with biocytin filling. *Neuroscience*, 74(4):1009–1018. 78

Driscoll, L. N., Duncker, L., and Harvey, C. D. (2022). Representational drift: Emerging theories for continual learning and experimental future directions. *Current Opinion in Neurobiology*, 76:102609. 41

Eichenbaum, H. (2014). Time cells in the hippocampus: a new dimension for mapping memories. *Nature Reviews Neuroscience*, 15(11):732–744. 10, 18, 42

Eichenbaum, H. (2017). The role of the hippocampus in navigation is memory. *Journal of neurophysiology*, 117(4):1785–1796. 6

Emiliani, V., Entcheva, E., Hedrich, R., Hegemann, P., Konrad, K. R., Lüscher, C., Mahn, M., Pan, Z.-H., Sims, R. R., Vierock, J., et al. (2022). Optogenetics for light control of biological systems. *Nature Reviews Methods Primers*, 2(1):55. 103

Feldman, D. E. (2012). The spike-timing dependence of plasticity. *Neuron*, 75(4):556–571. 19

Funahashi, S., Bruce, C. J., and Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of neurophysiology*, 61(2):331–349. 8, 65, 78

Fuster, J. M. and Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, 173(3997):652–654. 8, 65, 77

Fyhn, M., Molden, S., Witter, M. P., Moser, E. I., and Moser, M.-B. (2004). Spatial representation in the entorhinal cortex. *Science*, 305(5688):1258–1264. 19

Goddard, G. V. (1967). Development of epileptic seizures through brain stimulation at low intensity. *Nature*, 214(5092):1020–1021. 104

Grienberger, C., Giovannucci, A., Zeiger, W., and Portera-Cailliau, C. (2022). Two-photon calcium imaging of neuronal activity. *Nature Reviews Methods Primers*, 2(1):67. 11

Grienberger, C. and Konnerth, A. (2012). Imaging calcium in neurons. *Neuron*, 73(5):862–885. 11

Grienberger, C. and Magee, J. C. (2022). Entorhinal cortex directs learning-related changes in ca1 representations. *Nature*, 611(7936):554–562. 14, 18, 19, 20, 26, 39, 41

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. *Nature neuroscience*, 20(3):417–426. 14, 18, 39

Hafting, T., Fyhn, M., Molden, S., Moser, M.-B., and Moser, E. I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, 436(7052):801–806. 6, 19, 87

Hainmueller, T. and Bartos, M. (2018). Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, 558(7709):292–296. 15

Hampson, R. E., Simeral, J. D., and Deadwyler, S. A. (1999). Distribution of spatial and nonspatial information in dorsal hippocampus. *Nature*, 402(6762):610–614. 18

Hebb, D. O. (2005). *The organization of behavior: A neuropsychological theory.* Psychology press. 13

Herculano-Houzel, S. (2009). The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in human neuroscience*, 3:857. 3

Holtmaat, A. and Caroni, P. (2016). Functional and structural underpinnings of neuronal assembly formation in learning. *Nature neuroscience*, 19(12):1553–1562. 42

Howard, J. D., Plailly, J., Grueschow, M., Haynes, J.-D., and Gottfried, J. A. (2009). Odor quality coding and categorization in human posterior piriform cortex. *Nature neuroscience*, 12(7):932–938. 78

Igarashi, K. M., Lu, L., Colgin, L. L., Moser, M.-B., and Moser, E. I. (2014). Coordination of entorhinal–hippocampal ensemble activity during associative learning. *Nature*, 510(7503):143–147. 12, 19, 27, 42, 78, 87

Inagaki, H. K., Fontolan, L., Romani, S., and Svoboda, K. (2019). Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature*, 566(7743):212–217. 65, 78

Inagaki, H. K., Inagaki, M., Romani, S., and Svoboda, K. (2018). Low-dimensional and monotonic preparatory activity in mouse anterior lateral motor cortex. *Journal of Neuroscience*, 38(17):4163–4185. 65, 78

Itskov, V., Curto, C., Pastalkova, E., and Buzsáki, G. (2011). Cell assembly sequences arising from spike threshold adaptation keep track of time in the hippocampus. *Journal of Neuroscience*, 31(8):2828–2834. 78

Kappel, D., Habenschuss, S., Legenstein, R., and Maass, W. (2015). Network plasticity as bayesian inference. *PLoS computational biology*, 11(11):e1004485. 41

Kepecs, A. and Fishell, G. (2014). Interneuron cell types are fit to function. *Nature*, 505(7483):318–326. 40

Kerr, K. M., Agster, K. L., Furtak, S. C., and Burwell, R. D. (2007). Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. *Hippocampus*, 17(9):697–708. 87

Klausberger, T. and Somogyi, P. (2008). Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science*, 321(5885):53–57. 40

Knierim, J. J. (2015). The hippocampus. *Current Biology*, 25(23):R1116–R1121. 4, 87

Knierim, J. J., Neunuebel, J. P., and Deshmukh, S. S. (2014). Functional correlates of the lateral and medial entorhinal cortex: objects, path integration and local–global reference frames. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635):20130369. 7, 19, 87

Knowles, W. D. and Schwartzkroin, P. A. (1981). Local circuit synaptic interactions in hippocampal brain slices. *Journal of Neuroscience*, 1(3):318–322. 78

Kraus, B. J., Brandon, M. P., Robinson, R. J., Connerney, M. A., Hasselmo, M. E., and Eichenbaum, H. (2015). During running in place, grid cells integrate elapsed time and distance run. *Neuron*, 88(3):578–589. 78

Kropff, E., Carmichael, J. E., Moser, M.-B., and Moser, E. I. (2015). Speed cells in the medial entorhinal cortex. *Nature*, 523(7561):419–424. 6

Ku, S.-p., Nakamura, N. H., Maingret, N., Mahnke, L., Yoshida, M., and Sauvage, M. M. (2017). Regional specific evidence for memory-load dependent activity in the dorsal subiculum and the lateral entorhinal cortex. *Frontiers in systems neuroscience*, 11:51. 7

Kubota, K. and Niki, H. (1971). Prefrontal cortical unit activity and delayed alternation performance in monkeys. *Journal of neurophysiology*, 34(3):337–347. 8, 65, 78

Kuruvilla, M. V. and Ainge, J. A. (2017). Lateral entorhinal cortex lesions impair local spatial frameworks. *Frontiers in Systems Neuroscience*, 11:30. 7

Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338. 18

Le Duigou, C., Simonnet, J., Teleñczuk, M. T., Fricker, D., and Miles, R. (2014). Recurrent synapses and circuits in the ca3 region of the hippocampus: an associative network. *Frontiers in cellular neuroscience*, 7:262. 4, 78

Lever, C., Burton, S., Jeewajee, A., O'Keefe, J., and Burgess, N. (2009). Boundary vector cells in the subiculum of the hippocampal formation. *Journal of Neuroscience*, 29(31):9771–9777. 6

Levy, W. and Steward, O. (1983). Temporal contiguity requirements for long-term associative potentiation/depression in the hippocampus. *Neuroscience*, 8(4):791–797. 13

Li, X.-G., Somogyi, P., Ylinen, A., and Buzsáki, G. (1994). The hippocampal ca3 network: an in vivo intracellular labeling study. *Journal of comparative neurology*, 339(2):181–208. 4, 78

Li, Y., Xu, J., Liu, Y., Zhu, J., Liu, N., Zeng, W., Huang, N., Rasch, M. J., Jiang, H., Gu, X., et al. (2017). A distinct entorhinal cortex to hippocampal ca1 direct circuit for olfactory associative learning. *Nature neuroscience*, 20(4):559–570. 12, 19, 27, 42, 78, 87, 88, 96

Liu, D., Gu, X., Zhu, J., Zhang, X., Han, Z., Yan, W., Cheng, Q., Hao, J., Fan, H., Hou, R., et al. (2014). Medial prefrontal activity during delay period contributes to learning of a working memory task. *Science*, 346(6208):458–463. 21, 67

Liu, J. K. and Buonomano, D. V. (2009). Embedding multiple trajectories in simulated recurrent neural networks in a self-organizing manner. *Journal of Neuroscience*, 29(42):13172–13181. 42, 66, 77

Lu, L., Leutgeb, J. K., Tsao, A., Henriksen, E. J., Leutgeb, S., Barnes, C. A., Witter, M. P., Moser, M.-B., and Moser, E. I. (2013). Impaired hippocampal rate coding after lesions of the lateral entorhinal cortex. *Nature neuroscience*, 16(8):1085–1093. 7, 87

Ma, M., de Souza, F. S., Futia, G. L., Anderson, S. R., Riguero, J., Tollin, D., Gentile-Polese, A., Platt, J. P., Steinke, K., Hiratani, N., et al. (2024). Sequential activity of ca1 hippocampal cells constitutes a temporal memory map for associative learning in mice. *Current Biology*, 34(4):841–854. 65

MacDonald, C. J., Carrow, S., Place, R., and Eichenbaum, H. (2013). Distinct hippocampal time cell sequences represent odor memories in immobilized rats. *Journal of Neuroscience*, 33(36):14607–14616. 9, 78

MacDonald, C. J., Lepage, K. Q., Eden, U. T., and Eichenbaum, H. (2011). Hippocampal "time cells" bridge the gap in memory for discontiguous events. *Neuron*, 71(4):737–749. 9, 10, 18, 42, 78

Magee, J. C. and Grienberger, C. (2020). Synaptic plasticity forms and functions. *Annual review of neuroscience*, 43:95–117. 14, 18, 19, 39, 41

Magnus, C. J., Lee, P. H., Bonaventura, J., Zemla, R., Gomez, J. L., Ramirez, M. H., Hu, X., Galvan, A., Basu, J., Michaelides, M., et al. (2019). Ultrapotent chemogenetics for research and potential clinical applications. *Science*, 364(6436):eaav5282. 26

Mahn, M., Gibor, L., Patil, P., Cohen-Kashi Malina, K., Oring, S., Printz, Y., Levy, R., Lampl, I., and Yizhar, O. (2018). High-efficiency optogenetic silencing with soma-targeted anion-conducting channelrhodopsins. *Nature communications*, 9(1):4125. 103

Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215. 13

Masse, N. Y., Yang, G. R., Song, H. F., Wang, X.-J., and Freedman, D. J. (2019). Circuit mechanisms for the maintenance and manipulation of information in working memory. *Nature neuroscience*, 22(7):1159–1167. 9

Mayford, M., Siegelbaum, S. A., and Kandel, E. R. (2012). Synapses and memory storage. *Cold Spring Harbor perspectives in biology*, 4(6):a005751. 19

McKenzie, S. and Eichenbaum, H. (2011). Consolidation and reconsolidation: two lives of memories? *Neuron*, 71(2):224–233. 5

Messier, J. E., Chen, H., Cai, Z.-L., and Xue, M. (2018). Targeting light-gated chloride channels to neuronal somatodendritic domain reduces their excitatory effect in the axon. *Elife*, 7:e38506. 104, 107

Micou, C. and O'Leary, T. (2023). Representational drift as a window into neural and behavioural plasticity. *Current opinion in neurobiology*, 81:102746. 15, 41

Milner, B., Corkin, S., and Teuber, H.-L. (1968). Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of hm. *Neuropsychologia*, 6(3):215–234. 1

Milstein, A. D., Bloss, E. B., Apostolides, P. F., Vaidya, S. P., Dilly, G. A., Zemelman, B. V., and Magee, J. C. (2015). Inhibitory gating of input comparison in the ca1 microcircuit. *Neuron*, 87(6):1274–1289. 40

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:e73046. 14, 18, 19, 39, 41

Minerbi, A., Kahana, R., Goldfeld, L., Kaufman, M., Marom, S., and Ziv, N. E. (2009). Long-term relationships between synaptic tenacity, synaptic remodeling, and network activity. *PLoS biology*, 7(6):e1000136. 41

Morris, R. G., Anderson, E., Lynch, G. S., and Baudry, M. (1986). Selective impairment of learning and blockade of long-term potentiation by an n-methyl-d-aspartate receptor antagonist, ap5. *Nature*, 319(6056):774–776. 6

Moser, E. I., Moser, M.-B., and McNaughton, B. L. (2017). Spatial representation in the hippocampal formation: a history. *Nature neuroscience*, 20(11):1448–1464. 18

Murray, J. M. and Escola, G. S. (2017). Learning multiple variable-speed sequences in striatum via cortical tutoring. *Elife*, 6:e26084. 66, 77

Nakazawa, K., McHugh, T. J., Wilson, M. A., and Tonegawa, S. (2004). Nmda receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience*, 5(5):361–372. 6

Narayanan, N. S. (2016). Ramping activity is a cortical mechanism of temporal control of action. *Current opinion in behavioral sciences*, 8:226–230. 65

Nobre, A. C. and Van Ede, F. (2018). Anticipated moments: temporal structure in attention. *Nature Reviews Neuroscience*, 19(1):34–48. 10

O'Keefe, J. and Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat. *Brain research*. 6, 18

Olson, I. R., Page, K., Moore, K. S., Chatterjee, A., and Verfaellie, M. (2006). Working memory for conjunctions relies on the medial temporal lobe. *Journal of Neuroscience*, 26(17):4596–4601. 91

Pachitariu, M., Stringer, C., Dipoppa, M., Schröder, S., Rossi, L. F., Dalgleish, H., Carandini, M., and Harris, K. D. (2017). Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv*. 21, 30, 48, 84

Pastalkova, E., Itskov, V., Amarasingham, A., and Buzsaki, G. (2008). Internally generated cell assembly sequences in the rat hippocampus. *Science*, 321(5894):1322–1327. 9, 10, 18, 42, 65, 78

Posner, M. I. and Konick, A. F. (1966). Short-term retention of visual and kinesthetic information. *Organizational Behavior and Human Performance*, 1(1):71–86. 3

Qin, S., Farashahi, S., Lipshutz, D., Sengupta, A. M., Chklovskii, D. B., and Pehlevan, C. (2023). Coordinated drift of receptive fields in hebbian/anti-hebbian network models during noisy representation learning. *Nature Neuroscience*, 26(2):339–349. 41

Racine, R. J. (1972). Modification of seizure activity by electrical stimulation: I. after-discharge threshold. *Electroencephalography and clinical neurophysiology*, 32(3):269–279. 105

Rajan, K., Harvey, C. D., and Tank, D. W. (2016). Recurrent network models of sequence generation and memory. *Neuron*, 90(1):128–142. 65, 77

Riley, M. R. and Constantinidis, C. (2016). Role of prefrontal persistent activity in working memory. *Frontiers in systems neuroscience*, 9:181. 78

Rule, M. E., O'Leary, T., and Harvey, C. D. (2019). Causes and consequences of representational drift. *Current opinion in neurobiology*, 58:141–147. 15, 41

Salz, D. M., Tiganj, Z., Khasnabish, S., Kohley, A., Sheehan, D., Howard, M. W., and Eichenbaum, H. (2016). Time cells in hippocampal area ca3. *Journal of Neuroscience*, 36(28):7476–7484. 42

Sargolini, F., Fyhn, M., Hafting, T., McNaughton, B. L., Witter, M. P., Moser, M.-B., and Moser, E. I. (2006). Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science*, 312(5774):758–762. 19

Save, E. and Sargolini, F. (2017). Disentangling the role of the mec and lec in the processing of spatial and non-spatial information: contribution of lesion studies. *Frontiers in systems neuroscience*, 11:81. 7, 87

Sheintuch, L., Rubin, A., Brande-Eilat, N., Geva, N., Sadeh, N., Pinchasof, O., and Ziv, Y. (2017). Tracking the same neurons across multiple days in ca2+ imaging data. *Cell reports*, 21(4):1102–1115. 38, 48

Silva, A. J., Paylor, R., Wehner, J. M., and Tonegawa, S. (1992). Impaired spatial learning in α-calcium-calmodulin kinase ii mutant mice. *Science*, 257(5067):206–211. 6

Song, Y., Wang, T., Cai, P., Mondal, S. K., and Sahoo, J. P. (2023). A comprehensive survey of few-shot learning: Evolution, applications, challenges, and opportunities. *ACM Computing Surveys*, 55(13s):1–40. 18

Squire, L. R. (1992). Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. *Psychological review*, 99(2):195. 18

Squire, L. R. (2009). The legacy of patient hm for neuroscience. *Neuron*, 61(1):6–9. 1, 2, 3

Squire, L. R. and Zola-Morgan, S. (1991). The medial temporal lobe memory system. *Science*, 253(5026):1380–1386. 4

Stokes, M. G. (2015). 'activity-silent'working memory in prefrontal cortex: a dynamic coding framework. *Trends in cognitive sciences*, 19(7):394–405. 9

Strauch, C. and Manahan-Vaughan, D. (2020). Orchestration of hippocampal information encoding by the piriform cortex. *Cerebral Cortex*, 30(1):135–147. 78

Takahashi, H. and Magee, J. C. (2009). Pathway interactions and synaptic plasticity in the dendritic tuft regions of ca1 pyramidal neurons. *Neuron*, 62(1):102–111. 26

Taube, J. S., Muller, R. U., and Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. i. description and quantitative analysis. *Journal of Neuroscience*, 10(2):420–435. 6

Taxidis, J., Madruga, B., Melin, M. D., Lin, M. Z., and Golshani, P. (2023). Voltage imaging reveals that hippocampal interneurons tune memory-encoding pyramidal sequences. *bioRxiv*, pages 2023–04. 79

Taxidis, J., Pnevmatikakis, E. A., Dorian, C. C., Mylavarapu, A. L., Arora, J. S., Samadian, K. D., Hoffberg, E. A., and Golshani, P. (2020). Differential emergence and stability of sensory and temporal representations in context-specific hippocampal sequences. *Neuron*, 108(5):984–998. 9, 11, 12, 15, 18, 20, 21, 36, 44, 52, 65, 66, 67, 70, 78, 81, 85, 96, 100

Thom, M. (2014). Hippocampal sclerosis in epilepsy: a neuropathology review. *Neuropathology and applied neurobiology*, 40(5):520–543. 104

Tsao, A., Sugar, J., Lu, L., Wang, C., Knierim, J. J., Moser, M.-B., and Moser, E. I. (2018). Integrating time from experience in the lateral entorhinal cortex. *Nature*, 561(7721):57–62. 7, 78, 87

Umbach, G., Kantak, P., Jacobs, J., Kahana, M., Pfeiffer, B. E., Sperling, M., and Lega, B. (2020). Time cells in the human hippocampus and entorhinal cortex support episodic memory. *Proceedings of the National Academy of Sciences*, 117(45):28463–28474. 78

Vaidya, S. P., Chitwood, R. A., and Magee, J. C. (2023). The formation of an expanding memory representation in the hippocampus. *bioRxiv*. 42

Van Strien, N., Cappaert, N., and Witter, M. (2009). The anatomy of memory: an interactive overview of the parahippocampal–hippocampal network. *Nature reviews neuroscience*, 10(4):272–282. 19, 30

Vangkilde, S., Petersen, A., and Bundesen, C. (2013). Temporal expectancy in the context of a theory of visual attention. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1628):20130054. 66

Varga, C., Lee, S. Y., and Soltesz, I. (2010). Target-selective gabaergic control of entorhinal cortex output. *Nature neuroscience*, 13(7):822–824. 5

Watt, A. J., Van Rossum, M. C., MacLeod, K. M., Nelson, S. B., and Turrigiano, G. G. (2000). Activity coregulates quantal ampa and nmda currents at neocortical synapses. *Neuron*, 26(3):659–670. 13

Wood, E. R., Dudchenko, P. A., and Eichenbaum, H. (1999). The global record of memory in hippocampal neuronal activity. *Nature*, 397(6720):613–616. 18

Woods, N. I., Stefanini, F., Apodaca-Montano, D. L., Tan, I. M., Biane, J. S., and Kheirbek, M. A. (2020). The dentate gyrus classifies cortical representations of learned stimuli. *Neuron*, 107(1):173–184. 12, 19, 27, 42, 78, 87

Xiao, K., Li, Y., Chitwood, R. A., and Magee, J. C. (2023). A critical role for camkii in behavioral timescale synaptic plasticity in hippocampal ca1 pyramidal neurons. *Science Advances*, 9(36):eadi3088. 14, 18, 39

Yang, S., Yang, S., Moreira, T., Hoffman, G., Carlson, G. C., Bender, K. J., Alger, B. E., and Tang, C.-M. (2014). Interlamellar ca1 network in the hippocampus. *Proceedings of the National Academy of Sciences*, 111(35):12919–12924. 78

Yasumatsu, N., Matsuzaki, M., Miyazaki, T., Noguchi, J., and Kasai, H. (2008). Principles of long-term dynamics of dendritic spines. *Journal of Neuroscience*, 28(50):13592–13608. 41

Zemla, R. and Basu, J. (2017). Hippocampal function in rodents. *Current opinion in neurobiology*, 43:187–197. 5, 87

Zhang, Y. J., Lee, J. Y., and Igarashi, K. M. (2024). Circuit dynamics of the olfactory pathway during olfactory learning. *Frontiers in Neural Circuits*, 18:1437575. 12, 19, 27, 42, 78, 87

Zhao, X., Wang, Y., Spruston, N., and Magee, J. C. (2020). Membrane potential dynamics underlying context-dependent sensory responses in the hippocampus. *Nature neuroscience*, 23(7):881–891. 14, 18, 39

Zhou, S., Masmanidis, S. C., and Buonomano, D. V. (2020). Neural sequences as an optimal dynamical regime for the readout of time. *Neuron*, 108(4):651–658. 65

Zhou, S., Seay, M., Taxidis, J., Golshani, P., and Buonomano, D. V. (2023). Multiplexing working memory and time in the trajectories of neural networks. *Nature Human Behaviour*, 7(7):1170–1184. 10, 65, 66, 68, 77, 79

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., Gamal, A. E., and Schnitzer, M. J. (2013). Long-term dynamics of ca1 hippocampal place codes. *Nature neuroscience*, 16(3):264–266. 15