

THE UNIVERSITY OF CHICAGO

DYNAMIC HEMISPHERIC ASYMMETRY IN THE HIPPOCAMPUS THROUGHOUT
MEMORY FORMATION AND RETRIEVAL

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE BIOLOGICAL SCIENCES
AND THE PRITZKER SCHOOL OF MEDICINE
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

COMMITTEE ON COMPUTATIONAL NEUROSCIENCE

BY
ANQI JIANG

CHICAGO, ILLINOIS

JUNE 2025

Copyright © 2025 by Anqi Jiang

All Rights Reserved

This dissertation is dedicated to my grandma 杜素珍

Contents

List of Figures	vi
Acknowledgments	viii
Abstract	1
1. INTRODUCTION	3
2. A PREPROCESSING TOOLBOX FOR 2-PHOTON IMAGING	20
3. A LEARNING-DEPENDENT SHIFT IN CA3 INPUT DOMINANCE FROM LEFT TO RIGHT DETERMINES THE EVOLUTION OF CA1 SPATIAL MAPS	47
4. CONCLUSION AND FUTURE DIRECTIONS	89
Bibliography	100

List of Figures

Figure 1.1: Illustration of trisynaptic circuit	7
Figure 1.2: Illustration of CA1 layers and circuitry.	11
Figure 2.1: Imaging CA3 axons in CA1	27
Figure 2.2: Example time-series fluorescent traces from CA3 axons in CA1 before and after smoothing	27
Figure 2.3 Frequency band-based axon ROI selection.	31
Figure 2.4 Identification of movement artifacts in time-series traces	33
Figure 2.5 Baseline correction and peak detection	36
Figure 2.6 Comparison of clustering methods used for grouping ROIs.	40
Figure 2.7 SUBPREP pipeline summary.	42
Figure 3.1 Anticipatory licking as a behavioral readout of spatial learning and familiarization to a novel environment.	51
Figure 3.2 CA1 stabilization occurs within the first ~10 laps in a novel environment.	53
Figure 3.3 Optogenetic inhibition of left and right CA3 axons reduces CA1 pyramidal cell activity.	56
Figure 3.4 CA1 low-dimensional neural manifolds and spatial decoding reveal a dynamic contribution of left and right CA3 inputs that switches between initial learning and post-learning phases	59
Figure 3.5 The influence of left and right CA3 inputs on the emergence, reliability, and gain of CA1 place fields during initial learning and post-learning phases	

Figure 3.6 Recordings of CA3 axons reveal a switch from stronger right CA3 axon activity during the early phase to stronger left CA3 axon activity during the later phase.	64
Figure 3.7 CA3 Optogenetic inhibition did not alter the mean running speed of animals in novel environments.	68
Figure 3.8 Right optogenetic inhibition significantly increased the distance to centroid in the same example animal as Figure 3.4, representing a less stereotypical neural activity to location on the track under opto.	72
Figure 3.9 Left and right CA3 input inhibition similarly delayed the emergence of CA1 place cells, regardless of emergence criteria.	74
Figure 3.10 Left and Right CA3 axon place field features.	76

Acknowledgments

I am grateful to my advisor and mentor Dr. Mark Sheffield for his unwavering support and encouragement over the past six years. From the very beginning, Mark believed in me and made me feel welcomed and capable, even when I was unsure of myself. His ability to foster meaningful discussion and his openness to questions created an environment where I felt safe to learn, to fail, and to grow. I am extremely grateful for all his efforts and ways he invested in me growing as a scientist and person. He taught me to think critically about experiments, supported me through roadblocks, and celebrated every success, big or small. I feel incredibly lucky to have had his mentorship throughout this journey.

I would also like to thank the past and present members of the Sheffield lab—Seetha Krishnan, Antoine Madar, Douglas Goodsmith, Mikayla Voglewede, Can Dong, Chad Heer, Heather Macomber, Yuhung “Tommy” Chiu, Ariana Tortolani, Julliana Ramirez-Matias, Yizhi “Cherry” Wang, Bryan Garcia, Chery Cherian, Bennett Scott, Roma Shah—for have not only given me numerous insightful feedbacks along the way, but also made working in the lab truly enjoyable. I want to give special thanks to Seetha, Can, and Chad for teaching me and troubleshooting with me on how to do mouse surgeries and acquire good imaging planes. Seetha, Antoine, and Doug for all the fun and invaluable science discussion and project feedbacks that not only shaped my project but also challenged me to be a better scientist. Heather and Ariana for unwavering positivity and support through difficult times. Roma for supporting my project with her hard work.

I am also deeply thankful to my thesis committee members—Dr. Leslie Kay, Dr. Jason MacLean, and Dr. Narayanan “Bobby” Kasthuri—for their invaluable insights and thoughtful

discussions. It has been a true privilege to learn from them. Their guidance continually challenged me to think more broadly and deeply about science, and has profoundly enriched my research.

Outside the lab, I would like to thank the friends who have supported me through this journey. From weekly board games night to late night snacks, they are my rock who stood by me through every big and small events in life, high and low. Their kindness, resilience, warmth, intelligence, and curiosity continually remind me how lucky I am to be here and to have them in my life.

Finally, none of this would have been possible without the love and support of my family and people whom I consider as family. I would not be able pursue my dreams in science and live a life true to myself without their support and encouragement throughout life. I am especially thankful to Mary Kaye and John, who have become family to me in every sense of the word. Their kindness, wisdom, and unwavering support carried me through one of the most difficult times in my life. More than just support, they've been my role models—showing me how to live with resilience, empathy, and grace. Their impact on my life is profound, and I will always be grateful for the love and inspiration they have given me.

ABSTRACT

The hippocampus is critical for forming and stabilizing spatial memories. Within the classic trisynaptic circuit inside the hippocampus, CA1 is the output region of the hippocampus and place cells in CA1 are considered to collectively support spatial representations. While the contribution of unilateral or bilateral CA3 inputs to CA1 (Schaffer collateral pathway) have been well characterized, less is known about how inputs from the left and right CA3 differentially shape spatial representation in CA1. Most models of hippocampal function treat CA3 as a unified region and overlook the possibility of hemispheric specialization within its projections. Recent work has pointed to lateralized features of hippocampal circuitry primarily at the molecular and behavioral levels, yet the functional relevance of this hemispheric asymmetry among place cells *in vivo* remains poorly understood.

In this dissertation, I examine how left and right CA3 projections differentially contribute to the development and stabilization of spatial representations in right CA1 (CA1R) during learning. Using two-photon calcium imaging and optogenetic inhibition in head-fixed mice navigating a virtual environment, I characterize how CA1R spatial maps evolve over experience and how CA3 inputs support this process across distinct phases of learning. I find that CA1R maps emerge upon immediate novel exposure but are initially inaccurate. Over repeated laps, spatial maps gradually improve and stabilize after ~10 laps (early-phase), with later laps marking stability (late-phase). In the early phase, both CA3 inputs contribute to place field formation, but right CA3 inputs predominantly drive high-amplitude, reliable fields that support the

development of accurate spatial representations. In the late phase, left CA3 inputs become more prominent, supporting the maintenance of stable and reliable fields.

Complementary recordings of CA3 axonal activity within CA1R further reveal this dynamic hemispheric shift. Right CA3 axons exhibit elevated activity during the early phase, consistent with a role in supporting novel experience encoding. As learning continues and the environment becomes familiar, left CA3 axons show increased activity, aligning with their involvement in maintaining stable representations. These findings reveal a dynamic, experience-dependent shift in hemispheric contributions to spatial coding in CA1, moving from right-dominant during early learning to left-dominant during later stabilization.

Taken together, these results suggest that CA3 inputs are functionally lateralized and that their roles in spatial representation evolve across phases of learning. In Chapter 1, I introduce the motivation for this work and review prior literature on hippocampal lateralization and familiarization in novel environments. In Chapter 2, I describe the computational methods I developed to detect and functionally analyze the activity of CA3 inputs with axonal resolution. In Chapter 3, I present evidence for hemispheric differences in the dynamics and influence of CA3 projections on CA1 spatial representations and how these differences switch between learning and post-learning phases of experiences. This work provides new insight into how dynamic lateralized hippocampal inputs support the evolution of spatial representations and lays the foundation for future work on experience-dependent memory processing.

CHAPTER 1

INTRODUCTION

Learning and memory are essential functions of the brain. The ability to learn from experience and retrieve information when needed is crucial for adaptive behavior and survival by avoiding danger and maximizing rewards. Beyond just survival, in humans, this cognitive capacity is also central to our sense of self, shaping our perception of the world and interaction with the environment. When the brain's memory systems are impaired in neurological disorders such as Alzheimer's disease (Masurkar 2018; Fjell et al. 2014), individuals experience cognitive deficits associated with memory loss and troubles with everyday tasks. Gaining insights into how memories are formed, consolidated, and retrieved is critical for comprehending brain functions and serves as a foundational step to developing effective treatments to memory disorders.

Over decades of memory research, memory can be divided into short-term memory and long-term memory (Norris 2017; Cowan 2008; Andersen et al. 2006). Short-term memory transiently holds information relevant to the immediate tasks. The duration is estimated to be on the order of seconds. Long-term memory involves the retention of information over extended periods potentially indefinitely. Short-term memory can be transferred to long-term memory. Long-term memory can be further divided into non-declarative memory and declarative memory. Non-declarative memory is acquired and used unconsciously whereas declarative memory refers to all memories that can be consciously recalled. Declarative memory depends on the hippocampus and is the focus of this thesis.

1.1 What does the hippocampus do?

The understanding that the hippocampus is required for the formation of new declarative memory was only widely recognized after a series of studies of the patient H.M.. Patient H.M. suffered from epilepsy and underwent surgical removal of bilateral medial temporal lobes that included bilateral hippocampi (Scoville and Milner 1957) as the treatment. After the surgery, H.M. had epilepsy under control but suffered from inability to form new declarative memories. Interestingly, both procedural memory and declarative memory that existed well before his surgery were undisturbed. The early descriptions of H.M. indicated the importance of the hippocampus to forming new declarative memory and marked the beginning of the modern era memory research (Squire 2009).

Inspired by the patient studies, researchers began recording from rat hippocampus to investigate what memories looked like at the single cell level. They initially anticipated that these cells would encode specific aspects of experience, such as visual or sensory details from the environment (O'Keefe 2014). However, surprisingly, they discovered “place cells” that fired consistently at specific locations of a given environment (O'Keefe and Dostrovsky 1971). These cells provided the first physiological basis for spatial representation. Based on the finding, they proposed that the function of the hippocampus is to provide an internal representation of space in the brain through collection of place cells, termed cognitive map theory (Tolman 1948), implying that spatial cognition is at the core of memory. The profound implications of this theory, along with key findings from the parahippocampal formation, were later recognized with a Nobel Prize “for their discoveries of cells that constitute a positioning system in the brain”.

Indeed, place cells are strongly modulated by space and cover the entire path as animals traverse in environments (Dong, Madar, and Sheffield 2021; Krishnan et al. 2022; J. S. Lee et al. 2020). Furthermore, optogenetic excitation of place cells that represent reward locations can

strongly bias reward-related behavior, even when they are activated at locations far from the reward, providing evidence for a causal role of place cells in supporting spatial navigation and memory (Robinson et al. 2020). Nevertheless, place cells encode not just spatial information, but are also shown to be strongly modulated by time (MacDonald et al. 2011; Whittington et al. 2020), sensory cues like auditory tones and odor (Jeffery 2007; Moita et al. 2003; Best, White, and Minai 2001), internal states (Krishnan et al. 2022; Flavell et al. 2022; Fenton 2024), novelty (Kitanishi et al. 2015; Gómez-Ocádiz et al. 2022; Lever et al. 2010; Larkin et al. 2014), rewards (Krishnan et al. 2022; Nieh et al. 2021; J. S. Lee et al. 2020) etc. Beyond encoding specific features of the physical space or supporting general spatial navigation, hippocampal neurons can also form discrete firing fields in response to specific sound frequencies in a stationary non-spatial auditory task (Aronov, Nevers, and Tank 2017). Interestingly, these “frequency fields” also span the entire auditory task space, showing similar properties to place fields. These findings suggest that the hippocampus employs common mechanisms to encode various task-relevant variables, supporting cognitive processes beyond spatial navigation. Spatially tuned place cells therefore may be an example of memory cells, reflecting a special case of memory processes. Extending from just an internal representation of space, the hippocampus provides a representation of episodic memories, reflecting the general experience of “where”, “what”, and “when”. A growing body of research on hippocampal neural physiology supports the view that the hippocampus provides an underlying scaffolding that binds the task-relevant and -irrelevant information together to allow for memories to be sustained, updated, consolidated, and retrieved as animals experience the world (Geva et al. 2023; Keinath, Mosser, and Brandon 2022; Ziv et al. 2013; Klur et al. 2009; R. P. Kesner 2007; Nakazawa et al. 2002; Guo et al. 2024; Gava et al. 2021).

1.2 The anatomy of the hippocampus

The hippocampus is a cortical structure deep in the medial temporal lobe with a seahorse shape. As part of the archicortex, the hippocampus is evolutionarily older than the neocortex and conserved across animal species including mammals, birds, and reptiles, supporting memory functions in all these animal groups (T. A. Allen and Fortin 2013; Andersen et al. 2006; West 1990). It has a very different anatomical structure from the neocortex. For instance, the mammalian neocortex has multiple functional areas, each generally contains six layers of neurons; the hippocampus is divided into two regions, the *dentate gyrus* (DG) and the *cornus ammonis* (CA), each having only three layers of cells. CA can be further divided into CA3, CA2, and CA1, which all are interconnected. In a classic textbook style illustration of hippocampal circuit, the information in the hippocampus is predominately unidirectional through the trisynaptic circuit. Although by contemporary standards the trisynaptic circuit is oversimplified (Knierim 2015), it has been historically important and represents major pathways. The entorhinal cortex (EC) is the main source of inputs to the hippocampus through the perforant path to synapse with the DG neurons. DG neurons project to CA3 and CA2 through mossy fibers. CA3 and CA2 project to CA1 through Schaffer collaterals. CA1 neurons as the output of the hippocampus project to the subiculum and entorhinal cortex.

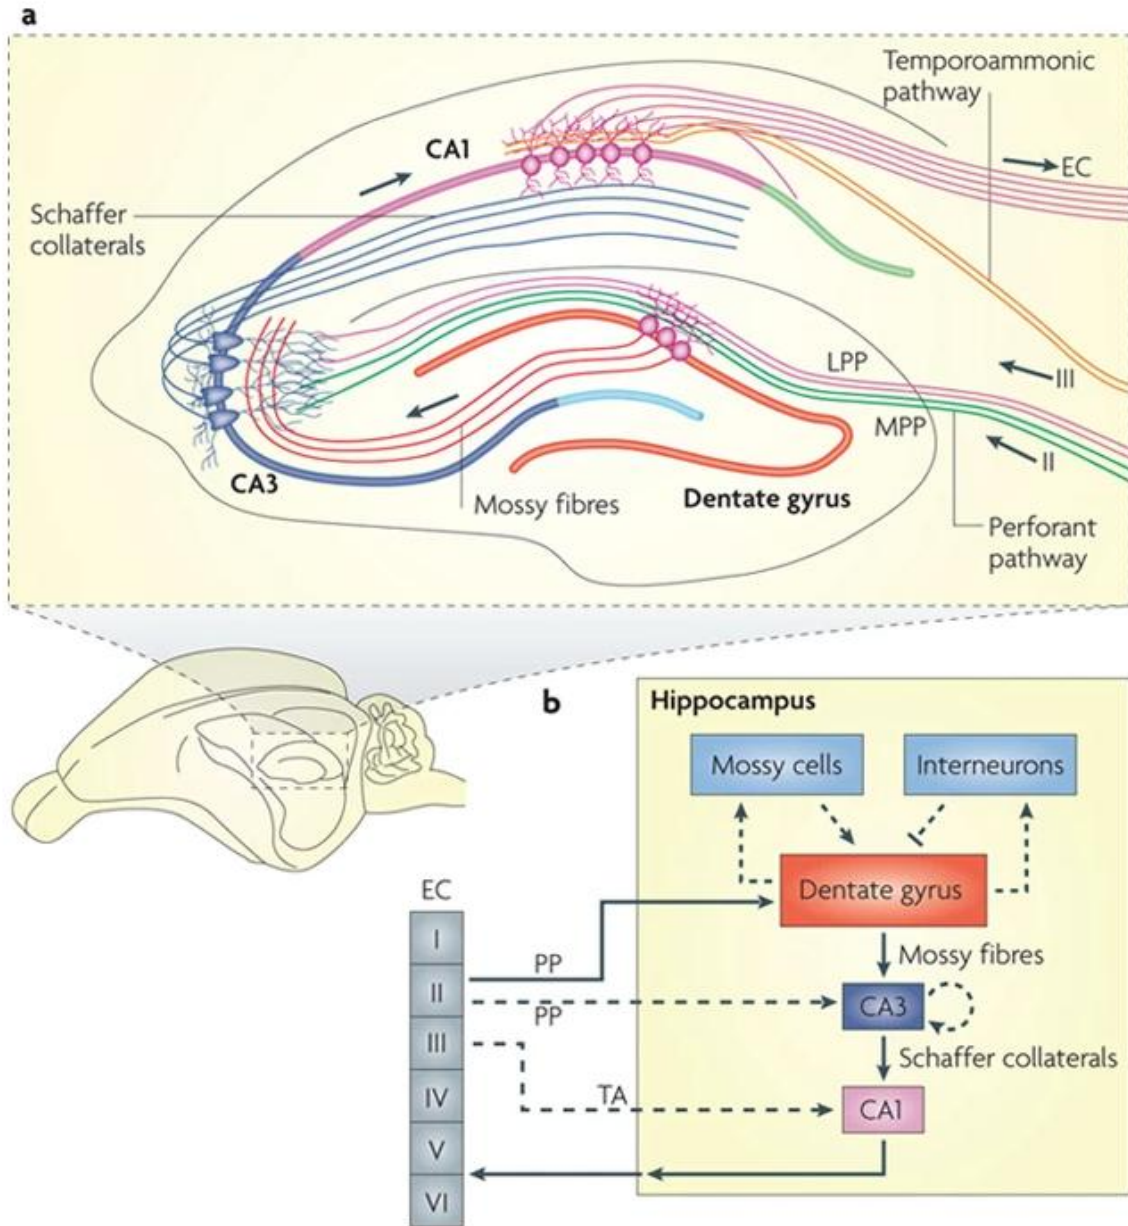


Figure 1.1: Illustration of trisynaptic circuit. (a) Depiction of major excitatory trisynaptic circuit (EC-DG-CA3-CA1-EC). LPP: lateral perforant pathway, MPP: medial perforant pathway. (b) Diagram of hippocampal network. Trisynaptic circuit in solid arrows. CA3 and CA1 also receive direct input from EC. The dentate granule cells also project to the mossy cells in the hilus and hilar interneurons, which send excitatory and inhibitory projections, respectively, back to the granule cells. TA: temporoammonic pathway. (Deng, Aimone, and Gage 2010)

The trisynaptic circuit is traditionally recognized for its major excitatory projections within the hippocampus, serving as a fundamental pathway for the flow of information from the dentate gyrus to CA1. However, this circuit is oversimplified and ignores interhemispheric connections within the hippocampus. In particular, the CA3 region not only project massively to bilateral CA1, it also projects to itself in the same or opposite hemisphere. In rat hippocampus, it is estimated that between 30% and 75% of synapses are formed between pairs of CA3 pyramidal neurons (X.-G. Li et al. 1994) and 0.92% of CA3 neurons within the population are synaptically connected, a rate significantly higher than in a randomly connected network (Guzman et al., 2016). Theorists propose that the robust autoassociative connectivity renders CA3 especially well-suited for *pattern completion*, a computational process through which a complete memory can be retrieved from partial or noisy inputs (Raymond P. Kesner and Rolls 2015; H. Lee, GoodSmith, and Knierim 2020). For example, how we can still recognize the same park in the context of different seasons with different sceneries. Inspired by the biological structure of CA3, the Hopfield network provided a theoretical model for achieving pattern completion, and Hopfield himself was later awarded the Nobel prize of physics in 2024 for this work.

The Hopfield network is at the foundation of today's advancements in artificial intelligence. In the context of hippocampal function, it provides a framework for understanding how neural circuits, particularly in the CA3 region, can retrieve complete memory representations from partial or degraded inputs. However, a key limitation of this model is its restricted storage capacity—only a limited number of patterns can be stored before the network begins to experience memory interference (H. Lee, GoodSmith, and Knierim 2020). This includes issues such as newer memories overwriting older ones or the retrieval of incorrect, overlapping patterns, a phenomenon known as catastrophic interference. To overcome this

challenge, theorists have proposed that *orthogonalization* must occur prior to pattern completion to enhance the fidelity of memory storage and retrieval. The DG is thought to carry out this function through a process termed *pattern separation*, in which it transforms similar sensory inputs from the EC into sparse and decorrelated activity patterns that are then relayed to CA3 (Yassa and Stark 2011; McNaughton and Nadel 1990; Cayco-Gajic and Silver 2019).

Experimental evidence from studies of DG granule cells has further supported the role of the DG in facilitating pattern separation (Madar, Ewell, and Jones 2019; GoodSmith et al. 2019, 2017; Kim et al. 2023; Leutgeb et al. 2007).

Within the trisynaptic circuit, DG is considered as the critical region for pattern separation and CA3 for pattern completion, CA1 is the output region of hippocampus and thus viewed as a key integrator and translator of information to the neocortex (H. Lee, GoodSmith, and Knierim 2020). CA1 receives converging inputs from CA3 via the Schaffer collaterals and from the entorhinal cortex via the temporoammonic pathway, allowing it to integrate signals before transmitting them outward.

In addition to receiving strong excitatory inputs from CA3 and the entorhinal cortex, CA1 also integrates inhibitory inputs from a diverse array of interneuron subtypes. These interneurons, including parvalbumin-expressing (PV+) basket cells and somatostatin-positive (SOM+) oriens-lacunosum moleculare (O-LM) cells, play critical roles in regulating excitatory output and shaping experience-dependent learning through local circuit computations (Jeong and Singer 2022; K. Allen and Monyer 2015; Pedrosa and Clopath 2020). Beyond these canonical circuit elements, neuromodulatory inputs to CA1 also make essential contributions to spatial representation and learning by modulating the excitability and synaptic plasticity of CA1 pyramidal neurons (Takeuchi et al. 2016; Duzskiewicz et al. 2019). These diverse inputs and

their complex interactions enable CA1 to perform functions that go beyond simple feedforward integration. Instead, CA1 engages in dynamic internal processing, allowing for adaptive computation and the refinement of spatial memory representations in response to experience (Cohen, Bolstad, and Lee 2017; Geiller, Priestley, and Losonczy 2023; Bilash et al. 2023; Rolotti et al. 2022; Adaikkan et al. 2024; Pedrosa and Clopath 2020).

To summarize, anatomically, information flows from the entorhinal cortex to DG (perforant path), DG to CA3 (mossy fibers), CA3 to CA1 (Schaffer collaterals), and finally from CA1 back to the entorhinal cortex and other cortical areas, completing the hippocampal-cortical loop. This organization underlies the hippocampus's ability to transform and relay episodic and spatial information with high specificity and fidelity.

1.3 Hemispheric asymmetry in CA3-CA1 synapses

While the trisynaptic circuit provides a foundational framework for understanding hippocampal information processing, it is oversimplified and does not account for significant hemispheric differences in hippocampal structure and function. In particular, evidence reveals that CA3–CA1 synapses exhibit marked hemispheric asymmetries in their molecular composition, synaptic physiology, and functional contributions to memory (Shipton et al. 2014; Kohl et al. 2011; Shinohara et al. 2008; Kawakami et al. 2003; El-Gaby et al. 2016, 2021; Guan et al. 2021; Song et al. 2020; Klur et al. 2009). These findings challenge the assumption of a uniform hippocampal circuit and point to lateralized specializations that may play distinct roles in learning and memory.

CA1 can be divided into four layers: *stratum oriens* (*s.o.*), *stratum pyramidale* (*s.p.*), *stratum radiatum* (*s.r.*), and *stratum lacunosum moleculare* (*s.l.m.*). Pyramidal cells in CA1 have

cell bodies densely packed in the s.p. layer with basal dendrites extending into the s.o. and apical dendrites into the s.r. and s.l.m.. The majority of excitatory inputs from bilateral CA3 synapses with CA1 in the s.o. and s.r.. Interestingly, tracing studies found that at the right dorsal CA1, more of the right CA3 inputs innervate at s.r. as compared to more of the left CA3 inputs innervate at s.o. (the ratio between left and right at s.r.: 1:2 and s.o.: 10:7) (Tao et al. 2021; Shinohara et al. 2012). Such innervation patterns suggest different dendritic computation may occur at the different layers (Ashby et al. 2021; Wright, Hedrick, and Komiyama 2025).




Figure 1.2: Illustration of CA1 layers and circuitry. A. CA1 in relation to other subregions of hippocampus. B. CA1 pyramidal neuron excitatory (triangle) and inhibitory (circle) inputs to different parts of its apical dendrite in SR, SLM and basal dendrite in SO. Distal apical dendrite receives direct input from entorhinal cortex (EC) and nucleus reuniens of thalamus (nRT). Indirect EC input arrives at the proximal apical dendrite, via DG and CA3, or at the basal dendrite, from CA2. Each pathway can elicit inhibition in feedforward fashion. (Masurkar 2018)

In addition to differences in the laminar innervation patterns from the left and right CA3 to CA1, interestingly, studies also showed that the CA3-CA1 synapses differ in size, shape, and glutamate receptor subunits depending on the laterality of CA3 hemispheric origin (Shinohara et al. 2008; Kawakami et al. 2003). Specifically, CA1 spines receiving input from the left CA3—

regardless of whether they are located in the left or right CA1—are smaller in size (typically thin spines) and exhibit higher expression of the NMDA receptor subunit GluN2B (also known as NR2B). In contrast, CA1 spines receiving input from the right CA3 tend to be larger (mushroom-shaped spines), express higher levels of the AMPA receptor subunit GluA1, and have lower levels of GluN2B. GluA1 is an AMPA receptor subunit, crucial for fast excitatory neurotransmission. GluN2B is a NMDA receptor subunit, known for supporting synaptic plasticity. Notably, spines with richer GluN2B density have a greater capacity for long-term potential (LTP) than those with lower GluN2B density (Shipton and Paulsen 2014). These findings suggest that while CA1 synapses receiving right CA3 input may support stronger baseline transmission due to larger spine size and higher GluA1 density, synapses receiving input from the left CA3 may possess a greater potential for LTP.

Consistent with this hypothesis, electrophysiology studies demonstrated a hemispheric asymmetry at CA3-CA1 synapses on LTP capacity (Shipton et al. 2014; Kohl et al. 2011). In particular, using high-frequency stimulation and spike-timing dependent plasticity protocols, LTP was only inducible at CA1 synapses receiving input from the left CA3, but not from the right CA3. These findings support the idea that LTP capacity at CA3–CA1 synapses is determined by the hemispheric origin of CA3 input. Moreover, the inability to induce LTP at right CA3–CA1 synapses indicates that these synapses may already operate near a close to saturation level of synaptic strength, potentially reflecting a trade-off between stable transmission and plastic potential.

Interestingly, CA3-CA1 synapses are initially rich in GluN2B subunits at birth with no hemispheric asymmetry observed (Shipton and Paulsen 2014; El-Gaby, Shipton, and Paulsen 2015). During development, as the brain matures, a shift from GluN2B to GluN2A occurs at the

right CA3-CA1, whereas left CA3–CA1 synapses maintain a high GluN2B density, preserving a more "immature," plastic state (El-Gaby, Shipton, and Paulsen 2015). The exact mechanisms driving this hemispheric divergence remain unknown. One hypothesis proposes that this molecular asymmetry reflects a functional specialization, whereby right CA3–CA1 synapses support rapid, one-trial learning through preconfigured networks built during development and left CA3–CA1 synapses enable more flexible, plastic encoding of new information that can be modified and integrated into existing representation (El-Gaby, Shipton, and Paulsen 2015).

Behavioral studies have provided partial support for this theory of hemispheric specialization in hippocampal function. In a long-term spatial learning task, mice were trained to locate a reward in a Y maze. Optogenetic inhibition of left CA3 during learning impaired task acquisition after 11 days of training, whereas right CA3 inhibition had no such effect, with mice displaying behavioral performance comparable to controls (Shipton et al. 2014). A follow-up study further localized this effect by selectively inhibiting left CA3 inputs to bilateral CA1, which produced similar impairments in learning, thereby narrowing the critical site of action to the CA3–CA1 synapses (El-Gaby et al. 2016). Together, these findings highlight the essential role of left CA3 input during spatial learning and support the theory that repeated training allows the left hippocampus, via synaptic plasticity mechanisms, to assume responsibility for spatial memory, while right CA3 input becomes less necessary over time.

However, this theory has been challenged by findings from a separate behavioral study. Consistent with the optogenetic results, pharmacological inactivation of the left hippocampus during learning impaired performance in the Morris water maze (Klur et al. 2009). Contrary to the proposed theory, however, inactivation of the right hippocampus during memory retrieval

also resulted in performance deficits. This finding suggests that the right hippocampus is necessary for recall, even in well-learned tasks. Based on the results, the authors proposed an information transfer theory that posits a hemispheric specialization with a leftward bias for memory encoding and a rightward bias for memory retrieval. According to this framework, interhemispheric communication is critical for successful task performance, as both hemispheres contribute at distinct stages of learning. However, evidence from a study using split-brain mice—where interhemispheric communication was disrupted through lesions to the ventral hippocampal commissure and corpus callosum, combined with monocular visual deprivation—complicates this view. Despite the forced reliance on a unilateral hippocampus, these mice were still capable of acquiring a spatial memory task with repeated training (Shinohara, Hosoya, Yamasaki, et al., 2012), raising questions about the necessity of interhemispheric transfer for spatial learning.

In summary, prior molecular and synaptic physiology studies have demonstrated clear hemispheric asymmetries at CA3–CA1 synapses, with distinct properties depending on the left or right hemispheric origin of CA3 (Shipton et al. 2014; Kohl et al. 2011; Kawakami et al. 2003; Shinohara et al. 2008, 2012). However, a unifying theory that integrates findings across molecular, synaptic, and behavioral levels to fully explain these lateralized functions remains elusive. Behaviorally, the left hippocampus has been consistently implicated during learning over repeated training, while the role of the right hippocampus remains more ambiguous. Addressing this gap will require more evidence from cellular and systems levels, especially in vivo recording across different stages of learning. In particular, the hemispheric contributions of CA3 on CA1 place cells remain unclear. Such work will be essential for clarifying the dynamic

and potentially complementary contributions of the left and right hippocampus in spatial memory processing.

1.4 The process of familiarization in a novel environment

The initial exposure to a novel environment causes distinct neural dynamics in the hippocampus. This brief period—often limited to the first few minutes of exploration—represents a critical time window during which the impact of neural activity appears amplified with disproportionately strong influence on the establishment of spatial representations (Cohen, Bolstad, and Lee 2017; Ashby et al. 2021; Kitanishi et al. 2015; Larkin et al. 2014; Tamboli et al. 2024; Priestley et al. 2022; Madar et al. 2025; S. Li et al. 2003).

Familiarization to a novel environment unfolds rapidly and engages mechanisms of incidental spatial learning. During this early phase in a novel environment, animals engage in exploratory behavior, often measured as a change in motor activity (e.g. distance traveled, rearing, etc.). Such behavior are believed to reflect a process of acquiring critical information about their surroundings—including potential food sources and threats—and constructing an internal representation of space, or “cognitive map” (Leussis and Bolivar 2006). As exposure continues and the environment becomes more familiar, exploratory behavior gradually diminishes.

Notably, during this early phase of exploration, CA1 exhibits a rapid surge in place field formation and a gradual organization into a stable spatial code. Place fields can emerge as early as the first lap through a novel environment (Dong, Madar, and Sheffield 2021), suggesting that CA1 neurons are capable of generating a spatial representation immediately upon exposure. Collectively, these early-forming place cells may offer a first-pass spatial map that can guide

downstream processing. As animals continue to explore the same environment, more place cells are formed and integrated into a coherent spatial representation. Interestingly, the presynaptic inputs driving place cell activity show increased amplitude during initial novel exploration, but this elevation is not sustained after place fields emerge or the environment becomes familiar (Cohen, Bolstad, and Lee 2017). Moreover, the spatial tuning of these inputs—initially inhomogeneous across novel environments—gradually become increasingly similar in both within and outside of the place fields. These observations highlight that the earliest moments to novelty creates a unique window with amplified impact, enabling rapid and robust spatial encoding.

This transformation may be orchestrated by interneuron dynamics. In the early phase of a novel environment, reduced activity in dendritic-targeting SOM+ interneurons enhances excitability in CA1 pyramidal dendrites, facilitating robust synaptic potentiation (Sheffield, Adoff, and Dombeck 2017; Geiller et al. 2022). Dendritic spikes are prevalent during this period, further promoting plasticity. As the environment becomes familiar, dendritic-targeting interneuron activity gradually increases, dendritic spike prevalence decreases, and inhibitory circuits develop inverse spatial tuning relative to excitatory place cells (Geiller et al. 2022). This shift marks a transition into a more stabilized representational state.

Notably, in contrast to the reduced activity observed in dendrite-targeting interneurons, soma-targeting (PV+) interneurons exhibit a transient increase in activity upon exposure to a novel environment (Sheffield, Adoff, and Dombeck 2017). PV+ interneurons in CA1 provide fast-spiking perisomatic inhibition that, in concert with excitatory signals, plays a critical role in regulating coordinated network events such as sharp wave ripples (SWRs) and gamma oscillations (Hall and Wang 2023; Colgin and Moser 2010; K. Allen and Monyer 2015;

Schlingloff et al. 2014). These oscillations are proposed to play important roles in memory consolidation and inter-region communication. Consistent with elevated PV+ interneuron activity, CA1 exhibits an increased occurrence rate of SWRs during novel exploration (Larkin et al. 2014). Furthermore, during the first two minutes of novel exposure, CA1 pyramidal neurons show enhanced phase-locked firing to slow gamma oscillations when compared to familiar environments (Kitanishi et al. 2015). This enhanced phase locking is dependent on the AMPA receptor subunit GluA1. Genetic disruption of GluA1 in CA1 impairs phase-locked firing to slow gamma oscillations and alters place field properties, including increased field size and reduced in-field firing rate during initial novel exposure. These findings suggest that GluA1-dependent synaptic plasticity facilitates phase-locked firing and enables the rapid development of spatially tuned activity in CA1 place cells during early learning. As discussed in chapter 1.3 on hemispheric asymmetry, GluA1 subunit has a higher expression level on CA1 dendritic spines that innervate with right CA3, implicating a potentially greater role for right CA3–CA1 synapses in the early stages of spatial learning.

In parallel with changes in local inhibition and oscillatory events, during initial novel exposure, neuromodulatory inputs from locus coeruleus (LC) in CA1 are rapidly elevated and decline with experience (Heer and Sheffield 2024). LC inputs are considered to serve as a novelty detector and drive novelty induced memory enhancement (Takeuchi et al. 2016). While primarily associated with norepinephrine release, the LC also co-releases dopamine in the hippocampus. Exposure to a novel environment promotes the induction of long-term potentiation (LTP) at CA1 synapses by lowering the threshold for LTP induction (S. Li et al. 2003). This facilitation is dependent on dopamine D1/D5 receptor activation and is specific to novel, but not familiar, environments—suggesting dopamine release from LC as a potential mechanism

underlying novelty-facilitated learning. In addition, in CA3, LC input has been shown to be essential for enabling rapid one-trial learning in a novel context (Wagatsuma et al. 2018). Although the precise mechanisms through which LC activity facilitates place cell formation and supports spatial learning remain unclear, these inputs represent another critical component of the hippocampal response to novelty. By modulating excitability and plasticity during this early phase, LC input likely works in concert with local inhibitory dynamics to establish the conditions necessary for rapid spatial encoding (Pedrosa and Clopath 2020).

Moreover, behavioral timescale synaptic plasticity (BTSP) is most prominent early during learning in a novel environment and declines with increasing familiarity (Madar et al. 2025; Priestley et al. 2022). BTSP is a non-Hebbian plasticity mechanism that enables drastic modification in synaptic weights with single or a few trials and facilitates rapid place fields formation (Bittner et al. 2017). In CA1, precisely timed convergent inputs from the entorhinal cortex (EC) and CA3 can trigger dendritic plateau potentials, which in turn activate BTSP mechanisms sufficient to induce the emergence of new place fields (Bittner et al. 2015). For place fields that have already emerged, BTSP-like mechanisms can still exert powerful effects by shifting the location and modulating the amplitude of their spatial firing (Madar et al. 2025; Priestley et al. 2022). Thus, although BTSP-triggering events are relatively rare, they produce highly impactful synaptic modifications. Altogether, the interplay of novelty-driven neuromodulation, local inhibition dynamics, and BTSP may create a transient period of heightened plasticity that supports the rapid encoding of novel spatial experiences. While much of this work has focused on mechanisms within CA1, the upstream contributions from CA3 are less well understood. In particular, how CA3 inputs help orchestrate the shift from early to late phases of learning—and whether these inputs differ by hemisphere—remains an open question.

To address this, my thesis investigates how left and right CA3 inputs differentially shape experience-dependent CA1 spatial representations across distinct phases of memory formation. Specifically, I used two-photon calcium imaging to record right CA1 place cell activity while selectively inhibiting left or right CA3 inputs at key time points during and after learning: (1) during the initial novel environment exploration, when behavioral measures indicate active learning and CA1 representations are still forming and refining; and (2) after learning, when exploration behavior plateaus and CA1 representations have stabilized. To gain further insight into any hemispheric asymmetries from left and right CA3 to CA1, I also recorded CA3 axon activity in CA1 to track how input dynamics from each hemisphere evolve over learning, offering insight into their respective roles in modulating CA1 encoding. In Chapter 2, I will describe the methods used to preprocess CA3 axon data in order to extract real physiological signals, and in Chapter 3, I will present the findings. Together, this work reveals a differential involvement of left and right CA3 inputs to CA1 that switches during distinct stages of learning.

CHAPTER 2

A PREPROCESSING TOOLBOX FOR 2-PHOTON SUBCELLULAR CALCIUM IMAGING

ABSTRACT:

Recording the spiking activity from subcellular compartments of neurons such as axons and dendrites during mouse behavior with 2-photon calcium imaging is increasingly common yet remains challenging due to low signal-to-noise, inaccurate region-of-interest (ROI) identification, movement artifacts, and difficulty in grouping ROIs from the same neuron. To address these issues, we present a computationally efficient pre-processing pipeline for subcellular signal detection, movement artifact identification, and ROI grouping. For subcellular signal detection, we capture the frequency profile of calcium transient dynamics by applying Fast Fourier Transform (FFT) on smoothed time-series calcium traces collected from axon ROIs. We then apply band-pass filtering methods (e.g. 0.05 to 0.12 Hz) to select ROIs that contain frequencies that match the power band of transients. To remove motion artifacts from z-plane movement, we apply Principal Component Analysis on all calcium traces and use a Bottom-Up Segmentation change-point detection model on the first principal component. After removing movement artifacts, we further identify calcium transients from noise by analyzing their prominence and duration. Finally, ROIs with high activity correlation are grouped using hierarchical or k-means clustering. Using axon ROIs in the CA1 region, we confirm that both clustering methods effectively determine the optimal number of clusters in pairwise correlation matrices, yielding similar groupings to “ground truth” data. Our approach provides a guideline for standardizing the

extraction of physiological signals from subcellular compartments during rodent behavior with 2-photon calcium imaging.

Significance Statement

The SUBPREP pipeline is specifically designed to process calcium imaging data from axons, dendrites, and other subcellular structures, which pose unique challenges due to their low signal-to-noise ratios, susceptibility to movement artifacts, and complex morphologies. This pipeline enables researchers to extract reliable physiological signals from noisy datasets, making it a significant improvement over manual or neural network-based approaches. By addressing key bottlenecks in subcellular imaging, this toolbox facilitates the standardization of preprocessing workflows, which is critical for reproducibility and cross-study comparisons. These improvements are especially timely as subcellular calcium imaging becomes increasingly common in neuroscience research. SUBPREP's potential for broad adoption, coupled with its ability to uncover profound biological insights, makes it a valuable contribution to in vivo imaging.

INTRODUCTION

Two-photon calcium imaging enables the recording of activity in subcellular neuronal structures such as dendrites, spines, and axons during behavior. Recent studies have demonstrated the feasibility of using this approach to investigate how specific sets of inputs transmit information to particular regions and how this information is integrated at postsynaptic sites (Bilash et al., 2023; Bowler & Losonczy, 2022; Cichon & Gan, 2015; Fay et al., 2024; Fenton, 2024; Garcia-Garcia et al., 2024; Grienberger et al., 2022; Heer & Sheffield, 2024; Kaifosh et al., 2014; Kaufman et al., 2020; Krishnan et al., 2022; Lovrić et al., 2014; Nikbakht et al., 2024; Pachitariu et al., 2017; Paukert & Bergles, 2012). This methodology requires meticulous analysis of calcium signals from these small structures, which are highly susceptible to movement artifacts and signal noise. Additionally, the intricate morphology of axonal and dendritic structures complicates the identification of segments belonging to the same neuron within the imaging plane. Addressing these challenges is essential for accurately interpreting the data collected in these experiments, which are increasingly being adopted by many laboratories. Standardizing the extraction of subcellular calcium signals through a widely accessible processing pipeline would facilitate direct comparisons across different studies.

Suite2P is the most popular software used for the initial identification of regions of interest (ROI) within a field of view (FOV) (Pachitariu et al., 2017). Suite2P can also be used for the initial identification of ROIs from a FOV containing subcellular structures. However, axon and dendritic ROIs are relatively small and are therefore subject to more noise from movement artifacts than somas and have fewer pixels to extract signals from. Additional steps are therefore required following initial ROI identification by Suite2P to ensure ROIs contain real physiological signals

that can then be analyzed. Various methods have been employed for the selection of active axon ROIs, such as manual labeling by visual inspection (Bilash et al., 2023) or using convolutional neural network classifiers (Krishnan et al., 2022). Manual labeling by looking through time-series movies for fluorescent axon or dendritic signals, while straightforward, is impractical for FOVs with many ROIs and is prone to user bias. Neural network classifiers, though effective, require extensive ground truth data for training, which is labor-intensive and computationally demanding, as well as being limited to one particular dataset, thus not generalizable.

Dual-color imaging offers another approach to identifying active subcellular ROIs by using a static fluorophore alongside an activity-dependent one (Cichon & Gan, 2015; Kaufman et al., 2020). For instance, imaging calcium dynamics with GCaMP in a green channel and using mCherry in a red channel can help distinguish real calcium signals from noise. However, differences in fluorophore behavior and excitation/emission spectra means that subtracting the static channel from the active channel may not remove all noise and can introduce new signal artifacts. Employing dual lasers to optimize excitation could help with this but further complicates the process and the addition of a second laser beam increases the probability of photo-damage to the tissue being imaged (Podgorski & Ranganathan, 2016). This is not to say a static channel is not useful, as it can help with motion correction of time-series FOVs in the X and Y axis and can help identify periods when the FOV may shift in the Z-axis, but simultaneous 2-channel imaging needs careful implementation and consideration.

Once subcellular ROIs are selected, identifying and correcting z-plane movement artifacts is critical. Slight brain movements during behavior can cause significant changes in the imaged structures, leading to false-positive calcium signals and this issue is amplified when imaging

subcellular structures which are much smaller than somas (Robert et al., 2023). Methods such as measuring pixel correlations within an ROI over time have been used to detect these artifacts, but they are computationally intensive and heavily dependent on the initial ROI selection (Ratigan et al., 2023; Rolotti et al., 2022).

Recording from small structures like thin axons also presents a low signal-to-noise ratio (SNR) challenge due to the limited number of pixels available (Sheffield & Dombeck, 2015). Identifying which subcellular ROIs belong to the same neuron can improve SNR by providing more pixels for averaging. Further, because single axons and dendritic branches can come in and out of the FOV, the same neural structure can be represented by disconnected ROIs that are far apart in the FOV. If ROIs belonging to the same neuron are not grouped, they will be overrepresented in the data and can confound results. However, resolving which ROIs are part of the same structure remains an issue.

We present a preprocessing pipeline for identifying subcellular ROIs with real physiological signals, removing motion artifacts, and grouping ROIs belonging to the same neuron. We apply Fast Fourier Transform (FFT) on smoothed calcium traces to identify the frequency band of transients, followed by band-pass filtering to select ROIs matching the power band of transients. Principal Component Analysis (PCA) and Bottom-Up Segmentation change point detection are then used to remove motion artifacts. Transient peak detection based on prominence and duration identifies real calcium transients. Finally, we employ hierarchical or k-means clustering to group correlated ROIs, and validate this approach with “ground truth” axon clusters. Our method provides a guideline for analyzing axonal and dendritic activity using two-

photon calcium imaging, aiding in the standardization of physiological signal extraction from subcellular structures during behavior.

METHODS

Ethics statement

All surgical and experimental procedures were approved Institutional Animal Care and Use Committee (IACUC) guidelines (protocol number: 72508). The IACUC consists of 8 scientists, 3 veterinarians, one non-scientist, and 2 public members.

Experimental model and subject details

We used 10–20 week old C57BL/6-Tg(Grik4-cre)G32-4Stl/J mice (23–33 g). Mice were individually housed in a reverse 12-hour light/dark cycle and behavioral experiments were conducted during the animal's dark cycle. Example animals for illustration purposes were one female and one male.

Injection protocol

Mice were anesthetized and I.P. injected with 0.5ml of saline and 0.5ml of meloxicam. A small craniotomy was made over the right hippocampus CA3 region (2.0mm lateral, 1.7mm caudal respectively of Bregma). An axon-targeted genetically-encoded calcium indicator, AAV9-axon-GCaMP6s-P2A-mRuby3 (pAAV-hSynapsin1-axon-GCaMP6s-P2A-mRuby3 was a gift from Lin Tian Addgene viral prep # 112005-AAV9; <http://n2t.net/addgene:112005>; RRID:Addgene_112005) was injected (~50 nL at a depth of 1.9 mm below the surface of the dura) using a beveled glass micropipette. Right after the viral injection, the site was covered with dental

cement (Metabond, Parkell Corporation) and a metal head plate (Atlas Tool and Die Works). After ~4-8 weeks, water scheduling (0.8-1ml per day) began. One week after the start of water scheduling, mice went through a hippocampal cannula implantation surgery over the right CA1 region (1.7mm lateral, -2.3mm caudal of Bregma). During the surgery, a head plate and head ring were attached to the cannula window for head-fixation and to house the microscope objective and block out ambient light. Post-surgery, mice had ad lib water, until a week later when they returned to water scheduling. Expression of axon-GCaMP6s in CA3 axons reached a steady state ~6-12 weeks after the virus was injected, as monitored through 2P imaging.

Two-photon imaging

Imaging was done using a laser scanning two-photon microscope (Neurolabware). Using a 8 kHz resonant scanner, images were collected at a frame rate of 15.49 Hz with unidirectional scanning through a 16×/0.8 NA/3 mm WD water immersion objective (MRP07220, Nikon). axon-GCaMP6s was excited at 920 nm with a femtosecond-pulsed laser (Insight DS+ Dual, Spectra-Physics) and emitted fluorescence was collected using two GaAsP PMTs (H11706, Hamamatsu). The average power of the laser measured after the objective ranged between 60–100 mW, and was kept constant across days of imaging. Time-series images were collected through Scanbox (Neurolabware) and a PicoScope Oscilloscope was used to synchronize frame acquisition timing with behavior.

Image processing and ROI identification

Time-series images were preprocessed using Suite2p (Pachitariu, M., et al. , 2017). Movement artifacts in the X and Y-axis were corrected using rigid and non-rigid transformations

based on the red mRuby channel. Regions of interest (ROIs) were also defined in the green channel using Suite2p. All identified ROIs were fed into SUBPREP for subsequent ROI selection. Raw fluorescence traces from ROIs were first smoothed using a Savitzky-Golay filter.




Fig. 1 Imaging CA3 axons in CA1

Left: Schematic representation of CA3 axonal imaging in CA1 Stratum Oriens. Right: Example CA1 field of view containing CA3 axons. Image is the max projection from 5000 frames.

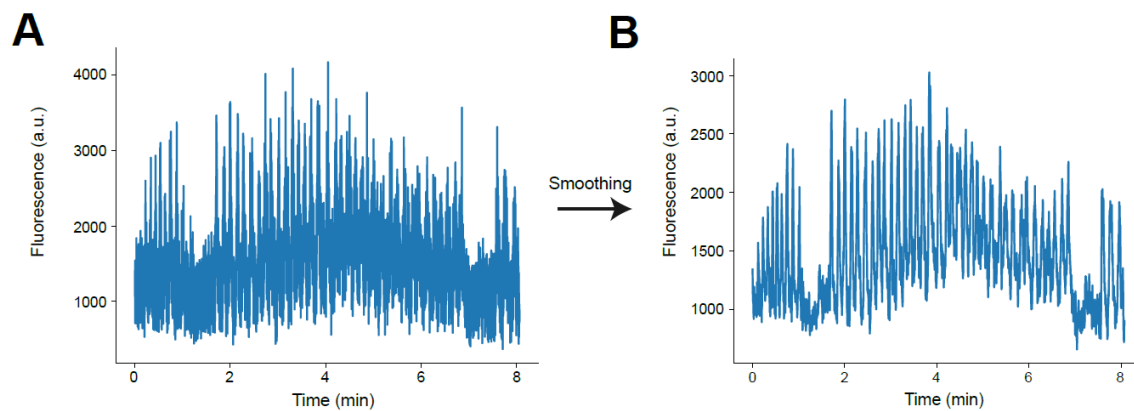


Fig. 2: Example time-series fluorescent traces from CA3 axons in CA1 before and after smoothing

A. Example raw fluorescent time-series trace from an ROI identified by Suite2P.

B. The same trace after smoothing using the Savitzky-Golay filter.

RESULTS

Frequency-based approach for transient identification and subcellular ROI selection

In selecting subcellular ROIs that contain real calcium transients, we developed a frequency-based ROI selection pipeline. The key idea underlying our selection procedure was that the rise-to-peak and decay phase of real transients have unique frequency domains. In particular, the rise of the transient occurs within a faster frequency band than the decay of the transient. Therefore, we designed a filter that captures both the rise and the decay of transients. To determine the key frequencies for the rise and the decay, we calculated the frequency bands using example ROIs with clearly visible transients and high SNR and compared them to the frequency bands from all ROIs. Because the vast majority of ROIs from Suite2P (85–95%, depending on the dataset) contain only noise, example ROIs with clear transients returned higher power in the frequency bands that reflect the rise and decay of transients compared to the mean of all ROIs. The frequency band with elevated power from example ROIs with clear transients became our frequency band of interest. If the filtered signal from an ROI returned high power in our frequency band of interest, the ROI was selected as “likely to contain transients.” By contrast, if the filtered signal did not contain sufficient power (see below) in our frequency of interest, the ROI was marked as “not likely to contain transients.” Note that the specific frequency band containing real transients will depend on which calcium indicator is being imaged (e.g., GCaMP6s versus GCaMP8f) and which

subcellular structures are being imaged (e.g., axons versus dendrites). The key frequency band will therefore need to be determined for each user depending on the specifics of the dataset.

To quantify the power within the key frequency band among all subcellular ROIs, we applied a band-pass filter using the key frequency range on smoothed time-series calcium traces using the `filtfilt` function in Python. We utilized a third-order Butterworth filter, a type of infinite impulse response filter known for its maximally flat frequency response within the passband. By selecting this filter design, we ensured effective attenuation of frequencies outside the passband while minimizing phase distortion within the passband. First, a low-pass filter with a cutoff frequency of 0.13 Hz (or the higher bound of key frequency bands from example ROIs) was applied to the green channel signal to remove high-frequency noise. Then, a high-pass filter with a cutoff frequency of 0.03 Hz (or the lower bound of key frequency bands from example ROIs) was applied to remove low-frequency noise (e.g., Z-plane drift). The `filtfilt` function employed a zero-phase digital filtering technique, which applied the filter forward and then backward to eliminate phase distortion, resulting in a zero-phase delay. Our procedure ensured that the filtered signal retained its temporal integrity. By using `filtfilt`, the signal was effectively band-pass filtered—removing both high-frequency noise and low-frequency drift—while preserving the original signal’s phase and shape.

Once the signal had been band-pass filtered, the power within the frequency band of interest (e.g., 0.03 Hz to 0.13 Hz) was calculated. We then calculated the *normalized power* of the signal, such that the power of interest was divided by the total power of the signal. An advantage of using normalized power is that the range is bounded (0 to 1), which enabled comparisons across different traces and animals, while absolute power is unbounded and depends on the signal length.

If the power within this frequency band exceeded a predefined threshold (e.g., 0.3 of the normalized power), the ROI was defined as containing transients and was selected. A higher threshold (e.g., 0.5) returned fewer ROIs but with higher SNR. This approach enabled the identification and isolation of signals with significant activity within the specified frequency range, facilitating the extraction of ROIs with transient-like patterns.

Fig. 3 shows two example ROIs from our dataset. The green ROI in **Fig. 3A** had a relatively low power in the frequency band of interest shown in **Fig. 3B** (0.182), suggesting this ROI is unlikely to contain transients. Contrarily, the black axon in **Fig. 3A**, which shows clear and obvious transients, had a higher power in the frequency range of interest (**Fig. 3B**, black; 0.454). Applying this approach to all ROIs from an example mouse, the distribution of the normalized power is shown in **Fig. 3C**. The ROIs selected as “likely to contain transient” were the ones with power above 0.3 (red dotted line) on the right side of the histogram that were kept for the next steps of the analysis, and the ROIs marked as “not likely to have transient” were shown on the left of the dotted line that were not used in the remainder of the analysis pipeline.

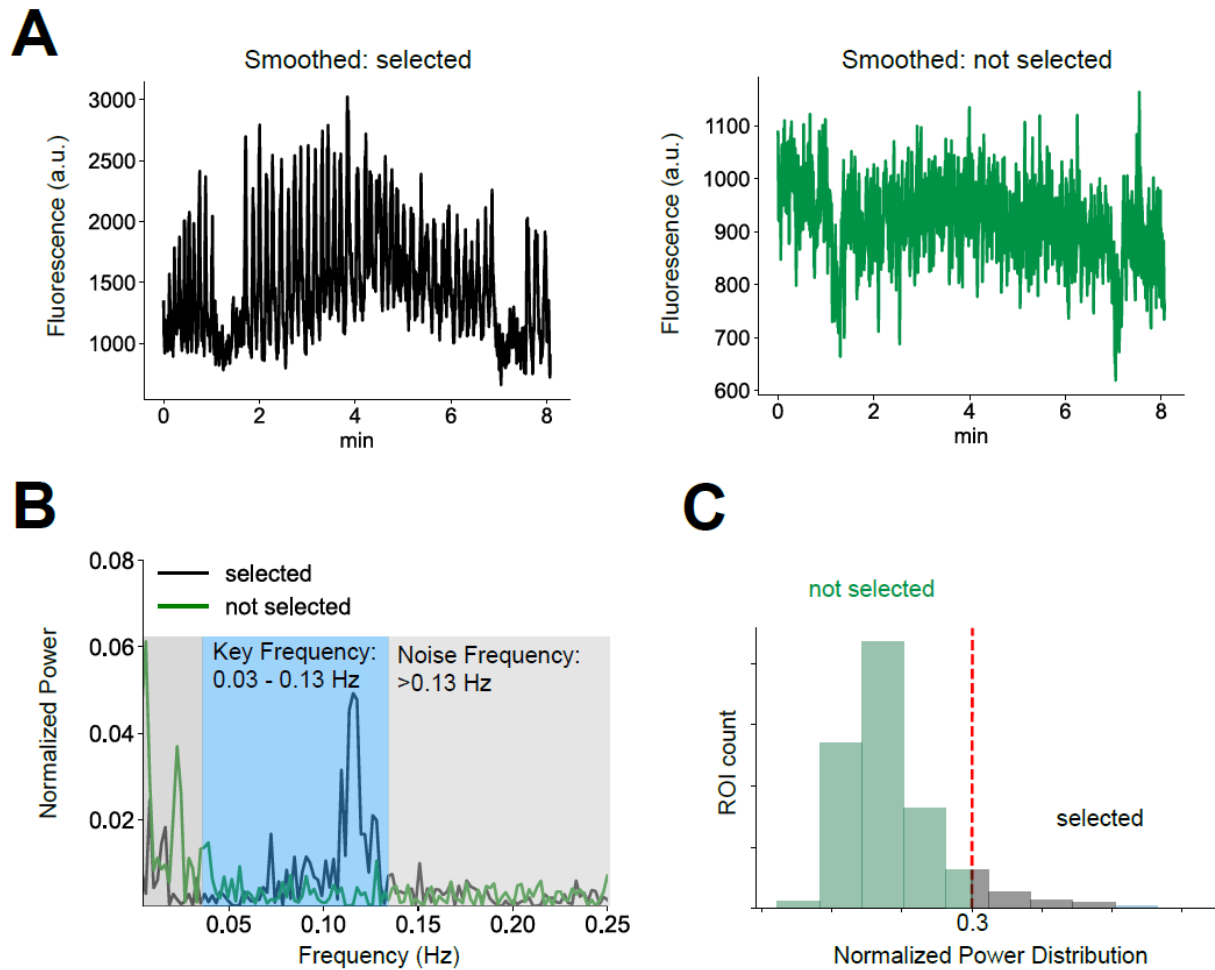


Fig. 3 Frequency band-based axon ROI selection.

A. Examples of selected (black) and discarded (green) ROIs in time domain after smoothing.

B. Same ROIs (first row in A.) in frequency domain. The x-axis displays the frequency of the FFT transformation, and the y-axis displays the normalized power. Note, the black ROI has a higher power in the frequency band of interest (shaded in blue) and lower power in the noise frequency (shaded in gray) while the discarded ROI is the opposite.

C. The distribution of normalized band power within the frequency band of interest (0.03-0.13 Hz) for all example ROIs from a single FOV. The ROIs with lower power in the frequency band interest (below 0.3; dashed line) were discarded, and the ROIs with higher power in the band of interest (above 0.3) were selected.

Time series anomaly detection for motion artifact identification

The previous step focuses on identifying transient-like patterns in our ROIs. However, the remaining ROIs might still contain traces that suffer from motion artifacts. X- and Y-axis movement are easily dealt with by motion correction algorithms like Suite2P. However, Z-axis movements may remain and are difficult to detect and remove. Typically, when a FOV moves along the Z-axis, many, if not all, ROIs are affected simultaneously. The effect of movement could be different in each ROI within the FOV (e.g. an ROI could drift in or out of the Z-plane). However, all affected ROIs would show synchronous abrupt changes (either an increase or decrease in signal) in fluorescence during a Z-axis shift. Thus, Z-axis shifts could be detected across the population of ROIs within a FOV by a sudden increase in coactivity strength and variance between ROIs. The goal here is to identify and remove periods of Z-axis shifts.

To achieve this, the first step is to z-score transform all ROI time-series traces to make them comparable. Real calcium transients should be stereotyped (as described above) and their occurrence should vary in time between ROIs, but z-axis shifts should cause synchronous signals that vary in terms of their shape (e.g. upwards or downwards). The distinct dynamics of real transients versus z-axis-induced signals, and the prediction that z-axis shifts will contribute to the largest variance in the moments shifting occurs compared to the rest of the time, should allow us to detect shifts using Principal Component Analysis (PCA) (**Fig. 4A**). We used the first principal component, which captures the largest variances in the data.

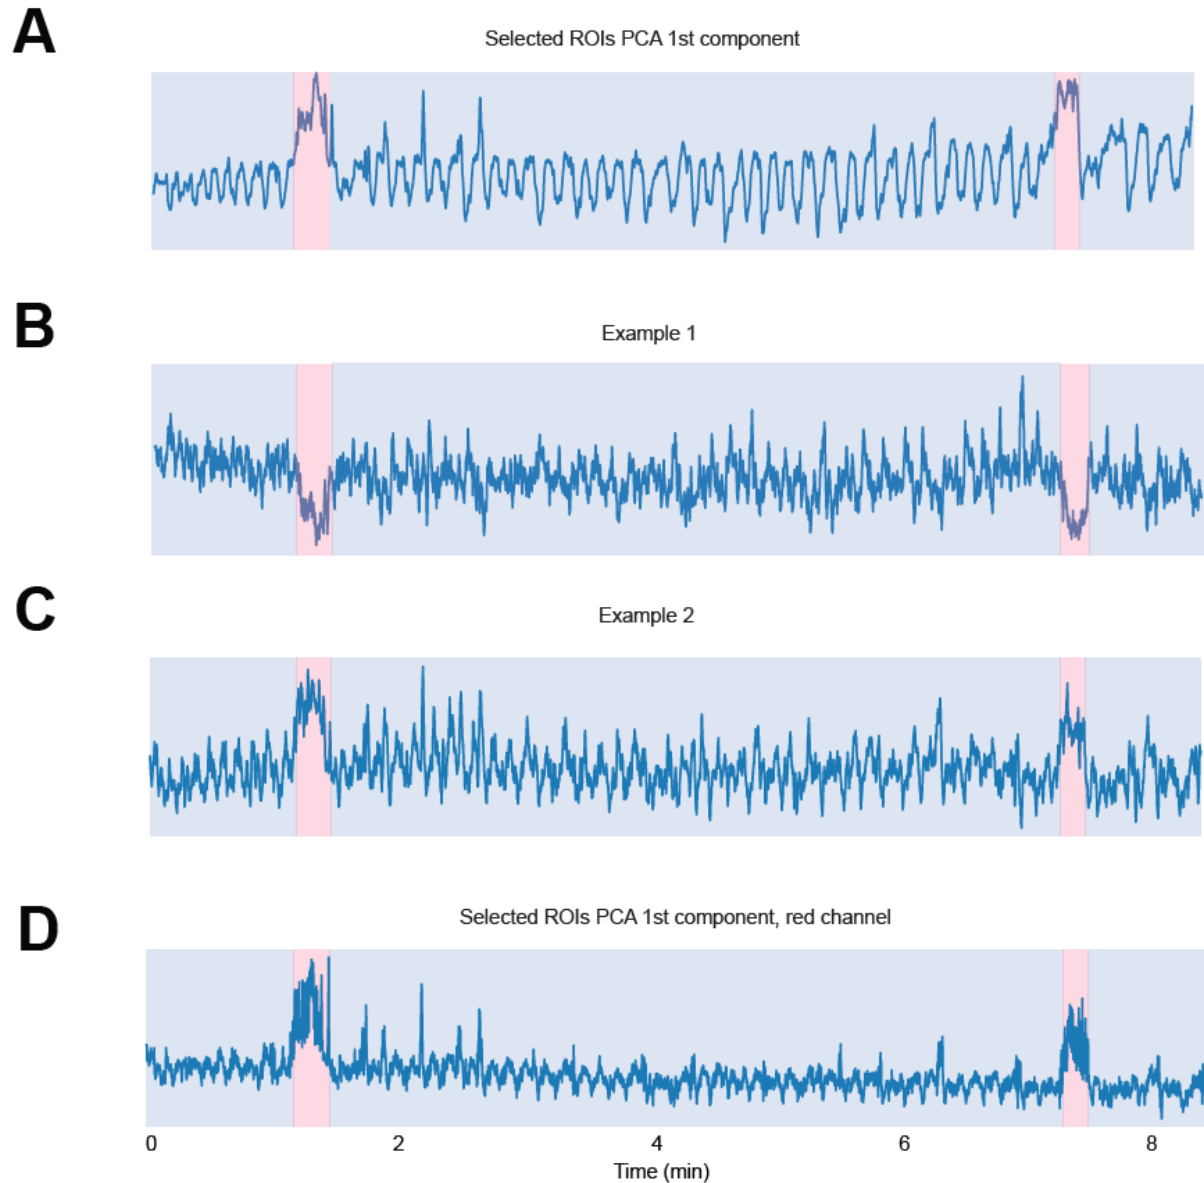


Fig. 4 Identification of movement artifacts in time-series traces

A. First principal component on z-scored ROI activities. All selected ROI activities were first standardized to z-score and then dimension reduced to the first component of PCA. Bottom-Up algorithm was used to segment and detect anomalies (identified anomaly period shaded in red).

B & C. Identified periods for potential movement artifacts were super-imposed on individual ROI activities for visual inspection. Two example selected ROI standardized activities are plotted and red shaded indicate potential periods of movement artifacts. All activities during the red shaded time period can then be ignored or removed from further data analysis. Note that the two example cells had activities in different directions but synchronous in time.

D. Identified time periods for movement artifacts on red channel as further verification.

After applying PCA on z-scored ROI traces, we used a Bottom-up algorithm to segment and identify the time periods where a likely shift occurred (Truong et al., 2020). Specifically, Bottom-up works by separating the time series data into a large number of very small segments of equal length and merging similar segments iteratively to bigger segments until met by a stopping criterion (Keogh et al., 2001). It is especially suitable for signals that contain complex and non-linear patterns with outliers and potentially multiple level shift changes, calculated by the algorithm. Using this method, we could detect shifts in the level of the baseline signal while filtering out noise and outliers. Each level shift change detected by the algorithm in the signal reflects a potential z movement shift artifact in the FOV (**Fig. 4A**).

The bottom-up segmentation algorithm requires the number of change points it will detect to be input by the user. Using a larger number will lead to more periods detected by the algorithm and potentially more false positives and vice versa when using a smaller number. **Fig.4A** shows an example first PC which we visually inspected for potential Z-shifts. We used $n=4$ as the input for the bottom-up segmentation algorithm, which accordingly identified 2 periods of potential Z-shift. The user input at this stage requires consideration of neural dynamics of the imaged population. For instance, synchronous activity among large populations of ROIs can cause a level shift change in the first PC space. For regions of brain that have highly correlated activity, a static red fluorophore is needed to distinguish between movement artifact and synchronous activity. Our dataset was recorded from hippocampal CA3 axons during a spatial navigation task. While it is possible for groups of axons to be co-active, especially near to the rewards in the hippocampus (Dupret et al., 2010; Dragoi & Buzsáki, 2006), the synchronous activity is unlikely to sustain for

a long time (beyond a second or two). Thus, any prolonged level shift change in the PC space is highly likely to be driven by z-axis shift. Using a large number of n as input to the algorithm and then validating the results by measuring the duration of level-shift change reduces the possibility of having false positives.

In addition, a static red fluorophore could help to distinguish between population co-activity and z-axis shifts. The same z-score standardization and PCA can be applied to the red channel. **Fig. 4D** shows the PCA from the red channel from the same ROIs used in **Fig. 4A**. The red channel shows the same periods of elevated activity in the PC space identified in the green channel, confirming they are most likely caused by z-axis shifts. We can further validate the shifts using individual ROI activities (described below).

After finding periods of potential movement artifacts using PCA and Bottom-up algorithm, to visually confirm the validity of results, we underwent another manual visual inspection of the time-series traces from individual ROIs with labeled potential movement artifacts (**Fig. 4B and C**). The two example ROIs illustrate that the movement artifact identified using the PCA method could be caused by z shifts from individual ROIs to be in the negative (**Fig. 4B**) or positive (**Fig. 4C**) direction, but the activities were synchronized during the labeled period, thus fitting our model of z-axis-induced signals. Multiple ROIs showing synchronized level shift in activity during the labeled period confirmed our findings of z-axis shift timing from PCA and Bottom-up algorithm. The frames during z-axis shifts can then be discarded across all ROIs within the FOV. Alternatively, if the FOV has shifted to a new Z-plane for a long time, our algorithm also allows for splitting the dataset into new sets of ROIs defined at the two planes to minimize data loss.

Baseline correction and peak detection

To remove slow drifts in fluorescence caused by bleaching and scale the F signal to baseline, slow time scale changes in the fluorescence traces were removed by examining the distribution of fluorescence in a ~20 second interval around each sample time point and subtracting the 8% percentile value. The resulting baseline corrected $\Delta F/F$ time-series traces were used to identify significant calcium transients. The initial steps in our pre-processing pipeline already detected ROIs with likely transients and removed motion artifacts. The following step was to further identify real transients from potential noise. To do this, peaks in the trace with a minimum amplitude of 0.12 $\Delta F/F$, duration of 0.5 s, and prominence of 0.1 $\Delta F/F$ were identified. These regions we selected as “real” calcium transients and were included in the subsequent analysis.




Fig. 5 Baseline correction and peak detection

All ROI activities were scaled to baseline to create $\Delta F/F$ traces. Baseline corrected $\Delta F/F$ traces across time were generated using sliding window of ~20 seconds and the 8% percentile value within the sliding window was subtracted from each timepoint. Across the baseline corrected $\Delta F/F$ traces, peaks were calculated using minimum height, distance, and prominence values. Of all potential peaks, only detected peaks (black) were kept and all other activities (green) were forced to 0.

Clustering of the selected ROIs

Due to the density, coverage, and highly branched nature of axons and dendrites that can move in and out of the 2-photon imaging plane, multiple independent ROIs within a FOV could belong to the same cell. As a result, ROIs belonging to the same cell should have highly similar patterns of activity. Previously, one of the biggest challenges for axon clustering is a lack of ground truth labels to evaluate the performance of the clustering method. Failure to group axons or dendrites that branch from the same cell leads to overrepresentation of a single cell's activity. Conversely, incorrect grouping of independent ROIs leads to loss of data and brings in artificially generated noise. In this section, we aim to compare two different clustering algorithms to combine axons and their corresponding performances when compared to ground truth labels.

Assuming that a high zero-phase cross-correlation between two ROIs reflects co-activation, we first performed a maximum correlation analysis across all ROIs over the entire time series (correlation matrix shape #ROIs x #ROIs). If the maximum correlation between a certain ROI with all the remaining ROIs exceeded a specific threshold ($r = 0.8$, see below for details), the ROI was likely to be part of a larger cluster of ROIs that belonged to the same axon. By contrast, if the maximum correlation between an ROI with all other ROIs was low, then the ROI was not likely to part of the same axon.

For the rest of the ROIs that could originate from the same axon, we calculated the pairwise correlation between ROIs and compared two different clustering methods, namely k-means and hierarchical clustering. To test whether the clustering methods were doing a good job, we compared their ROI clustering results to “ground truth data”, i.e. ROI groupings that were highly likely from the same axon. To obtain ground truth data in our experimental setup, animals ran

through two different virtual reality environments while imaging hippocampal CA3 axons in the CA1. A feature of hippocampal cellular activity is their spatial firing and their ability to remap their spatial firing fields upon exposure to a novel environment, meaning that while it is possible for distinct place cells to have similar spatial firing within one environment, it is highly unlikely for them to also fire in a similar pattern in more than one environment (Fenton, 2024). Thus, if the pairwise correlation between ROIs was consistently above 0.7 across two distinct environments, we considered those pairs to belong in the same axon group. This served as our “ground truth data” in which our clustering methods could be compared.

We evaluated the similarity between the ground truth data and our assigned clusters using two different metrics: the Adjusted Mutual Information (AMI) score and the Silhouette score. While the AMI score measured the similarity between the true cluster labels (i.e., > 0.7 correlations between two ROIs in both of our virtual environments) and the labels assigned by the clustering algorithm (e.g., k-means or hierarchical clustering), adjusting for chance, the Silhouette score evaluated the quality of the clusters themselves without the need for ground truth labels. The AMI score ranged from 0 to 1, where a score of 1 indicated perfect agreement between true and predicted cluster labels. The Silhouette score ranged from -1 to 1, with higher values indicating better clustering results. This metric assessed the separation of clusters based solely on the clustering output produced by clustering algorithms. Specifically, a high Silhouette score indicates clear separation between cluster groups, that objects within a cluster were close to each other and far away from objects in other clusters. A high performance in either metric suggests that the clustering results were reliable, and the number of clusters that maximize these metrics represents the best-fit grouping assignments for the remaining ROIs. Because not every axon dataset will have ground truth labels to calculate AMI, we compared the consistency between AMI and

Silhouette scores to test if Silhouette score is sufficient to lead to a similar clustering result without the need for ground truth data.

To estimate the number of clusters, K number of clusters was examined from 2 to 50% the of the selected ROIs. For both example mice, Silhouette scores and AMI scores converged (contained a maximum number that represent best performance) and reached optimal K for clustering (see **Fig. 6A**, left panel). In addition, both performance metrics (AMI and Silhouette) returned a similar number K, which suggests that hierarchical clustering was efficient with or without true labels of clusters.

In addition to hierarchical clustering, the k-means algorithm can also be used to identify clusters of high correlations between pairs of ROIs. Like hierarchical clustering, k-means uses the full correlation matrix between all pairs of ROIs as input (citation needed here). Our main goal was to identify the optimal number of clusters, k, that best explains the covariance structure of the ROIs. We applied the k-means algorithm with a predefined number of clusters, ranging from 2 to half the number of ROIs (n). The algorithm partitions the samples into k clusters based on similarity, iteratively assigning each sample to the nearest cluster centroid and updating centroids to minimize the within-cluster sum of squares. This process continues until convergence, resulting in clusters where correlations between ROIs are maximized within each cluster and minimized between clusters. Notably, the optimal number of clusters identified by k-means closely matched the results from hierarchical clustering. Therefore, we conclude that both methods reliably estimate optimal clustering.

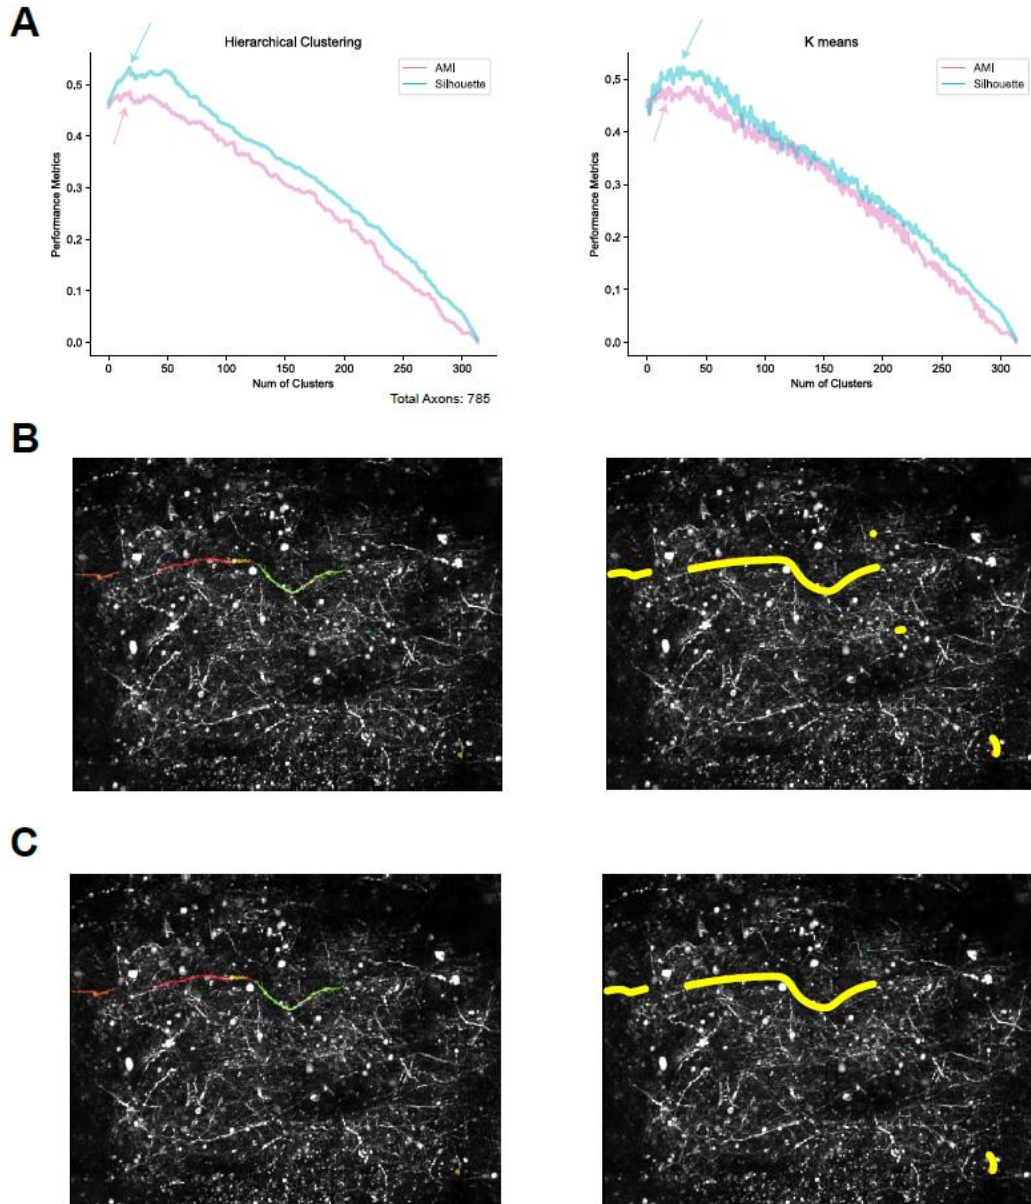


Fig. 6 Comparison of clustering methods used for grouping ROIs.

A. Clustering performance metrics with hierarchical clustering and K-means clustering. The left panel shows that hierarchical clustering converged to a similar optimal number of clusters (peak pointed by the arrows) for both performance metrics, Silhouette scores and AMI. The right panel showed that K-means clustering converged to a similar optimal number of clusters for both performance metrics, Silhouette scores and AMI.

B. Example grouping from hierarchical clustering transposed to original FOV. Left: each color represents an individual ROI in the group. Right: all ROIs in the same group highlighted in yellow

C. Same grouping from “ground truth” data.

Summary of the SUBPREP

A sample paradigm of our SUBPREP pipeline is displayed in **Fig. 7**, with all the significant steps and an example ROI on the right side of the flowchart. The first step, pre-processing, was applied to all ROIs and the goal was to smooth the data across time. After pre-processing, the ROIs with higher power in the key frequency band for transients were selected with band-pass filtering methods (0.03 to 0.13 Hz). Then, to identify movement artifacts, we standardized and used PCA to reduce the dimensions in the population activity. We used a Bottom-up algorithm to identify anomalies in signal of the reduced PC space. We confirmed the results by imposing the labeled movement artifact periods onto individual ROIs and visually inspected their activity. Confirmed periods of movement artifact were removed from all ROIs in the FOV. The remaining ROI signals, highly likely to contain transients from stable FOVs, were identified by thresholding signal prominence which allowed detection of transient peaks. Lastly, to identify the potential clustering of the ROIs belonging to the same axon, we performed two methods, hierarchical clustering and k-means clustering. Both methods were effective in identifying the optimal number of clusters in the pairwise ROI correlation matrices and resulted in similar grouping compared to the ground truth labels.

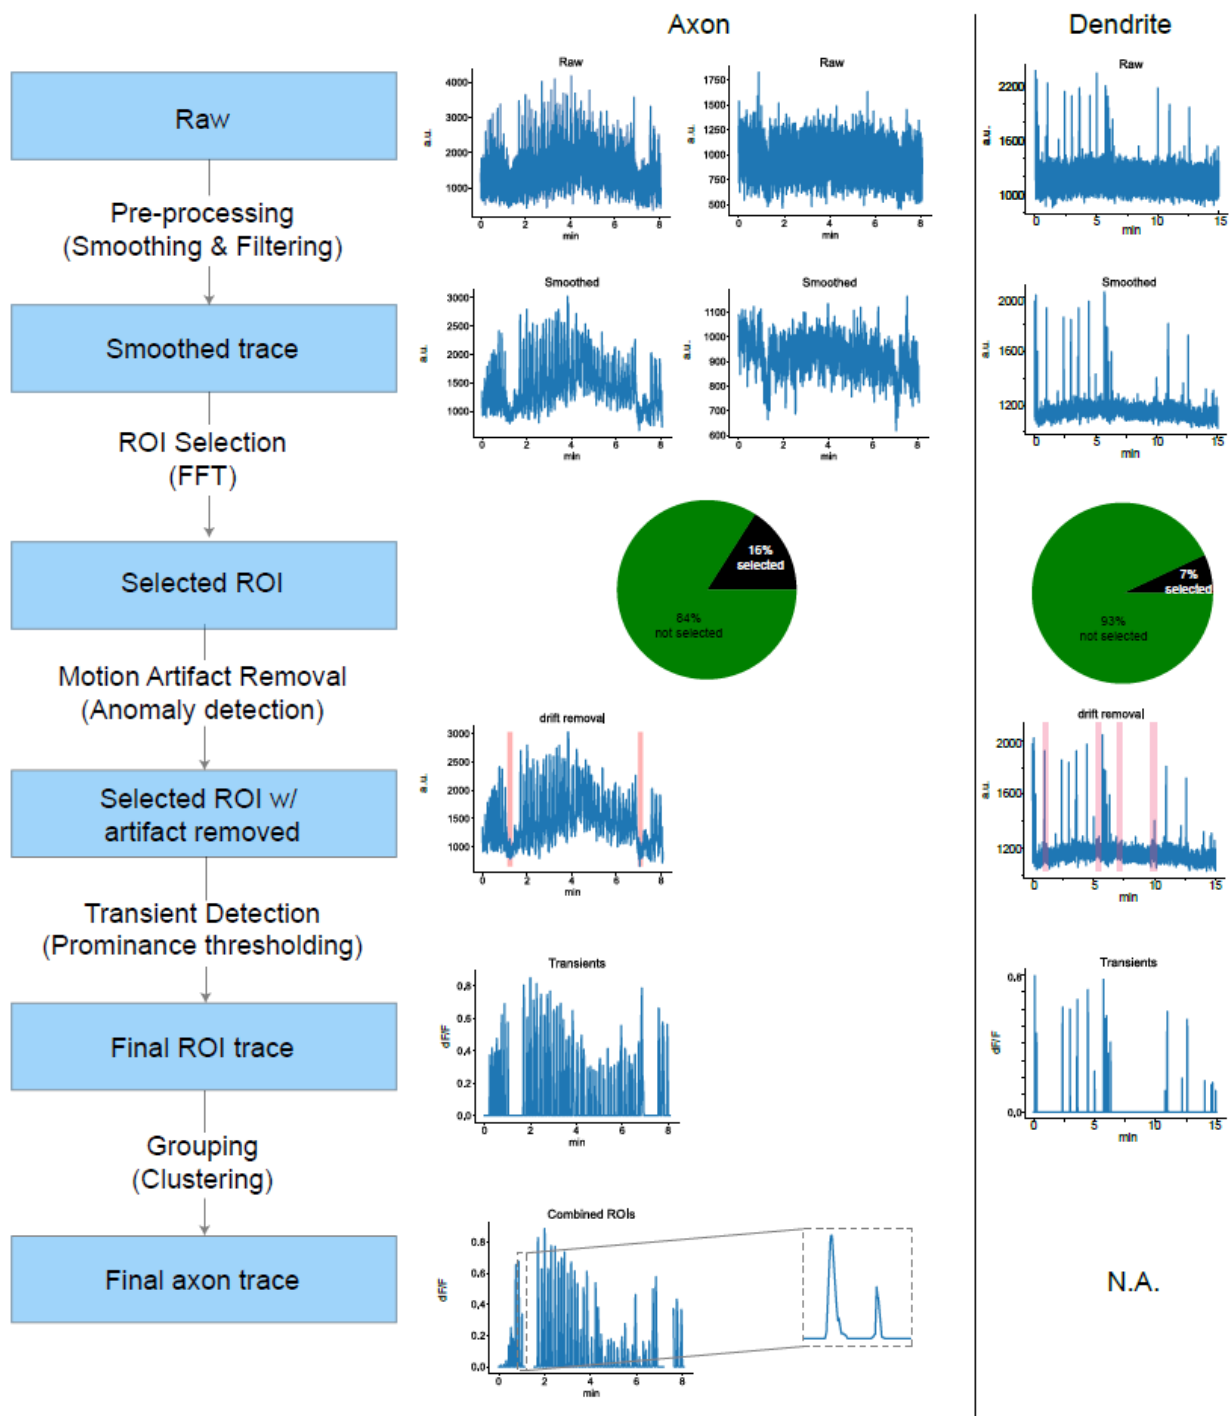


Fig. 7 SUBPREP pipeline summary.

Axon (middle) and dendrite (right) example traces are shown being processed as they move through the SUBPREP pipeline.

DISCUSSION

In this study, we developed and validated a computationally efficient preprocessing pipeline for subcellular signal detection, movement artifact identification, and ROI grouping in two-photon calcium imaging data. Our approach addresses several key challenges associated with analyzing subcellular structures, such as low signal-to-noise ratio (SNR), movement artifacts, and difficulty in grouping subcellular ROIs from the same neuron. Through the application of frequency-based filtering, Principal Component Analysis (PCA) and level change detection for motion artifact identification, and clustering algorithms for ROI grouping, our pipeline provides a robust framework for extracting physiological signals from subcellular structures during behavior.

Suite2P and SIMA, designed for detecting somatic activity, provide robust frameworks for initial subcellular data processing (Pachitariu et al., 2017; Kaifosh et al., 2014), but additional steps are required as subcellular datasets are inherently noisy and sensitive to Z-plane movements because of their small size relative to somas. Our processing pipeline offers a systematic approach to sub-select ROIs from the initial pool identified by Suite2P by leveraging the unique frequency domains of calcium transient dynamics. By employing Fast Fourier Transform (FFT) and band-pass filtering, we were able to isolate ROIs likely to contain true calcium transients, allowing accurate signal detection in noisy subcellular datasets. Other methods have been used for this step, such as manual ROI selection and convolutional neural network (CNN)-based approaches, but these are either labor-intensive and prone to bias, or require extensive ground truth data for training (Dupret et al., 2010; Dragoi & Buzsáki, 2006; Soltanian-Zadeh et al., 2019). A second, “static”, fluorescence channel has often been used to subtract noise from an “active” channel, usually from GCaMP signals, but this can introduce artifacts as noise is not perfectly mirrored across channels.

For instance, PMTs are a source of noise and separate PMTs are used for 2 channel imaging (Chen et al., 2013), and the fluorophores themselves are distinct sources of noise. Our approach offers an unbiased and computationally cheap method for subcellular ROI identification that does not require dual-color imaging, thereby simplifying experimental design and reducing potential artifacts from channel subtraction.

Accurately identifying periods of Z-plane shifts is always important but is amplified when imaging subcellular ROIs versus somas because axons and dendrites are relatively small and have elaborate structures that can be found in and out of the imaging plane in complex branching patterns (Chen et al., 2013). A small movement in Z can result in a distinct set of subcellular ROIs being imaged. After identifying ROIs with potential transients, our pipeline effectively identifies Z-plane movement artifacts by utilizing PCA and a Bottom-Up Segmentation algorithm to detect artifacts based on the first principal component. This allows for the analysis of data collected from a stable imaging plane by excluding time points collected from different planes. Additionally, if a prolonged Z-shift to a new imaging plane occurs, the new Z-plane can be identified and used for data analysis of a distinct set of subcellular ROIs during this period. This limits data loss from technically challenging subcellular imaging experiments. Importantly, our pipeline is able to detect and remove (or separately analyze) only abrupt shifts of the entire Z-plane. It is not designed to deal with within field movement artefacts or rhythmic movements caused by breathing or the animal's heart rate. Indeed, a stable baseline FOV is a requirement in order to detect rare and abrupt changes across the entire FOV.

Despite the strengths of our approach, several limitations remain. For instance, our method's reliance on frequency-based filtering requires careful consideration of the smoothing

filters applied prior to FFT, as these can impact the key frequency band of real transients and thus the selection of ROIs. Additionally, while our PCA-based motion artifact identification effectively detects Z-plane shifts, it still requires user input to define the number of change points, which introduces potential bias. Automating this process or developing more sophisticated algorithms to reduce user dependency could further improve the robustness of the pipeline. Furthermore, an appropriate selection of frequency bands is needed to best extract the signal of interest in calcium imaging data.

Our toolbox could work with or without existing calcium imaging software such as Suite2P. Despite that SUBPREP did not require Suite2P to function, we believe that future work should integrate our SUBPREP toolbox into Suite2P. This would provide multiple benefits for the user. One benefit would be that the user interface and real-time visual displays in Suite2P could be utilized for SUBPREP to show the user which ROIs in the FOV contain real transients or which segments of the time-series movies have potential motion artifacts. It would also streamline analysis allowing one software package that can be used for cellular and subcellular preprocessing. Other future work should focus on expanding the validation of our approach across different brain regions and experimental conditions will be crucial for generalizing its applicability. To further reduce high-frequency noise sources such as PMT shot noise and electrical artifacts, future iterations of our pipeline could incorporate state-of-the-art denoising algorithms (e.g., DeepInterpolation; Lecoq et al., 2021) applied to motion-corrected image stacks prior to fluorescence signal extraction. While our current approach relies on motion correction and low-pass filtering to suppress such noise, the application of deep learning-based denoising methods has the potential to significantly reshape the power spectrum and improve the fidelity of neural activity signals.

SUBPREP offers a standardized approach for analyzing axonal and dendritic activity by simplifying the process of ROI selection, artifact removal, and ROI grouping, providing a practical and effective solution to several key challenges in the field of sub-cellular imaging. We hope our publicly available processing pipeline, SUBPREP, is used by the field to help standardize the extraction of physiological signals from subcellular structures, enabling more accurate and reproducible studies of subcellular activity across labs. This would lead to more robust insights into the mechanisms of information transmission and integration in the brain, ultimately advancing our understanding of neural circuitry and its role in behavior.

DATA AVAILABILITY

Raw imaging data is large and not feasible for upload to an online repository but is available upon request to the lead contact. Processed source data for all figures and associated statistical analysis are provided will be provided in the final version of the paper.

CODE AVAILABILITY

SUBPREP software used for subcellular data analysis is available on github <https://github.com/anqijiang/subprep>

CHAPTER 3

A LEARNING-DEPENDENT SHIFT IN CA3 INPUT DOMINANCE FROM LEFT TO RIGHT DETERMINES THE EVOLUTION OF CA1 SPATIAL MAPS

3.1 Abstract

CA1 place fields underlie spatial maps supporting memory, yet how bilateral CA3 inputs shape these maps during learning remains unclear. Using two-photon calcium imaging and optogenetic inhibition in head-fixed mice navigating a virtual track, we examine left and right CA3 projections to right CA1 (CA1_R) as they learn about a novel environment. CA1_R spatial maps are initially inaccurate but gradually improve and stabilize after ~10 laps (early-phase), with later laps marking stability (late-phase). In the early-phase, both CA3 inputs drive CA1_R field emergence, but right CA3 inputs predominantly drive high-amplitude, reliable fields that contribute to developing maps. Later, left CA3 inputs predominantly maintain reliable fields and stable CA1_R maps. Axonal recordings reveal heightened right CA3 activity early, shifting to increased left CA3 activity later. These findings demonstrate a shift from right-to-left CA3 input dominance that mediates the progression from early map development during learning to later stabilization in CA1_R.

Introduction

The hippocampus is essential for spatial and episodic memory, with dorsal CA1 place cells forming the neural basis of a spatial representation (Eichenbaum et al. 1999; Robinson et al. 2020). These neurons fire at specific locations within an environment, collectively generating a cognitive map of space (O'Keefe and Dostrovsky 1971). Understanding how these representations emerge and evolve over time is critical for uncovering the neural mechanisms underlying memory formation, retrieval, and updating.

Familiarization to a novel environment is a form of incidental spatial learning that occurs within a few minutes of exploration (Leussis and Bolivar 2006; Wilkinson et al. 2006; Bevins, Koznarova, and Armiger 2001). During this period, CA1 exhibits an elevated increase in place field formation (Dong, Madar, and Sheffield 2021), and at the population level, neural activity gradually evolves into a spatially structured state that is relatively stable (Berners-Lee et al. 2022; Cohen, Bolstad, and Lee 2017; Chiu et al. 2023; Zemla et al. 2022). Associated with this period, dendritic-targeting interneurons in CA1 show reduced activity that enhances dendritic spike prevalence that can boost synaptic potentiation (Sheffield, Adoff, and Dombeck 2017; Geiller et al. 2022). Over the next few minutes and trials as animals become more familiar with their environment, interneuron activity gradually increases, dendritic spike frequency declines, and inhibitory circuits develop inverse spatial tuning relative to place cells (Sheffield, Adoff, and Dombeck 2017; Geiller et al. 2022). Moreover, behavioral timescale synaptic plasticity (BTSP) (Bittner et al. 2017, 2015)—facilitating rapid place field formation and shifting of place fields—is highest during the initial moments in a novel environment but gradually diminishes with familiarity (Madar et al. 2025; Priestley et al. 2022). Together, these observations support two distinct phases in CA1 during exploration of a novel environment: an early phase of spatial

representation formation and refinement, followed by a late phase of stabilized representations. But how inputs into CA1 contribute to these early and late phases is not clear.

Bilateral CA3 provides the primary excitatory input to CA1 and is essential for spatial learning (R. P. Kesner 2007; Nakazawa et al. 2002; X.-G. Li et al. 1994). Left and right CA3 inputs form synapses that exhibit molecular and physiological differences that may influence their functional contributions (Shipton et al. 2014; Kohl et al. 2011; Shinohara et al. 2008; El-Gaby et al. 2016, 2021; Song et al. 2020; Kawakami et al. 2003; Jordan 2020). Synapses from Left CA3 inputs onto bilateral CA1 pyramidal cells contain more GluN2B-enriched NMDARs, which facilitate synaptic potentiation, whereas right CA3 synapses have a higher density of AMPARs, supporting stronger baseline transmission (Shinohara et al. 2008; Kawakami et al. 2003). Electrophysiological studies suggest that long-term potentiation (LTP) is induced only at left CA3-CA1 synapses, while right CA3-CA1 synapses do not undergo LTP (Shipton et al. 2014; Kohl et al. 2011). Behaviorally, inhibiting left CA3 disrupts long-term spatial memory, whereas right CA3 inhibition has minimal effects (Shipton et al. 2014; El-Gaby et al. 2016; Kohl et al. 2011). These findings suggest that left and right CA3 inputs may differentially contribute to CA1 spatial dynamics across distinct memory phases.

To address this, we used two-photon calcium imaging to record right CA1 (CA1_R) place cell activity while selectively inhibiting left or right CA3 inputs at specific timepoints during and following learning: during initial novel exploration (first 3 mins) when behavioral measures show ongoing learning and CA1 representations are forming and refining; after 3 mins when behavior shows no further learning and CA1 representations have stabilized (Fig. 1, 2). Additionally, we recorded CA3 axon activity in CA1 to examine how left and right CA3 inputs

evolve over time during the early and later phases, providing insight into how they influence CA1_R dynamics.

Results

We trained mice to run in a 2-meter 1D linear track virtual environment (Fig. 1a). Mice received water reward by the end of the track and were teleported back to the beginning. In the familiar environment animals were trained in, we found that animals robustly and precisely licked at locations on the track prior to the reward site (Fig. 1b, c, green; Fig. 1d). In comparison, during the initial laps in novel environments, as the animals were learning, they ran at lower velocity and licked all over the track (Fig. 1b, middle; Fig. 1d right). Over a few laps, animals developed precise anticipatory licking that reached a plateau close to the behavior in the familiar environment by lap 8 (Fig. 1e). The development of precise anticipatory licking behavior reflects spatial learning as animals undergo familiarization to the environment (Krishnan et al. 2022). Once precise anticipatory licking is established, we consider animals to be familiar with the novel environment.

Using two-photon microscopy, we recorded calcium transients from pyramidal cell populations in CA1_R as animals were navigating in novel environments (Fig. 2a). Upon immediate novel exposure, we observed elevated neural activity (Fig. 2b-d) that lasted under 2 minutes and within the first 5-10 laps (Fig. 2b-d) as animals familiarized with the environment. We trained decoders to predict location on the track during and following this learning process. CA1 decoding accuracy using population activity (all cell; Fig. 2e) and place cell firing (place cells; Fig. 2f) were initially low but progressively and rapidly improved, plateauing within ~10

laps and matching decoder performance in the familiar environment (Fig. 2e, bottom inset). Importantly, both decoding accuracy (Fig. 2) and anticipatory licking behavior (Fig. 1) showed quick development and plateaued within the first ~10 laps (or 2 minutes) in the novel environment. Thus, we consider this early phase in a novel environment (the first 10 laps) to be the learning phase in which familiarization to the environment is taking place. Following this early phase, we consider animals to be familiar with the novel environment.

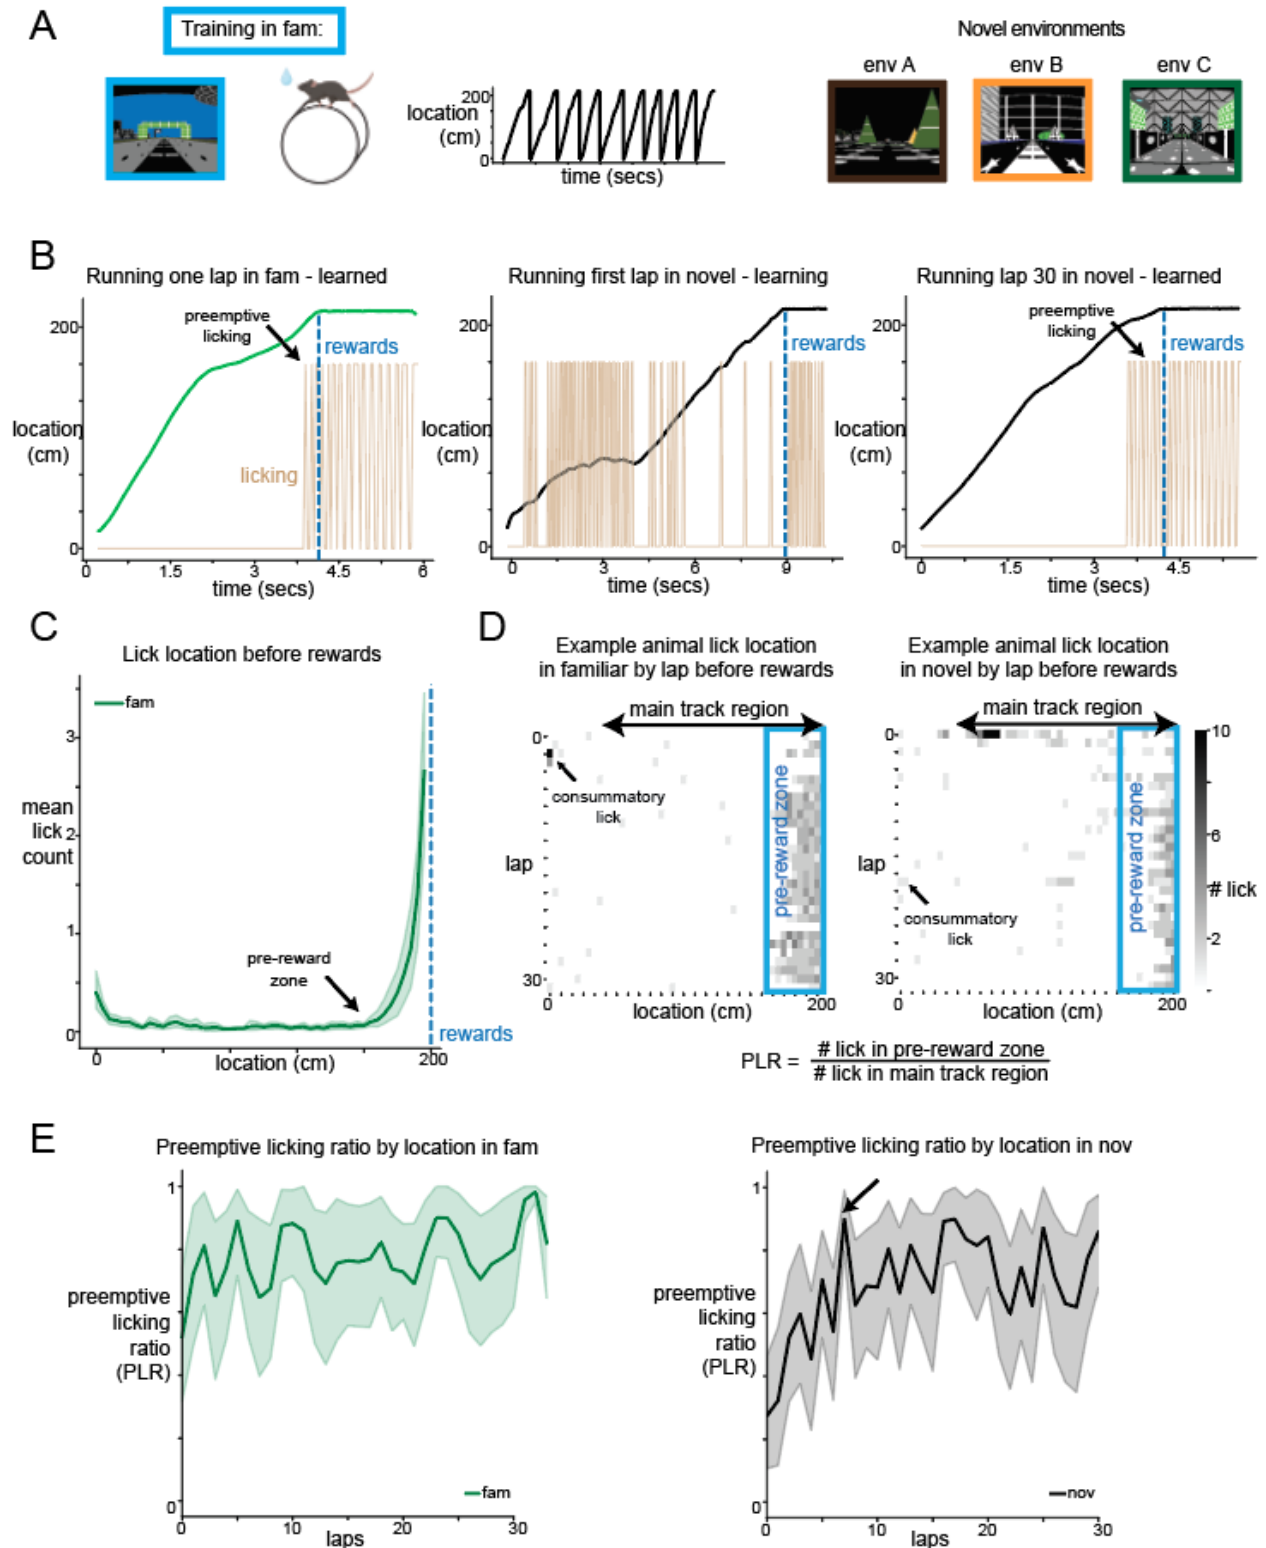


Fig1 Anticipatory licking as a behavioral readout of spatial learning and familiarization to a novel environment.

- A. Behavioral training setup: mouse running on the treadmill in a 2m track VR environment. left: familiar environment VR. right: 3 different novel VR environments.
- B. Example animal running one lap in left: familiar env (learned), middle: first lap in novel env (during learning), and right: lap 30 in novel env (learned). Over repeated laps, animals develop preemptive licking anticipating water rewards prior to reaching the rewards zone.
- C. Mean lick count by location on the track before reward delivery from all animals (n=12). The pre-reward zone (start of zone indicated by arrow) was defined as the start of elevated lick ratio in the familiar environment identified by the kneedle algorithm (see Methods).
- D. Example animal lick location on the track by lap in familiar (left) and novel (right) environments. Consummatory licks often happen at the beginning of the track as animals finish consuming water rewards from the previous lap. We defined the “main track region” to include track locations after consummatory licks and before reward delivery. Preemptive licking ratio (PLR) is defined as the ratio of number of licks in pre-reward zone over number of licks in the main track region.
- E. PLR over laps in familiar (left) and novel (right) environments averaged across all animals (n=12). Arrow points at the transition lap (lap 8) from continuous rising to stability in the novel environment identified by the kneedle algorithm.

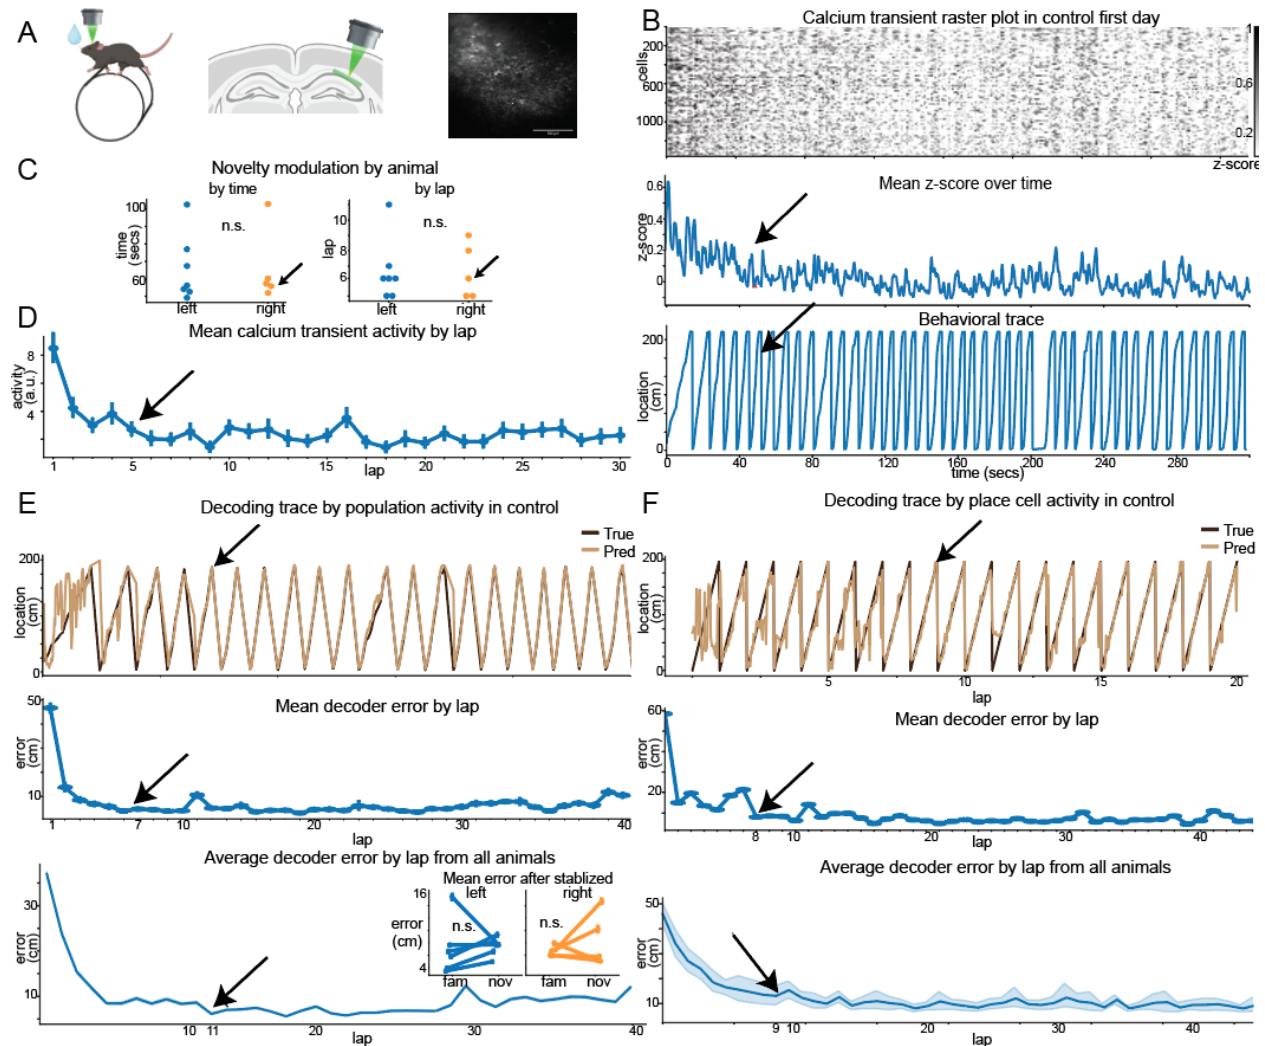


Fig 2. CA1 stabilization occurs within the first ~10 laps in a novel environment.

A. Imaging setup: in vivo recording of CA1 population activity as animals navigate in novel environments.

B. Z-score transformed neural activity of an example animal upon entering a novel environment. Top: Neural activity (see Methods) on control day 1, binned in ~160 ms intervals. Middle: Mean z-score over time. The arrow marks the transition point where novelty-driven elevated neural activity declined, identified using the kneedle algorithm (see methods) on smoothed (800 ms window) mean z-score data. Bottom: Behavioral trace of the same animal, with the arrow indicating the transition point calculated from the middle panel.

C. Novelty modulation time across animals, identified by the kneedle algorithm on smoothed mean z-score data, expressed in seconds and laps. Arrows highlight the example animal from A. No significant difference in time or laps between left and right groups (Mann-Whitney U test, laps: $p = 1$; time: $p = 0.88$)

D. Mean calcium transient activity across all cells per lap of the example animal in A., resulted in a comparable transition lap as A.

E. Decoder performance predicting location on the track trained using population activity dimension-reduced by CEBRA and a KNN regressor (see Methods). Top: Decoder-predicted

(light) vs. true (dark) location for the example animal. Middle: Mean decoder error per lap for the same animal, calculated as the mean absolute difference between true and predicted location. Bottom: Mean decoder error per lap averaged across animals. The arrow marks the transition lap where decoder error plateaued, identified using the kneedle algorithm. Inset: Mean decoder error in laps after the transition lap in novel vs. familiar conditions. No significant differences were found (Wilcoxon test, left: $p = 0.44$; right: $p = 0.81$).

F. Same as D., but using only place cell activity and Bayesian decoder (see Methods) resulting in similar transition laps as D.

To examine the contribution of bilateral CA3 inputs to spatial learning at the early familiarization and late familiarized phases in novel environments, we optogenetically inhibited left or right CA3 at these early and late phases while recording CA1 cell population. To achieve this, we injected cre-dependent opsin eOPN3 (Mahn et al. 2021) in either left or right CA3a (referred as CA3 from here on) of Grik4-cre mice (Fig. 3a). These mice have restricted cre expression in CA3 pyramidal cells (Nakazawa et al. 2002). In the same animals we also expressed GCaMP6f in dorsal CA1_R pyramidal cells. Using two-photon microscopy, we recorded calcium transients from pyramidal cell populations in CA1_R and delivered light pulses to dorsal CA1_R through the same objective to optogenetically inhibit either left or right CA3 inputs to dorsal CA1_R. We designed an inhibition and imaging protocol that contained two optogenetic conditions and a control condition. Each mouse experienced all 3 conditions, each one on a different day and each condition associated with a different novel environment. Also, In each mouse the same fields of view (FOVs) were tracked across all 3 conditions (Fig. 3b, c). The conditions: “Opto early” provided inhibition immediately upon entering into a novel environment and remained on for 3 minutes, “Opto late” provided inhibition after 3 minutes in a different novel environment on a different day and remained on for the subsequent 3 minutes. “Control” provided no optogenetic inhibition but the early and late phases were time-matched to the opto conditions for comparison. We selected a 3-minute optogenetic inhibition for the “opto

early” condition to span the period when familiarization to the novel environment is taking place (Fig. 1, 2). Importantly, our opto early condition provided inhibition throughout this early learning phase. In contrast, the opto late period was applied only after familiarization had occurred, i.e. once animals had developed stable and precise anticipatory licking behavior (Fig. 1d) and CA1 decoding accuracy and neural activity had stabilized (Fig. 2a-f).

Animals entered the inhibition and imaging protocol after passing behavioral criteria of running >2 laps per minute for 3 consecutive days and confirmed efficacy of opto inhibition by checking opsin expression in CA1 and testing inhibition of CA1 activity while animals are in a dark environment. On each day animals ran through the familiar environment first and then one of three novel environments paired with either the opto-early, opto-late, or the control condition (Fig. 3b and c). The combination and order of novel environments and opto conditions were counterbalanced between animals. Animals momentarily slowed down during the first lap in a novel environment compared to laps in the familiar environment, but did not show differences in running during optogenetic inhibition compared to control (Supp Fig. 1), suggesting that any changes in neural dynamics we found were not caused by changes in running speed in response to optogenetic inhibition.

We first examined the effects of optogenetic inhibition of left and right CA3 inputs on CA1 population activity. For each animal, we binarized calcium transients from all pyramidal cells per 160 ms time-bins and counted the number of transients in each time bin (Fig. 3c). We calculated the normalized transient frequency for each cell (calculated by transient frequency during opto normalized by transient frequency across the entire session) in the opto condition and compared to the time-matched control condition in the same animal. All animals showed

reduced transient frequency upon optogenetic inhibition (Fig. 3d), confirming that optogenetic inhibition worked on all animals.

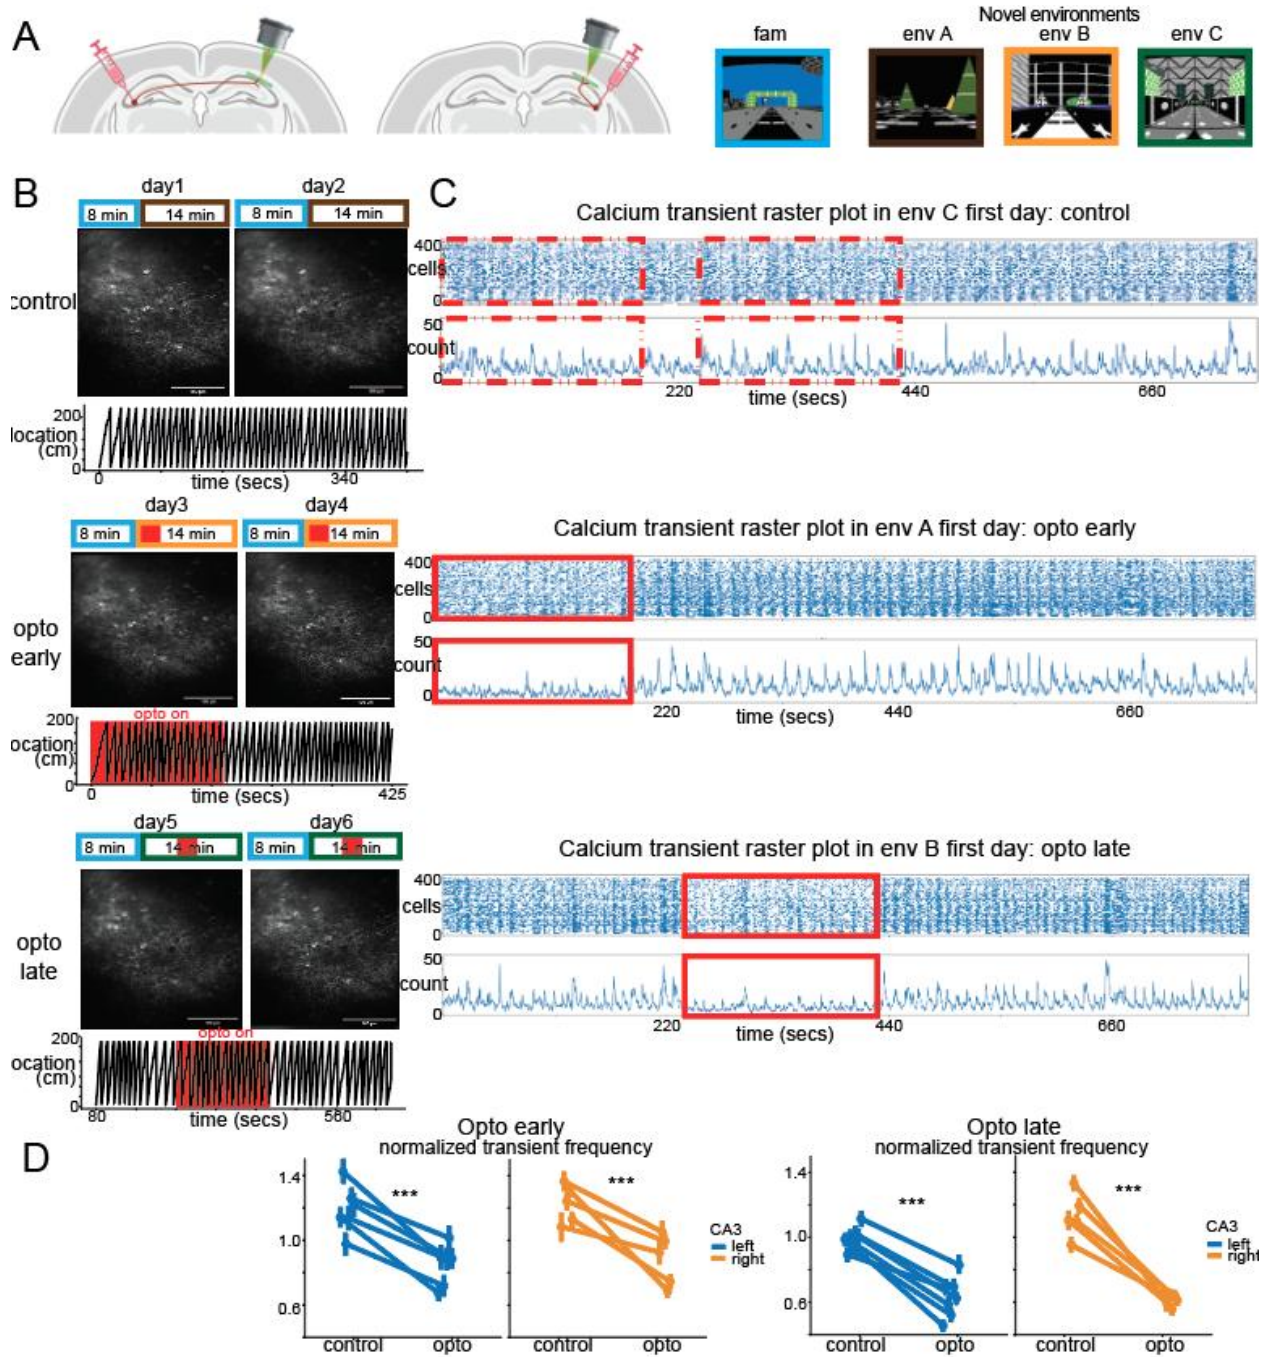


Fig3 Optogenetic inhibition of left and right CA3 axons reduces CA1 pyramidal cell activity.

- A. Left: Viral injection protocol. GCaMP6f is expressed in CA1_R for imaging of CA1 pyramidal cells (green) and eOPN3 is expressed in left or right CA3 for axon inhibition in CA1_R (red). Right: The VR environments (familiar and 3 novel environments).
- B. The 6-day recording protocol. The same FOV was recorded each day (example from one mouse shown). The environments and time in the environments are indicated above the FOVs by the color-coded boxes that match the environments in A. Example behavior traces are shown below the FOVs and the opto-on periods are indicated in red.
- C. Example animal calcium transient raster plot from the same population of CA1 pyramidal cells in the 3 different novel environments. Each dot in the raster plot indicates a single calcium transient. Below each raster plot is the number of cells with a calcium transient per 160ms time-bin. Solid red box: opto-on periods. Dashed red box: No-opto time-matched control period.
- D. Opto-inhibition reduces normalized transient frequency in both opto early and late conditions in every animal. Normalized transient frequency is calculated by dividing transient frequency during opto on period (solid or dashed red box) by transient frequency during the entire time. Wilcoxon signed-rank test done on all cells within each individual animal with Bonferroni correction. $p < 0.001$ for all 11 animals in opto early and opto late.

Right CA3 Inputs Underlie Early CA1 Spatial Refinement, Whereas Left CA3 Inputs Sustain Late Stability

Besides facilitating the firing activity of individual CA1 pyramidal cells, bilateral CA3 inputs have also shown to contribute to CA1 population spatial coding (Guan et al. 2021; Davoudi and Foster 2019; Zutshi et al. 2022). We investigated if left and right CA3 inputs made distinct contributions to CA1_R spatial coding during early and late phases following novel environment exposure. Multiple studies have identified neural manifolds within CA1 population activity that encode both task-relevant and irrelevant information using various dimensionality reduction techniques (Guo et al. 2024; Yang et al. 2024; Nieh et al. 2021; Aery Jones and Giocomo 2023). To uncover low-dimensional trajectories of CA1 population activity and the effect of optogenetic inhibition of CA3 inputs, we applied CEBRA (McKenzie et al. 2021), a nonlinear dimensionality reduction algorithm that links neural dynamics to behavioral data. The resulting manifolds preserved the topology of the 1D VR environment and exhibited clear linear separation in latent space (Fig. 4a, left & middle), consistent with prior findings on structured spatial representations

in CA1 (Guo et al. 2024; Yang et al. 2024). During optogenetic inhibition (red-coded; Fig. 4a, right), the projected manifold spread out farther from the centroid compared to the time-matched control condition (Supp. 2a, see methods), suggesting a potential disruption in the low-dimensional neural representation of the environment.

To test this, we assessed whether location on the track during opto-on periods could be accurately decoded using a latent embedding learned during opto-off periods. Under time-matched control conditions, location decoding accuracy was high (Fig. 4b, c). In the control condition, decoder performance improved rapidly over the first few traversals (Fig. 2e, f, 4c), resembling a spatial learning process in which CA1 representations emerge and refine. In contrast, in the late-phase in the control condition, decoder performance was stable across laps, suggesting CA1 representations had already stabilized with familiarity. However, optogenetic inhibition significantly impaired decoding accuracy during both early and late phases (Fig. 4c). Interestingly, right CA3 inhibition had a significantly stronger effect than left CA3 inhibition on reducing decoder performance during the early phase, whereas left CA3 inhibition had a significantly greater impact than right CA3 inhibition during the late phase (Fig. 4d). These results indicate that right CA3 inputs contribute more to the development of CA1 representations during early novel exposure, while left CA3 inputs contribute more once representations have stabilized.

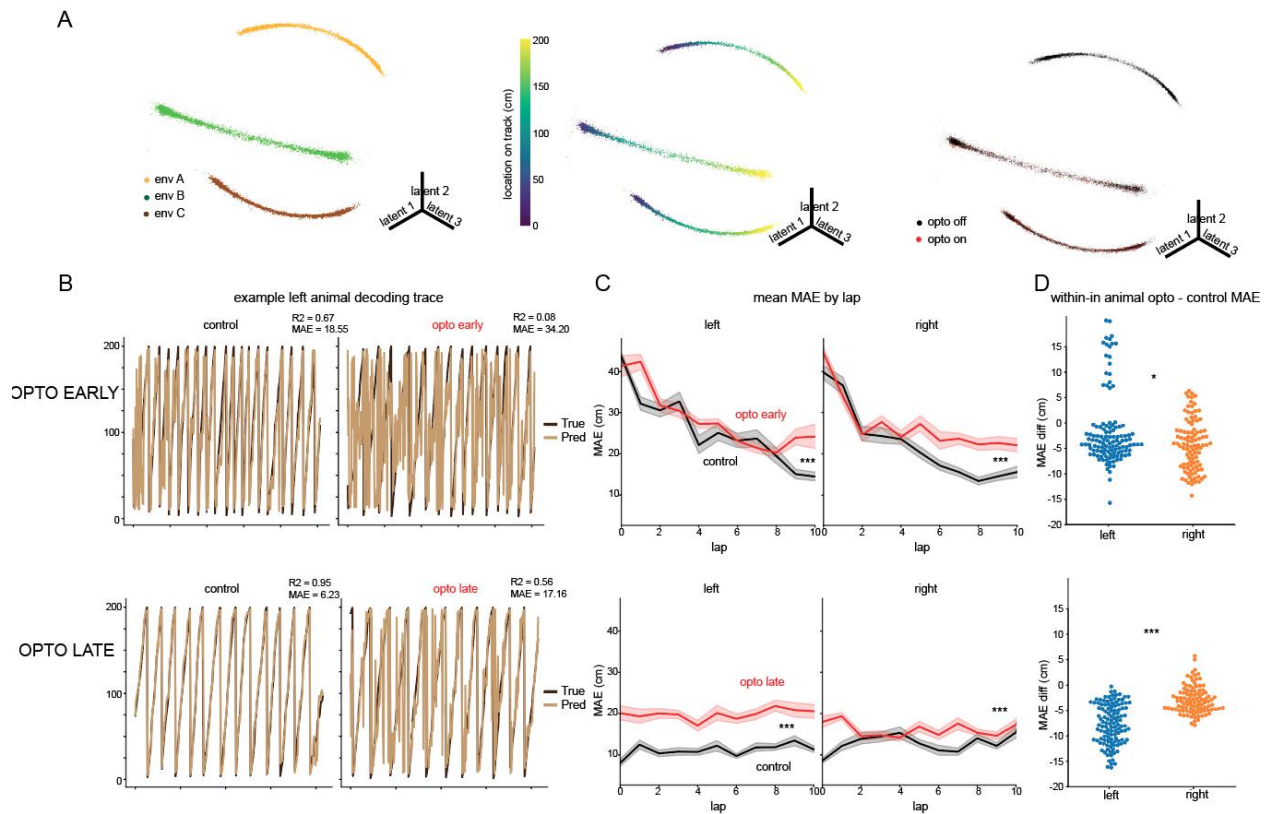


Fig4 CA1 low-dimensional neural manifolds and spatial decoding reveal a dynamic contribution of left and right CA3 inputs that switches between initial learning and post-learning phases

A. Neural manifold color coded by environment (left), location on the track (middle), opto condition (right). Each point corresponds to the low dimensional representation of binarized population activity per 160ms.

B. True running trajectory (dark) vs. decoded running trajectory (light) from latent embeddings in time matched control condition and opto condition during different phases.

C. Decoder error per lap represented by mean absolute error (MAE) in cm averaged across animals. Both left and right opto-inhibition significantly increased MAE during opto-early (top) and opto-late (bottom) phases. Wilcoxon signed rank test comparing within animal MAE between opto and control: $p < 0.001$ for all conditions.

D. Time-matched within-animal decoder MAE difference between opto and control condition. Each dot represents the MAE difference between opto and control in a lap in C. Below 0 reflects that the decoder performed worse under opto than control within the same animal on the same lap. Top: opto-early, right opto-inhibition disrupted decoder performance significantly more than left opto-inhibition. Mann Whitney U, $P = 0.02$. Bottom: opto-late, left opto-inhibition significantly disrupts decoder performance more than right inhibition. Mann Whitney U, $P < 0.001$.

Right CA3 inputs control early CA1 place cell dynamics whereas left CA3 inputs dominate later

To investigate how individual place cells contribute to the disruptions in population coding described above, we analyzed lap-by-lap dynamics of CA1 place cells in response to optogenetic inhibition of left and right CA3 inputs. Since inhibiting either input reduced CA1 pyramidal cell firing frequency (Fig. 4d), we assessed whether in-field firing within place fields, critical for spatial coding (McKenzie et al. 2021; Robinson et al. 2020), was affected.

We first detected place cells with significant place fields based on their activity across all laps during the session. We then assessed place field reliability, a measure of how often a place cell fires in its place field across laps. Specifically, we quantified the number of laps with an in-field calcium transient within a 6-lap sliding window for each place field (Fig. 5a). During the early phase, inhibiting either left or right CA3 inputs significantly reduced place field reliability (Fig. 5b, c) but right CA3 inhibition had a greater impact than the left (Fig. 5b, inset). However, left CA3 inhibition more strongly affected reliability than the right during the later phase (Fig. 5c, inset).

Reliability could be influenced by both newly emerging place fields—characterized by abrupt, step-like increases in reliability as they transition from absent or sporadic firing to consistent in-field activity (Fig. 5d)—and by the lap-lap reliability of already established place fields (Fig. 5g). To disentangle these contributions, we directly tested whether left versus right CA3 inhibition differentially affected CA1 place field emergence.

We defined the emergence lap as the first lap in which subsequent in-field firing became reliable, requiring at least 3 out of 6 consecutive laps to contain in-field activity. Upon novel

environment exposure, CA1 place fields typically emerge rapidly within the first few laps (Dong, Madar, and Sheffield 2021). Consistent with this, many CA1 place fields formed quickly under control conditions (Fig. 5e, black lines). However, optogenetic inhibition of either left or right CA3 inputs during initial novel exposure significantly delayed place field emergence (Fig. 5e, upper panel) and reduced the total number of fields formed (Fig. 5e, lower panel), with no significant difference between left and right CA3 inhibition (Supp Fig. 3b) and this effect was not due to our definition of emergence (Supp Fig. 3a, Supp Fig. 3b). This indicates that the stronger impact of right CA3 inhibition on place field reliability is not due to a greater disruption of place field emergence.

To further investigate the impact of CA3 inhibition on place field reliability, we examined whether already emerged place fields became less reliable over successive laps under inhibition (Fig. 5g). Following the formation of place fields, reliability gradually declined in the control group (Fig. 5h, black lines), potentially due to reported inhibitory dynamics in CA1 that show initially low inhibition in novel environments that gradually increase during the early phase ^{12,13}. We found that during the early phase, inhibiting the right CA3 inputs reduced place field reliability, with this effect becoming more pronounced over laps and reaching statistical significance towards the end of the early phase (Fig. 5h). In contrast, left CA3 inhibition had no significant effect on place field reliability at any point during the early phase (Fig. 5h). These results indicate that the stronger impact of right CA3 inhibition on place field reliability is specifically due to decreased reliability of emerged place fields rather than a delay in their emergence.

In contrast, during the later phase, both left and right CA3 inhibition significantly reduced place field reliability (Fig. 5i), with left CA3 inhibition having a significantly stronger effect than right CA3 inhibition throughout the phase (Fig. 5i, inset). Additionally, left CA3 inhibition during the later phase also impaired new place field emergence, whereas right CA3 inhibition did not (Fig. 5f). These findings indicate that while both left and right CA3 inputs are necessary for the rapid emergence of CA1 place fields during initial novel exposure, right CA3 primarily regulates place field reliability early on, whereas left CA3 plays a greater role in maintaining reliability later and continues to drive new place field emergence during the later phase.

The reliability analysis assessed whether a place cell fired on a given lap (a binary measure) but did not account for variations in firing rate within place fields. Place cells modulate their firing rates in response to contextual changes (rate remapping), reflecting both task-relevant and irrelevant cognitive processes (Keinath, Mosser, and Brandon 2022; Ji and Wilson 2008). As a result, CA1 place field gain carries important information that can influence downstream circuits (Kitanishi et al. 2015; Butola et al. 2025). To examine whether left and right CA3 inputs differentially regulate place field gain across memory phases, we quantified the area under the curve (AUC) during the rise-to-peak period of calcium transients on each lap as a proxy for firing rate (Fig. 5j). To facilitate comparisons across cells and experimental conditions, we z-score transformed within-field transient amplitudes, focusing on relative changes rather than absolute values.

Optogenetic inhibition of either left or right CA3 inputs significantly reduced place field gain across all opto conditions (Fig. 5k, Fig. 5l), consistent with prior findings from bilateral

CA3 inhibition (Davoudi and Foster 2019). However, the effects varied by learning phase: right CA3 inhibition caused a greater reduction in in-field firing during early novel exposure, whereas left CA3 inhibition had a stronger effect during the later phase. These findings suggest that right CA3 plays a dominant role in early gain determination, while left CA3 takes over in the later phase to determine gain.

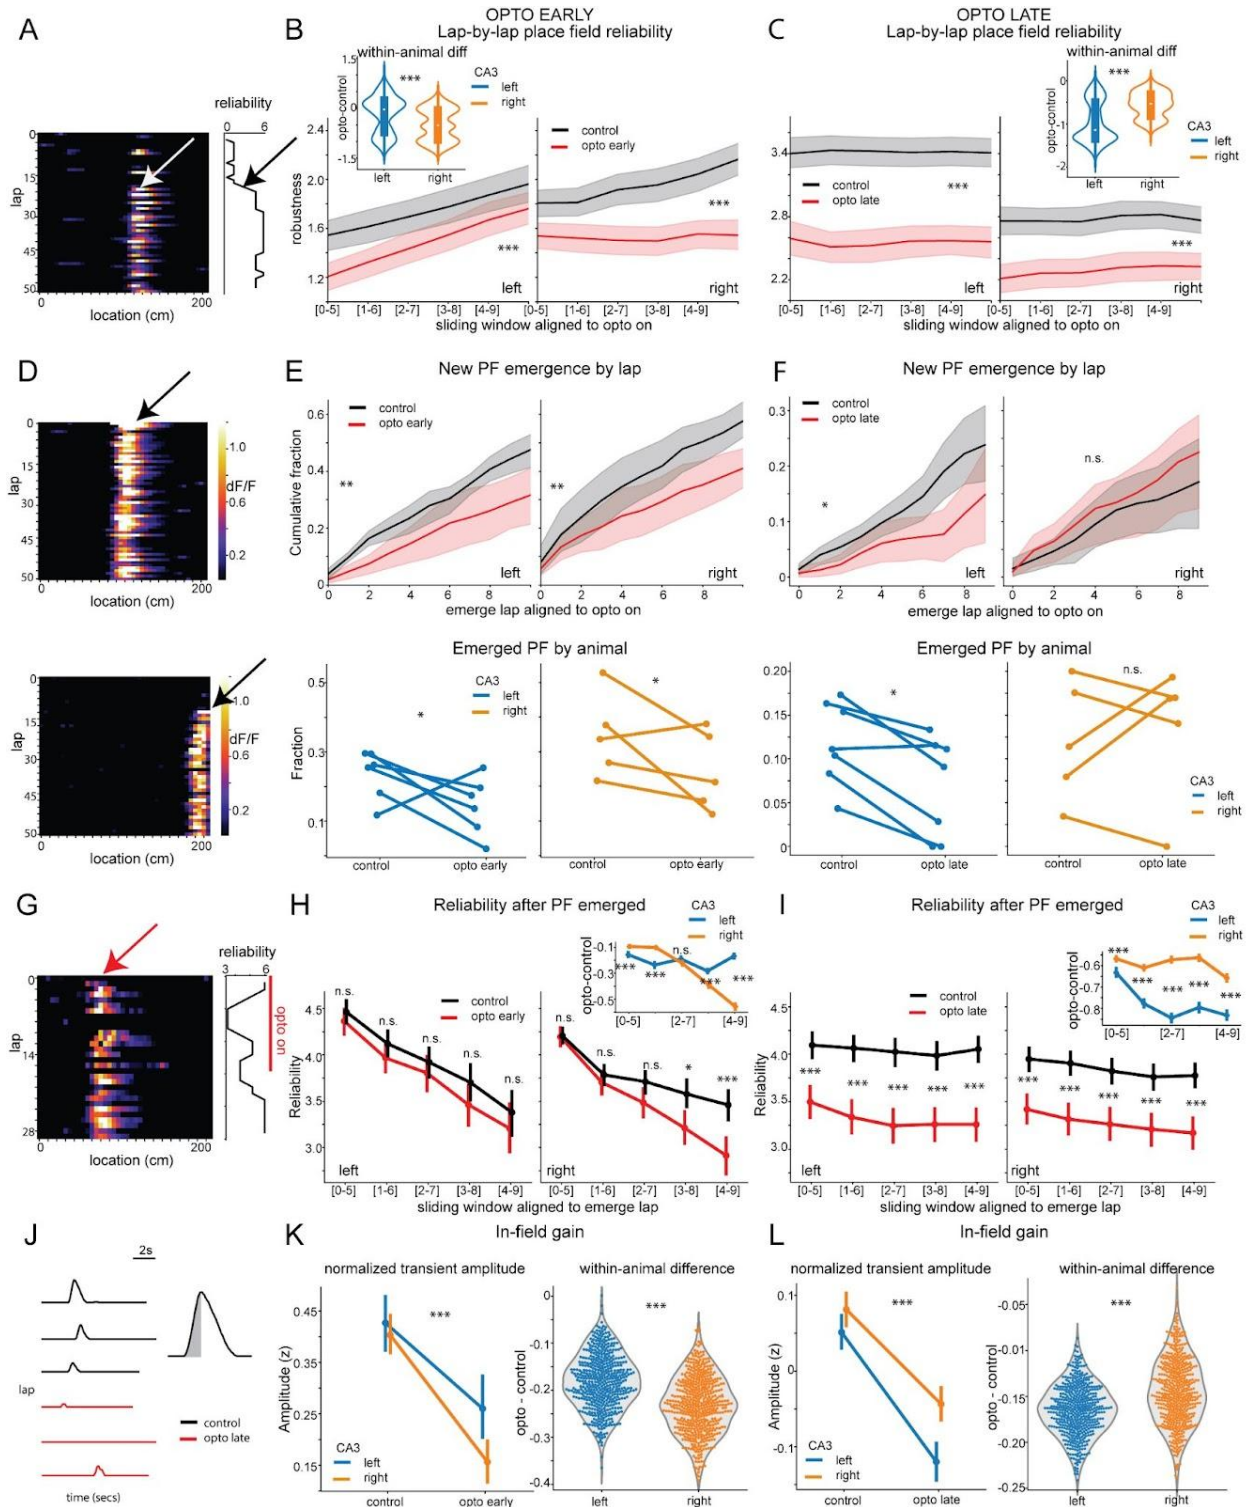


Fig5 The influence of left and right CA3 inputs on the emergence, reliability, and gain of CA1 place fields during initial learning and post-learning phases

- A. Example place field activity and its reliability measured over a 6 lap sliding window. 0 is inactive and 6 is always firing in-field for 6 laps. The cell only began consistently firing within-field after lap 18 (pointed by the white arrow), represented by a step-like increase in reliability (black arrow).
- B. Average in-field firing reliability across all place fields for first 10 laps under optogenetic inhibition compared to within-animal lap-matched control. Both left and right CA3 input inhibition significantly reduced reliability (Mann-Whitney U test $p < 0.001$). Inset: Bootstrapped data for within animal difference distribution. Above 0: place fields are more reliable in opto versus control. Below 0: more reliable in control. Right CA3 inputs inhibition reduced place field reliability significantly more than left CA3 input inhibition during opto-early (Mann-Whitney U test $p < 0.001$).
- C. Same as B for opto-late. Both left and right CA3 inputs inhibition significantly reduced in-field firing reliability (Mann-Whitney U test $p < 0.001$ between opto and control for both left and right). Inset: Left CA3 inputs inhibition significantly reduced place fields reliability more than right CA3 input inhibition in opto-late (Mann-Whitney U test $p < 0.001$).
- D. Two example place fields; one emerged on lap 1 (top) and one on lap 18 (bottom). Emergence lap is defined as the first lap in which 3 out of 6 subsequent laps fired.
- E. Top: Cumulative histogram showing the proportion of place fields that have emerged as a function of laps during opto-early (red) versus control (black). Both left and right CA3 inhibition significantly delayed the emergence of CA1 place fields. Wilcoxon Signed-Rank Test left: $p = 0.026$, right: $p = 0.003$. Bottom: Within-animal comparison of fraction of place fields emerged over laps under opto-inhibition and control. Both left and right CA3 inhibition significantly decreased the number of fields that emerged (linear mixed effects model, main effect of opto $p = 0.005$, interaction term opto x CA3 $p = 0.97$, CA3 $p = 0.062$).
- F. Same as E for comparing opto-late (red) with control (black). Top: Only left CA3 inhibition significantly delayed the emergence of place fields. Wilcoxon Signed-Rank Test left: $p = 0.014$; right: $p = 0.143$. Bottom: Only left CA3 inhibition significantly reduced the fraction of fields formed under opto compared to lap-matched control. (linear mixed effects model, main effect of opto $p = 0.004$ and interaction term opto x CA3 $p = 0.017$, CA3 $p = 0.923$. Wilcoxon signed-rank test left: $p = 0.031$; right: $p = 0.19$).
- G. Example place field showing emergence lap (red arrow) and reliability (right).
- H. Reliability of place fields after emergence for opto early (red) and control (black). Left CA3 input inhibition did not affect reliability: 5 lap windows: $p = 1.2, 0.8, 1.5, 0.35, 1.7$. Right CA3 input inhibition significantly reduced reliability in the last two sliding windows: 5 lap windows: $p = 4.6, 2.1, 0.11, 0.011, < 0.001$. Inset: Bootstrapped distribution within-animal difference in mean reliability between opto and control (see methods). Right CA3 inhibition reduced reliability significantly more than left in the last two sliding windows. 5 lap windows from left to right: $p < 0.001, < 0.001, 0.15, < 0.001, < 0.001$. All tests used Mann-Whitney U test with Bonferroni correction (*5).
- I. Same as H for opto late. Both left and right CA3 inhibition significantly reduced place field reliability: $p < 0.001$ for all comparisons. Inset: Left CA3 inhibition reduced reliability significantly more than right CA3 inhibition ($p < 0.001$ for all comparisons).
- J. Calcium transients from an example place field without optogenetic inhibition (black) and with optogenetic inhibition (red). Transient magnitude is calculated as the area under the curve from rise to peak (shaded area).

K. Left: Normalized in-field transient magnitudes in control vs. opto early from all transients pooled together across animals. Both left and right inhibition significantly reduced in-field gain (Mann-Whitney U test left: $p < 0.001$; right: $p < 0.001$). Right: Bootstrapped distribution for within-animal difference in mean z-score transformed in-field transient magnitudes between control and opto early. Right CA3 input inhibition reduced in-field transient magnitudes significantly more than left CA3 input inhibition (two sample t-test, $p < 0.001$).

L. Right: Same as K for opto late. Left: Both left and right CA3 inhibition significantly reduced transient magnitudes (left: $p < 0.001$; right: $p < 0.001$). Right: Left CA3 inhibition reduced in-field transient magnitudes significantly more than right CA3 input inhibition (two sample t-test, $p < 0.001$).

Lateralized CA3 axon dynamics underlie phase-specific contributions to CA1 during initial learning and post-learning phases.

To investigate whether differences in left and right CA3 input dynamics underlie their distinct roles in CA1 during early and late phases of novel environment exposure, we recorded the activity of left or right CA3 axons in right CA1. Using Grik4-Cre mice, we injected Cre-dependent axon-GCaMP6s into either left or right CA3 and performed in vivo two-photon calcium imaging of CA3 axons in the stratum oriens of CA1 during spatial navigation in a familiar environment and two novel environments over three days (Fig. 6a). All subsequent analyses focused on the first day of exposure to the two novel environments combined.

We first analyzed the spatial coding properties of individual CA3 axons from the left and right and found axon place fields that tiled the entire track (Supp Fig. 4A). The emergence dynamics of these place fields were the same in the left and right axons (Supp Fig. 4B) and their place field properties were similar (although some small differences were observed; Supp Fig. 4C). We then quantified the overall activity of left and right axons to see if they provided different levels of excitation to CA1. To match the experimental optogenetic inhibition protocol, we defined the first 10 laps of novel exposure as the early phase of learning and laps 20–30 as

the late phase following familiarization. Axon transients were detected following the approach used in the *subprep* package (Jiang, Zhao, and Sheffield 2024), and transient magnitude was quantified using the area under the curve (AUC) during the rise-to-peak period as a proxy for spiking activity (Fig. 6c). To assess overall axon activity within each phase, we normalized transient activity by summing all transient magnitudes and dividing by 10 laps. Right CA3 axons exhibited significantly higher activity than left CA3 axons during the early phase. However, following familiarization, right CA3 activity declined and fell below left CA3 activity, which remained stable between the early and the late phase (Fig. 6d). These findings indicate that right CA3 axons are more active early in a novel environment when familiarization is taking place, whereas left CA3 axons become relatively more active than right axons following familiarization. This shift in relative activity likely underlies the transition from right CA3 dominance in controlling CA1 spatial representations during learning to left CA3 dominance following learning when representations are relatively stable.

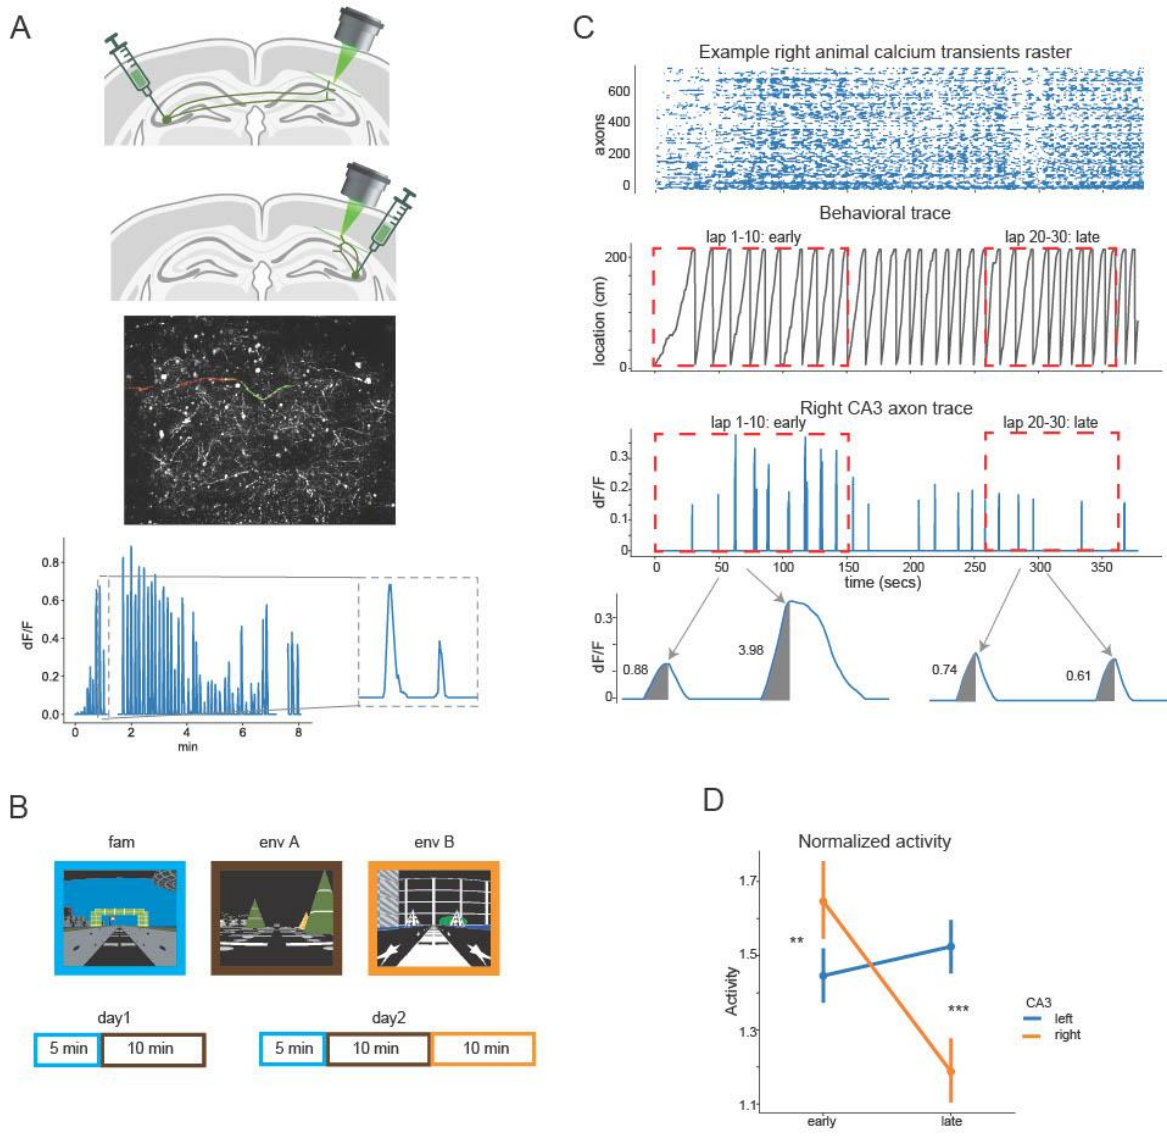


Fig6 Recordings of CA3 axons reveal a switch from stronger right CA3 axon activity during the early phase to stronger left CA3 axon activity during the later phase.

A. Left: Cre-dependent axon gcamp6s was injected to Grik4-cre mice in left/right CA3. Middle: Example FOV of axon population recording of left CA3 axons in right CA1 at S.O. Right: Example axon calcium activity over time after preprocessing.

B. Example VR environments. After reaching for the same behavioral criteria, animals went through 3 days of recording in familiar and novel environments.

C. Top: Example calcium transient raster of all axons in env A. Middle: Behavioral trace. First 10 laps are defined as early. Lap 20-30 are defined as late. Bottom: Example right CA3 axon calcium transient over time. Normalized axon activity in the early or late phase is

calculated as the rise-to-peak AUC (shaded area for zoomed in example transients) of all transients divided by the number of laps in the early or late phase.

D. Normalized transients activity across all axons in early and late phases. early: $p=0.006$; late: $p<0.001$, Mann whitney U test.

Discussion

Our results reveal that left and right CA3 inputs contribute to CA1 spatial representations in distinct, time-dependent ways during familiarization to a novel environment - a form of incidental learning. Initially, when animals are exposed to a novel environment, both inputs are required for the rapid emergence of CA1 place fields. However, once these fields form, right CA3 inputs predominantly enhance early spatial coding by boosting place field reliability, amplitude, and overall decoding accuracy. Once the environment becomes familiar, left CA3 inputs assume a more dominant role, maintaining stable spatial representations by preserving field reliability and gain and by supporting the emergence of new fields. This functional transition is mirrored by axonal recordings in CA1: right CA3 axons are more active during early exposure, while left CA3 axons become relatively more active following familiarity. Together, these findings demonstrate a dynamic, temporally structured division of labor between left and right CA3 inputs in shaping CA1 spatial maps during and following learning.

Manifold analysis revealed that during initial exploration, inhibiting right CA3 inputs significantly reduced spatial decoding accuracy, indicating that right CA3 is crucial for rapidly establishing a high-fidelity spatial map. This early role is driven by right CA3's capacity to quickly generate reliable, high-gain place fields when memory formation dominates, ensuring that new spatial experiences are effectively encoded (Nakazawa et al. 2002; Rebola, Carta, and Mulle 2017). Since CA1 is the primary hippocampal output, its accurate spatial representations during novel exploration and memory formation likely play a critical role in driving synaptic

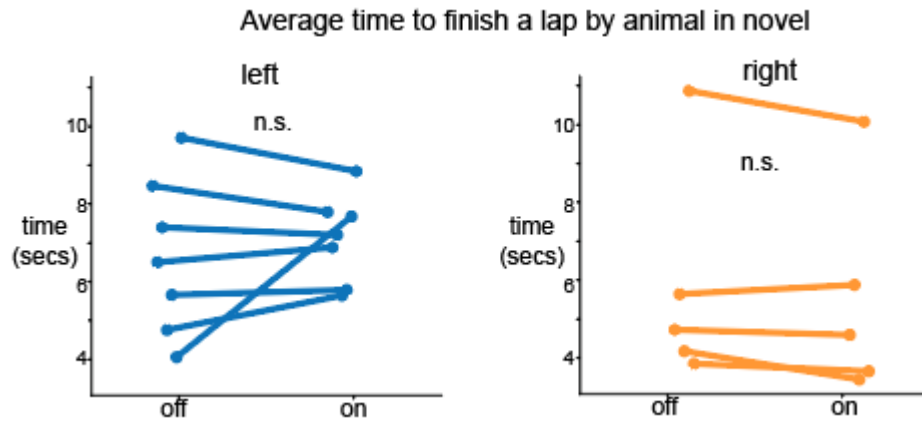
plasticity in downstream brain regions (Butola et al. 2025; Robinson et al. 2020). At the single cell level, high-gain, reliable place fields may promote robust synaptic plasticity at CA1 output synapses, supporting the formation of memory engrams (Kitamura et al. 2017; Zaki and Cai 2025). Disrupting these dynamics through right CA3 inhibition results in less consistent CA1 output (lower gain, unreliable place fields), leading downstream neurons to receive variable or misaligned inputs that could impair memory formation.

Following familiarization once animals precisely predict the reward location in the novel environment, the role of CA3 inputs shifts. During this later phase, left CA3 inputs assume greater control, as inhibition of left CA3 led to a more pronounced reduction in spatial decoding accuracy. This suggests that left CA3 is critical for refining and maintaining CA1 spatial maps once initial memory formation has occurred and the hippocampus operates in a memory retrieval and updating mode (familiarity requires memory retrieval) (Shipton et al. 2014; El-Gaby et al. 2016). While CA1 representations naturally drift over time—a process thought to support memory updating—excessive drift may impair memory retrieval (Geva et al. 2023; Keinath, Mosser, and Brandon 2022; Khatib et al. 2023; Krishnan and Sheffield 2023; Delamare et al. 2023). Our findings indicate that left CA3 inputs help stabilize the CA1 population code by maintaining reliable, high-gain place fields, which is essential for ensuring that downstream neurons receive a consistent hippocampal output during memory retrieval. Simultaneously, left CA3 also promotes the emergence of new place fields, allowing for ongoing memory updating as new experiences are integrated.

Two-photon calcium imaging of CA3 axons in the right CA1 revealed that right CA3 axons exhibit elevated activity during the early learning phase of novel exposure—precisely

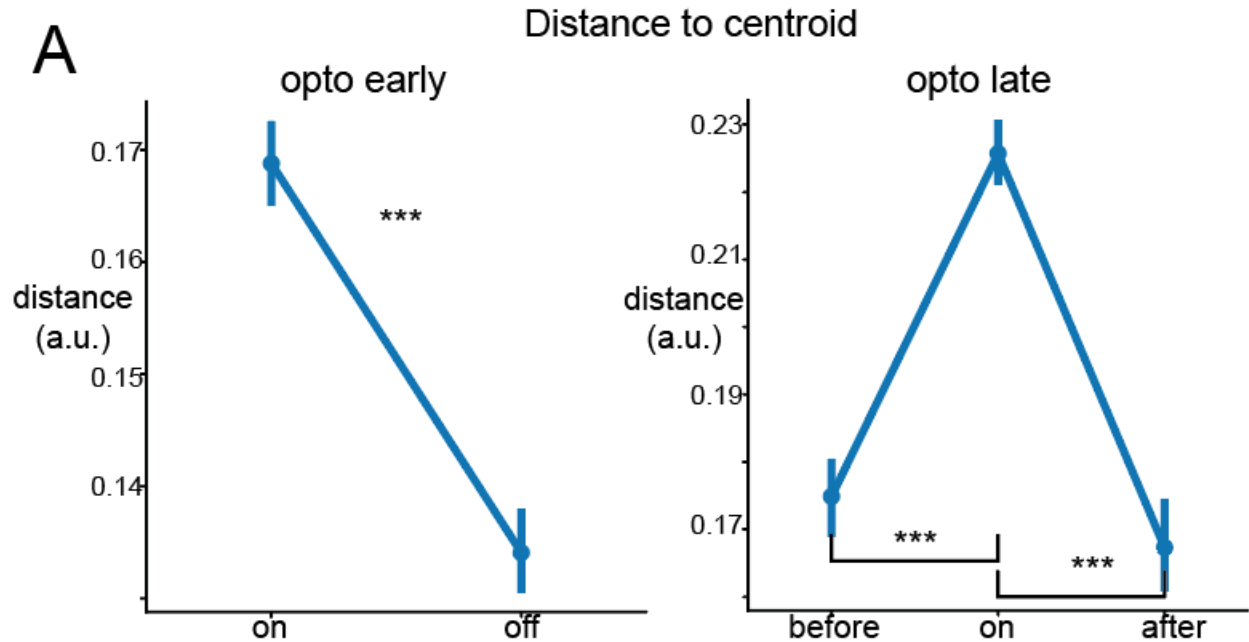
when CA1 place fields rapidly emerge with high gain and reliability. This early surge in right CA3 activity likely drives the rapid encoding process necessary for effective memory formation by ensuring that CA1 outputs are reliable and capable of consistently driving downstream targets. In contrast, as the environment becomes familiar, right CA3 activity greatly declines while left CA3 axon activity remains stable and becomes higher than right CA3 axon activity. This shift aligns with our CA1 findings, where left CA3 inputs become essential for maintaining and refining spatial maps, supporting memory retrieval and memory updating. Thus, the dynamic axonal activity patterns we observed explain how the temporal balance of CA3 inputs orchestrates the transition from rapid encoding during novel exposure (via right CA3) to the stable, adaptive representation of space once the environment is familiar (via left CA3).

Our findings extend previous molecular work by linking synaptic properties to functional roles in spatial learning. Prior studies showed that right CA3 axons form strong, AMPAR-rich synapses on large mushroom spines, whereas left CA3 axons innervate smaller, GluN2B-enriched spines with greater plasticity potential (Shinohara et al. 2008; Kawakami et al. 2003). Consistent with these distinctions, we observed that during early novel exploration, right CA3 inputs are highly active and drive the rapid emergence of reliable, high-gain CA1 place fields, providing robust initial encoding. As the environment becomes familiar, left CA3 inputs increase in activity, suggesting that synaptic plasticity potentiates these initially weaker synapses to support maintenance and refinement of spatial maps. This coordinated transition—right CA3 driving rapid encoding and left CA3 sustaining stabilization and ongoing updating—aligns with previous findings that left CA3 inputs are essential for learning over repeated trials (Shipton et al. 2014; El-Gaby et al. 2016, 2021), further underscoring the distinct, experience-dependent contributions of CA3 pathways to memory processing.



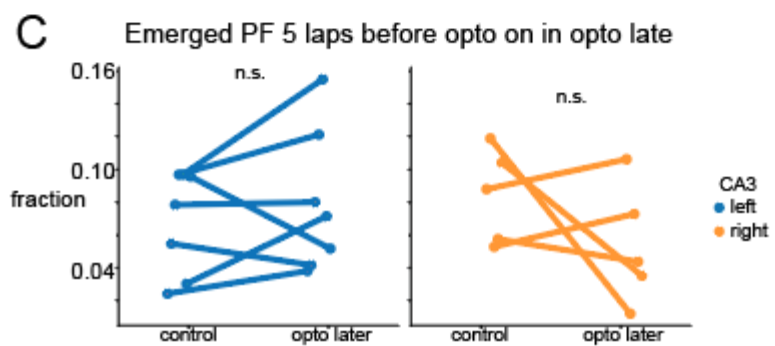
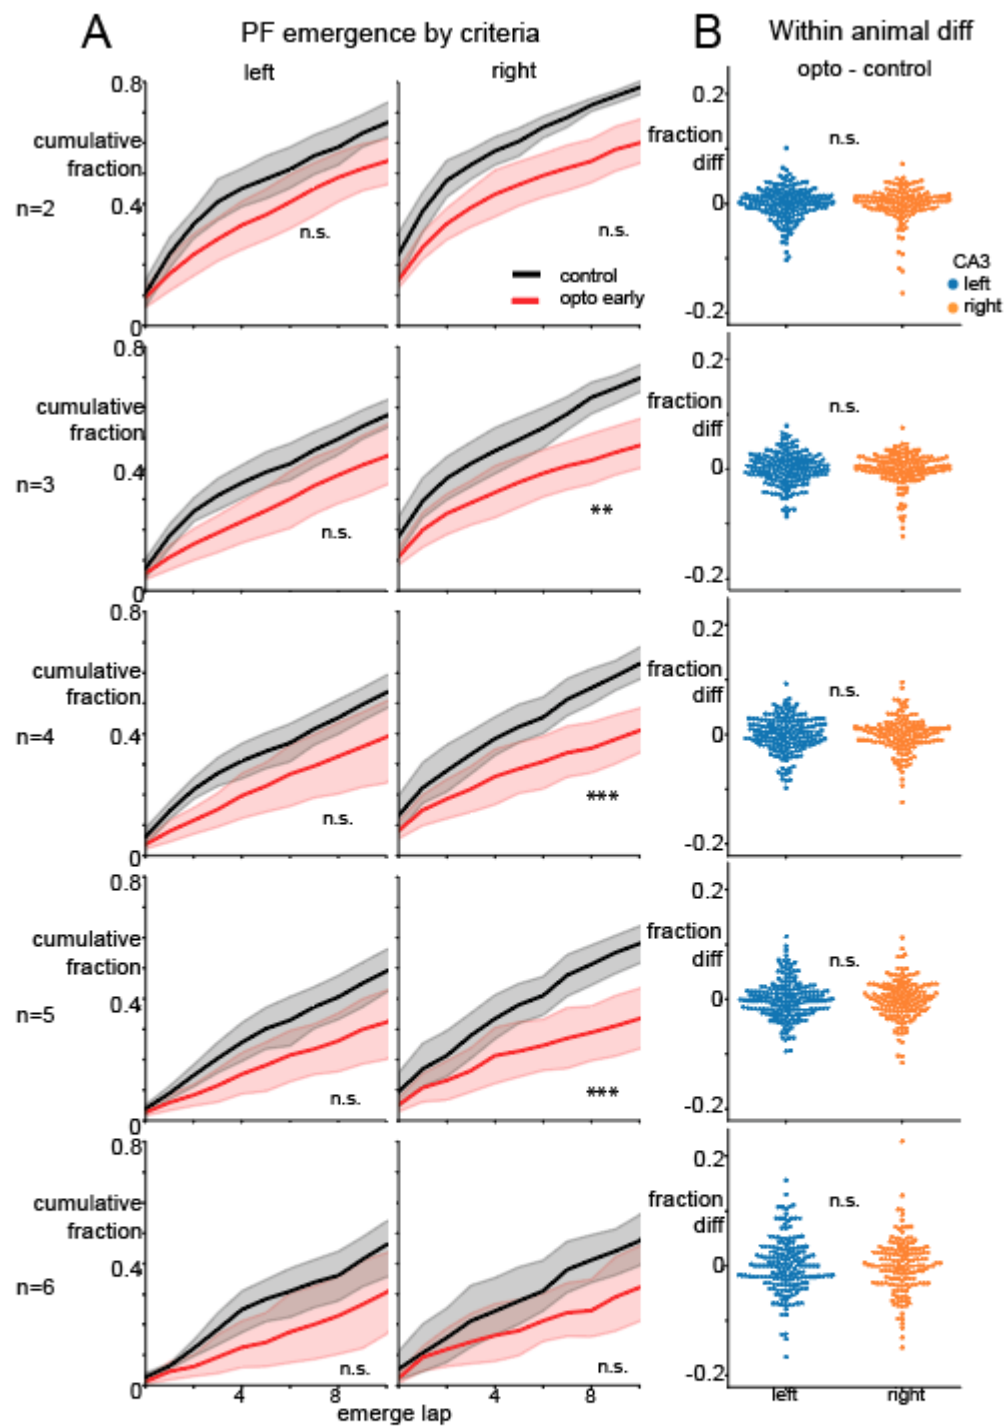
Supp1. CA3 Optogenetic inhibition did not alter the mean running speed of animals in novel environments.

Average time elapsed (in seconds) for running one lap under optogenetic inhibition and before turning on opto across all mice. Frames recorded within the teleportation window were not included. Wilcoxon test between opto off and on, left: $p=0.69$; Right: $p=0.31$. Between off and on pooled between all animals: Wilcoxon test $p=0.79$.



Supp2. Right optogenetic inhibition significantly increased the distance to centroid in the same example animal as Fig4, representing a less stereotypical neural activity to location on the track under opto.

A. Distance to centroid in opto early and opto late. Same example animal as (Fig4. a). See methods for calculating distance to centroid. Left: Optogenetic inhibition significantly increased distance to centroid than off in opto early. (Mann whitney u stats= 152935071, $p < 0.001$). Right: Turning on optogenetic inhibition significantly increased distance to centroid than before and after. (Before vs. on: $p < 0.001$; after vs. on: $p < 0.001$).

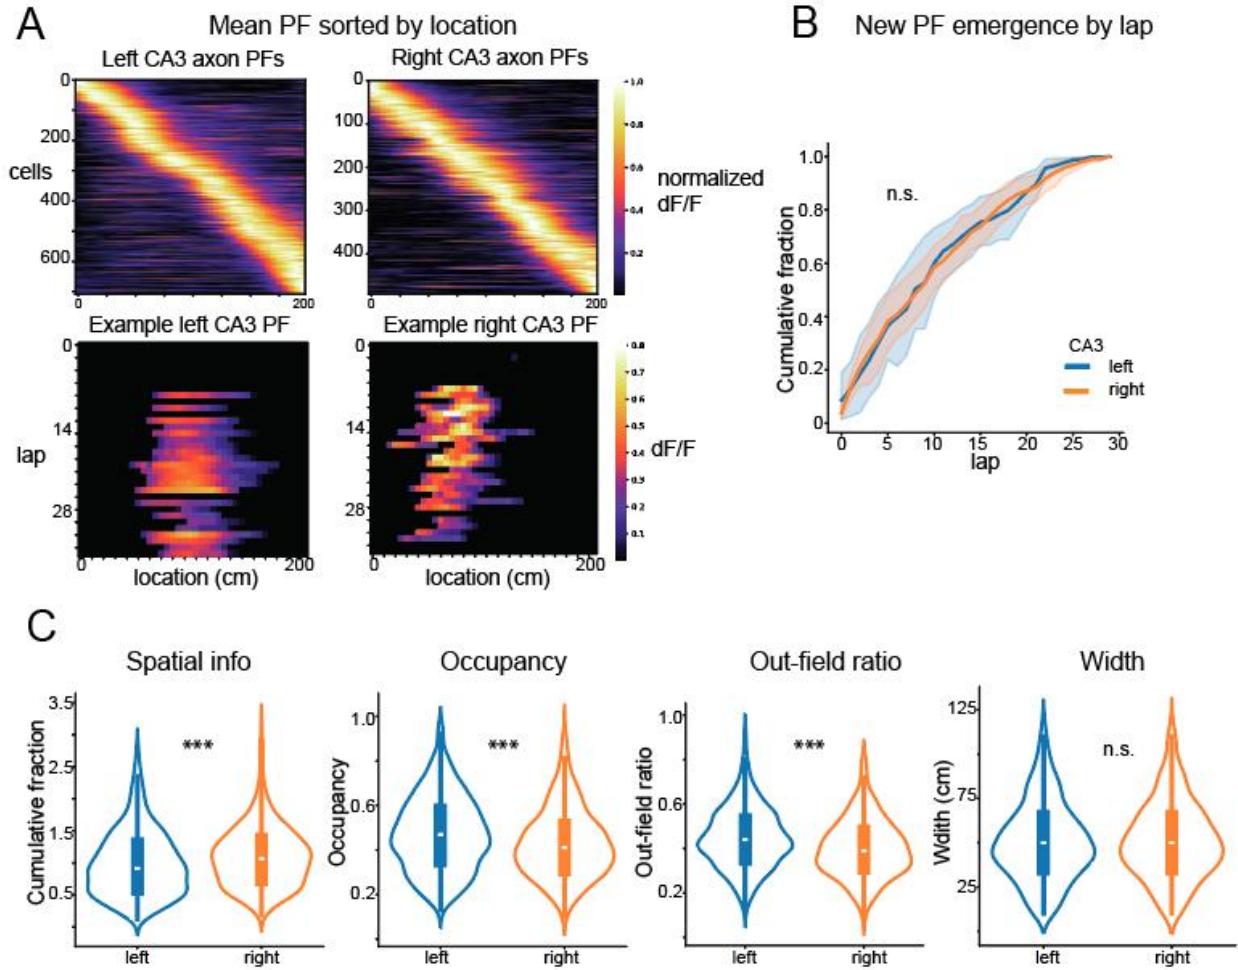


Supp3. Left and right CA3 input inhibition similarly delayed the emergence of CA1 place cells, regardless of emergence criteria.

A. Cumulative histogram of fractions of place fields emerged by lap number comparing opto early (red) with control (black) for first 10 laps under optogenetic inhibition by difference emergence criteria. Emergence lap was defined as the first lap in N out of 6 consecutive laps of within-field firing. P values reported were after Bonferroni correction (*10). N=2, left: Wilcoxon signed-rank stats=701, $p = 1.2$; right: stats= 361.5, $p=0.07$. N=3, left: stats=612, $p=0.26$; right: stats=262, $p<0.01$; N=4, left: stats=588, $p=0.16$; right: stats=209, $p<0.001$; N=5, left: stats=543, $p=0.062$; right: stats=190, $p<0.001$; N=6, left: stats=506, $p=0.11$; right: stats=371, $p=0.26$.

B. Within-animal emergence fraction difference between control and opto. Each dot represented one lap: below 0 meant the lap had less fraction of PF emerged under opto than that under lap-matched control, which reflected that optogenetic inhibition delayed the emergence of place fields in that lap. No significant difference found between left and right under any criteria of N. P values reported were after Bonferroni correction (*5). N=2, Mann whitney U stats=1653, $p=1.40$; N=3, Mann whitney U stats=1702, $p=1.23$; N=4, T stats = 1.29, $p=1$; N=5, T stats=1.74, $p=0.41$; N=6, T stats=0.24, $p=4.05$.

C. Fractions of place fields emerged 5 laps before turning on optogenetic in opto late by animal. No significant differences were found between control and opto in the left or right group. (linear mixed effects model, no significant effects, CA3 $p=0.99$, opto $p=0.46$, CA3 x opto $p=0.087$. Wilcoxon signed-rank test left: $p=0.375$, right: $p=0.625$.)



Supp4. Left and Right CA3 axon place field features.

A. Example axon place fields and mean place fields map sorted by center of mass location.

B. No significant difference found between left and right CA3 axon place fields in emergence in novel environments.

C. Left and right CA3 axon place fields feature comparison. Right CA3 place fields contained significantly more spatial information. Mann Whitney U stats=450693, $p < 0.001$. Left CA3 axon place fields had significantly longer occupancy than right CA3 axon place fields. Mann Whitney U stats=607665, $p < 0.001$. Left CA3 axon place fields had significantly more out-field firing ratio than right CA3 axon place fields. Mann Whitney U stats=633514, $p < 0.001$. Left and right CA3 axon place fields were not significantly different from each other. Mann Whitney U stats=514315, $p = 0.38$

Methods

Subjects. All experimental and surgical procedures were in accordance with the University of Chicago Animal Care and Use Committee guidelines. For this study, 11-15 week old C57BL/6-Tg(Grik4-cre)G32-4Stl/J mice were individually housed in a reverse 12 hour light/dark cycle in controlled temperature of ~20 Celsius and ~50% humidity animal facility. Both male and female mice were included in the experiments. All training and experiments were conducted during the animal's dark cycle. The entire experimental protocol, from the initial injection to the end of the experiments, would last 4-6 months.

Injection protocol. Mice were anesthetized and injected with 0.5ml of saline and 0.5ml of meloxicam. A small craniotomy was made over the left or right hippocampus CA3 region (2.0mm lateral, 1.7mm or -1.7mm caudal respectively of Bregma). An Cre-dependent opsin, pAAV-hSyn1-SIO-eOPN3-mScarlet-WPRE (Addgene #125713), was injected into CA3 (~100nl at a depth of 1.9mm below the dura). Right after the viral injection, the site was covered with dental cement (Metabond, Parkell Corporation) and a metal head plate (Atlas Tool and Die Works). After at least 2 weeks, another small craniotomy was made over the right CA1 region (1.7mm lateral, -2.3mm caudal of Bregma), where a genetically encoded calcium indicator, AAV1-CamKII-GCaMP6f (Addgene #100834) was injected using a beveled glass micropipette. Water restriction (0.8-1ml per day) began after a week of the initial craniotomy. One week after the start of water restriction, mice went through another surgery to implant a hippocampal cannula window at the right CA1 region. During the surgery, head plate and head ring were attached to the cannula window to house the microscope objective and block out ambient light.

Post-surgery, mice were back to ad lib with water, until a week later when they returned to the same water restriction schedule. Expression of eOPN3 in CA3 terminals and GCaMP6f in pyramidal cells at CA1 was checked every week starting 4 weeks after injection.

For axon imaging, the same injection protocol was followed except for a different virus. The genetically-encoded calcium indicator, AAV9-axon-GCaMP6s-P2A-mRuby3 (Addgene #112005) was injected into the left or right CA3 of Grik4-Cre mice (200 nL at a depth of 1.9 mm below the surface of the dura). Expression was checked every week starting 10 weeks after injection.

Behavioral training. Animals were trained to run on a treadmill in a virtual reality (VR) environment by five LED screens. All VR environments (one training environment, named as familiar environment, and three novel environments env A, B, C) were created using VIRMEn. Each environment had a two-meter linear track with unique detailed distal and proximal 3D visual cues. At the end of the track in all environments, 4 microliter water rewards were delivered and a short VR pause of 1.5s was implemented to allow for water consumption before “teleported” back to the beginning of the track. During training, mice were placed in the familiar environment for ~30 minutes each day to learn to run and lick the water rewards. Mice had to run at least 2 laps/min for 5 days in a row to proceed to imaging.

Two-photon imaging and optogenetic inhibition. Imaging was done using a laser scanning two-photon microscope (Neurolabware). The microscope consisted of an 8 KHz resonant

scanning module (Thorlabs), a 16×/0.8 NA/3 mm WD water immersion objective (MRP07220, Nikon). Images were collected at a frame rate of 30Hz with bidirection. GCaMP6f was excited at 920 nm with a femtosecond-pulsed two-photon laser (Insight DS + Dual, Spectra-Physics) and the fluorescence was collected using a GaAsP PMT (H11706, Hamamatsu). Average laser power measured at the objective was ~80mw each day and kept constant between days of imaging in the same animal. Sampling rate was 31Hz with bidirectional scanning. Opto-stimulation consisted of 625nm light delivered by a LED (Thorlabs) through the microscope objective, with pulse width of 30ms (during which 2P-imaging was stopped), pulse margin of 5.25ms (during which 2P-imaging was performed), frequency of 28.37Hz, and ~3 minutes of duration. A shutter stopped imaging during an opto-pulse but scanning continued at the same speed. The asynchrony between opto-stimulation and scanning resulted in missing lines for some frames. Each region of interest (ROI) had a slightly different effective sampling rate from the varying number of missing scan lines. The lowest effective sampling rate for a given ROI was ~10Hz during opto-stimulation. The PicoScope Oscilloscope (PICO4824, Pico Technology) collected the signal from the microscope to synchronize frame acquisition timing with behavior. Mice behavior including treadmill running speed, position, and licking was collected by Picoscope Oscilloscope and synchronized with the imaging for analysis.

For axon imaging, the same setup was used except for the following adjustments: A reduced sampling rate of 15Hz with unidirectional scanning. Average laser power measured at the objective was ~90mw each day. Axon-GCaMP6s was excited at 920nm and mRuby was excited at 1040nm. Emitted fluorescence was collected using two GaAsP PMTs.

Imaging sessions. Once mice met the behavioral criteria described above, they underwent test imaging sessions to assess the expression and functionality of eOPN3 and GCaMP6f. eOPN3 expression was confirmed in the basal and apical layers of CA1 using the red anatomical channel. GCaMP6f expression was verified based on the following criteria: absence of resting fluorescence in the nucleus, fast transient kinetics, and normal soma morphology.

During each test imaging session, mice ran for ~10 minutes in a dark environment (no visual cues and rewards) followed by ~30 minutes in the familiar environment. Optogenetic inhibition was delivered for 3 minutes once the mouse had been running for at least 3 minutes in either environment. The 6-day imaging protocol on the same field of view was initiated only after confirming a clear reduction in CA1 pyramidal cell activity during optogenetic inhibition in both the dark and familiar environments (Fig. 3d).

During the 6-day imaging protocol, each imaging day consisted of navigation in the familiar environment (10 min) before switching to a novel environment A, B or C (20 min). Each novel environment was paired with an optogenetic inhibition condition: opto-early, opto-late, or control condition (no opto delivered). Each pairing ran for 2 consecutive days. The order and association of novel environments and optogenetic conditions was counterbalanced between animals. On opto-early days, optogenetic inhibition was delivered as soon as the mouse entered the novel environment and lasted for ~3 minutes. On opto-late days, optogenetic inhibition was delivered for ~3 minutes after the mouse had completed at least 15 laps in the novel environment. No optogenetic inhibition was delivered in control condition. Imaging was recorded from the same field of view (FOV) throughout the protocol. After each imaging

session, multiple averaged FOVs were saved as references for alignment on subsequent days. Data from sessions with noticeable misalignment between days were discarded.

For axon imaging, mice followed the same behavioral criteria, and axon GCaMP and mRuby expression were confirmed before proceeding. Each day, mice ran in the familiar environment (~5 minutes) before transitioning to a novel environment (~10 minutes). To increase axon collection per animal, recordings were performed in two novel environments. Axon activity from the first day exposure in both novel environments was combined for analysis.

Image processing and ROI selection. Imaging planes acquired on different days were combined and preprocessed using Suite2p (Pachitariu et al. 2016). Time-series images went through two times of rigid and non-rigid transformations in Suite2p to remove movement artifacts. Motion corrected videos were visually assessed to ensure the absence of drifts in z-direction. Datasets with visible z-drifts were discarded. ROIs were identified by Suite2p and manually inspected for accuracy. For each ROI, baseline corrected $\Delta F/F$ traces were generated within each day, filtered for significant calcium transients as reported before. When generating baseline corrected $\Delta F/F$ traces for a given ROI, missing frames in optogenetic conditions were linearly interpolated using the *fillmissing* function in MATLAB. To match optogenetic conditions, two control datasets were generated: one downsampled at opto-early inhibition periods and the other at opto-late inhibition periods.

For axon imaging, ROIs were first defined using Suite2p and then preprocessed using default parameters in package *subprep* (Jiang, Zhao, and Sheffield 2024). A red mRuby channel was recorded simultaneously to axon GCaMP for motion correction and motion artifacts removal.

Transient Raster and Normalized Transient Frequency. To ensure behavioral consistency, neural data during stationary periods was excluded. Control data was downsampled during the same frames as opto-on periods in opto-early and opto-late conditions, ensuring a matched sampling rate during optogenetic inhibition. The data was then binned into ~160 ms time bins and z-score transformed (described as “z-score transformed neural data” in the study). Time bins with a z-score > 1 were classified as active and counted toward the number of active transients. Transient frequency for each pyramidal cell was calculated as the number of active transients divided by the condition duration. Normalized transient frequency was computed as the transient frequency during opto-on (or time-matched control) divided by the transient frequency of the entire session.

Low-Dimensional Manifold Visualization with CEBRA. CEBRA was applied on the z-score transformed neural data using environment labels and track location as behavioral variables, with each point in the low-dimensional manifold representing population activity within a single time bin.

To compute embedding centroids for each environment, the 2 m track was divided into 40 spatial bins (5 cm each). Within each bin, the centroid was defined as the average position of all points from the same lap in the latent space. Distance to centroid was calculated as the Euclidean

distance between each point and its corresponding centroid based on track location. To ensure equal representation of time bins across conditions, opto-off time bins were randomly downsampled to match the number of opto-on time bins. The average distance to centroid within each opto condition was computed after 20 shuffles.

Spatial Location Decoding with CEBRA. Similar to manifold visualization, z-score transformed neural data was used. In each shuffle, 80% of the opto-off period was randomly selected to generate the low-dimensional embedding, which was used to train a k-nearest neighbor (KNN) regressor. The model was then tested on 80% of the opto-on period projected embedding. Each animal underwent 20 shuffles, and hyperparameters for both the KNN regressor and latent embedding were optimized via grid search. Within-animal controls were time-matched to the opto-on period in both opto-early and opto-late conditions. Decoding accuracy at each time bin was measured using mean absolute error (MAE), calculated as the absolute difference (in cm) between predicted and true locations. To examine the temporal evolution of spatial decoding accuracy, mean MAE was averaged across all time bins within each lap. To assess the impact of optogenetic inhibition on decoder performance, the within-animal difference in mean MAE between opto and time-matched control conditions was computed for each shuffle (Fig. 4d).

Spatial Location Decoding with Bayesian Decoder. The 200 cm track was divided into 40 spatial bins (5 cm each). Place cell activity was binarized based on spatial bins and used to train and test a Gaussian Naive Bayes decoder. A cross-validation strategy was applied over 500 shuffles, with 80% of laps randomly selected for training and the remaining 20% for testing in

each shuffle. Decoding error was calculated as the absolute difference between predicted and true location bins (e.g., a 2-bin difference corresponds to a 10 cm error). Within each animal, mean decoder error was averaged across all locations within the same lap over all shuffles.

Licking Behavior. Licking data was recorded using a capacitive sensor on the waterspout. In familiar environments, well-trained mice exhibited increased licking in the pre-reward zone, the region just before reward delivery (Fig. 1). To identify this zone, we averaged lick counts by location across all animals in the familiar environment and applied the kneedle algorithm to detect the transition to elevated anticipatory licking. The pre-reward zone extended from this transition point to just before reward delivery and was consistently applied across all animals and environments. Similarly, the main track zone was determined using the kneedle algorithm to identify the shift from elevated consummatory licking to stable no licking in the familiar environment. Preemptive licking ratio is a lap-by-lap measurement of the ratio between lick count in the pre-reward zone and main track zone.

Kneedle Algorithm to Determine Novelty Modulation Time. Implemented and adapted from the Kneedle algorithm (Satopaa et al. 2011), this procedure was applied to each time series (e.g. Bayesian decoder error by lap, transient raster count over time etc). Following preprocessing (smoothing as described and normalized), a linear trend was fitted between the first and last time points of the signal. For each intermediate time point, the perpendicular distance to this line was computed. The time point with the maximum deviation—i.e., the “elbow” point—was defined as the novelty modulation time, marking the transition from a phase of rapid decline to a stable period.

Place Field Selection. We used the peak method (Grijseels et al. 2021) with additional criteria to account for noise in the dataset. Potential place fields (PFs) were identified as contiguous regions on the mean tuning curve with a minimum amplitude of $0.1 \Delta F/F$ and exceeding 40% of the difference between peak $\Delta F/F$ and baseline values. At least 30% of laps had to contain a significant transient within the PF boundaries, and PF width was constrained to 6%–50% of the track length. Each PF had to be active within-field in at least 3 of 6 consecutive laps (see emergence lap below), but only one instance of this criterion being met was required, defined as the emergence lap. After passing all those criteria, neural activity was randomly shuffled 600 times by shifting fluorescence transients in time while preserving their calcium dynamics. For each shuffle, peak and bottom fluorescence values were averaged across the fluorescence map. Cells with true peak fluorescence in the top 1% and bottom fluorescence in the bottom 1% were classified as place cells.

Emergence lap. For each cell, the first lap in n ($n=3$ if not otherwise specified) out of 6 consecutive laps of within-field firing is defined as an emergence lap.

Reliability. For each place cell, reliability is defined as the number of laps with significant in-field transient over a 6-lap sliding window.

Transient Amplitude. Neural activity traces were smoothed using a centered rolling average with a ~ 320 ms time window (minimum ~ 100 ms per window) to reduce high-frequency noise.

Peaks were then detected using *scipy.find_peaks* with the following criteria: Minimum prominence of $0.1\Delta F/F$, minimum peak height of $0.12\Delta F/F$, minimum interpeak interval of ~ 320 ms, peak width of ~ 160 ms to ~ 3.88 s. Transient amplitude was defined as the area under the curve from transient onset to peak.

Normalized axon activity. Calcium transient amplitudes were calculated as above but with lower prominence ($0.05 \Delta F/F$) and minimum peak height ($0.1 \Delta F/F$). Normalized axon activity was defined as the sum of all transient amplitudes divided by the number of laps (10 laps in early, 15 laps in late), capturing both transient frequency and amplitude.

Statistics and Reproducibility. Error bars in all figures represent 95% confidence intervals of a bootstrap-generated difference (5,000 resamples). Data distributions were assessed for normality using the Shapiro–Wilk test. If normally distributed, a paired or unpaired Student’s t-test was used where applicable. For non-normal distributions, a paired Wilcoxon signed-rank test or an unpaired Mann–Whitney U-test was applied. For samples with six or fewer data points, only non-parametric tests were used. Multiple comparisons were corrected using Bonferroni post hoc adjustments. $p < 0.05$ was chosen to indicate statistical significance. P values presented in figures are as follows: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$, n.s. not significant.

To compare delta-delta differences (Fig. 4d, 5b inset, 5c inset, 5h inset, 5i inset, 5k, 5l) and assess whether left or right inhibition had a greater effect, a within-animal subsampled bootstrap difference distribution was generated. For each shuffle, an equal number of subsamples were randomly selected from the control and opto conditions within the same animal, and their mean

difference was computed. This process was repeated 500 times per animal to construct a bootstrap distribution of control–opto differences. Within-animal differences from the same inhibition group were pooled for left vs. right comparisons, and statistical significance was tested accordingly.

A linear mixed effects model (Fig. 5E bottom, 5F bottom, Supp. Fig. 3c) was implemented in MATLAB using *fitlme* with the formula: $\text{EmergePercent} \sim \text{CA3} * \text{opto} + (1|\text{mouse})$. This model assesses the effects of optogenetic manipulation, CA3 side, and their interaction while accounting for variability between animals.

Acknowledgements. We thank Leslie Kay, PhD and Seetha Krishnan, PhD for feedback on the manuscript and all members of the Sheffield lab for valuable discussions that helped shape the manuscript. We thank Roma Shah for assistance in training animals.

Funding. This work was supported by The Whitehall Foundation, The Searle Scholars Program, The Sloan Foundation, The University of Chicago Institute for Neuroscience start-up funds and the National Institute for Health (1DP2NS111657-01, 1RF1NS127123-01, 1R21NS128822-01, 1R01MH136274-01) awarded to M.S.

Author contributions. A.J. and M.S. conceived and designed the experiments. A.J. collected the data. A.J. analyzed the data. A.J. and M.S. interpreted the data and wrote the manuscript. M.S. supervised the research and obtained funding.

Declaration of interests. The authors declare no competing interests.

Data and code availability. The original codes have been deposited at https://github.com/anqijiang/opto_analysis and will be publicly available on the date of publication. Raw imaging data is large and not feasible for upload to an online repository but is available upon request to the lead contact at sheffield@uchicago.edu. Processed source data for all figures and associated statistical analysis will be uploaded to Box and provided in the final version of the paper. Any additional information required to reanalyze the data reported in this paper is available from the lead contact upon request.

CHAPTER 4

CONCLUSION AND FUTURE DIRECTIONS

4.1 Summary of Findings

Chapter 2: A preprocessing toolbox for 2-photon subcellular calcium imaging

1. Challenges of preprocessing for 2-photon subcellular calcium imaging include ROI detection, movement artifact identification, and ROI grouping for axon data.
2. We developed and validated a computationally efficient preprocessing pipeline for subcellular calcium imaging preprocessing. Through the application of frequency-based filtering, Principal Component Analysis (PCA) and level change detection for motion artifact identification, and clustering algorithms for ROI grouping, our pipeline provides a robust framework for extracting physiological signals from subcellular structures during behavior.

Chapter 3: A learning-dependent shift in CA3 input dominance from left to right CA3 determines the evolution of CA1 spatial maps

1. CA1R spatial maps develop with experience in novel environments and then stabilize. There is an initial phase of gradual improvement in spatial decoding (first 10 laps: early phase) that plateaus and is stable after ~10 laps (late phase). In parallel, anticipatory licking for a reward within the novel environment develops and becomes precise over

~10 laps, reflecting the animal's learning about the environment in the early phase (familiarization).

2. Both left and right CA3 projections are essential for forming CA1R place fields during the early phase. Once formed, right CA3 inputs predominantly drive CA1R place-field reliability, amplitude, and population spatial decoding accuracy during the early phase. During the late phase, left CA3 input dominates, supporting place field reliability, precise population spatial decoding, and new place field emergence. Right CA3 axons are more active during the early phase, shifting to left CA3 dominance as the environment becomes more familiar at the late phase.
3. Right CA3 supports the development of accurate spatial representations during learning, while left CA3 engages more when spatial representations have stabilized - revealing a greater role in memory retrieval.

4.2 Discussion and Future Directions

Updated hypothesis on the specialized functions of hippocampal lateralization

As discussed in the introduction, the "lateralized plasticity hypothesis" (Jordan 2020), proposed by the Paulsen lab, suggests that the left CA3 specializes in acquiring and storing new memories through its enhanced synaptic plasticity, while the right CA3 supports rapid spatial representation of novel environments via stable, developmentally preconfigured transmission (El-Gaby, Shipton, and Paulsen 2015). While my data align with a role for right CA3 in supporting rapid spatial representations during novel experiences, I propose a revised model in

which both left and right CA3 contribute to CA1 spatial representation in distinct, experience-dependent ways during familiarization to a novel environment.

Specifically, both hemispheric inputs are necessary for the rapid emergence of CA1 place fields. Once place fields have formed, right CA3 inputs facilitate the rapid development and refinement of spatial representations during early phases of novel exposure. In contrast, left CA3 inputs contribute to the stabilization and updating of spatial maps by maintaining field reliability and gain, and by integrating newly emerging fields as the environment is familiarized. This proposed dynamic division of labor is supported not only by axonal imaging data showing a shift in activity from right to left CA3 during familiarization, but also by prior literature demonstrating hemispheric asymmetries in synchronized network events (Kitanishi et al. 2015; Benito et al. 2016; Guan et al. 2021) and in synaptic plasticity mechanisms (Shipton et al. 2014; Kohl et al. 2011; S. Li et al. 2003; Madar et al. 2025; Priestley et al. 2022; Kitanishi et al. 2015).

In contrast to the lateralized plasticity hypothesis—which posits that right CA3's role is largely restricted to generating a fast but inflexible representation during immediate novelty and becomes dispensable with continued learning—my findings extend the functional relevance of right CA3 beyond just the initial exposure. I show that right CA3 contributes to rapid plasticity-based learning during the first ~10 laps (~2 minutes) of novel environment exploration. While plasticity is not required for the initial formation of all place fields (Sheffield, Adoff, and Dombeck 2017), it is crucial for improving spatial coding within already formed fields. The original hypothesis is based on slice physiology studies

showing that only left CA3 inputs could induce LTP at CA1 using HFS and STDP protocols (Shipton et al. 2014; Kohl et al. 2011). However, the inability to induce LTP under specific protocols does not imply a complete absence of plastic potential. Early novel exposure likely presents a unique window with heightened susceptibility to alternative plasticity rules, such as behavioral time-scale synaptic plasticity (BTSP) (Madar et al. 2025; Priestley et al. 2022) and neuromodulatory facilitation (Heer and Sheffield 2024; S. Li et al. 2003), which may enable LTP induction even from right CA3 inputs.

Additional support for a specialized role of right CA3 during early learning comes from studies on hippocampal oscillations. Gamma oscillations, which occur during awake behavior, are thought to promote interregional communication through precise temporal coordination (Colgin and Moser 2010). Phase-locked firing of CA1 pyramidal cells to gamma oscillations likely enhances the convergence and impact of their outputs to downstream regions. Notably, during initial novel exposure, CA1 neurons exhibit transiently increased phase-locked firing to slow gamma oscillations (Kitanishi et al. 2015). This enhancement depends on GluA1, an AMPAR subunit that is more densely expressed on CA1 dendritic spines receiving right CA3 inputs (Shinohara et al. 2008), and has been associated with the rapid improvement of spatial coding in CA1 place fields during early learning of novel environments (Kitanishi et al. 2015).

Consistent with these findings, my data demonstrate a more prominent role for right CA3 inputs in enhancing place field gain and reliability during the early phase of learning. The specialized involvement of GluA1 and enhanced gamma

phase-locking likely represents a key mechanism by which right CA3 facilitates the rapid refinement of spatial representations. Thus, even if right CA3 inputs are less effective at inducing NMDAR-dependent plasticity (Shipton et al. 2014; Kohl et al. 2011), they may nevertheless play a dominant role in supporting AMPAR-dependent plasticity and temporally coordinated coding during initial novel experience.

Further evidence for the specialized influence of right CA3 comes from studies in anesthetized rats, where right CA3 inputs were shown to induce stronger gamma-paced firing and exert greater excitatory drive on CA1 neurons compared to left CA3 (Benito et al. 2016). Complementing this, slow gamma coherence between CA3 and CA1 transiently increases during exploration of novel objects (Trimper et al. 2017), suggesting that enhanced synchrony between these regions may play a critical role in coordinating spike timing and engaging synaptic plasticity mechanisms to support memory encoding. Moreover, CA3 inputs has been implicated in mediating interhemispheric coordination of ensemble and ripple activity in CA1, without which the cell assembly activity, SWRs, and replays in CA1 become poorer in quality and more lateralized within the hemispheres (Guan et al. 2021). Together, these findings support the idea that right CA3 inputs are more likely to synchronize more effectively with CA1 networks and exert a larger influence on population-level activity during the early stages of novel experience.

Taken together, these findings support a revised model in which right CA3 enables rapid, experience-dependent plasticity and synchronization during initial

novel exposure. After the environment is familiarized, my data suggest a more prominent role of left CA3 in maintaining spatial representation and supporting the emergence of new fields over continued learning (see below for detailed discussion on the topic of continuous learning). Although in my experimental setup, learning occurs rapidly during early novel exposure—evidenced by the establishment of anticipatory licking behavior and stabilization of the spatial population code—the later “post-learning” phase likely retains capacity for ongoing plasticity and learning, consistent with prior behavioral studies (Shipton et al. 2014; Kohl et al. 2011; Klur et al. 2009; El-Gaby et al. 2016).

I propose that the distinction between early and late phases of novel exposure is modulated by the experience of novelty itself—during which specialized plasticity rules and network events occur (Kitanishi et al. 2015; S. Li et al. 2003; Heer and Sheffield 2024; Madar et al. 2025; Priestley et al. 2022), — rather than by the specific timepoint at which the task is learned, which depends on the task difficulty and cognitive demands. When a hippocampal-dependent learning task is embedded within a novel environment, this model predicts that learning-related neural dynamics—such as refinement of spatial representations, as observed in my data—will also be temporally modulated by early versus late phases of novelty exposure. For tasks requiring extended training beyond the period of initial novelty, the model further predicts that right CA3 would still play a dominant role early on, rapidly enhancing spatial representations through novelty-specific plasticity, but would be insufficient for full task acquisition. After the initial novelty phase, left CA3 would gradually assume a leading role,

supporting continued learning through slower, more conventional forms of synaptic plasticity.

This hypothesis is based on the assumption that novelty modulation and the process of learning are independently regulated by distinct pathways and mechanisms—a concept that itself warrants further investigation. For instance, if an animal is exploring a maze with particularly complex or hard-to-learn spatial configurations, would the novelty modulations persist longer to facilitate better, more effective learning? It is unknown if novelty and learning may operate on partially dissociable timelines with different underlying neural mechanisms.

The idea that continuous learning and memory updating is occurring even “post-learning” is consistent with recent work on representational drift, which highlights continuous turnover in place cell identity and shifting spatial activity patterns over time (Geva et al. 2023; Keinath, Mosser, and Brandon 2022; Khatib et al. 2023; Delamare et al. 2023; Krishnan and Sheffield 2023). While left CA3 inputs, but not right CA3 inputs, indeed support the emergence of new place fields over continued learning, indicating its role in ongoing learning beyond mere maintenance of already learned spatial representations, it is unknown if left CA3 actively drives these experience-dependent representational drift dynamics, either stabilizing or facilitating drift in the CA1 code. This possibility warrants further investigation in future studies.

Lateralization of CA3 inputs during phases of learning

As discussed in Chapter 3, while both left and right CA3 inputs are essential for rapid formation of CA1 place fields, once formed, right CA3 inputs dominate the early familiarization phase (within the first ~10 laps) of CA1 place field refinement whereas left CA3 inputs drive the maintenance of CA1 place fields when the environment becomes familiar. Although we measured the anticipatory licking prior to rewards as a behavioral readout of learning, a more complicated behavioral design that directly depends on successful spatial learning would be needed for further establishing a causal role between the contribution of lateralized CA3 inputs and behavior during the process of familiarization.

Prior research has consistently reported a left-hemisphere dominance in hippocampus-dependent spatial learning tasks (Shipton et al. 2014; El-Gaby et al. 2016, 2021; Klur et al. 2009). However, these studies employed pharmacological or optogenetic inhibition protocols that involved continuous inactivation of left or right CA3 throughout the entire training period, which spanned multiple days. According to my proposed model, right CA3 contributes most strongly during the initial novel exposure phase, while left CA3 gradually assumes a dominant role during the later phase of learning. The inhibition protocols used in these prior studies would therefore have encompassed both early and late phases of learning, as defined in this thesis.

Given the extended training durations and the complexity of the tasks in these studies, it is likely that most of the learning occurred after the novelty-driven modulation had subsided. Under my revised framework, behavioral impairments from early right CA3 inhibition might be compensated over time by the

progressive engagement of left CA3 inputs. In contrast, inhibition of left CA3 during the later phase—when it plays a primary role in supporting continued learning—would result in more persistent deficits that cannot be compensated for by right CA3 inputs. Therefore, the greater behavioral impairment observed with left CA3 inhibition in these studies may reflect the dominance of left CA3 during the prolonged learning phase, rather than an exclusive role in all phases of learning.

This interpretation suggests that the apparent discrepancy between my findings—emphasizing a critical early role for right CA3—and prior studies—highlighting left CA3 necessity—can be reconciled by differences in task complexity and training duration. I hypothesize that in tasks requiring extended training, selective inhibition of left CA3 during the later phase would lead to more pronounced behavioral impairments than inhibition of right CA3, consistent with the dynamics hemispheric specialization proposed in this thesis.

In contrast, I hypothesize that for tasks that require rapid representation, right CA3 inputs inhibition would reduce behavioral performance more than left CA3 inputs inhibition. Prior research found CA3 essential for rapid one-trial learning in novel environments (Wagatsuma et al. 2018; Nakazawa et al. 2002; Nakashiba et al. 2008). Bilateral inhibition of CA3 significantly reduced the amount of freezing in a one-trial contextual fear conditioning task (Nakashiba et al. 2008). Given the dominance of right CA3 inputs during the initial novelty exposure phase that I report in this thesis, I expect that in a similar contextual fear

conditioning task that right CA3 inputs inhibition during the early phase would lead to worse behavioral performance than left CA3 input inhibition.

Broader impact and future directions

This work reveals a dynamic hemispheric asymmetry in the hippocampus throughout memory formation and retrieval. Rapid encoding is critical for survival, the ability of animals to form immediate spatial memories in unfamiliar environments has significant evolutionary advantage, highlighting the importance of right CA3's early dominance. By providing a new framework for memory encoding as a novelty modulated, lateralized process, this model opens new directions for understanding how learning unfolds across time and brain regions. A key question for future studies is what controls the ability to dynamically shift the influence of left and right CA3. How does the brain perform such a shift in a learning-dependent manner? Neuromodulatory systems like the locus coeruleus may play a crucial role in coordinating this transition by linking novelty, arousal, and plasticity. Uncovering these mechanisms will offer deeper insight into memory flexibility and its disruption in disease.

Bibliography

- Adaikkan, Chinnakkaruppan, Justin Joseph, Georgios Foustoukos, Jun Wang, Denis Polygalov, Roman Boehringer, Steven J. Middleton, Arthur J. Y. Huang, Li-Huei Tsai, and Thomas J. McHugh. 2024. “Silencing CA1 Pyramidal Cells Output Reveals the Role of Feedback Inhibition in Hippocampal Oscillations.” *Nature Communications* 15 (1): 2190.
- Aery Jones, Emily A., and Lisa M. Giocomo. 2023. “Neural Ensembles in Navigation: From Single Cells to Population Codes.” *Current Opinion in Neurobiology* 78 (102665): 102665.
- Allen, Kevin, and Hannah Monyer. 2015. “Interneuron Control of Hippocampal Oscillations.” *Current Opinion in Neurobiology* 31 (April): 81–87.
- Allen, Timothy A., and Norbert J. Fortin. 2013. “The Evolution of Episodic Memory.” *Proceedings of the National Academy of Sciences of the United States of America* 110 Suppl 2 (supplement_2): 10379–86.
- Andersen, Per, Richard Morris, David Amaral, Tim Bliss, and John O’Keefe, eds. 2006. *The Hippocampus Book*. Oxford Neuroscience Series. New York, NY: Oxford University Press.
- Aronov, Dmitriy, Rhino Nevers, and David W. Tank. 2017. “Mapping of a Non-Spatial Dimension by the Hippocampal-Entorhinal Circuit.” *Nature* 543 (7647): 719–22.
- Ashby, Donovan M., Stan B. Floresco, Anthony G. Phillips, Alexander McGirr, Jeremy K. Seamans, and Yu Tian Wang. 2021. “LTD Is Involved in the Formation and Maintenance of Rat Hippocampal CA1 Place-Cell Fields.” *Nature Communications* 12 (1): 100.
- Benito, Nuria, Gonzalo Martín-Vázquez, Julia Makarova, Valeri A. Makarov, and Oscar Herreras. 2016. “The Right Hippocampus Leads the Bilateral Integration of Gamma-Parsed Lateralized Information.” *ELife* 5 (September). <https://doi.org/10.7554/eLife.16658>.
- Berners-Lee, Alice, Ting Feng, Delia Silva, Xiaojing Wu, Ellen R. Ambrose, Brad E. Pfeiffer, and David J. Foster. 2022. “Hippocampal Replays Appear after a Single Experience and Incorporate Greater Detail with More Experience.” *Neuron* 110 (11): 1829-1842.e5.
- Best, P. J., A. M. White, and A. Minai. 2001. “Spatial Processing in the Brain: The Activity of Hippocampal Place Cells.” *Annual Review of Neuroscience* 24 (1): 459–86.
- Bevins, Rick A., Jana Koznarova, and Tracy J. Armiger. 2001. “Environmental Familiarization in Rats: Differential Effects of Acute and Chronic Nicotine.” *Neurobiology of Learning and Memory* 75 (1): 63–76.
- Bilash, Olesia M., Spyridon Chavlis, Cara D. Johnson, Panayiota Poirazi, and Jayeeta Basu. 2023. “Lateral Entorhinal Cortex Inputs Modulate Hippocampal Dendritic Excitability by Recruiting a Local Disinhibitory Microcircuit.” *Cell Reports* 42 (1): 111962.
- Bittner, Katie C., Christine Grienberger, Sachin P. Vaidya, Aaron D. Milstein, John J. Macklin, Junghyup Suh, Susumu Tonegawa, and Jeffrey C. Magee. 2015. “Conjunctive Input Processing Drives Feature Selectivity in Hippocampal CA1 Neurons.” *Nature Neuroscience* 18 (8): 1133–42.

- Bittner, Katie C., Aaron D. Milstein, Christine Grienberger, Sandro Romani, and Jeffrey C. Magee. 2017. "Behavioral Time Scale Synaptic Plasticity Underlies CA1 Place Fields." *Science (New York, N.Y.)* 357 (6355): 1033–36.
- Butola, Tanvi, Melissa Hernández-Frausto, Stefan Blankvoort, Marcus Sandbukt Flatset, Lulu Peng, Ariel Hairston, Cara Deanna Johnson, et al. 2025. "Hippocampus Shapes Entorhinal Cortical Output through a Direct Feedback Circuit." *Nature Neuroscience*, February. <https://doi.org/10.1038/s41593-025-01883-9>.
- Cayco-Gajic, N. Alex, and R. Angus Silver. 2019. "Re-Evaluating Circuit Mechanisms Underlying Pattern Separation." *Neuron* 101 (4): 584–602.
- Chiu, Yuhung, Can Dong, Seetha Krishnan, and Mark E. J. Sheffield. 2023. "The Precision of Place Fields Governs Their Fate across Epochs of Experience." *ENeuro* 10 (12): ENEURO.0261-23.2023.
- Cohen, Jeremy D., Mark Bolstad, and Albert K. Lee. 2017. "Experience-Dependent Shaping of Hippocampal CA1 Intracellular Activity in Novel and Familiar Environments." *ELife* 6 (July): e23040.
- Colgin, Laura Lee, and Edvard I. Moser. 2010. "Gamma Oscillations in the Hippocampus." *Physiology (Bethesda, Md.)* 25 (5): 319–29.
- Cowan, Nelson. 2008. "What Are the Differences between Long-Term, Short-Term, and Working Memory?" *Progress in Brain Research* 169: 323–38.
- Davoudi, Heydar, and David J. Foster. 2019. "Acute Silencing of Hippocampal CA3 Reveals a Dominant Role in Place Field Responses." *Nature Neuroscience* 22 (3): 337–42.
- Delamare, Geoffroy, Yosif Zaki, Denise J. Cai, and Claudia Clopath. 2023. "Drift of Neural Ensembles Driven by Slow Fluctuations of Intrinsic Excitability." <https://doi.org/10.7554/elife.88053.2>.
- Deng, Wei, James B. Aimone, and Fred H. Gage. 2010. "New Neurons and New Memories: How Does Adult Hippocampal Neurogenesis Affect Learning and Memory?" *Nature Reviews. Neuroscience* 11 (5): 339–50.
- Dong, Can, Antoine D. Madar, and Mark E. J. Sheffield. 2021. "Distinct Place Cell Dynamics in CA1 and CA3 Encode Experience in New Environments." *Nature Communications* 12 (1). <https://doi.org/10.1038/s41467-021-23260-3>.
- Duszkiewicz, Adrian J., Colin G. McNamara, Tomonori Takeuchi, and Lisa Genzel. 2019. "Novelty and Dopaminergic Modulation of Memory Persistence: A Tale of Two Systems." *Trends in Neurosciences* 42 (2): 102–14.
- Eichenbaum, Howard, Paul Dudchenko, Emma Wood, Matthew Shapiro, and Heikki Tanila. 1999. "The Hippocampus, Memory, and Place Cells." *Neuron* 23 (2): 209–26.
- El-Gaby, Mohamady, Hayley M. Reeve, Vítor Lopes-dos-Santos, Natalia Campo-Urriza, Pavel V. Perestenko, Alexander Morley, Lauren A. M. Strickland, István P. Lukács, Ole Paulsen, and David Dupret. 2021. "An Emergent Neural Coactivity Code for Dynamic Memory." *Nature Neuroscience* 24 (5): 694–704.
- El-Gaby, Mohamady, Olivia A. Shipton, and Ole Paulsen. 2015. "Synaptic Plasticity and Memory: New Insights from Hippocampal Left-Right Asymmetries: New Insights from Hippocampal Left-Right Asymmetries." *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry* 21 (5): 490–502.
- El-Gaby, Mohamady, Yu Zhang, Konstantin Wolf, Christof J. Schwiening, Ole Paulsen, and Olivia A. Shipton. 2016. "Archaerhodopsin Selectively and Reversibly Silences Synaptic Transmission through Altered PH." *Cell Reports* 16 (8): 2259–68.

- Fenton, André A. 2024. “Remapping Revisited: How the Hippocampus Represents Different Spaces.” *Nature Reviews. Neuroscience* 25 (6): 428–48.
- Fjell, Anders M., Linda McEvoy, Dominic Holland, Anders M. Dale, and Kristine B. Walhovd. 2014. “What Is Normal in Normal Aging? Effects of Aging, Amyloid and Alzheimer’s Disease on the Cerebral Cortex and the Hippocampus.” *Progress in Neurobiology* 117 (June): 20–40.
- Flavell, Steven W., Nadine Gogolla, Matthew Lovett-Barron, and Moriel Zelikowsky. 2022. “The Emergence and Influence of Internal States.” *Neuron* 110 (16): 2545–70.
- Gava, Giuseppe P., Stephen B. McHugh, Laura Lefèvre, Vítor Lopes-Dos-Santos, Stéphanie Trouche, Mohamady El-Gaby, Simon R. Schultz, and David Dupret. 2021. “Integrating New Memories into the Hippocampal Network Activity Space.” *Nature Neuroscience* 24 (3): 326–30.
- Geiller, Tristan, James B. Priestley, and Attila Losonczy. 2023. “A Local Circuit-Basis for Spatial Navigation and Memory Processes in Hippocampal Area CA1.” *Current Opinion in Neurobiology* 79 (102701): 102701.
- Geiller, Tristan, Sadra Sadeh, Sebastian V. Rolotti, Heike Blockus, Bert Vancura, Adrian Negrean, Andrew J. Murray, et al. 2022. “Local Circuit Amplification of Spatial Selectivity in the Hippocampus.” *Nature* 601 (7891): 105–9.
- Geva, Nitzan, Daniel Deitch, Alon Rubin, and Yaniv Ziv. 2023. “Time and Experience Differentially Affect Distinct Aspects of Hippocampal Representational Drift.” *Neuron* 111 (15): 2357-2366.e5.
- Gómez-Ocádiz, Ruy, Massimiliano Trippa, Chun-Lei Zhang, Lorenzo Posani, Simona Cocco, Rémi Monasson, and Christoph Schmidt-Hieber. 2022. “A Synaptic Signal for Novelty Processing in the Hippocampus.” *Nature Communications* 13 (1): 4122.
- GoodSmith, Douglas, Xiaojing Chen, Cheng Wang, Sang Hoon Kim, Hongjun Song, Andrea Burgalossi, Kimberly M. Christian, and James J. Knierim. 2017. “Spatial Representations of Granule Cells and Mossy Cells of the Dentate Gyrus.” *Neuron* 93 (3): 677-690.e5.
- GoodSmith, Douglas, Heekyung Lee, Joshua P. Neunuebel, Hongjun Song, and J. Knierim. 2019. “Dentate Gyrus Mossy Cells Share a Role in Pattern Separation with Dentate Granule Cells and Proximal CA3 Pyramidal Cells.” *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 39 (48): 9570–84.
- Grijseels, Dori M., Kira Shaw, Caswell Barry, and Catherine N. Hall. 2021. “Choice of Method of Place Cell Classification Determines the Population of Cells Identified.” *PLoS Computational Biology* 17 (7): e1008835.
- Guan, Hefei, Steven J. Middleton, Takafumi Inoue, and Thomas J. McHugh. 2021. “Lateralization of CA1 Assemblies in the Absence of CA3 Input.” *Nature Communications* 12 (1). <https://doi.org/10.1038/s41467-021-26389-3>.
- Guo, Wei, Jie J. Zhang, Jonathan P. Newman, and Matthew A. Wilson. 2024. “Latent Learning Drives Sleep-Dependent Plasticity in Distinct CA1 Subpopulations.” *Cell Reports* 43 (12): 115028.
- Hall, Arron F., and Dong V. Wang. 2023. “The Two Tales of Hippocampal Sharp-Wave Ripple Content: The Rigid and the Plastic.” *Progress in Neurobiology* 221 (102396): 102396.
- Heer, Chad, and Mark Sheffield. 2024. “Distinct Catecholaminergic Pathways Projecting to Hippocampal CA1 Transmit Contrasting Signals during Navigation in Familiar and Novel Environments.” *ELife* 13 (November). <https://doi.org/10.7554/eLife.95213>.

- Jeffery, Kathryn J. 2007. "Integration of the Sensory Inputs to Place Cells: What, Where, Why, and How?" *Hippocampus* 17 (9): 775–85.
- Jeong, Nuri, and Annabelle C. Singer. 2022. "Learning from Inhibition: Functional Roles of Hippocampal CA1 Inhibition in Spatial Learning and Memory." *Current Opinion in Neurobiology* 76 (102604): 102604.
- Ji, Daoyun, and Matthew A. Wilson. 2008. "Firing Rate Dynamics in the Hippocampus Induced by Trajectory Learning." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 28 (18): 4679–89.
- Jiang, Anqi, Chong Zhao, and Mark Sheffield. 2024. "A Preprocessing Toolbox for 2-Photon Subcellular Calcium Imaging." *BioRxiv*. <https://doi.org/10.1101/2024.10.04.616737>.
- Jordan, Jake T. 2020. "The Rodent Hippocampus as a Bilateral Structure: A Review of Hemispheric Lateralization." *Hippocampus* 30 (3): 278–92.
- Kawakami, Ryosuke, Yoshiaki Shinohara, Yuichiro Kato, Hiroyuki Sugiyama, Ryuichi Shigemoto, and Isao Ito. 2003. "Asymmetrical Allocation of NMDA Receptor E2 Subunits in Hippocampal Circuitry." *Science (New York, N.Y.)* 300 (5621): 990–94.
- Keinath, Alexandra T., Coralie-Anne Mosser, and Mark P. Brandon. 2022. "The Representation of Context in Mouse Hippocampus Is Preserved despite Neural Drift." *Nature Communications* 13 (1). <https://doi.org/10.1038/s41467-022-30198-7>.
- Kesner, R. P. 2007. "Behavioral Functions of the CA3 Subregion of the Hippocampus." *Learning & Memory (Cold Spring Harbor, N.Y.)* 14 (11): 771–81.
- Kesner, Raymond P., and Edmund T. Rolls. 2015. "A Computational Theory of Hippocampal Function, and Tests of the Theory: New Developments." *Neuroscience and Biobehavioral Reviews* 48 (January): 92–147.
- Khatib, Dorgham, Aviv Ratzon, Mariell Sellevoll, Omri Barak, Genela Morris, and Dori Derdikman. 2023. "Active Experience, Not Time, Determines within-Day Representational Drift in Dorsal CA1." *Neuron* 111 (15): 2348-2356.e4.
- Kim, Sang Hoon, Douglas GoodSmith, Stephanie J. Temme, Fumika Moriya, Guo-Li Ming, Kimberly M. Christian, Hongjun Song, and James J. Knierim. 2023. "Global Remapping in Granule Cells and Mossy Cells of the Mouse Dentate Gyrus." *Cell Reports* 42 (4): 112334.
- Kitamura, Takashi, Sachie K. Ogawa, Dheeraj S. Roy, Teruhiro Okuyama, Mark D. Morrissey, Lillian M. Smith, Roger L. Redondo, and Susumu Tonegawa. 2017. "Engrams and Circuits Crucial for Systems Consolidation of a Memory." *Science (New York, N.Y.)* 356 (6333): 73–78.
- Kitanishi, Takuma, Sakiko Ujita, Mehdi Fallahnezhad, Naomi Kitanishi, Yuji Ikegaya, and Ayumu Tashiro. 2015. "Novelty-Induced Phase-Locked Firing to Slow Gamma Oscillations in the Hippocampus: Requirement of Synaptic Plasticity." *Neuron* 86 (5): 1265–76.
- Klur, Sandra, Christophe Muller, Anne Pereira de Vasconcelos, Theresa Ballard, Joëlle Lopez, Rodrigue Galani, Ulrich Certa, and Jean-Christophe Cassel. 2009. "Hippocampal-Dependent Spatial Memory Functions Might Be Lateralized in Rats: An Approach Combining Gene Expression Profiling and Reversible Inactivation." *Hippocampus* 19 (9): 800–816.
- Knierim, James J. 2015. "The Hippocampus." *Current Biology: CB* 25 (23): R1116-21.
- Kohl, Michael M., Olivia A. Shipton, Robert M. Deacon, J. Nicholas P. Rawlins, Karl Deisseroth, and Ole Paulsen. 2011. "Hemisphere-Specific Optogenetic Stimulation

- Reveals Left-Right Asymmetry of Hippocampal Plasticity.” *Nature Neuroscience* 14 (11): 1413–15.
- Krishnan, Seetha, Chad Heer, Chery Cherian, and Mark E. J. Sheffield. 2022. “Reward Expectation Extinction Restructures and Degrades CA1 Spatial Maps through Loss of a Dopaminergic Reward Proximity Signal.” *Nature Communications* 13 (1). <https://doi.org/10.1038/s41467-022-34465-5>.
- Krishnan, Seetha, and Mark E. J. Sheffield. 2023. “Reward Expectation Reduces Representational Drift in the Hippocampus.” *BioRxiv*. <https://doi.org/10.1101/2023.12.21.572809>.
- Larkin, Margaret Carr, Christine Lykken, Lynne D. Tye, Jeanette Graham Wickelgren, and Loren M. Frank. 2014. “Hippocampal Output Area CA1 Broadcasts a Generalized Novelty Signal during an Object-Place Recognition Task: Generalized Novelty Signals in The Hippocampus.” *Hippocampus* 24 (7): 773–83.
- Lee, Heekyung, Douglas GoodSmith, and James J. Knierim. 2020. “Parallel Processing Streams in the Hippocampus.” *Current Opinion in Neurobiology* 64 (October): 127–34.
- Lee, Jae Sung, John J. Briguglio, Jeremy D. Cohen, Sandro Romani, and Albert K. Lee. 2020. “The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value.” *Cell* 183 (3): 620-635.e22.
- Leussis, M., and V. Bolivar. 2006. “Habituation in Rodents: A Review of Behavior, Neurobiology, and Genetics.” *Neuroscience and Biobehavioral Reviews* 30 (7): 1045–64.
- Leutgeb, Jill K., Stefan Leutgeb, May-Britt Moser, and Edvard I. Moser. 2007. “Pattern Separation in the Dentate Gyrus and CA3 of the Hippocampus.” *Science (New York, N.Y.)* 315 (5814): 961–66.
- Lever, Colin, Stephen Burton, Ali Jeewajee, Thomas J. Wills, Francesca Cacucci, Neil Burgess, and John O’Keefe. 2010. “Environmental Novelty Elicits a Later Theta Phase of Firing in CA1 but Not Subiculum.” *Hippocampus* 20 (2): 229–34.
- Li, Shaomin, William K. Cullen, Roger Anwyl, and Michael J. Rowan. 2003. “Dopamine-Dependent Facilitation of LTP Induction in Hippocampal CA1 by Exposure to Spatial Novelty.” *Nature Neuroscience* 6 (5): 526–31.
- Li, X-G, P. Somogyi, A. Ylinen, and G. Buzsáki. 1994. “The Hippocampal CA3 Network: An in Vivo Intracellular Labeling Study.” *The Journal of Comparative Neurology* 339 (2): 181–208.
- MacDonald, Christopher J., Kyle Q. Lepage, Uri T. Eden, and Howard Eichenbaum. 2011. “Hippocampal ‘Time Cells’ Bridge the Gap in Memory for Discontiguous Events.” *Neuron* 71 (4): 737–49.
- Madar, Antoine D., Laura A. Ewell, and Mathew V. Jones. 2019. “Pattern Separation of Spiketrains in Hippocampal Neurons.” *Scientific Reports* 9 (1): 5282.
- Madar, Antoine D., Anqi Jiang, Can Dong, and Mark E. J. Sheffield. 2025. “Synaptic Plasticity Rules Driving Representational Shifting in the Hippocampus.” *Nature Neuroscience*, March. <https://doi.org/10.1038/s41593-025-01894-6>.
- Mahn, Mathias, Inbar Saraf-Sinik, Pritish Patil, Mauro Pulin, Eyal Bitton, Nikolaos Karalis, Felicitas Bruentgens, et al. 2021. “Efficient Optogenetic Silencing of Neurotransmitter Release with a Mosquito Rhodopsin.” *Neuron* 109 (10): 1621-1635.e8.
- Masurkar, Arjun V. 2018. “Towards a Circuit-Level Understanding of Hippocampal CA1 Dysfunction in Alzheimer’s Disease across Anatomical Axes.” *Journal of Alzheimer’s Disease and Parkinsonism* 8 (1): 412–412.

- McKenzie, Sam, Roman Huszár, Daniel F. English, Kanghwan Kim, Fletcher Christensen, Euisik Yoon, and György Buzsáki. 2021. “Preexisting Hippocampal Network Dynamics Constrain Optogenetically Induced Place Fields.” *Neuron* 109 (6): 1040-1054.e7.
- McNaughton, B., and L. Nadel. 1990. “Hebb-Marr Networks and the Neurobiological Representation of Action in Space.” *Neuroscience and Connectionist Theory*. 405: 1–63.
- Moita, Marta A. P., Svetlana Rosis, Yu Zhou, Joseph E. LeDoux, and Hugh T. Blair. 2003. “Hippocampal Place Cells Acquire Location-Specific Responses to the Conditioned Stimulus during Auditory Fear Conditioning.” *Neuron* 37 (3): 485–97.
- Nakashiba, Toshiaki, Jennie Z. Young, Thomas J. McHugh, Derek L. Buhl, and Susumu Tonegawa. 2008. “Transgenic Inhibition of Synaptic Transmission Reveals Role of CA3 Output in Hippocampal Learning.” *Science (New York, N.Y.)* 319 (5867): 1260–64.
- Nakazawa, Kazu, Michael C. Quirk, Raymond A. Chitwood, Masahiko Watanabe, Mark F. Yeckel, Linus D. Sun, Akira Kato, et al. 2002. “Requirement for Hippocampal CA3 NMDA Receptors in Associative Memory Recall.” *Science (New York, N.Y.)* 297 (5579): 211–18.
- Nieh, Edward H., Manuel Schottdorf, Nicolas W. Freeman, Ryan J. Low, Sam Lewallen, Sue Ann Koay, Lucas Pinto, Jeffrey L. Gauthier, Carlos D. Brody, and David W. Tank. 2021. “Geometry of Abstract Learned Knowledge in the Hippocampus.” *Nature* 595 (7865): 80–84.
- Norris, Dennis. 2017. “Short-Term Memory and Long-Term Memory Are Still Different.” *Psychological Bulletin* 143 (9): 992–1009.
- O’Keefe, J. 2014. “Spatial Cells in the Hippocampal Formation.” <https://www.nobelprize.org/uploads/2018/06/okeefe-lecture.pdf>.
- O’Keefe, J., and J. Dostrovsky. 1971. “The Hippocampus as a Spatial Map. Preliminary Evidence from Unit Activity in the Freely-Moving Rat.” *Brain Research* 34 (1): 171–75.
- Pachitariu, Marius, Carsen Stringer, Mario Dipoppa, Sylvia Schröder, L. Federico Rossi, Henry Dalglish, Matteo Carandini, and Kenneth D. Harris. 2016. “Suite2p: Beyond 10,000 Neurons with Standard Two-Photon Microscopy.” *BioRxiv*. <https://doi.org/10.1101/061507>.
- Pedrosa, Victor, and Claudia Clopath. 2020. “The Interplay between Somatic and Dendritic Inhibition Promotes the Emergence and Stabilization of Place Fields.” *PLoS Computational Biology* 16 (7): e1007955.
- Priestley, James B., John C. Bowler, Sebi V. Rolotti, Stefano Fusi, and Attila Losonczy. 2022. “Signatures of Rapid Plasticity in Hippocampal CA1 Representations during Novel Experiences.” *Neuron* 110 (12): 1978-1992.e6.
- Rebola, Nelson, Mario Carta, and Christophe Mulle. 2017. “Operation and Plasticity of Hippocampal CA3 Circuits: Implications for Memory Encoding.” *Nature Reviews Neuroscience* 18 (4): 208–20.
- Robinson, Nick T. M., Lucie A. L. Descamps, Lloyd E. Russell, Moritz O. Buchholz, Brendan A. Bicknell, Georgy K. Antonov, Joanna Y. N. Lau, Rebecca Nutbrown, Christoph Schmidt-Hieber, and Michael Häusser. 2020. “Targeted Activation of Hippocampal Place Cells Drives Memory-Guided Spatial Behavior.” *Cell* 183 (6): 1586-1599.e10.
- Rolotti, Sebi V., Mohsin S. Ahmed, Miklos Szoboszlay, Tristan Geiller, Adrian Negrean, Heike Blockus, Kevin C. Gonzalez, et al. 2022. “Local Feedback Inhibition Tightly Controls Rapid Formation of Hippocampal Place Fields.” *Neuron* 110 (5): 783-794.e6.

- Satopaa, Ville, Jeannie Albrecht, David Irwin, and Barath Raghavan. 2011. "Finding a 'Kneedle' in a Haystack: Detecting Knee Points in System Behavior." In *2011 31st International Conference on Distributed Computing Systems Workshops*. IEEE. <https://doi.org/10.1109/icdcs.2011.20>.
- Schlingloff, Dániel, Szabolcs Káli, Tamás F. Freund, Norbert Hájos, and Attila I. Gulyás. 2014. "Mechanisms of Sharp Wave Initiation and Ripple Generation." *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience* 34 (34): 11385–98.
- Scoville, W. B., and B. Milner. 1957. "Loss of Recent Memory after Bilateral Hippocampal Lesions." *Journal of Neurology, Neurosurgery, and Psychiatry* 20 (1): 11–21.
- Sheffield, Mark E. J., Michael D. Adoff, and Daniel A. Dombeck. 2017. "Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation." *Neuron* 96 (2): 490-504.e5.
- Shinohara, Yoshiaki, Hajime Hirase, Masahiko Watanabe, Makoto Itakura, Masami Takahashi, and Ryuichi Shigemoto. 2008. "Left-Right Asymmetry of the Hippocampal Synapses with Differential Subunit Allocation of Glutamate Receptors." *Proceedings of the National Academy of Sciences of the United States of America* 105 (49): 19498–503.
- Shinohara, Yoshiaki, Aki Hosoya, Kazuko Yahagi, Alex S. Ferecskó, Kunio Yaguchi, Attila Sík, Makoto Itakura, Masami Takahashi, and Hajime Hirase. 2012. "Hippocampal CA3 and CA2 Have Distinct Bilateral Innervation Patterns to CA1 in Rodents: Rodent Hippocampal CA3 and CA2 Projection to CA1." *The European Journal of Neuroscience* 35 (5): 702–10.
- Shipton, Olivia A., Mohamady El-Gaby, John Apergis-Schoute, Karl Deisseroth, David M. Bannerman, Ole Paulsen, and Michael M. Kohl. 2014. "Left-Right Dissociation of Hippocampal Memory Processes in Mice." *Proceedings of the National Academy of Sciences of the United States of America* 111 (42): 15238–43.
- Shipton, Olivia A., and Ole Paulsen. 2014. "GluN2A and GluN2B Subunit-Containing NMDA Receptors in Hippocampal Plasticity." *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* 369 (1633): 20130163.
- Song, Da, Deheng Wang, Qinghu Yang, Tianyi Yan, Zhe Wang, Yan Yan, Juan Zhao, et al. 2020. "The Lateralization of Left Hippocampal CA3 during the Retrieval of Spatial Working Memory." *Nature Communications* 11 (1). <https://doi.org/10.1038/s41467-020-16698-4>.
- Squire, Larry R. 2009. "The Legacy of Patient H.M. for Neuroscience." *Neuron* 61 (1): 6–9.
- Takeuchi, Tomonori, Adrian J. Duzkiewicz, Alex Sonneborn, Patrick A. Spooner, Miwako Yamasaki, Masahiko Watanabe, Caroline C. Smith, et al. 2016. "Locus Coeruleus and Dopaminergic Consolidation of Everyday Memory." *Nature* 537 (7620): 357–62.
- Tamboli, Suhel, Sanjay Singh, Dimitry Topolnik, Mohamed El Amine Barkat, Risna Radhakrishnan, Alexandre Guet-McCreight, and Lisa Topolnik. 2024. "Mouse Hippocampal CA1 VIP Interneurons Detect Novelty in the Environment and Support Recognition Memory." *Cell Reports* 43 (4): 114115.
- Tao, Sijue, Yihang Wang, Jundan Peng, Yang Zhao, Xiaobin He, Xuefeng Yu, Qing Liu, Sen Jin, and Fuqiang Xu. 2021. "Whole-Brain Mapping the Direct Inputs of Dorsal and Ventral CA1 Projection Neurons." *Frontiers in Neural Circuits* 15 (April): 643230.
- Tolman, E. C. 1948. "Cognitive Maps in Rats and Men." *Psychological Review* 55 (4): 189–208.

- Trimper, John B., Claire R. Galloway, Andrew C. Jones, Kaavya Mandi, and Joseph R. Manns. 2017. “Gamma Oscillations in Rat Hippocampal Subregions Dentate Gyrus, CA3, CA1, and Subiculum Underlie Associative Memory Encoding.” *Cell Reports* 21 (9): 2419–32.
- Wagatsuma, Akiko, Teruhiro Okuyama, Chen Sun, Lillian M. Smith, Kuniya Abe, and Susumu Tonegawa. 2018. “Locus Coeruleus Input to Hippocampal CA3 Drives Single-Trial Learning of a Novel Context.” *Proceedings of the National Academy of Sciences of the United States of America* 115 (2): E310–16.
- West, M. 1990. “Stereological Studies of the Hippocampus: A Comparison of the Hippocampal Subdivisions of Diverse Species Including Hedgehogs, Laboratory Rodents, Wild Mice and Men.” *Progress in Brain Research* 83: 13–36.
- Whittington, James C. R., Timothy H. Muller, Shirley Mark, Guifen Chen, Caswell Barry, Neil Burgess, and Timothy E. J. Behrens. 2020. “The Tolman-Eichenbaum Machine: Unifying Space and Relational Memory through Generalization in the Hippocampal Formation.” *Cell* 183 (5): 1249-1263.e23.
- Wilkinson, Jamie L., Laura Herrman, Matthew I. Palmatier, and Rick A. Bevins. 2006. “Rats’ Novel Object Interaction as a Measure of Environmental Familiarity.” *Learning and Motivation* 37 (2): 131–48.
- Wright, William J., Nathan G. Hedrick, and Takaki Komiyama. 2025. “Distinct Synaptic Plasticity Rules Operate across Dendritic Compartments in Vivo during Learning.” *Science (New York, N.Y.)* 388 (6744): 322–28.
- Yang, Wannan, Chen Sun, Roman Huszár, Thomas Hainmueller, Kirill Kiselev, and György Buzsáki. 2024. “Selection of Experience for Memory by Hippocampal Sharp Wave Ripples.” *Science (New York, N.Y.)* 383 (6690): 1478–83.
- Yassa, Michael A., and Craig E. L. Stark. 2011. “Pattern Separation in the Hippocampus.” *Trends in Neurosciences* 34 (10): 515–25.
- Zaki, Yosif, and Denise J. Cai. 2025. “Memory Engram Stability and Flexibility.” *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology* 50 (1): 285–93.
- Zemla, Roland, Jason J. Moore, Maya D. Hopkins, and Jayeeta Basu. 2022. “Task-Selective Place Cells Show Behaviorally Driven Dynamics during Learning and Stability during Memory Recall.” *Cell Reports* 41 (8): 111700.
- Ziv, Yaniv, Laurie D. Burns, Eric D. Cocker, Elizabeth O. Hamel, Kunal K. Ghosh, Lacey J. Kitch, Abbas El Gamal, and Mark J. Schnitzer. 2013. “Long-Term Dynamics of CA1 Hippocampal Place Codes.” *Nature Neuroscience* 16 (3): 264–66.
- Zutshi, Ipshita, Manuel Valero, Antonio Fernández-Ruiz, and György Buzsáki. 2022. “Extrinsic Control and Intrinsic Computation in the Hippocampal CA1 Circuit.” *Neuron* 110 (4): 658-673.e5.