THE UNIVERSITY OF CHICAGO

MEMORIES OF AVERSIVE EXPERIENCES: INSIGHTS FROM THE THALAMIC NUCLEUS REUNIENS TO HIPPOCAMPAL CA1 PATHWAY AND CA1 NEURONAL DYNAMICS

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 NEUROBIOLOGY

BY

HEATHER CASSIDY RATIGAN

CHICAGO, ILLINOIS

DECEMBER 2023

This dissertation is dedicated to my parents, Linda Travis Macomber and Brian Lee Macomber, my dog Emmy, my spouse, Timothy Raymond Justice Ratigan, and my cherished friends, who I am fortunate to say are too numerous to list. I could not have done this without each and every one of you.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF SUPPLEMENTARY FIGURES

ACKNOWLEDGMENTS

I am immensely grateful to my advisor Mark Sheffield for all of the guidance and support you have provided over the last five years. Reading your paper as an undergraduate on dendritic contributions to hippocampal representations was immensely inspiring, and I distinctly remember presenting it in a seminar full of excitement about this interesting new research. From our very first conversation at 2017 SfN before I had even applied to or been accepted into the program, you were willing to talk for hours about science, life, and everything in between. After you welcomed me into the lab you were willing to take a chance on my wild ideas about nucleus reuniens, a brain region neither of us knew much of anything about. You let me develop a vague idea of 'what does this weird projection even do' into a fully-fledged project, with complete support that I could figure it out. Our conversations have helped to greatly guide my scientific and personal development, and your insights have consistently made me a better scientist. You have fostered an incredible lab environment of scientists who truly care about one another, and that kind of care starts from the top.

On that note, I'd like to thank all the current and former members of the Sheffield lab I have interacted with throughout this project, including Seetha Krishnan, Antoine Madar, Douglas Goodsmith, Chad Heer, Anqi Jiang, Yuhung 'Tommy' Chiu, Bryan Garcia, Chery Cherian, Shai Smith, Valerie Barreto, Phoebe Cao, Roma Shah, Anthony James Jr,, Sophia Vann-Adibe, Denisse Morales Rodriguez, and Mohamed Fawaz. I have been fortunate to be in such a selfless lab full of thoughtful, brilliant, and above all *kind* scientists who have freely given their valuable time, insightful thoughts, and an ear to commiserate with over these last five years. I'd like to give special thanks to Chad for teaching me everything you knew about axonal imaging, Seetha for tirelessly guiding me through learning your CFC protocol and 2p imaging, Anqi for our fruitful coding collaborations, Tommy for your contributions to replay analysis, Cherry Wang for taking up the mantle of continuing this work down whatever path you find most fruitful, and

extra especially to my two undergraduate assistants, Shai Smith and Valerie Barreto, for your incredibly tireless work on this project over the years. I couldn't have asked for better labmates. Science is a team sport, and I couldn't have finished this project without all of you.

I additionally would like to thank my thesis committee members: Dr. Murray Sherman, Dr. Matt Kaufman, and Dr. David Freedman. Each of you has given your time freely to provide insightful suggestions along the way. To Murray, for including me in lab meetings and both broad and specific discussions about the importance of oft-disregarded thalamic regions. To Matt, for helping me think rigorously about how to extract meaningful and interpretable insights from my data, and for graciously letting us borrow reagents and techniques from your lab. To Dave, for your insights on how all the pieces can combine into a bigger picture, and for your broader thoughts on science and life. It has been a great privilege to learn from each and every one of you, and I am a better scientist for the effort you have each put into molding me.

PhDs are rooted in a place, and I would like to thank departmental, institute, and university staff, especially Elena Rizzo and Stephanie Thomas, for troubleshooting bureaucracy and keeping everything organized and running smoothly. I would like to acknowledge and thank members of the UChicago imaging core for their contributions to confocal imaging in this thesis, and members of the ARC veterinary and animal care staff for their tireless, difficult, and often under-acknowledged work keeping our experimental animals happy and healthy.

I have been so grateful to have an incredible community of friends and found family surrounding me throughout my scientific journey. To our neighbors, for making our home in Hyde Park into a true community. To Rajeev Emirelli, Katie Blume, and Marissa Applegate, my fearless graduate students from afar, for always being willing to talk about science or life in ways that only people going through the same experience can understand. To Deepti Varathan, for every hour spent in our dorm studying in college, for officiating Tim and I's wedding, for every

moment of support. To Katherine Frain for always being willing to proofread every grant and paper and important document I've ever written (this one included), for being a voice of reason outside of science, and for all the creativity and fun. To Nikki Lyons, the challenges you have faced to stay a fearless scientist are nothing short of inspiring, and your unflagging faith in my ability to do this has stayed with me. You three are my rocks, my strongest cheerleaders, and my heart from college and beyond. To Sasha Warren for countless conversations that got me through this PhD, and for countless hours of extremely necessary fun, and a special shoutout to James Zak for last-minute R help to rescue the statistics in Chapter 2. To Ross Rad, for the last five years of friendship started at interviews and still going strong. Every moment of sitting on the back porch after a long day, every show watched quietly hanging out, and the very critical task of co-parenting our clingy dogs together, there's not enough words in the english language to describe what your friendship has meant to me - you're the brother I never had. On that note, thank you to Emmy, the best emotional support dog in the universe, for all of the well-earned cuddles and walks - while you can't read, I still feel the need to thank you here nonetheless.

I have been so incredibly blessed to have not only an incredible found family, but to have an amazing family both by blood and marriage. Thank you to my parents, my father Brian Lee Macomber, and mother Linda Travis Macomber. Without your support in uncountable ways material and immaterial, I could not have made it anywhere near this point. To my dad, thank you for always fostering my curiosity, teaching me about space and science, and always supporting every choice I make. All those hours watching PBS, playing board games, and reading science fiction helped mold my curious developing mind into the scientist I am today. I hope to have inherited your curiosity and charisma in equal doses. To my mom, thank you for your tireless efforts to give me every possible opportunity, often at personal sacrifice, and for letting me go wide and afar in search of the best possible education. Your tireless support in raising me has not gone unnoticed. Every summer camp, every museum visit, every trip near and

afar, every girls-in-STEM program, made it possible for me to chase my ambitions and complete the neurobiology PhD I so audaciously declared I wanted to get when I was 12. You never failed to believe I could do it, and moved heaven and earth to make it possible. Thank you for all the good memories. Thank you to Jennifer Topping and Terry Ratigan, my in-laws, and to Joanna and Lucas Ratigan, my siblings-in-law, for graciously welcoming me into your home and your family, for cheering me on, and for so much support from near and afar over all these years.

Lastly, I don't have words to express how much my spouse, Timothy Raymond Justice Ratigan, means to me. For nearly a decade, you have been my best friend, my safe harbor in a storm, and the person I most want to be around all of the time. Your brilliant, selfless, and kind heart is unlike anyone I know, and shines through in every moment with everyone who is lucky enough to know you. My life has been made truly incredible by virtue of you being by my side, and you are my home, no matter where we are. Your unwavering, unconditional support in ways material and immaterial too numerous to count throughout every hard moment in college and throughout every hard moment in this PhD have been instrumental to me getting to this point. For every time you walked Emmy when I was working late in the lab, for every time you offered a shoulder to cry on when an experiment failed, for your unwavering belief that I am brilliant and capable, even when I didn't believe it myself. This dissertation, in particular, would not exist without all the uncountable hours spent bouncing ideas off of you, without your help with tricky coding problems, and getting your expert revisionary eye on writing. Thank you, for everything.

ABSTRACT

Being able to flexibly remember aversive events is a critical neural function for mammals. However, aversive memories can have a negative impact on animal well-being if they are not restricted to the appropriate context, or not updated with new information over time. The hippocampus is a key brain region for both forming and updating aversive memories –specifically, the dorsal section of hippocampal area CA1 is crucial for integrating spatial and non-spatial information into coherent contextual memories. My doctoral work focuses on understanding the hippocampal circuits underlying the formation, recall, and extinction of aversive memories.

In Chapter 2, I reveal the contextual fear-suppressive role of the projection from the thalamic subregion nucleus reuniens (NR) to dorsal CA1. I implemented a novel head-restrained virtual reality contextual fear conditioning paradigm (VR-CFC) where I administered mild tail shocks in one context (shocked) but not the other (neutral). I showed that inactivation of the NR-CA1 pathway following VR-CFC increases fearful freezing, induces fear generalization to a neutral context, and delays extinction of fearful responses. Using in vivo sub-cellular calcium imaging, I found that NR-axons become selectively tuned to fearful freezing only after VR-CFC.

In Chapter 3, I demonstrate how the NR-CA1 pathway impacts population dynamics in excitatory pyramidal neuron somata in dorsal CA1 throughout VR-CFC. I reveal that with the NR-CA1 pathway intact, CA1 somata are preferentially active in the shocked context after shocks, globally remap their place fields to more flexible post-fear extinction maps, tune to both shocks and to freezing epochs, and increase both the frequency and scale of simultaneous neural activations. In contrast, I show that inhibiting the NR-CA1 pathway prevents preferential activity

in the shocked context post-shocks by over-activating somata, interferes with global remapping by retaining more stable post-extinction maps, increases freeze-tuning, and disrupts simultaneous neural activations. Lastly, in Chapter 4, I investigated the differential activity of the basal and apical dendritic arbors through novelty and reward manipulation. I reveal functional segregation between predominantly spatial representations in the basal dendrites and predominantly non-spatial reward and novelty-modulated representations in the apical dendrites, paving the way for future investigations of dendritic mechanisms underlying somatic activity changes in VR-CFC.

Overall, this dissertation provides both methodological and experimental leaps forward in our understanding of how the hippocampus flexibly encodes aversive experiences by highlighting the critical role of projections from thalamic nucleus reuniens to hippocampal dorsal CA1. I pioneered a novel contextual fear conditioning paradigm in virtual reality that reliably elicits context-dependent freezing behaviors similar to non-head-fixed behavior. I imaged two distal subcellular components for the first time – NR axons and dorsal CA1 distal apical tuft dendrites – and developed a new technique for identifying population-level synchronized activity in dorsal CA1 somata. Our findings paint a picture of hippocampal dorsal CA1 as a critical region for memory that both encodes a variety of salient features to aversive contextual memories and flexibly updates fear memories to extinguish fear responses through modulating inputs from thalamic nucleus reuniens.

CHAPTER 1

INTRODUCTION

Learning, remembering, and flexibly altering memories of threatening experiences is a critical evolutionarily-conserved behavior. An organism's survival often depends on appropriately responding to threats in its environment, a behavior heavily reliant on accessing and correctly applying underlying neural encodings of previously negative experiences. These memories need to be stable enough to be recalled across the organism's lifetime, but also be flexible enough to incorporate new information and generalize to non-identical but similarly threatening situations, all while not occupying too much space in the brain. Understanding this process of how the brain transforms disparate sensory information into flexible, stable, and compact memories has been a focus of neuroscience research for decades.

One form of memory that is essential for recollecting aversive events is episodic memory, in which key components of event timing, location, and context combine into a singular memory. These episodic memories are, by partial definition, contextual memories, as they require encoding contextual representation. Several brain regions are necessary to form contextual memories, perhaps chief among them the hippocampal formation (HPC). However, the recall and modification of aversive memories requires the activation of a widespread, coordinated network of brain regions, including but by no means limited to the amygdala (AMY), medial prefrontal cortex (mPFC), and nucleus reuniens of the thalamus (NR).

This dissertation will focus on how these regions interact to shape hippocampal activity throughout context-dependent aversive memory formation, retrieval, and suppression via subcellular and population-level processes. Since the experiments conducted in this dissertation

are entirely on mice, the work reviewed and discussed here will be primarily, albeit not exclusively, focused on studies conducted in rodent models.

What is Contextual Fear?

Understanding contextual fear requires first defining the term 'context.' A context is a set of continuous experienced conditions that define an event[1]. Contexts are typically composed of individual multisensory cues that the brain combines to create a singular, abstract representation of the surroundings involved in a salient event. It is important to note that the working definition of context here is a neural representation, *not* a representation in physical space. Forming a context requires the brain to combine incoming information about shapes, colors, smells, and textures with more abstract physical information about spatial location and temporal properties[2]. These aspects of a context rooted more firmly in the physical world are further imbued with internal information, such as psychological state and prior knowledge[3]. All contexts do not necessarily contain each of these components, and some contexts include additional components such as social, cultural, or further multi-sensory representations. However, all contexts are critically more than the sum of their parts. Despite being made of many specific elements, the context *itself* is represented distinctly from the representation of individual elements. This distinct encoding keeps contextual representations intact even if specific elements are changed[1].

As a more concrete example, let us examine a context in which a person approaches a red light at an intersection while driving a car on their way home from work. This experience's physical features include the car's location at an abstract level (country, city, cross-streets) and a sensory-detectable level (the location of this traffic signal compared to a block ahead). It includes the time in which the driver is approaching at an abstract level (time on the clock, day of the week, month, year), a sensory-detectable level (environmental luminance of the sun setting), and

an internal relativistic level (after work, before dinner). They include the smell of asphalt and exhaust, the sounds of honking and the car radio, objects made of various colors and shapes that form into cars, people, the signal light, street poles, sidewalks, and other details. Internal features include if the person driving is tired and stressed from a long work day, causing sensory information to be less clearly encoded than it otherwise might be. It includes learned cues, such as the knowledge that 'green' on the signal means they should continue through the intersection and 'red' to stop, while those colors in other contexts have other meanings. It includes learned responses, such as a loud honk, indicating that the light changed while the tired driver was inattentive. It is also a flexible representation; a sign could be removed, or a new house could be built on the corner, without fundamentally changing the contextual representation of the intersection. These features all combine to create the context against which a memory forms.

A key question then becomes why does the brain create these contexts? Contexts are critical to remember previous experiences, guide present behavior, and make decisions impacting the future. Contexts enable organisms to flexibly situate experiences by creating succinct, semi-abstract recollections. In the example above, the reader likely pictured a 'generic' intersection containing all the described features. However, each specific context pictured will vary from reader to reader based on different learned experiences - for example, the driver in this example would be on the left side of the car and the right side of the street for an American reader, while a British reader would reverse these details. The brain's capacity to form specific yet flexible contexts is how both readers can readily access a gestalt of this intersection context despite each individual's precise context containing different elements.

What information is meaningful is a question of great importance to the brain and those studying its processes. Organisms have a limited capacity to remember each experience in their

lifetimes and must prioritize memories of behaviorally salient experiences over more mundane experiences[4]. Therefore, we remember surprising, positive, and negative events more vibrantly than ordinary events[5]. Of particular relevance to survival are aversive experiences. To return to the intersection example, if the driver stops at the red light and then goes through again at the green light, the particular context surrounding this intersection will likely be forgotten.

However, suppose the driver stops at the red light, accelerates at the green light, and is hit by a car. In that case, the context suddenly gains salience as a location where a dangerous event happened. Imbuing this context with salience might guide future behavior in helpful ways, such as slowing down at this intersection and looking carefully at oncoming traffic before accelerating the next time the driver passes through. In real-world situations for humans and other organisms, correctly embedding fearful memories in their appropriate contexts is a crucial survival skill.

Fear Conditioning: A Brief History

To study how fearful experiences embed into a context, the two are linked together through 'conditioning.' The earliest examples of conditioning come from Ivan Pavlov in the 1920s, whose contributions to the field were so notable that certain forms of conditioning are still termed 'Pavlovian' today[6]. His paradigm trained dogs to pair a conditioned stimulus (CS) that was intrinsically meaningless, like a ringing bell, with an unconditioned stimulus, like food, that had a positive salience. After ringing the bell and then presenting food, dogs quickly learned to salivate after the presentation of the bell but before the presentation of food, showing that they had associated the US with the CS[7].

Researchers quickly applied the same principle to an aversive US, eliciting fear responses from creatures after CS presentation but before US presentation. The most infamous early example of fear conditioning is found in the 'Little Albert' study in 1920, in which a human

toddler with no innate fear of rats (the CS) but a substantial innate fear of loud noises (the US) was presented with a rat and a loud noise simultaneously, resulting in a long-lasting conditioned fear response (CR) to rat presentation alone[8]. It is important to note that while this experiment yielded a strong case of fear conditioning, its replication would not meet modern ethical standards.

A full discussion of what defines fear and the conscious experience of fear in non-human animals is outside this dissertation's scope. However, I will use a working definition of fear as a negative internal state that orchestrates psychological and behavioral defense responses to threats, real or perceived[9,10]. Following the 'Little Albert' experiment, researchers designed a paradigm pairing a context as the CS with the US of mild foot-shocks in rodent models[11–13]. Animals can acquire contextual fear memory in as little as a single shock, although they encode it more reliably through a series of shocks[14]. Fear memory retrieval is measured by subsequent elicitation of a CR in the CS in the absence of the US, most commonly freezing, defined as a complete cessation of bodily movement except breathing[13].

Through repeated re-exposure of subjects to the conditioned context, fear CRs can be extinguished, with freezing returning to pre-shock levels[7,15,16]. Fear extinction is a learned behavior that does not erase or overwrite the underlying fear memory. Instead, fear extinction is a newly-formed memory that competes with and, when successful, suppresses the original fear memory for control of animal behavior[7,17,18]. Evolutionarily, extinction is a critical component of the fear-learning process. To return to the intersection car accident example, if the driver never extinguished their fear memory of that accident, they would be paralyzed with fear and potentially unable to drive through that intersection in perpetuity. This fear memory would

interfere with their ability to successfully navigate to and from work, impairing their normal functioning.

Another modification of the fear memory retrieval process is generalization, a process through which a CR is linked to a context similar to, but not the same as, the CS[7,19–21]. Generalization enables fear learning to be flexible and apply to situations where defensively enacting fear behavior may be an appropriate response. For example, driving carefully through intersections that share some features with an intersection where a crash occurred may be an appropriate enactment of defensive fear behavior. In the face of a potential threat, it is generally evolutionarily advantageous to have high 'false positive' rates of fear expression, as the downsides of excess CR in a threatening situation are often lower - although not zero - when contrasted with the downsides of fearlessness.

However, every context should not induce a fear response, just as getting in a car alone should not induce a severe fear response after a fender-bender. Organisms must be able to constantly find a delicate balance, discriminating between fearful and safe contexts, to exhibit fear responses as long as they are salient and still extinguish them once they lose their utility, while regularly updating fear memories in a constantly-evolving, often-threatening world. Failure to find this balance can be catastrophic, as demonstrated by the devastation experienced by sufferers of neurological disorders that have inappropriate responses to perceived threats as their core pathology, chief among them post-traumatic stress disorder but also including a variety of anxiety disorders, bipolar disorder, and trauma-induced emotional dysregulation[22].

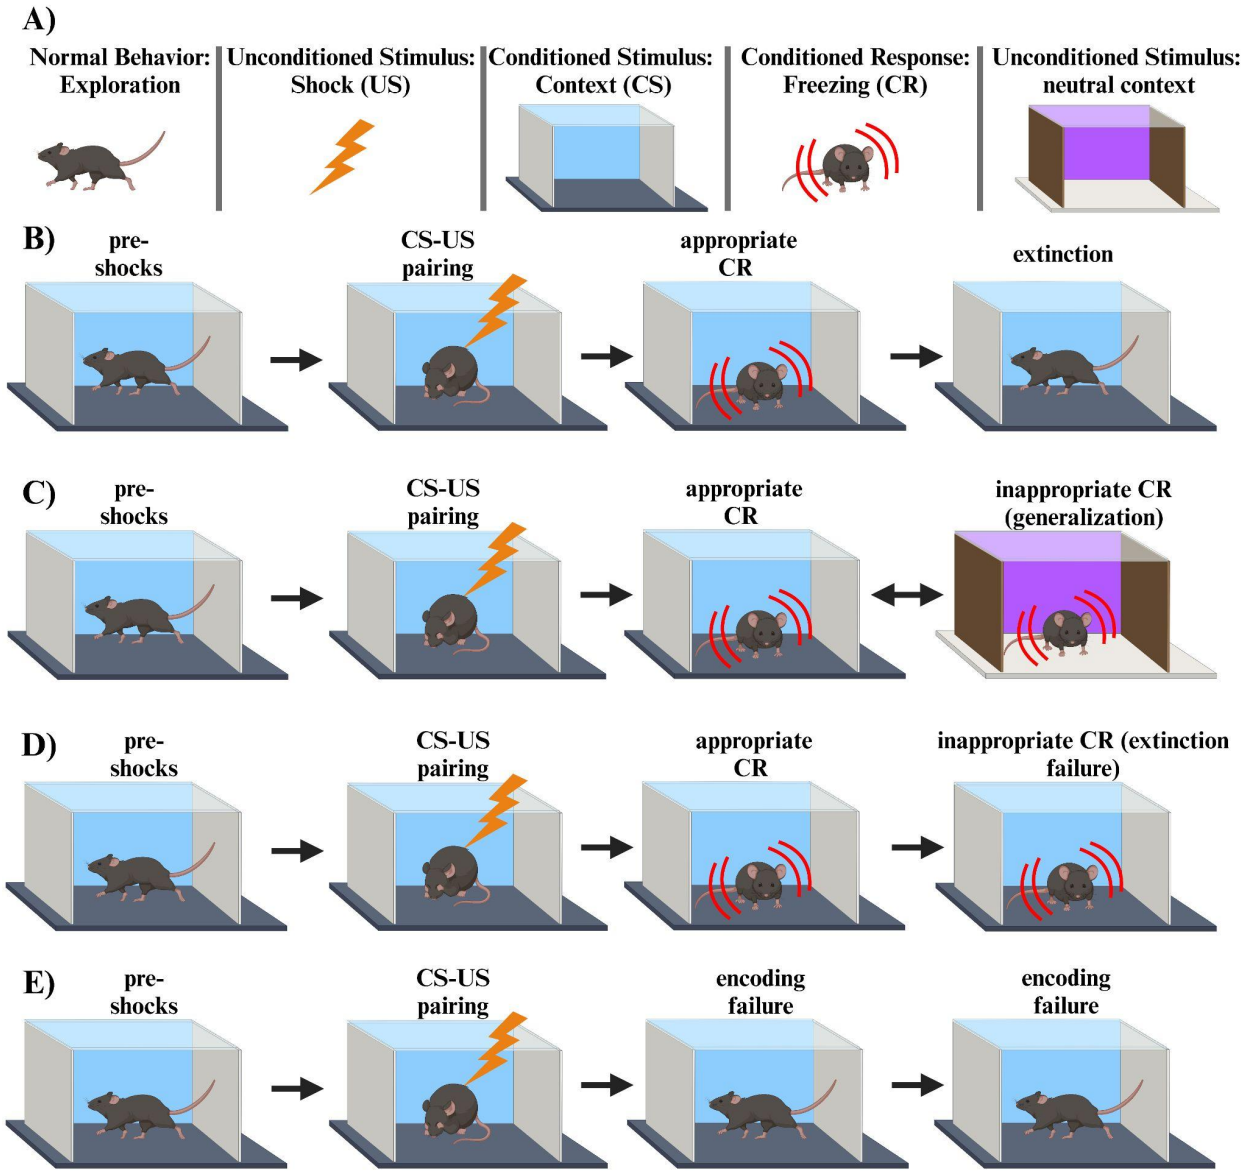

Figure 1.1: Schematic of Contextual Fear Conditioning.
A) Legend of components important in contextual fear conditioning. **B)** Contextual fear conditioning working in a typical sense. After pairing the context and shock, mice freeze in response to being placed back in the context. Over re-exposures to the context, mice extinguish the fear memory, stop freezing, and resume normal exploration. **C)** The mouse is placed in an unconditioned neutral context and freezes, an example of fear memory generalization. **D)** The mouse fails to extinguish the CR, even on repeated exposures. **E)** The mouse fails to encode the CS-US pairing, never exhibiting appropriate CRs. Made with BioRender[23].

The Contextual Fear Regulation Network

Over the years, dozens of cortical and subcortical brain regions have been identified as playing a role in forming, recalling, and regulating contextual fear memories[1,14,24–30]. Factors such

as the organism studied, the type of fear memory, and the stage of memory processing impact which of these dozens of brain regions are engaged in contextual fear memory. However, in the interest of brevity and a discussion of work with the highest relevance to the experimental work presented in this dissertation, I will focus on just four significant regions at the center of contextual fear research: the amygdala, medial prefrontal cortex, nucleus reuniens, and the hippocampus.

Amygdala: Fear Encoding

The amygdala is an integrated network of over a dozen distinct nuclei, categorized broadly into the central (CeA), medial (MeA), cortical (CoA), basomedial (BMA), and basolateral (BLA) regions, with excitatory glutamatergic neurons primarily located in the BLA[31]. In both early bilateral amygdala ablation studies in monkeys and later studies of amygdaloid complex damage in humans, subjects show a remarkable lack of fear of innately fearful stimuli, such as a loud abrupt tone or a snake-shaped object[32–36]. While the amygdala is implicated in a broader range of behaviors than fear and threat assessment, including in an increasing body of work on the amygdala's role in reward, its critical role in fear responses is still its best understood[37,38]. During classical fear conditioning, information about the CS - most frequently an auditory tone - and the US -most frequently a shock - directly converge on the BLA, sometimes onto the same neuron, altering its synaptic plasticity and forming the CS-US association[39]. It is generally accepted that NMDA-receptor-dependent long-term potentiation underlies this plasticity[40–42]. Contextual fear conditioning requires integrating cues into a contextual representation outside the BLA, which is then associated with the US within the BLA[43,44]. While often discussed together, the basal amygdala (BA) and lateral amygdala (LA) are morphologically and functionally distinct substructures, with the BA, in particular, necessary for

contextual fear encoding, mainly due to BA receiving preferential hippocampal input[45]. Following contextual fear memory encoding, subsequent presentation of the CS elicits robust reactivation of BA neurons. This response is subsequently modified by the CeA and routed onwards to output regions elsewhere in the brain[46]. While the amygdaloid complex generally and BA specifically are critical for contextual fear encoding and fear CR behaviors, it does not act alone. Inputs encoding the CS come to the amygdala through the hippocampal formation, and amygdalar activity is strongly regulated by the medial prefrontal cortex and by the nucleus reuniens, with both regions possessing strong bilateral connections to the BLA[42,47].

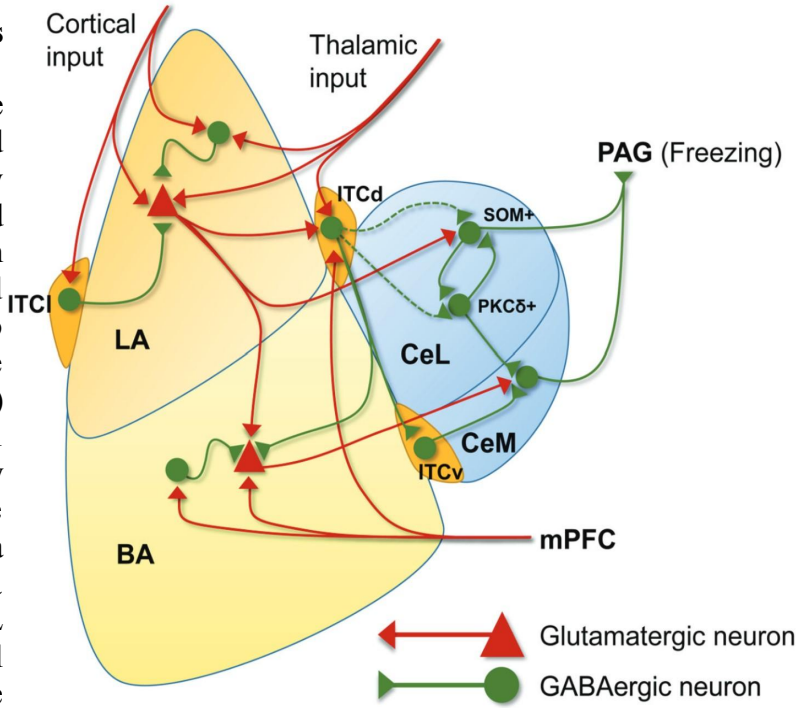

Figure 1.2: Amygdala circuits involved in fear learning. Fear-related sensory signals are delivered from cortical and thalamic inputs to the LA by monosynaptic excitation and disynaptic inhibition. Projection neurons in the LA send fear-related signals to downstream areas, such as the lateral central amygdala (CeL) or BA. CeL and medial central amygdala (CeM) inhibitory neurons project to the periaqueductal gray (PAG), a brain region that elicits a conditioned fear response. CeL projects to CeM. Dorsal intercalated cells (ITC) receive neural signals from the LA, thalamus, and medial prefrontal cortex. Dorsal ITC neurons are known to project to the CeL. Ventral ITC neurons receive inhibitory inputs from dorsal ITC, and mainly inhibit CeM neurons. Excitatory inputs from mPFC can activate both projection neurons and interneurons in the BA, as well as the neurons in the dorsal ITC. Adapted from Lee et al. 2017[48].

mPFC: Emotion Regulation

The rodent medial prefrontal cortex (mPFC) is a midline frontal cortical region comprised of three subdivisions: the anterior cingulate cortex (ACC), the prelimbic cortex (PL),

and the infralimbic cortex (IL)[49]. Unlike more strongly conserved regions in the brain, there is substantial discussion about whether rodent PFC generally can be considered analogous to non-human primate and human PFC and whether or not anatomical boundaries and general functions of rodent PFC subregions transfer adequately[50,51]. One aspect of PFC architecture that adds confusion is its lack of singular functional definition, with its excitatory activity exhibiting a wide diversity of responses to stereotyped inputs and a tendency towards multimodal high-dimensional neural responses over straightforward behavior-linked, easily-interpretable activity[52,53].

Since the role of mPFC in contextual fear conditioning is of particular relevance to this discussion, I will focus there, which also means a focus on PL and IL. Both regions receive varied excitatory inputs, most strongly from the mediodorsal nucleus of the thalamus (MD), BLA, HPC, and NR[54–60]. Both IL and PL strongly output to AMY, although these projections target different regions, with PL outputs centering more strongly on BA and IL outputs centering more strongly on LA, MeA, and CeA[61]. Both regions additionally project reciprocally to thalamic regions, including large projections to MD and NR, primarily from layers 1, 5, and 6, although we do not yet know if IL and PL target the same neurons within NR[62]. Interestingly, despite major innervation from the hippocampal formation in its inputs, neither IL nor PL project back to the hippocampal formation directly[47]. Both regions have bilateral connectivity with each other and their contralateral partners[63–65].

Overall evidence for mPFC's involvement in fear processing came first from ablation studies, in which rats with mPFC lesions showed typical conditional fear acquisition and formed long-lasting fear memories with normal CRs but showed deficits in extinction learning and prolonged timelines for fear extinction[66,67]. More selective PL lesions yielded outsized cued and

contextual fear CRs and extinction deficits[68]. Individual neurons in PL increase their spiking during aversive CSs in a way that is highly correlated with ongoing CRs, specifically freezing, and microstimulation of PL increases fear expression levels[69,70]. Consistent with lesion studies, pharmacological or optogenetic inactivation during contextual fear learning or memory retrieval also impairs CRs[71]. PL seems to play a much less important role during extinction, as inactivation of PL during extinction learning or retrieval does not impact the rodent's ability to learn to extinguish fear[72].

Meanwhile, early work targeting specifically IL showed that while contextual fear extinction learning was not impaired, its consolidation was, although subsequent work has shown mixed effects of IL inhibition before extinction[73–76]. Consistently, IL inactivation after extinction learning impairs both consolidation and retrieval [73,77,78]. IL neurons burst fire during extinction memory formation and fire preferentially to extinction memory retrieval[79,80]. Despite a focus here on differences, it is noteworthy that not all evidence agrees that PL and IL are performing strictly separate roles. Some studies have shown similar patterns of activity and decreased function in both regions during fear memory encoding, expression, extinction, and renewal, potentially due in part to similarities in their inputs, strong cross-regional connectivity, or more nuanced processing roles for these regions than present studies have captured[81]. Generally, mPFC is important for top-down regulation of contextual fear memory acquisition, expression, and extinction, with a segregation of roles between PL for contextual fear memory acquisition and expression and IL for contextual fear extinction consolidation and expression.

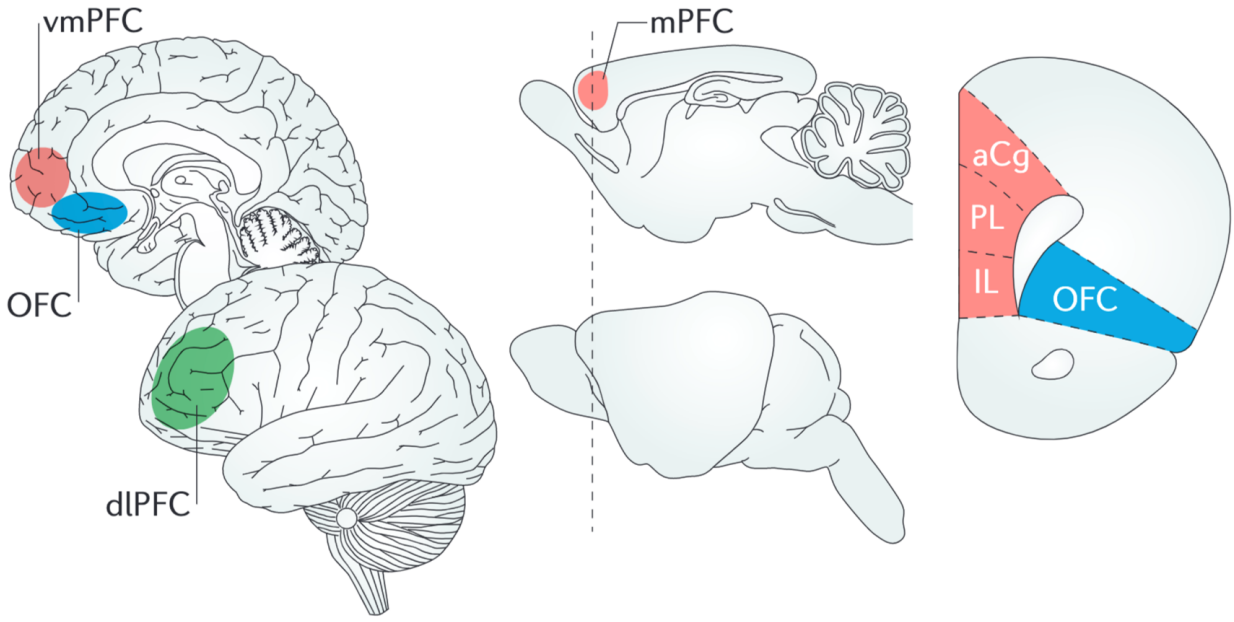

Figure 1.3: The prefrontal cortex
Left: Approximate regions with the three major segments of the prefrontal cortex (PFC) are shown in humans: the ventromedial PFC (vmPFC), orbitofrontal cortex (OFC), and dorsolateral PFC (dlPFC). Middle: Sagittal section of the mouse brain with medial PFC (mPFC) highlighted in red. Right: Coronal section through mPFC with the anterior cingulate cortex (aCg, labeled ACC elsewhere in this discussion, prelimbic cortex (PL), and infralimbic cortex (IL) labeled. The orbitofrontal cortex is shown in blue. No accepted mouse analog of dlPFC exists. Adapted from Eichenbaum 2017[82].

Nucleus Reuniens: mPFC-CA1 Communication Hub

The role of thalamic nuclei in contextual fear processing has received significantly less attention than many of the other brain regions discussed in this chapter due in large part to a persistent belief that the thalamus's main purpose is to 'relay' sensory and motor information between subcortical and cortical areas[83,84]. Historically, the thalamus is described as having specific and nonspecific nuclei, i.e., regions that serve a clear 'relay' function and nuclei whose function are unclear from basic anatomical explorations[85,86]. Those nonspecific nuclei are further divided roughly into two groups. First, nuclei that receive primary inputs from layer 5 pyramidal neurons in sensorimotor cortices and send that information elsewhere, often largely to the cortex, which are often associated with primary nuclei. Second are the midline and intralaminar nuclei,

largely defined by their anatomical location along the midline and their lack of anatomically clear directional connectivity between subcortical and/or cortical regions [86].

Nucleus reuniens is the largest of the midline ventral thalamic nuclei, located directly dorsal to the third ventricle and along the majority of the rostral-caudal axis of the thalamus [87]. Its makeup is entirely glutamatergic, with GABAergic modulation coming largely from the thalamic reticular nucleus (TRN) [87,88]. Within these glutamatergic neurons, subsets of neurons express either the calcium-binding proteins calretinin, calbindin, or both[89,90]. CB^+ neurons are especially relevant to discussions of thalamic function as they are considered 'matrix' cells, which generally lack the topographical sensory precision of 'core' cells found more frequently in primary thalamic nuclei[91]. Internal subdivisions are largely between the center of NR and its lateral 'wings', sometimes termed peri-reuniens but often grouped with NR proper[47]. Like investigations into most thalamic nuclei, researchers first examined NR's inputs and outputs for clues to its function. Herkenham's influential 1978 study was the first to unveil its massive projections to the mPFC and hippocampus[87]. Subsequent investigations found NR projections targeted the stratum lacunosum-moleculare (SLM), roughly the hippocampal equivalent of cortical layer 1, of the dorsoventral axis of hippocampal CA1 and lateral dorsoventral subiculum (Sub) specifically while sparing CA2, CA3, and the dentate gyrus[54,92,93]. The targets of these projections seem to be both excitatory dendrites and inhibitory dendrites/somata in the SLM, although there is some controversy on whether NR synapses only on inhibitory interneuron or directly on excitatory pyramidal neuron distal apical tuft dendrites[94–101].

The second major site of NR projections investigated is the mPFC, where it terminates densely in both layers 1 and layers 5/6 of all mPFC subregions, projections that further confirm the 'matrix' status of NR[102]. More recent investigations have found a substantial subpopulation of

5-10% of NR rostral pole neurons project both to CA1 and to mPFC[103,104]. Other important projections of note from NR are to the entorhinal cortex (EC) which is the major input region for the hippocampus, retrosplenial cortex, and BLA[105,106]. Major inputs to NR are largely from cortical and neocortical regions, arising from reciprocal connections with likely driving inputs from mPFC and CA1/Sub, as well as with entorhinal, retrosplenial, orbitomedial, insular, and perirhinal cortices[107]. NR also receives inputs from a variety of subcortical regions, including locus coeruleus, the dorsal raphe, lateral septum, the basal forebrain nuclei, and of particular importance to this discussion, AMY. However, these subcortical inputs are sparser and weaker than cortical inputs. They likely serve a modulatory function to NR processing, unlike the function of primary thalamic nuclei, which are driven by their subcortical inputs[107].

The unique placement of NR as the primary thalamic input to CA1 and as a massive modulatory thalamic input to mPFC, as well as its reciprocal connectivity with many brain regions crucial for processing contextual emotional memories, quickly drew interest to its function[47]. The circuit of connectivity between mPFC and CA1 is notable in particular, as these regions need to synchronize for various complex cognitive functions but do not have a bidirectional connection. Instead, both PL and IL receive projections from CA1, but do not return them[47]. Given this lack of bidirectional connectivity with each other alongside the strong bidirectional connectivity of both mPFC and CA1 with NR, NR has been hypothesized to be the communication hub altering and regulating cross-regional mPFC-CA1 communication[47,108]. Consistent with this idea, lesioning or inactivating NR impairs higher-order cognitive behaviors that require both mPFC and CA1 to be online and working together, including but not limited to spatial navigation, task switching, spatial decision-making, goal-directed behavior, behavioral flexibility, working memory, and - crucially - contextual fear conditioning, but not non-spatial

Pavlovian fear conditioning[47,108–110]. Another supporting piece of evidence for NR's role as a mediator of mPFC-CA1 activity is its well-documented functions in synchronizing oscillatory activity critical for memory acquisition and consolidation, coupling delta activity between regions, occurring directly before both CA1 and mPFC beta in non-spatial memory tasks, phase-locking hippocampal theta during memory-dependent spatial tasks, and burst-firing prior to CA1-mPFC gamma waves in anesthetized rodents[111–114].

While behavioral investigations of NR activity are still in their relative infancy, perhaps the best-studied behavior that requires NR to be active and functional is contextual fear conditioning (CFC). The first evidence for NR involvement in CFC comes from Xu and Südof in 2013, where the authors showed that virally-expressing tetanus toxin in NR prevents context-dependent learning during CFC in rodents. While both acquisition of contextual fear and post-encoding freezing responses were unaffected, as well as tone-based Pavlovian fear conditioning, NR ablation massively increased generalization to a novel, unpaired context[97]. They also showed that constitutive optogenetic activation of NR reduced generalization[97]. A multitude of recent studies have confirmed and expanded NR's role in fear conditioning, with optogenetic inactivation of either mPFC-NR projections or whole NR increasing the time mice spend freezing after contextual fear encoding, impairing extinction learning, and inducing fear generalization. Activation has the inverse effect[115–119]. Online optogenetic activation or inactivation of whole NR or the NR-BLA during remote contextual fear memory retrieval also respectively shortened or lengthened freezing epochs[120]. Interestingly, while NR is critical for acquiring contextual fear extinction, it is not necessary for later consolidation of extinction memories[121].

The role of NR in fear response specificity and extinction is further supported by how NR dysfunction manifests in neurological disorders and diseases. NR dysfunction is suggested in Alzheimer's disease, anxiety, depression, schizophrenia, and post-traumatic stress disorder (PTSD), all disorders with dysregulated contextual emotion processing [47]. The human analog of NR, the interthalamic adhesion (massa intermedia), shows hypoactivity in PTSD patients, neurofibrillary tangles that increase with severity in Alzheimer's disease, decreased connectivity in major depressive disorder, and selective cell death during Korsokoff's syndrome that correlates with increased emotional dysregulation and memory confabulation [122–130]. Overall, much work remains before a full picture of NR will emerge, but existing research indicates NR is the critical intersection of information synchrony across the contextual fear encoding network.

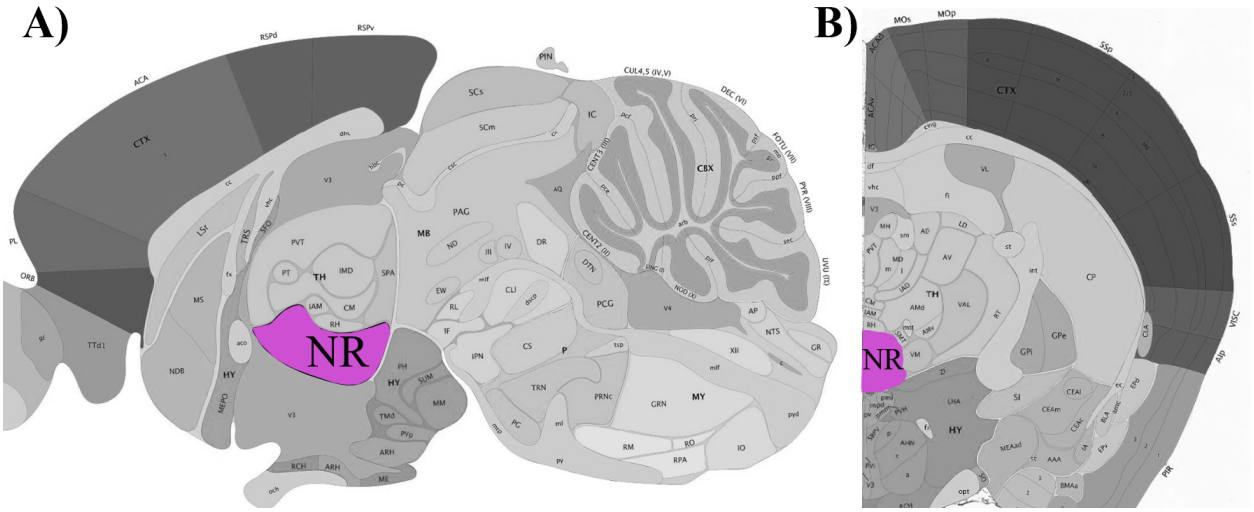

Figure 1.4 The Thalamic Nucleus Reuniens in Mice
A) Sagittal section of the mouse brain through the midline, with Nucleus reuniens (NR) highlighted purple. **B)** Coronal section of the anterior mouse brain with NR highlighted in purple. Adapted from the Allen Brain Atlas[131].

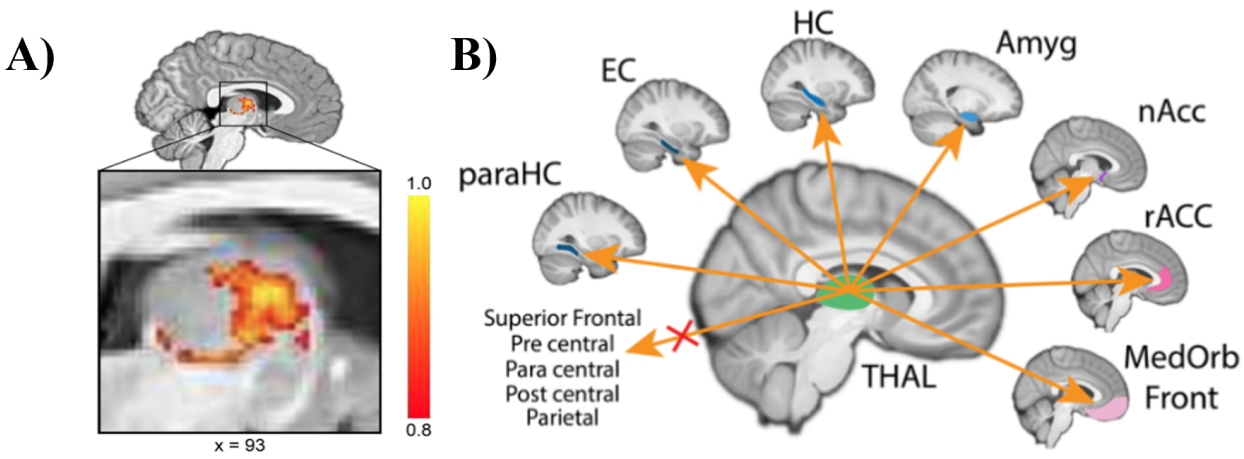

Figure 1.5: The Thalamic Nucleus Reuniens in Humans
A) Using K-means connectivity clustering, a region with similar connectivity to nucleus reuniens was identified in humans. When extracting the cluster with high connectivity to MTL, mPFC and nACC, and not control regions, a mask was obtained of the midline thalamus with over 80% overlap among all participants. **B)** Midline thalamus has high connectivity with mPFC, nACC and MTL, but little connectivity with other regions such as motor and somatosensory cortex (e.g., postcentral and precentral gyrus in the human brain. Adapted from Reeders 2023[130].

The Hippocampal Formation: Memories in Context

The hippocampal formation is one of the best-studied regions of the brain. Early investigations of the hippocampus came from the 'father of neuroscience' himself, Ramón y Cajal, who beautifully illustrated the largely self-contained significant connections within the hippocampal formation [132]. The hippocampus is an archicortical structure in the medial temporal lobe that is evolutionarily conserved in species as diverse as reptiles, birds, and mammals[133]. Unlike the classically six-layered neocortex, the evolutionarily-older mammalian hippocampus has only three layers. It receives its most impactful input from the medial entorhinal cortex (MEC) into the dentate gyrus (DG), which next travels through the cornus ammonis (CA) and finally outputs to other brain regions from the Sub. CA is further subdivided into CA3, which sends densely packed axons, i.e. the Schaffer collateral, to CA1. CA1 then sends information out via Sub[134]. A small zone, CA2, also exists between CA3 and CA1, which is important for

contextual social behavior encoding, but is not included in the traditional view of hippocampal information routing[135]. The general structure of the hippocampal information loop persists along the dorsal-ventral axis, although considerable heterogeneity exists between these regions[136]. Most research has been on excitatory neurons, either granule cells within DG or pyramidal neurons in CA/Sub, each forming a densely packed layer[137]. Quite unlike the neocortex, the hippocampus is a somewhat self-contained unidirectional 'loop,' a feature that has proved helpful in anatomic, physiologic, and behavioral investigations into the hippocampus.

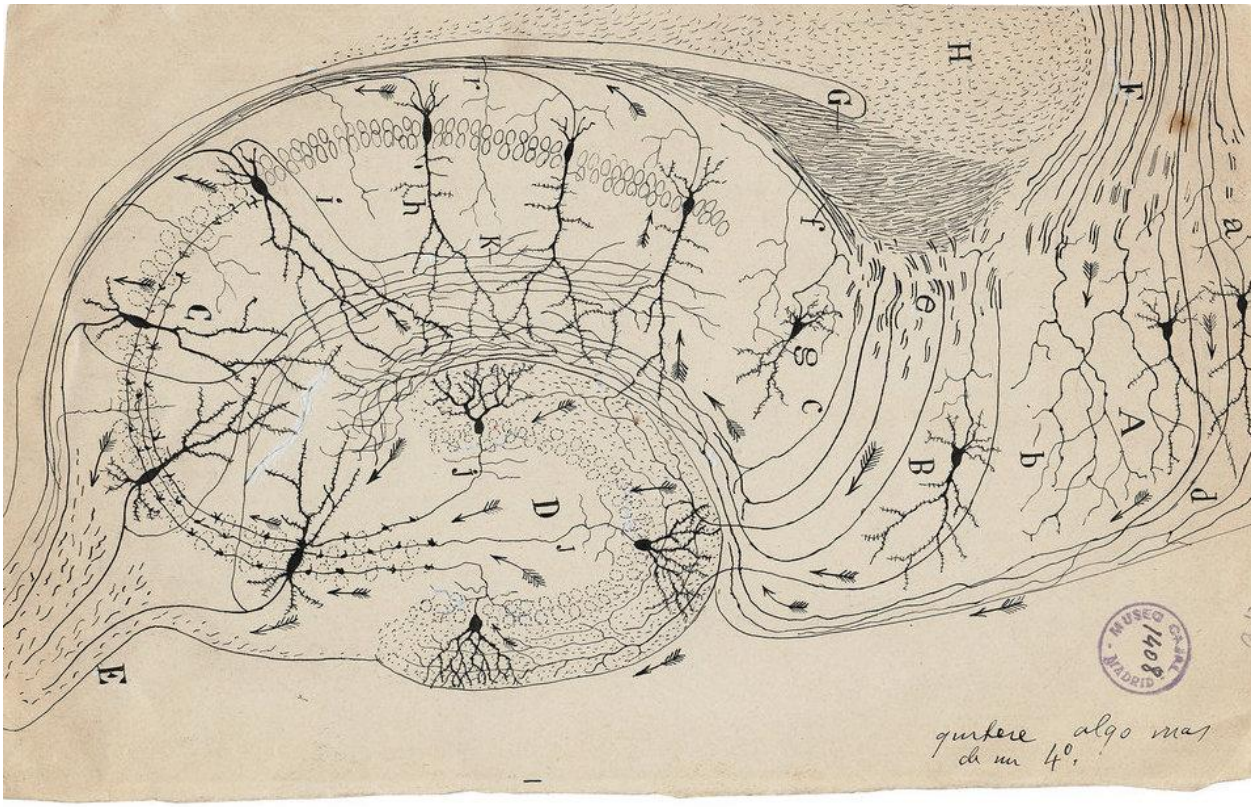

Figure 1.6: The Hippocampal Formation as Drawn by Ramón Y Cajal
An anatomical drawing of some major cell types, their connection, and the general anatomy in the hippocampus based on a silver stain. Arrows denote the at the time theorized direction of information transfer. Note especially drawings of pyramidal neurons labeled i, h, and k in CA1. Adapted from Cajal, 1911 [132]

Role in Place Encoding

The first evidence of HPC's role in memory came from patient H.M (Henry Molaison), whose severe temporal lobe epilepsy was 'treated' by the bilateral removal of a large segment of the medial temporal lobe - one containing the amygdala, some multimodal association cortical regions, and the entirety of the hippocampus - after which H.M. could no longer form new memories of events or learn declarative information[138]. Further research revealed the critical role of the hippocampus in encoding memories in spatial context with the discovery of place cells, hippocampal neurons that regularly fire at a specific location in a spatial context, which tile to form 'place maps'[139,140]. This work later shared the Nobel Prize in Medicine in 2014 with the discovery of space-encoding 'grid cells' in the medial entorhinal cortex[141]. These early hippocampal researchers argued that grid and place cells underlie the 'cognitive map,' memory representations that record routes, paths, and abstract maps of space first proposed as a cognitive model decades earlier by Tolman in 1948[142–144].

Substantial subsequent research has characterized the properties of hippocampal cognitive maps and, more broadly, the role of the hippocampus in spatial memory. Place maps rapidly form when mammals enter a novel context, stabilize over several minutes of exploration, and persist semi-stably across days to months[145–149]. Individual hippocampal neurons, as well as hippocampal populations, additionally encode a wide range of non-spatial cues relevant to contextual memory, including encoding the location of physical objects, sensory information such as sound, odor, and texture, and more abstract concepts like time, social valence, reward, reward expectation, and punishment[150–159]. Place cells also encode information about the past and future locations of both the subject themselves and of others. Tasks that require spatial

navigation based on previous experience rely critically on the hippocampus in ways simplistic spatial exploration does not[160].

Further reinforcing the mnemonic and predictive role of place cells is the phenomenon of hippocampal replay, the rapid reactivation of a sequence of past place cells in a trajectory[160]. First discovered during sleep and rest, replay has since been demonstrated as important during active behavioral tasks, especially those requiring animal decision-making[160–164]. Replays happen at about ~20x speed of actual experience and preferentially encode behaviorally relevant regions of a spatial environment. Replay events that occur while the subject is resting or sleeping, i.e. 'offline', are necessary for memory consolidation of new experiences[,163,165–169]. Not just the locale of the 'brain's GPS,' as the 2014 Nobel prize announcement asserts, the hippocampus seems to be broadly involved in encoding contextual representations crucial for episodic memory formation and consolidation.

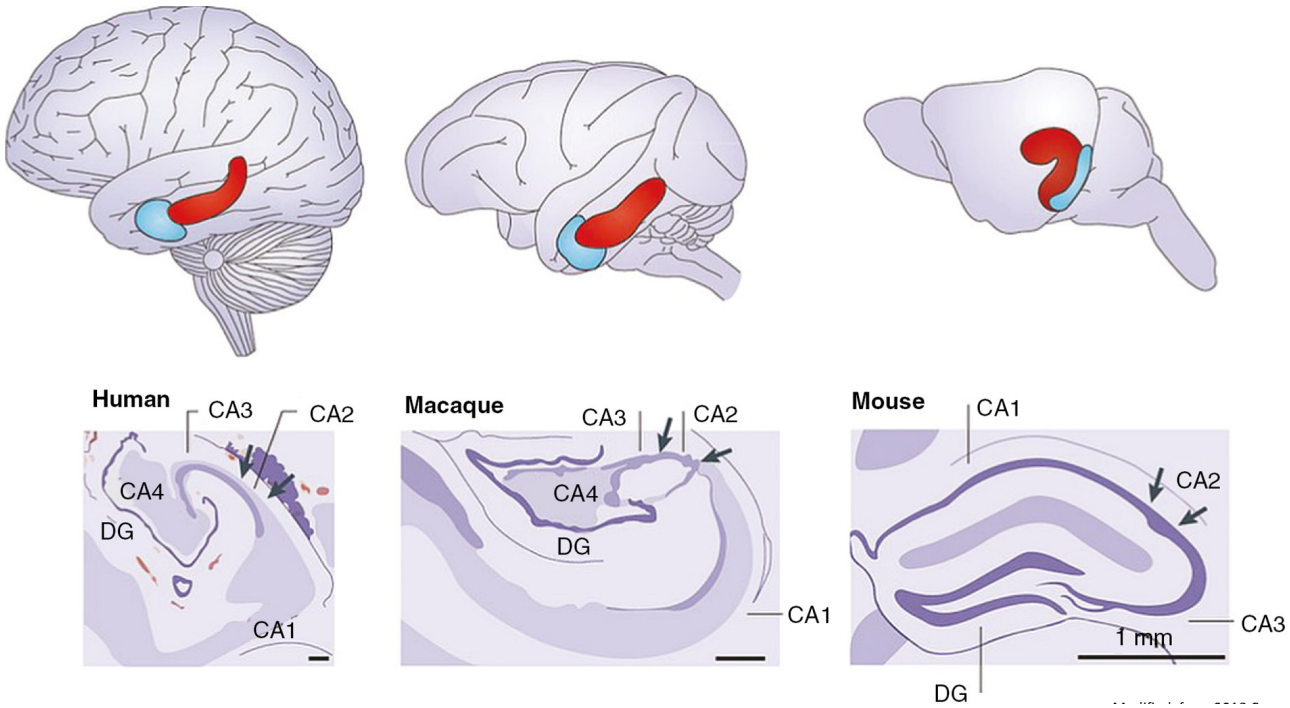

Figure 1.7: The Hippocampal Formation is conserved across species
Top: The hippocampus is shown in red and the entorhinal cortex shown in blue, outlined in the brains of mice, macaque monkeys and humans. Bottom: Subregions of the hippocampal formation across species are demonstrated in coronal sections. Figure is adapted from Strange et al., 2014 [170].

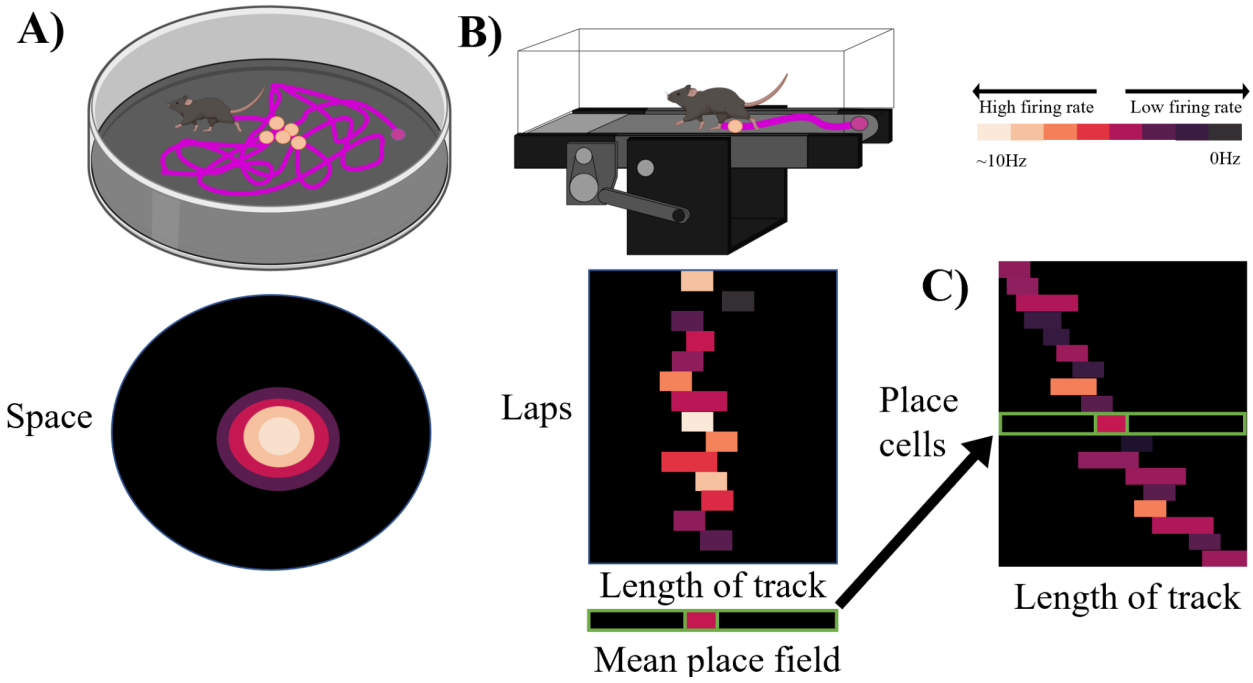

Figure 1.8: Place Cells and Maps
A) Schematic of O'Keefe et al. 1971 experiment. Top: rat navigated a 2D round open field freely while tetrodes recorded from hippocampal CA1. Their trajectory is shown in bright purple, firing of an example place cell in beige dots. Bottom: Schematic of an idealized average session of activity of a single neuron within CA1 with a stable place field, color-coded by Hz (see upper right corner for scale). **B)** Schematic of Harvey et al. 2009. Top: mouse running along a 1D virtual track. Their trajectory is shown in bright purple; the firing of an example place cell in beige dots. Bottom: Schematic of idealized lap-by-lap activity of a single neuron within CA1 with a stable place field color-coded by Hz. Below the lap-by-lap average is a mean place field for that session. **C)** Mean place field is joined with other place fields, showing place fields tiling the length of the track[140,171].

Role in Contextual Fear

Despite the 'contextual' core of contextual fear conditioning, extensive investigations into the hippocampus's role in CFC are largely removed from investigations of place maps. Lesioning or temporarily inactivating the hippocampus prevents the association of the contextual CS with the US[26,172–175]. Generally speaking, non-spatial Pavlovian fear associations do not need the hippocampus, while contextual fear associations do need hippocampal involvement in encoding, but there is some contention in the literature[28,172,173,176]. The point at which the hippocampus is

taken offline during a contextual fear conditioning paradigm impacts the resulting encoding deficit. Pre-training temporary inactivation or plasticity impairment will impair CS-US pairing, but permanent lesions with time for recovery yield no such impairment, indicating the existence of an alternate, non-hippocampal-dependent strategy for contextual CS-US pairing[28,177–183]. However, non-hippocampally encoded contextual CS-US fear memories show severe deficits in contextual discrimination, with significantly higher fear memory generalization to safe contexts[177,184]. The hippocampus is also critically involved in the extinction process, as inactivation prior to extinction learning prevents context-dependent fear memory extinction[174,185–187].

How, exactly, the hippocampus is encoding contextual fear has been a subject of great interest. Most research in this area has focused on the concept of hippocampal engrams - a sparse subset of neurons that are activated by an experience, modified by that experience, and then reactivated on recall[188]. The concept of engrams was first proposed by Richard Semon in the 1920s, with early behaviorist Karl Lashley then searching and failing to find evidence of engrams[189,190]. Confirmation of synaptic changes strengthening between co-active neurons gave evidence to Donald Hebb's 'fire together, wire together' doctrine and supported engram theory. H owever, it took the discovery of immediate early genes (IEGs) for true engrams to finally be identified[191–193]. IEGs, such as c-Fos, Arc, and Zif/268, are genes that only express after a neuron fires intensely[194–199].

By *post-hoc* staining or, later, by *in vivo* tagging and activating/silencing IEG-expressing neurons, researchers were able to find small subsets of hippocampal neurons - in the dozens to hundreds scale - that reliably activated in response to a specific context or task[200–203]. Subsequent optogenetic activation of these IEG-tagged neurons alone during CFC CS-US pairing can induce

contextual CSs without mice being in the conditioned context, and optogenetic inactivation of this population can cause early extinction [204–208]. 'False' context-specific fear memories can even be induced in DG and CA1 by artificially activating fear engrams from one context in another, unshocked safe context[205]. Curiously, the neurons that causally trigger freezing are found throughout the hippocampal formation, indicating a role for all hippocampal subregions in CFC encoding. Using the same techniques, separate ensembles for fear extinction learning were also identified in the hippocampus, indicating that typical contextual fear extinction learning exists as a semi-independent phenomenon from fear engram suppression[207].

Integrating Place and Memory

A key question is what relationship, if any, exists between these IEG-tagged hippocampal engrams and between hippocampal neurons that activate as place cells or in response to specific context clues. While these research avenues have largely existed independently, recent work has attempted to bridge the theoretical and experimental gap [209]. Researchers originally thought there would be a large overlap between a context-tuned place cell and an IEG-tagged engram cell. However, recent evidence has shown that these groups arise from at least partially different subpopulations, raising questions about using one method exclusively for identifying memory formation and storage neural mechanisms in the hippocampus[209–211]. While the hippocampal CFC literature has near-exclusively focused on IEG-tagging as the method of choice, a few recent studies have attempted to reconcile the place field (PF) literature with CFC engram findings.

Recordings *in vivo* have shown populations of synchronized co-active CA1 neurons emerge after contextual fear memory acquisition, activate during active navigation, and only partially are made up of place cells[212,213]. These non-place activated neurons identified in active navigation may overlap with IEG-tagged engram neurons found in previous work. Two recent

studies have also examined PF remapping and activity throughout CFC [159,214]. Kinsky et al 2023 found that lower levels of overlap between a neutral and shocked context in active dorsal CA1 pyramidal neurons of freely-behaving mice that had undergone CFC correlated with better context-selective CFC CRs, and that mice that associated the CR-US pairing without generalizing demonstrated higher levels of PF remapping and co-varying freezing-tuned neural activity. Arresting protein synthesis impaired this remapping, prevented context-selective CRs, and decreased co-variance of shock-selective neurons[214]. Similarly, Blair et al. 2023 showed that CA1 place cells in freely-behaving rats remapped following avoidance learning, and preferentially remapped when aversive experiences were remembered rather than artificially forgotten[159]. These results suggest that place cells, co-active ensembles, and engrams all play important, semi-overlapping roles in the encoding, expression, and extinction of contextual fear. However, much unknown information remains to untangle the relationships among these memory mechanisms.

Hippocampal CA1: What makes it special?

While engram and place cell analysis have been conducted across the hippocampus, the place cell literature has largely focused on the dorsal section of the hippocampal region CA1 (dCA1). Some of this preferential focus is attributable to the technical ease of recording and manipulating dCA1 in rodents compared to other subregions, with place cells first identified in dCA1[140]. Another region for the concentrated interest is that CA1, as well as Sub, serves as the primary output region of the hippocampus. In contrast, DG and CA3 serve as important pre-processing structures, but are less important for communication with the rest of the brain[134].

While PFs are found across the hippocampus, differences have emerged in how these PFs are formed, their stability, and how they encode the environment, with an important caveat that

PF research in CA1 is massively more developed than in DG or CA3. DG is largely considered a pattern separation region, which helps to form orthogonal representations of different contexts[215–220]. It contains two major cell types: granule cells that fire sparsely - with fewer than 10% forming singular PFs in any given spatial context - and mossy cells that fire more commonly in multiple locations across a context[219]. One study that tracked the same granule cells across days found DG PFs to be highly stable[221]. Place cells in CA3 are very common, at ~50% of the active pyramidal population, and have been shown to shift backward with experience, emerge relatively late during exposure to an environment with a low level of 'instant' forming place cells, and show medium levels of remapping[221–224].

Meanwhile, in CA1, place cells represent ~30% of active pyramidal neurons, but emerge instantly, 'shift' their representation backward, and experience representational drift - partial remapping in familiar contexts - far more often than place cells in CA3 or DG[221,223]. Pyramidal neurons in CA1 are especially dynamic in novel environments and have a higher frequency of co-active ensemble events during CFC than in CA3[213,223]. The instability of PF encoding, especially in CA1, has been a topic of much discussion, as one might think that the more stable a place map is, the moer stable a memory it forms of a context.

However, one major clue as to why this instability might be a feature and not a bug comes from the unique inputs to hippocampal CA1. Unlike CA3, which is largely self-contained to glutamatergic inputs from itself, contralateral CA3, and DG, CA1 receives not only contralateral CA1 and the major Schaffer collateral input from CA3 to its basal and somatic dendrites, but also receives major glutamatergic input from medial entorhinal cortex, lateral entorhinal cortex, and nucleus reuniens to its apical tuft dendrites, with ventral CA1 additionally receiving inputs from the BLA and olfactory regions[225,226]. In concert with various

neuromodulatory inputs, these additional major glutamatergic inputs to the apical tuft may serve as a significant modulator on valence, making CA1 the ideal place to form representations of especially salient novel, rewarding, or aversive experiences.

Dendrites in CA1 Computation

While somatic firing has been the best-studied principle of individual neurons, it is the subcellular input region of dendrites that enables individual neurons to massively multiply their computational power through spatial and electrical compartmentalization, leading some to propose dendritic segments as the foundational computational unit of the brain[227]. Dendrites contain multiple ion channels that produce different types of dendritic spikes (dSpikes) that vary from a millisecond to hundreds of milliseconds in duration, with the most influential being sodium, NMDA, and calcium (Ca^{2+}) spikes[228–231]. Dendrites possess both passive cable properties that enable attenuation of the signal based on distance from the soma and active signal boosting or attenuating processes that can lead to super and sublinear signal transforms, effectively modulating the importance of inputs to the somatic decision to fire[228,229,232–239] After firing takes place, information about that decision is then 'back-propagated' to the dendritic branches, imbuing individual dendrites with the knowledge of if their contributions were relevant to the spike, increasing the computational capacity of single neurons to include coincidence detection, gain modulation, and memory formation[231,233,240–242].

At a synaptic level, evidence for long-term potentiation and depression, synaptic plasticity mechanisms essential for memory formation, modification, and storage, was discovered initially in the Schaffer collateral CA3-CA1 connection, and NMDA-dependent modulation of the postsynaptic, i.e., dendritic, side is critical for a wide variety of learning

mechanisms[192,243–245]. At a circuitry level, many pathways in the brain are specifically oriented to synapse at precise points of the dendritic arbor, perhaps most strikingly in the archicortical six-layered neocortex and three-layered hippocampus, which precisely orient their pyramidal neurons to receive specified inputs in different layers[231]. Despite some species variation, the dendritic arbor of excitatory pyramidal neurons spans all layers of a hippocampal region and is largely made of three categories: basal dendrites branching out dorsally from the base of the soma, a long trunk dendrite extending ventrally that branches repeatedly into intermediate apical dendrites, and the distal apical tuft dendrites branching laterally from the trunk endpoints[246]. Single pyramidal neurons receive over 10,000 inputs along this dendritic arbor and often receive extraordinarily different inputs to its apical tuft region than to its basal[246,247].

The properties of dendritic events in dorsal CA1 pyramidal neurons have long been a subject of great interest, partially due to the region's importance in memory formation. Early evidence showed robust dendritic regenerative, i.e., supralinear, Ca^{2+} dSpikes in CA1 dendrites and spontaneous shorter sodium spikes[248–250]. Interestingly, early dCA1 dSpikes *in vivo* found evidence for Ca^{2+} dSpikes in the apical tuft to coincide with sharp-wave ripples indicative of hippocampal replay, indicating that they were contributing to the greater memory functions of the hippocampal network[250]. Inducing Ca^{2+} dSpikes in these distal dendrites *in vitro* can trigger robust burst spiking at the soma[251–253]. In particular, the dendritic-mediated integration of information from the SLM, where CA1 distal apical tuft dendrites receive modulatory inputs from a variety of brain regions, and the Schaffer collateral inputs from CA3, which largely carry spatial PF information, may hold the key to understanding the unique role CA1 plays in contextual memory.

To link dendritic activity to contextual memory requires *in vivo* investigations of hippocampal dendritic activity during navigation and learning behaviors, which until recently proved technically challenging due to the tiny 1-3uM wide structures of dendritic segments. The first evidence for *in vivo* branch spiking events cleverly used somatic plateaus to infer dSpikes in the arbor and found evidence that these dSpikes may be the foundation for PF formation[254]. The advent of high-quality Ca^{2+} sensors enabled directly *in vivo* dendritic recordings during navigation of familiar and novel environments in virtual reality, showing the critical role of basal dendritic spiking in PF formation[147]. Despite the importance of distal apical tuft dendritic spiking to coincidence detection and learning *in vivo* in neocortex, and the critical nature of distal Ca^{2+} dSpikes for CA1 burst to fire *in vitro,* to our knowledge, no investigation of these dendrites has been carried out in dCA1 *in vivo*, likely due to technical difficulties accessing these deep segments[255–260]. How spatial and non-spatial information is processed by and reflected in these distal apical tuft dendrites is of great interest to understanding the dendritic computational mechanisms underlying contextual memory.

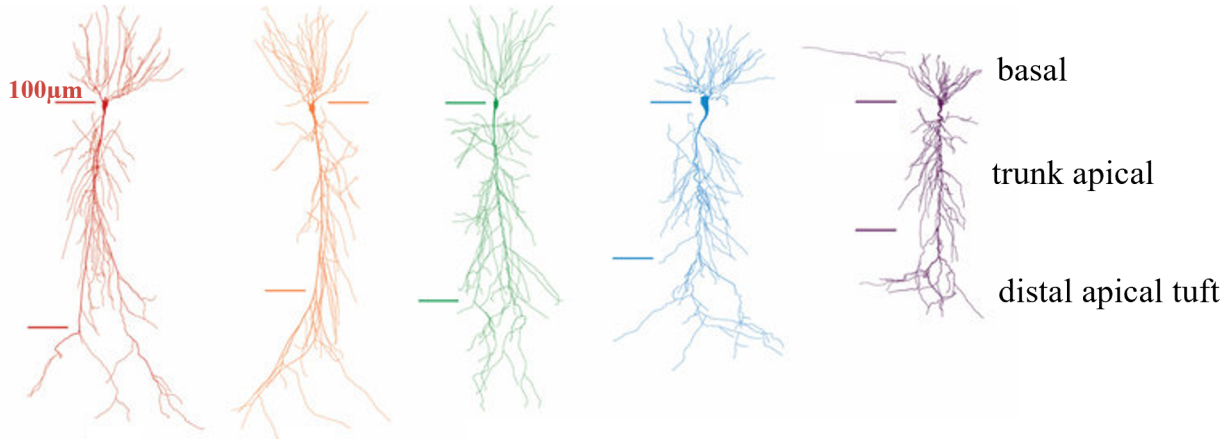

Figure 1.9: Dendritic Arbors of Mouse CA1 Pyramidal Neurons
Different colors denote different pyramidal neurons traced from mouse CA1 hippocampus. Lines indicate major sectioning of the arbor, the basal versus apical distinction at the soma (top), and the trunk apical versus the distal apical tuft distinction in the apical arbor (bottom). Adapted from Basak et al., 2020[261].

In Summary: Interconnected Circuits for Contextual Fear Memory

This section has discussed contextual fear, highlighting important brain regions necessary for contextual fear memory, including the amygdala, prefrontal cortex, nucleus reuniens of the thalamus, and the hippocampal formation with a particular focus on hippocampal CA1 and its dendritic computations. Of course, no brain region exists in a vacuum, and the fear regulation network is no exception. Because of the highly plastic nature of the brain, fear conditioning can activate, strengthen, or weaken the connections among these brain regions. Fear conditioning will strengthen the connection between the prelimbic cortex and amygdala, a projection necessary for fear expression[262–265]. In contrast, the connection from IL to the amygdala is essential for encoding fear extinction[76,266–268]. Evidence also supports that strengthening of the reciprocal pathway from the amygdala back to mPFC occurs after fear conditioning[269]. The IL additionally has a well-documented projection to GABAergic inhibitory neurons in CeA that directly inactivate BLA neurons involved in fear memory retrieval[262,270–272]. Adding to the contextual component of contextual fear conditioning is information from the hippocampus, as CA1 and Sub directly project strongly to both PL and IL with CA1-PL inputs shown to be critical to prevent fear memory generalization between contexts[273,274]. Hippocampal projections to the amygdala have a different role, one that is seemingly necessary to renew previously extinguished fear memories [61,185,275–277].

The lack of a return pathway from IL or PL to CA1 is largely ameliorated by the indirect pathway through NR, in which individual neurons receiving IL or PL driving input directly project into CA1[83]. The mPFC-NR pathway itself is critical for fear extinction learning and preventing fear memory generalization, while the NR-CA1 pathway of the circuit has not been elucidated *in vivo*[278,279]. *In vitro,* this pathway seems well-positioned to exert modulatory effects

on CA1 during CFC, as it induces strong excitatory depolarizations on pyramidal dendrites in the distal apical tuft while inducing robust firing in interneuron populations, positioning NR inputs to CA1 to modulate both feed-forward excitation and inhibition of CA1 activity[95,100]. The existence of co-projecting NR neurons to both CA1 and mPFC simultaneously suggest that this collateral pathway may be critical to regional synchrony, and further work on this subpopulation may yield insights into their role in CFC[60,103,104]. Another NR-mediated parallel pathway in the fear inhibition circuitry comes from the IL to NR to BLA pathway, each component critical for extinguishing remote fear memories[120]. Interestingly, while NR projects to both dorsal and ventral CA1 and Sub - albeit with stronger projections to vCA1 than dCA1 - the BLA seems to exclusively project to vCA1, which is differentially heavily-wired with other subcortical emotion processing regions[54,226].

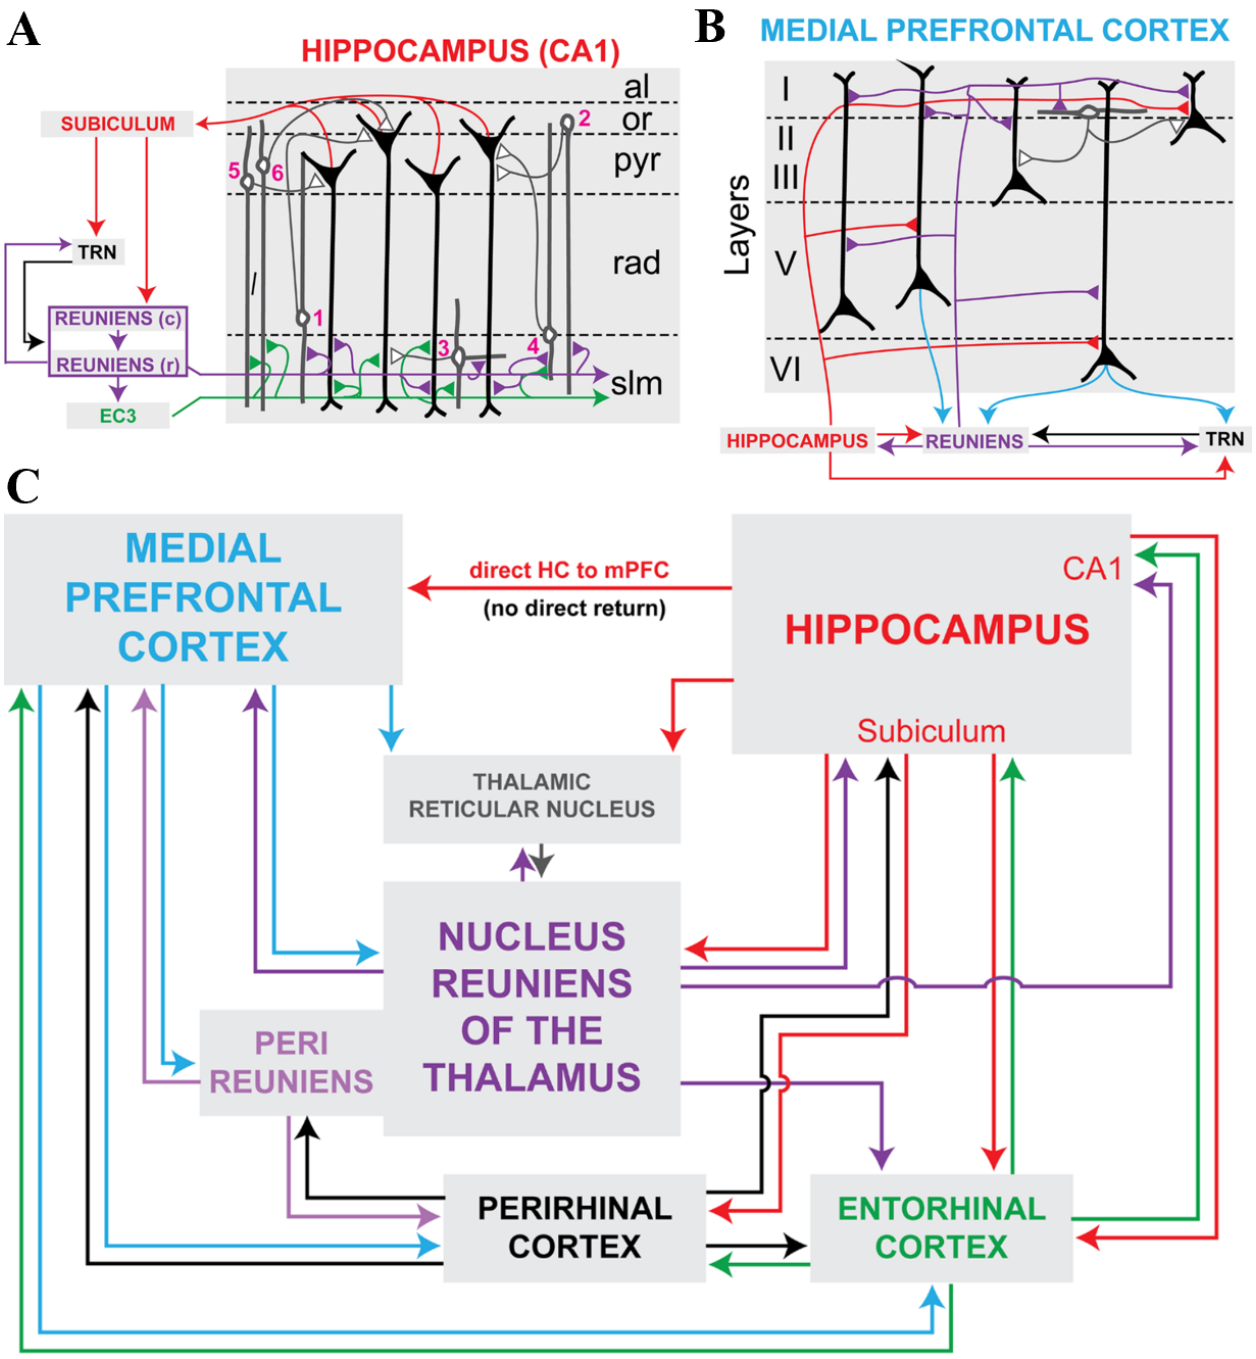

Figure 1.10: The mPFC-RE-HC System

A) Excitatory RE (purple lines) and EC (green lines) inputs in CA1, and a closed CA1→ subiculum → RE → CA1 circuit, including the thalamic reticular nucleus (TRN). The direct RE → CA1 input in SLM originates from the rostral part of RE; only a minor contribution arises from the caudal part. The output of CA1 (red lines) via the subiculum can be relayed back to caudal RE. In turn, caudal RE projects to rostral RE, which results in a di-synaptic (cRE → rRE → CA1) input in slm, thereby closing the loop. The direct EC → CA1 pathway, arising from EC layer III cells, overlaps with the RE input in SLM. RE → EC input may have the ability to modulate the activity level of EC layer III cells. Electrophysiological data support the view that

RE and EC inputs converge (at least partly) onto the same dendritic branch of a pyramidal cell in proximal SLM. RE has been shown to drive presumed Schaffer collaterals-associated cells (1) which are thought to inhibit pyramidal cells and other (unidentified) interneurons, and vertical oriens/alveus cells (2), mediating feedforward perisomatic inhibition of CA1 cells. EC is assumed to drive parvalbumin- positive basket cells (5) and chandelier or axo-axonic cells (6) providing feedforward somatic and axonal inhibition, respectively, of pyramidal cells; EC also drives neurogliaform cells in SLM (3). RE and EC inputs are proposed to converge on slm-neurogliaform cells (3) providing feedforward inhibition of pyramidal cells and other interneurons in SLM, and on a subclass of (presumed CCK/VIP) basket cells (4) located at the slm/radiatum border, providing (peri)somatic inhibition of CA1 cells. **B)** mPFC layers 1 and 5/6 contain a high density of RE fibers onto mostly unidentified pyramidal cells and interneurons. They overlap with fibers from CA1/subiculum (red lines); it is not known whether RE and CA1/subiculum inputs converge onto single neurons in mPFC. RE input drives presumed neurogliaform cells in layer 1, providing feedforward inhibition of pyramidal cells in layers 2/3 and other unidentified layer 1 interneurons. mPFC (light blue lines) provides direct excitatory input to RE as well as an indirect, pre-sumed inhibitory input via the TRN (black arrow); RE reciprocates the TRN input (purple arrow). **C)** Schematic representation of the connectivity in the mPFC–RE– HC system. The HC projects directly to the mPFC, but there are no direct return projections. RE (including periRE) thus serves as the main link completing an HC–mPFC loop as follows: HC → mPFC→ RE→ HC. Abbreviations: (al) alveus, (c) caudal, (EC3) entorhinal cortex layer III, (or) stratum oriens, (pyr) pyramidal cell layer, (r) rostral, (rad) stratum radiatum, (SLM) stratum lacunosum moleculare,(TRN) thalamic reticular nucleus (HC) hippocampus, (mPFC) medial prefrontal cortex, (periRE) peri-reuniens nucleus, (RE) the nucleus reuniens of the thalamus. Adapted from Dolleman-van der Weel et al., 2019[47].

While further research is critical on all aspects of this interconnected circuitry, I will focus much of this dissertation on further understanding the role of the NR-CA1 pathway in mediating contextual fear conditioning. Whether this pathway is involved in aversive learning, what information it sends to CA1 apical tuft dendrites, how this activity is reflected in the dendritic arbor, and how that translates to changes in CA1 population activity are all direct areas of investigation. In Chapter 2, I discuss the role of the NR-CA1 pathway in contextual fear memory suppression, extinction, and discrimination by using modern pathway-selective chemogenetic techniques to inactivate, or using 2-photon imaging to record from, the NR-CA1 pathway in dCA1 during a novel virtual reality contextual fear conditioning paradigm (VR-CFC). In Chapter 3, I discuss the role of the NR-CA1 pathway in modulating dCA1

excitatory pyramidal neuron population dynamics throughout CFC by imaging these populations across days with NR-CA1 intact, with NR-CA1 inactivation and without induction of contextual fear through shock delivery. In Chapter 4, I discuss the differential role of the dCA1 distal apical tuft dendrites, which to our knowledge have never been recorded from *in vivo* before, and dCA1 basal dendrites in encoding both neutral and aversive contextual information in a novelty and reward manipulation task. Altogether, this dissertation contributes significantly to our understanding of the role of dorsal CA1 hippocampal encoding of aversive experiences, with a significant focus on the understudied role of the NR-CA1 pathway.

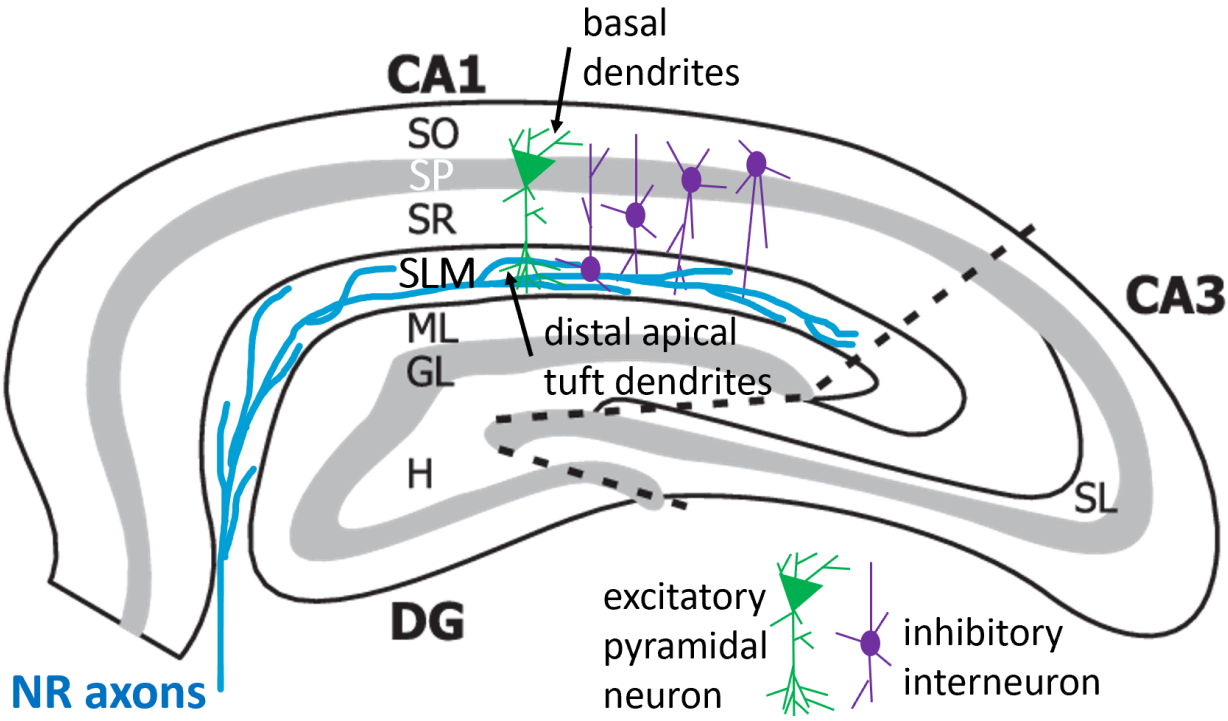

Figure 1.11: NR-axonal circuit interactions in CA1
Schematic representation of NR axonal input (blue) to the SLM of hippocampal CA1. Representative excitatory pyramidal neuron (green) and inhibitory interneurons with dendrites in SLM (purple) are shown. NR axons putatively synapse on all neurons with dendritic processes in SLM. Abbreviations: (CA1) Cornus Ammonis 1, (CA3) Cornus Ammonis 3, (DG) Dentate Gyrus, (NR) Nucleus Reuniens, (SO) Stratum Oriens, (SP) stratum pyramidale, (SR) stratum radiatum, (SLM) stratum lacunosum moleculare, (ML) molecular level, (GL) granular level, (H) hilus, (SL) stratum lucidum. Adapted from Ludkiewicz et al., 2002[280].

CHAPTER 2

A THALAMIC-HIPPOCAMPAL CA1 SIGNAL FOR CONTEXTUAL FEAR MEMORY SUPPRESSION, EXTINCTION, AND DISCRIMINATION

This chapter is reprinted with minor modifications for dissertation formatting continuity from **"A Thalamic-Hippocampal CA1 Signal for Contextual Fear Memory Suppression, Extinction, and Discrimination "** (Heather C. Ratigan, Seetha Krishnan, Shai Smith, Mark E.J. Sheffield, Accepted in Principle at Nature Communications, 2023), in which I was first author. The work is included with permission from all authors.

Abstract

Adaptive memory of fearful experiences is a conserved survival skill. To prevent excessive or inappropriate fear expression, fear memories must be suppressed and updated. If this does not occur, it can result in post-traumatic-stress-disorder or other anxiety disorders. Fear memories can be acquired through contextual fear conditioning (CFC) which relies on the hippocampus. The thalamic subregion nucleus reuniens (NR) is necessary for the extinction of contextual fear, and projects to hippocampal CA1. However, the NR-CA1 pathway has not been investigated during behavior, leaving its role in contextual fear memory unknown. We implement a novel head-restrained virtual reality (VR) CFC paradigm, demonstrating that male mice can acquire and extinguish context-dependent fear responses in virtual contexts. We show that inactivation of the NR-CA1 pathway following CFC increases fear responses. This was observed through longer freezing epochs, fear generalization to other contexts, and delayed fear extinction. Using *in vivo* sub-cellular imaging, we recorded from NR-axons innervating CA1 before and after VR-CFC. We found NR-axons become selectively tuned to freezing only after VR-CFC,

and this activity was well-predicted by an encoding model. We conclude that the NR-CA1 pathway actively suppresses fear by disrupting contextual fear memory retrieval in CA1 during fearful freezing behavior - a process that also limits fear generalization and promotes extinction.

Introduction

Flexibly encoding and retrieving memories of fearful events is a critically conserved survival behavior, as a single failure can be deadly. However, failing to suppress inappropriate fear responses can also have devastating consequences, manifesting as negative affective states in generalized anxiety disorder and post-traumatic stress disorder[281,282]. One way in which fear memories can be studied in the laboratory is through contextual fear conditioning (CFC), in which a spatial context, the conditioned stimulus (CS), is repeatedly paired with a noxious unconditioned stimulus (US), generally a mild shock[25,26,2831]. Freezing is a species-specific fear response, and a quantifiable readout of contextual fear memory retrieval (CFMR) of the learned association[11,284]. With continued exposure to the CS in the absence of the US, freezing generally decreases and exploratory behavior increases - a process termed fear extinction. Fear extinction occurs as animals learn that the context no longer predicts shocks[285–288].

Contextual fear in both mice and humans relies on coordinated brain regions including the medial prefrontal cortex (mPFC), thalamus, amygdala, and hippocampus[289]. The contextual component of these memories relies on the hippocampus, which retrieves and updates contextual fear memories[1,177,198,209,221,290,291]. Following CFC, experimental inhibition of a subset of hippocampal neurons tagged using immediate early genes active during CFC is sufficient to suppress CFMR[204,211,292]. This suggests that natural suppression of ongoing CFMR must involve a circuit that can modulate hippocampal activity. One potential source of this modulation is the

ventral midline thalamic subregion, nucleus reuniens (NR). Sometimes termed 'limbic thalamus' for its diverse set of inputs from limbic-related regions in the brainstem, hypothalamus, amygdala, basal forebrain, mPFC, entorhinal cortex (EC), and hippocampal subregion CA1, NR sits at the nexus of emotional regulation and serves as a major communication hub among these limbic-activated areas[47,59,102,107,108,293]. While mPFC does not have a direct excitatory projection to CA1, it does send a strong excitatory projection to NR[102,107].

The mPFC-NR projection and NR itself is necessary for both fear extinction and for preventing fear generalization to a neutral context, a process in which animals associate a non-shocked context with fear[97,115–118,279,294]. NR stimulation reduces contextual fear-induced immediate early gene expression in both mPFC and CA1[278,279]. While the roles of the mPFC-NR pathway and NR itself have been explored during CFMR, the role of the NR-CA1 pathway is unknown. We hypothesize that NR transmits a signal to CA1 to suppress ongoing CFMR, thereby reducing fear responses (freezing) and promoting exploratory behavior (movement).

To test our hypothesis, we used a chemogenetic approach to directly inhibit the NR-CA1 pathway, and 2-photon calcium imaging in head-restrained male mice to record NR-axons in CA1, before and after CFC. While CFC induction in head-restrained mice in virtual reality (VR) has been attempted, none to our knowledge have replicated the characteristic 'freezing' behavior of freely-moving mice in real-world CFC[295,296]. We therefore developed a new VR-based CFC paradigm (VR-CFC), using a conductive fabric to deliver mild tail shocks that induces context-dependent freezing. By combining VR-CFC, targeted chemogenetic NR-CA1 inhibition, and 2-photon NR-axonal calcium imaging, we were able to determine the role of the NR-CA1 pathway in CFMR, generalization, and extinction.

Results

Contextual Fear Conditioning and Extinction in Virtual Contexts

To ensure mice were comfortable with the VR setup before shocks were delivered, we trained water-restricted mice to run in a VR context for water rewards until they reached ~4 traversals of the context per minute, as previously discussed[223,297]. To avoid confounds from the water reward during CFC, these trained mice were then introduced to two novel VR contexts without a water reward. At this stage, the custom-designed conductive tailcoat was fitted to their tails (Fig. 2.1a, Methods: Behavior). Mice then spent ~5 minutes in each novel VR context, which allowed them to habituate to running with the tailcoat (Fig. 2.1b; Methods: Behavior).

Mice in all experimental conditions that continued to meet the criterion for movement (i.e. >4 traversals per minute) were advanced to the next stage the following day (46/79 mice), where they were re-exposed to both novel contexts for ~5 minutes each. During this period, mice demonstrated low levels of spontaneous freezing behavior (Supplementary Fig. 2.1f). Mice were then administered 6 mild 0.6 mA tail shocks through the tailcoat for a duration of 1 s each, 20-26 s apart, (Fig. 2.1b: day 0). These shocks were delivered at pseudorandom locations throughout one of the contexts ('shocked'), but not the other ('control'; Fig. 2.1h). Mice responded to each tail shock with an abrupt stereotyped increase in running speed, a behavioral validation of successful shock delivery (Fig. 2.1c). To test for CFMR and subsequent fear memory extinction, mice were then re-exposed for ~5 minutes each, in a pseudorandom order, to both the shocked and control contexts while wearing the tailcoat for the following three retrieval days (Fig. 2.1b: day 1-3). Elevated freezing behavior on post-shock days was interpreted as an expression of CFMR (Methods: Behavioral Parameters - Freezing).

In the shocked context on retrieval day 1, mice (N = 20) froze 32.7 ± 9.4% (95% CI) more than the pre-shocks baseline day, a statistically significant increase, while mice froze in the control context on average 8.7 ± 4.6% more than the baseline day, a non-significant increase (Fig. 2.1e). Comparing across contexts, mice increased their freezing significantly more in the shocked context compared to the control context, showing context specificity in CFMR (Supplementary Fig. 2.1d). This remained true on retrieval day 2, as mice continued to freeze at significantly elevated levels above baseline in the shocked context and compared to the control context (Fig. 2.1e and Supplementary Fig. 2.1d). By the third day of retrieval, mice froze at comparable levels between contexts (Supplementary Fig. 2.1d), although remained slightly elevated above baseline in the shocked context (Fig. 2.1e).

While freezing levels differed between contexts and days of retrieval, freezing position was distributed evenly across all track locations in both contexts on all retrieval days. This shows that mice associated fear with the entire context, and not specific locations along the track or near specific VR objects (Fig. 2.1i). As an additional control, a separate group of mice went through the same process but were never shocked in either context. These mice had similar baseline freezing levels to the VR-CFC mice, but did not have any significant differences in freezing levels on subsequent days compared to baseline or compared across contexts (Fig. 2.1f; Supplementary Fig. 2.1f). The low level of spontaneous freezing epochs in both the non-shocked control group and the pre-shocked contexts (i.e. before the delivery of any shocks) could potentially be caused by the lack of water reinforcement, the presence of the tail coat itself, or a temporary disinterest in running, and provides a useful within-mouse comparison to post-shock fear-evoked freezing.

To further quantify freezing behavior, we measured the duration of each individual freezing event (freezing epoch) and found that freezing epochs were longer in the shocked context at an average of 5.9 s compared to 4.6 s in the control context per freezing epoch on retrieval day 1 (Fig. 2.1g). Freezing epochs also remained significantly longer on day 2 in the shocked compared to the control context, however, they became similar by day 3 (Supplementary. Fig. 2.2b-d), corresponding with trends in the total time spent freezing. We additionally examined if VR-CFC impacted non-freezing running behavior, and showed that average non-freezing velocity in either context compared to baseline or between contexts was not statistically impacted by VR-CFC, indicating that VR-CFC selectively impacts freezing behavior (Fig. 2.1j). Our results show that VR-CFC produces robust CFMR that can be measured via context-specific increases in freezing, which is extinguished following ~3 days of re-exposure to the shocked context in the absence of additional shocks.

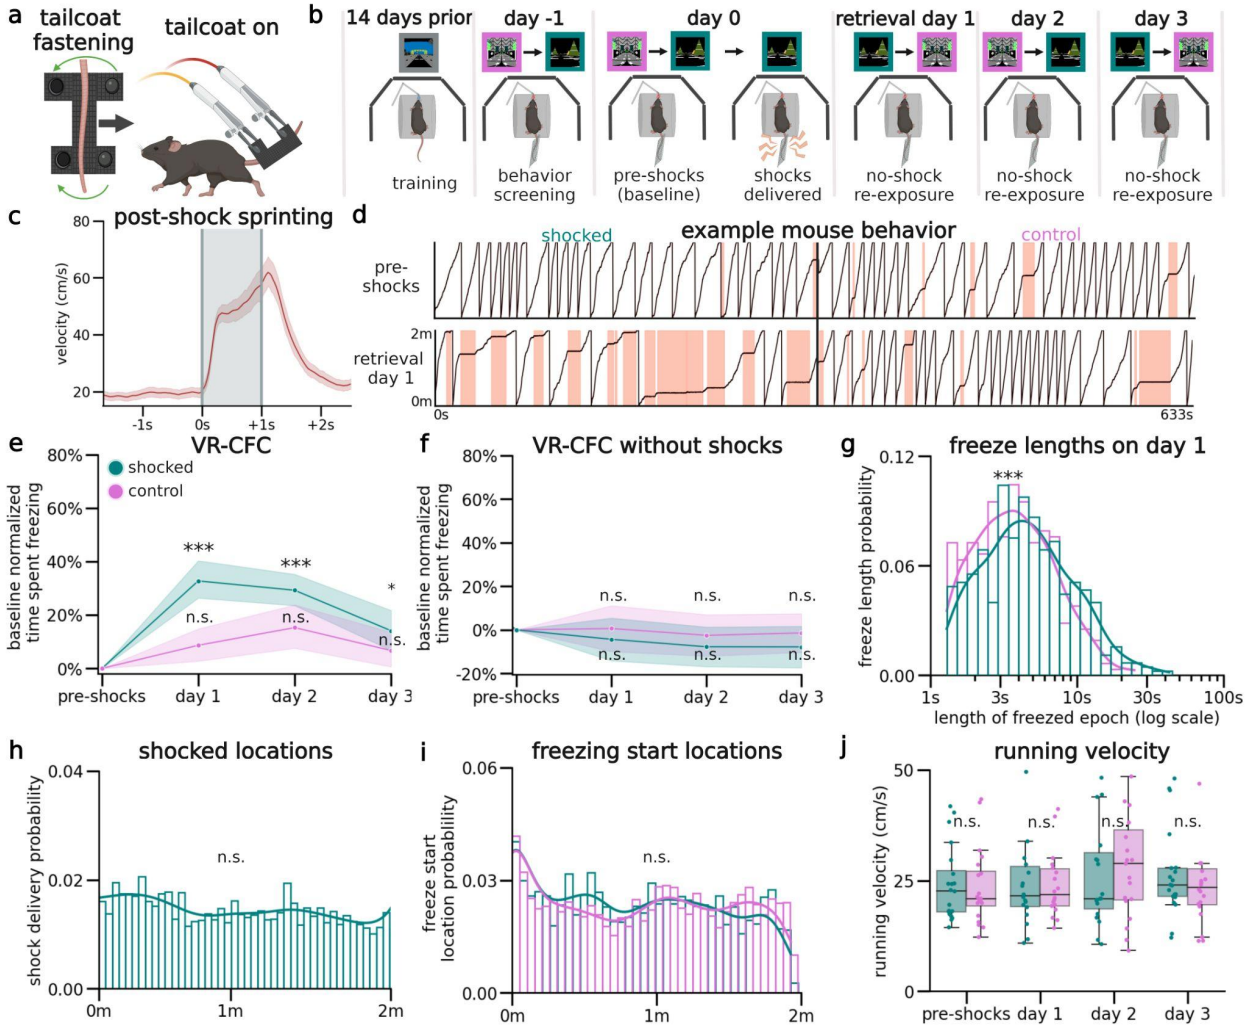

Figure 2.1. Virtual reality contextual fear conditioning induces fearful freezing responses

a, Schematic of tailcoat apparatus. Left: Tailcoat is custom-made out of conductive cloth with two metal snaps that secure the cloth around the tail. Right: The tail coat is wrapped around the mouse's tail and suspended using lightweight alligator clips forming a supportive 'hammock' structure connected to a device that generates electric shocks. The apparatus weighs less than 1.8 g, the weight of which is supported by the 'hammock' structure. **b,** After training, head-restrained mice were put in two unrewarded contexts. They received mild tail shocks in only one context (shocked), and not the other (control). They were re-exposed to both contexts in pseudo-random order for 3 retrieval days. **c,** Average animal velocity aligned to shock initiation and plotted across all shocks shows immediate sprinting behavior in response to the shocks (red line = mean, red shading = 95% CI). Grey shading indicates duration of shock. **d,** Example track position from a mouse on day 0 pre-shocks (Top), and on retrieval day 1 (Bottom), in the shocked (left) and control (right) contexts. Red shading indicates freezing epochs. **e,** Per mouse, we tested the difference in percent freezing from baseline pre-shocks within each context (N = 20 mice; CI = 95% shaded area). Shocked mice froze significantly more in the shocked context post-shocks

than pre-shocks on retrieval days 1-3 (Estimated Marginal Mean on ANOVA (EMM), P = day 1: 1.16e-8, day 2: 8.65e-8, day 3: 0.04). Shocked mice did not freeze more in the control context post-shocks (EMM, P = day 1: 0.76, day 2: 0.07, day 3: 1.) **f,** A separate group of mice that were not shocked on day 0 froze at baseline levels across all days in both conditions, (N = 7, EMM, Shocked context: P = 1). **g,** Kernel density estimates and density histogram of the length of individual freezing epochs show they were shorter in the control than the shocked context on day 1 (Mann-Whitney U, P = 1.20e-3). **h,** Density histogram of shock locations calculated from all mice show that shocks were administered evenly across the track (Kolmogorov-Smirnov, P = 0.91). **i,** Density histogram of locations where mice start individual freezing epochs across the virtual track indicate that mice freeze indistinguishably at all locations (Kolmogorov-Smirnov, P = 0.89), and at similar locations across contexts (Mann-Whitney U, P = 0.24). **j,** Mean running velocity in each mouse when not freezing (dots). Boxplot indicates median, 25-75th interquartile range, whiskers include all data points not determined to be outliers. We did not observe any differences between contexts across days (Student's T, P = pre-shocks: 0.60, day 1: 0.60, day 2: 0.30, day 3: 0.57).

Inhibition of the NR-CA1 Pathway during CFMR

To test the involvement of NR-CA1 projecting neurons in CFMR, we developed a designer receptor exclusively activated by designer drugs (DREADD) based inhibition paradigm[298–300] (Fig. 2.2a). We injected a Cre-expressing virus bilaterally in NR, and a retrograde Cre-dependent virus carrying the inhibitory G(i)-coupled DREADD receptor, hM4Di-DREADD, bilaterally in the SLM of dorsal CA1 where hippocampal-projecting NR-axons terminate. This led to the expression of DREADDs selectively in NR neurons projecting to CA1 and can be seen in NR axons in SLM of CA1 (Fig. 2.2b)[54]. This enabled us to intraperitoneally (IP) inject the hM4Di agonist, deschloroclozapine dihydrochloride (DCZ), before the first post-shock re-exposure to the contexts on retrieval day 1, therefore selectively inhibiting a subset of NR-CA1 projecting neurons during CFMR. DCZ was selected over CNO for its comparatively higher DREADD-selective binding at lower doses, reduced off-target effects due to reduced conversion into clozapine, and rapid onset kinetics.[300]

To ensure our injection paradigm and administration of DCZ did not alter context-dependent fear behavior, we had two DREADD-control groups. One group (N = 4) expressed mCherry in place of hM4Di, and received DCZ on retrieval day 1. A separate group (N = 4) expressed the hM4Di receptor, and received saline instead of DCZ on retrieval day 1 (Supplementary Fig. 2.1g). In both control groups, freezing behavior was similar to the experimental mice shown in Fig. 2.1e, and the groups were thus combined with other NR-CA1 uninhibited mice and termed the NR-CA1 intact group for analysis. We then examined the behavioral impact of inhibiting the NR-CA1 pathway on day 1, and on subsequent retrieval days 2 and 3 (Fig. 2.2e).

We found that in the shocked context on retrieval day 1, NR-CA1 inhibited mice (N = 9) spent 58.5 ± 10.4% more time freezing on day 1 under the influence of DCZ in the shocked context when compared to their pre-shock baseline freezing levels, with every mouse increasing their freezing level in amounts ranging from a minimum of 41.3% to a maximum of 84.6 (Fig. 2.2d). Freezing in the control context also significantly increased from baseline to 31.0 ± 7.9%. under the influence of DCZ (Fig. 2.2e). Freezing levels on day 2 remained significantly elevated in the shocked context on day 2 in the absence of DCZ at 48.3 ± 9.0%, while freezing in the control context fell to baseline levels at 14.8 ± 11.1% (Fig. 2.2e). Freezing levels on day 3 remained slightly elevated in the shocked context, at 23.4 ± 9.1% while remaining not significantly different from baseline in the control context at 8.44 ± 8.5%. Overall, NR-CA1 inhibition induced large increases in freezing levels in both contexts, which persisted beyond the DCZ delivery day in the shocked context.

NR-CA1 Inhibition Elevates Freezing, Reduces Context Discrimination, and Delays Extinction

To further test the impact of NR-CA1 inhibition, we directly compared freezing behavior in NR-CA1 intact mice to NR-CA1 inhibited mice (Fig. 2.3). We found that NR-CA1 inhibited mice under DCZ froze ~78% more in the shocked context (Fig. 2.3a) and ~256% more in the control context (Fig. 2.3b) compared to NR-CA1 intact mice on day 1. This elevated freezing response remained true for the shocked context on day 2, with a ~65% increase compared to intact mice, even though DCZ was no longer inhibiting the NR-CA1 pathway (Fig. 2.3a). Thus, inhibiting the NR-CA1 pathway on day 1 reduces the amount of extinction on day 2. However, in the control context on day 2, freezing levels returned to baseline and were not different between the intact and inhibited groups (Fig. 2.3b). By day 3, freezing levels in both the the shocked and control contexts in NR-CA1 inhibited mice were not statistically different from their NR-CA1 intact counterparts (Fig. 2.3a and b).

We next compared the freeze length distributions between NR-CA1 inhibited and NR-CA1 intact groups to see if an increase in average freeze length was driving the observed increased time spent freezing (Fig. 2.3c-d). Indeed this was the case, as average freeze lengths nearly tripled in NR-CA1 inhibited mice, increasing 173% from 5.9 s to 16.2 s in the shocked context on day 1 (Fig. 2.3c). A similar increase was observed in the control context on day 1, with average freeze lengths increasing 156% from 4.6 s to 11.8 s in NR-CA1 inhibited mice. These findings suggest that inhibiting the NR-CA1 pathway during CFMR increases the time spent in an ongoing fearful state of freezing by the lengthening individual freezing epochs. This suggests that the intact role of the NR-CA1 pathway suppresses CFMR in both appropriate (shocked) and inappropriate (unshocked control) contexts, to reduce fearful freezing.

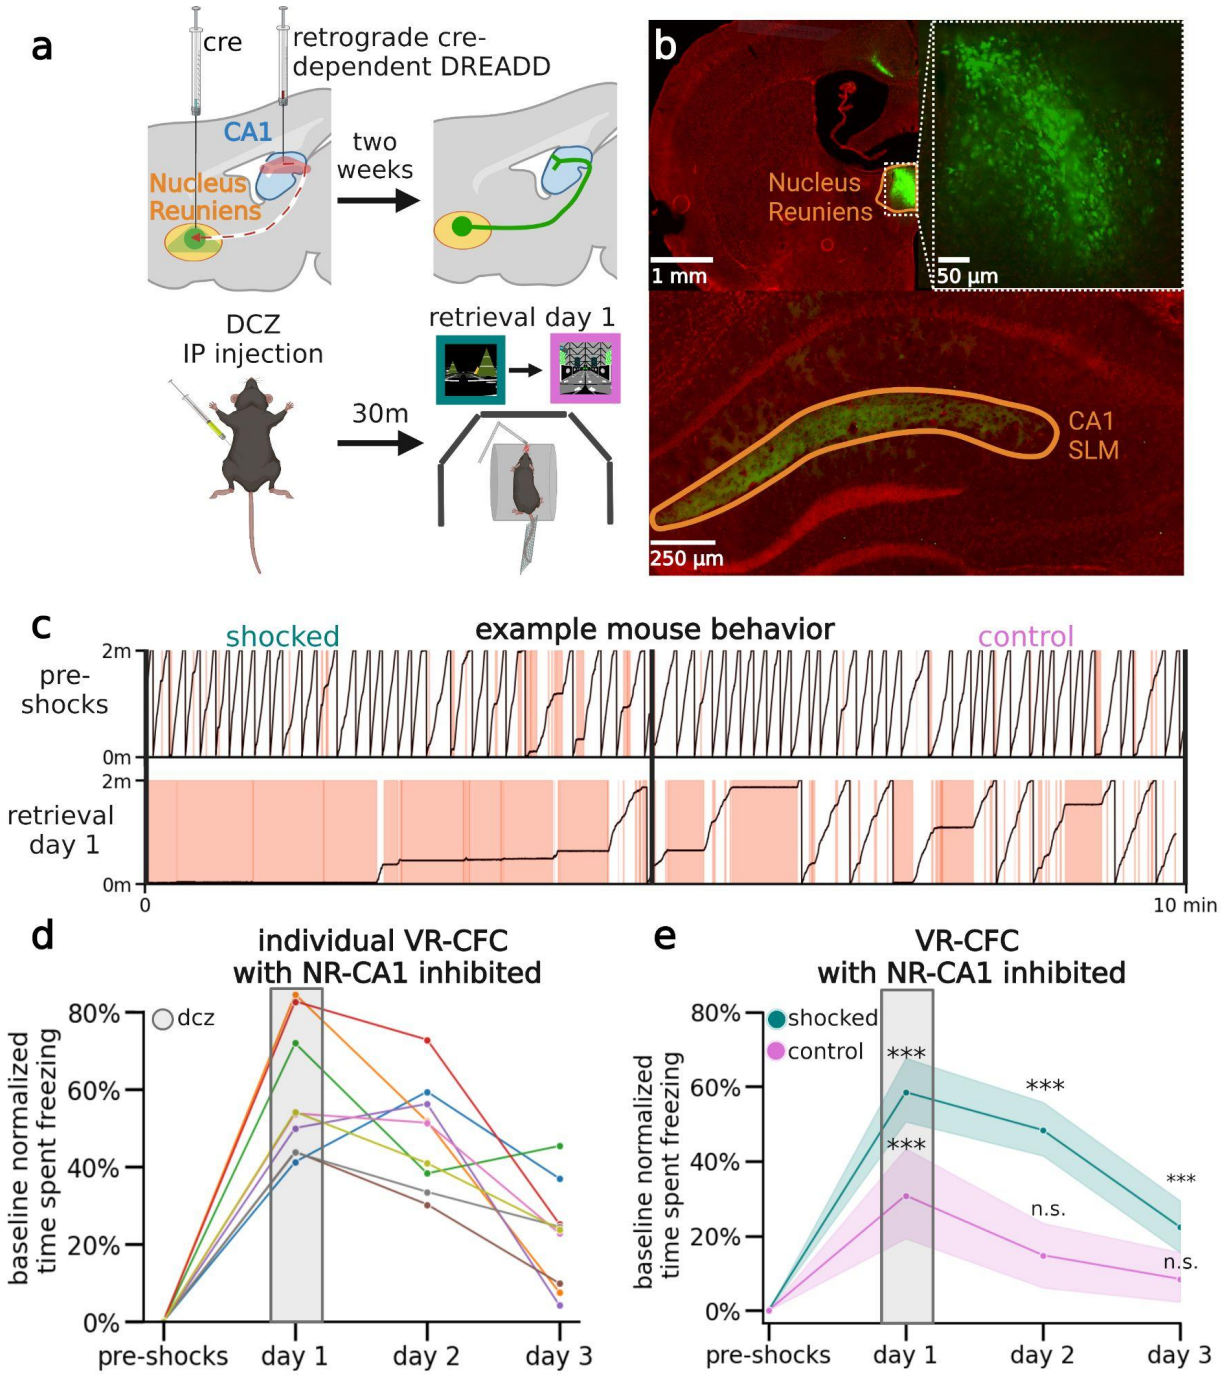

Figure 2.2. Nucleus reuniens-CA1 pathway inhibition one day following CFC

a, The NR-CA1 pathway was inhibited by first injecting retrograde CRE-dependent DREADDs in SLM of CA1 and CRE in NR and then injecting DCZ systemically 2 weeks later. Bottom: ~30 minutes before context re-exposure on retrieval day 1 (one day after CFC), mice received 0.1 mg/kg of the HM4d agonist DCZ. **b,** Confocal imaging of HM4di+Cre expression in NR and SLM of CA1 in a coronal section from an example mouse. HM4di expression is labeled in green,

DAPI in red. Top left: Section with Nucleus Reuniens (orange) showing DREADD expression selectivity to NR (green). Top right: NR section zoomed in to highlight expression in individual neurons. Bottom: CA1 hippocampal section showing DREADD expression in NR axons in CA1 restricted to the SLM. **c,** Example track position from a mouse on day 0 pre-shocks (Top), and on retrieval day 1 with NR-CA1 inhibited (Bottom), in the shocked (left) and control (right) contexts. Red shading indicates freezing epochs. **d,** Percent time freezing normalized to baseline for each mouse plotted individually across days in the shocked context. DCZ administration is shown highlighted in gray on retrieval day 1. **e,** Green line shows summary data of 2d showing mean time spent freezing across days with NR-CA1 inhibited on day 1. Red line shows time spent freezing in the same mice but in the control context (N = 9 mice). Freezing levels stay elevated in the shocked context on days 1-3, but are only elevated on day 1 in the control context. (Estimated Marginal Mean on ANOVA, Shocked P = day 1: 4.15e-10, day 2: 2.34e-9, day 3: 9.40e-3, Control P = day 1: 2.93e-4, day 2: 0.99, day 3: 1).

Given these findings, we asked if mice could still discriminate between the shocked and the control context after NR-CA1 inhibition. To do so, we calculated a discrimination index (DI)[97] which revealed a significant decrease in discrimination between the shocked and control contexts on day 1, going from 37.7 ± 8.1% in NR-CA1 intact mice to 14.2 ± 7.25% with NR-CA1 inhibited (Fig. 2.3e). No significant difference in the DI was observed on any other day. This result suggests that during inhibition of the NR-CA1 pathway fear-induced contextual discrimination is reduced (Fig. 2.3e).

We wanted to ensure that DCZ-induced increases in freezing were not due to a general decrease in movement. To do so, we exposed NR-CA1 inhibited and intact mice to a 'dark' context (devoid of any visual cues) for ~5 minutes after they were exposed to both the shocked and control contexts on retrieval days 1-3. In this dark context, mice quickly recovered their running behavior, with NR-CA1 inhibited mice freezing on average only 4.3 ± 3.8% of the time across all 3 days of retrieval. Both within and across-mice controls froze at comparably low levels in the dark context (on average under 5%; Fig. 2.3f). Therefore, neither DREADD inhibition, nor DCZ itself, impacted the mouse's ability to move, and the increase in freezing

behavior is specific to when mice are navigating in VR contexts. Our results thus indicate that the increase in freezing behavior in NR-CA1 inhibited mice is not due to motor impairment of DCZ. Our results overall indicate that the intact NR-CA1 pathway sends a potent fear suppression signal to CA1, critical for shortening the length of freezing epochs, limiting fear generalization to other contexts, and speeding up extinction. Inhibiting this fear suppression signal, therefore, induces a higher fear response by increasing the length of freezing epochs, increasing fear generalization, and decreasing extinction.

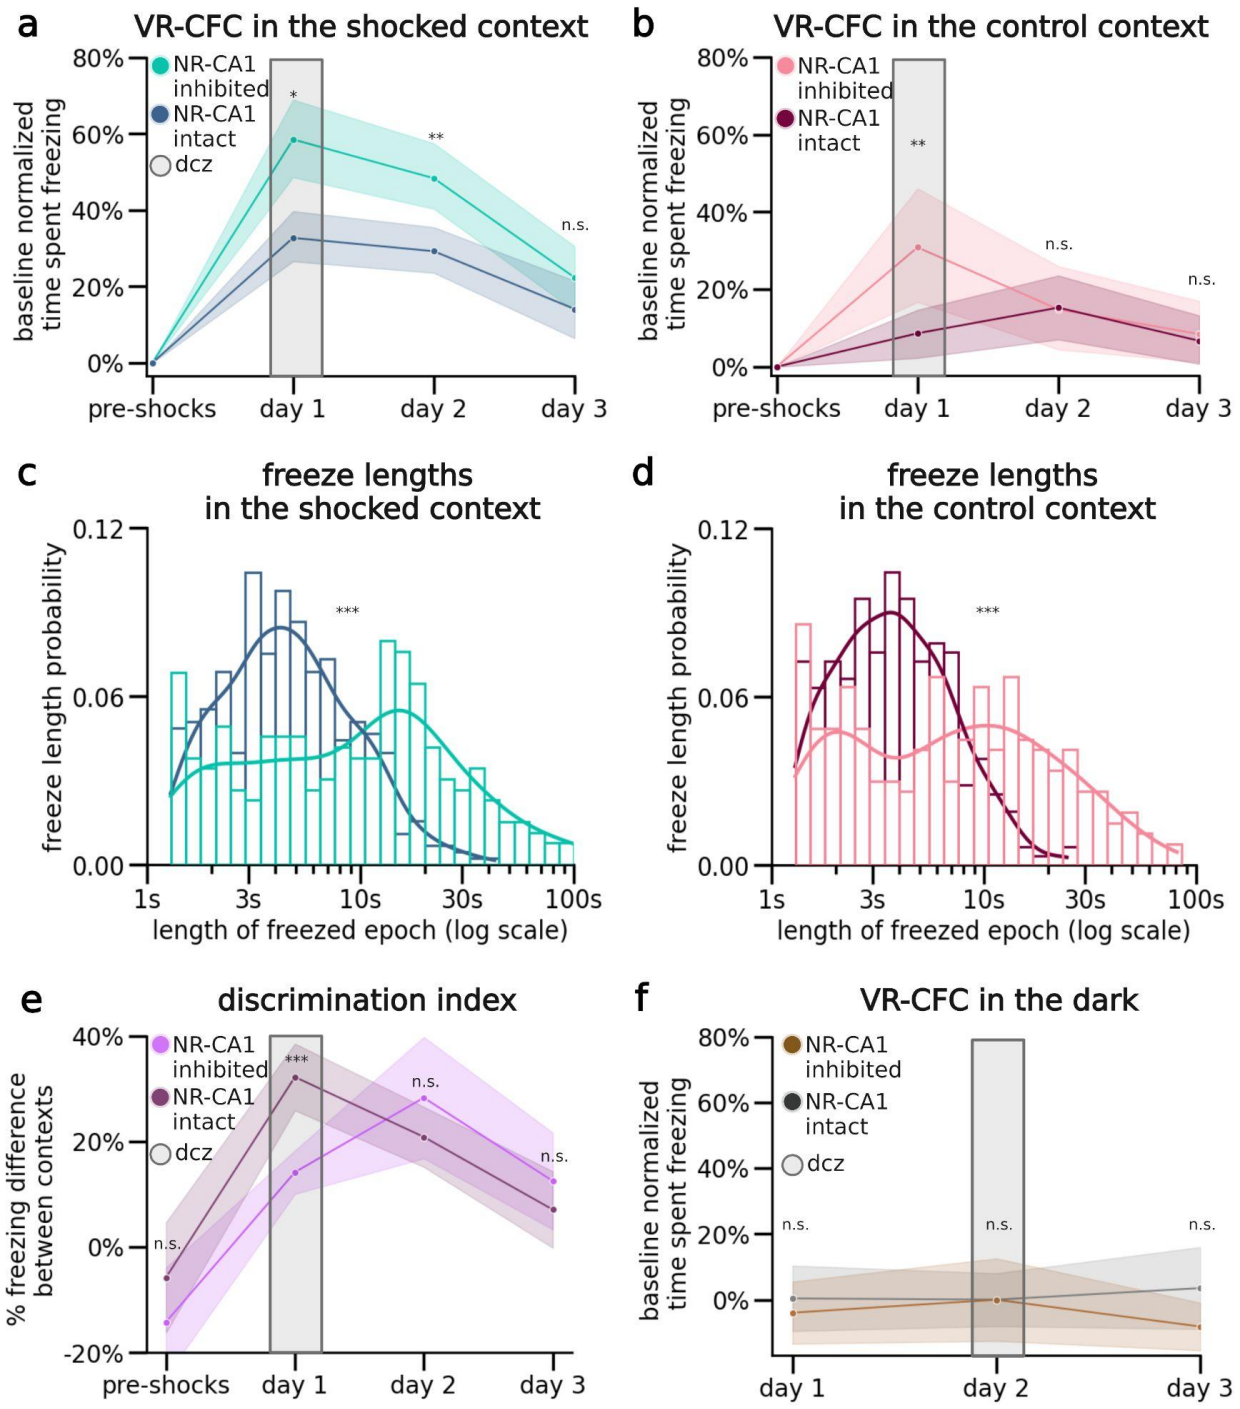

Figure 2.3. Nucleus reuniens-CA1 pathway inhibition increases freezing, reduces context discrimination, and delays extinction

a, Direct comparison of freezing behavior in NR-CA1 inhibited (same data as Fig. 2.1e: shocked context; blue) versus NR-CA1 intact (same data as Fig. 2.2e: shocked context; green) mice (Estimated Marginal Mean on ANOVA (EMM), P = day 1: 0.027, day 2: 1.63e-3, day 3: 1). **b,**

Same as a, but for the control context (EMM, P = day 1: 2.18e-3, day 2: 1, day 3: 1). **c,** The length of individual Freeze epochs in the shocked context with NR inhibition skewed longer compared to NR intact on day 1 (Mann-Whitney U, P = 2.49e-15). **d,** Same analysis as c, but for the control context (Mann-Whitney U, P = 5.49e-12). **e,** Discrimination index was calculated per day for both the NR-CA1 intact mice and NR-CA1 inhibited mice as (% time spent freezing in shocked context - % time spent freezing in control context)/total % time spent freezing in both contexts. NR-CA1 inhibition caused mice to discriminate less between the two contexts (Wilcoxon Rank Sum, P = pre-shocks: 0.67, day 1: 1.20e-3, day 2: 0.89, day 3: 0.98). **f**, On retrieval days 1-3, we additionally recorded in a 'dark' context - a dark VR with no visual cues - in NR-CA1 inhibited mice and NR intact mice for the same length of time as the other context exposures. Mice froze at consistently low levels in the dark context across days, with no difference in freezing levels between groups (EMM, P = day 1: 1, day 2: 1, day 3: 0.10).

NR-CA1 Axon Activity Becomes Tuned to Freezing Behavior Following CFC

Excitatory NR projections to the hippocampus are restricted to the stratum lacunosum-moleculare (SLM) of CA1 and subiculum; NR does not project to any other hippocampal subregions or layers[54,87,106]. Previous work stimulating this projection shows it depolarizes CA1 pyramidal neurons across the dorsal-ventral axis and induces firing in multiple interneuron subtypes with dendritic processes in SLM[94,99,100,301]. However, the activity of the NR-CA1 projection *in vivo* during behavior is unknown. To determine the information transmitted directly from the NR to CA1 during CFMR, we performed *in vivo* 2-photon Ca^{2+} imaging of NR-axons in SLM.

We injected an axon-targeted virus carrying axon-GCaMP6s into NR, followed by a cannula window over CA1 as previously described (Fig. 2.4a[223,297]). Expression in NR was confirmed via histological evaluation following the completion of experiments (Fig 2.4a; left). NR-axons could be observed in the SLM of CA1 both *post-hoc* (Fig 2.4a, middle) and during experiments under 2-photon imaging (Fig. 2.4a, right). We successfully recorded reliable GCaMP6s expression from 1 highly branching NR-axon per mouse (N = 10) in hippocampal CA1 during the VR-CFC paradigm (day 0 and retrieval days 1-3). We limited our analysis to a

putative single axon per animal since all identified axonal segments within the field of view with above-baseline activity were highly correlated, likely due to single, highly branched axons occupying the majority of our imaging field-of-view (see Methods: Image Processing and ROI Selection).

We found that NR axons switched their activity from untuned sparse activity (Fig. 2.4b; Top) pre-shocks, to activity highly selective for freezing epochs post-shocks, even after filtering for axons with detectable pre-shock activity (Fig. 2.4b; Bottom; same axon shown on both days). Since behavior necessarily changes following successful CFC, which induces more and longer freezing epochs, we needed to avoid potential confounds in comparing axon activity during dissimilar freezing epoch lengths before and after CFC. To do so, we quantified axon activity in three different ways.

First, we examined the mean normalized $\Delta f/f$ of peaks in both contexts during running (Fig. 2.4c, Control: Green, Shocked: Blue) and freezing (Fig. 2.4c, Control: Orange, Shocked: Red). Axons were normalized each day to near-maximum activity, controlling for any potential differences in amplitudes across days. This analysis shows that pre-shocks, NR-axonal activity in the to-be shocked context was similar between running and freezing epochs, with a slight preference for running epochs, which was slightly elevated in the control context (Fig. 2.4c). Conversely, post-shocks in both contexts and on all 3 retrieval days, we found that NR-axons had significantly higher mean peak activity during freezing compared to running (Fig. 2.4c).

Second, we binned freezing epochs into 1 s intervals based on their total length from 1-2 s to 6-7 s, and compared pre-shock to post-shocks activity within those bins, therefore comparing the NR-axonal activity of similar lengths of freezing both pre and post shocks (Fig. 2.4c, Supplementary Fig. 2.3e). Freezing epochs that were longer in length than 7 s were not used for

this analysis due to the low quantity of such epochs present pre-shock. Within each binned epoch, we trial-aligned activity to the freezing to running transition point (Fig. 2.4d, black center line). We then compared average NR-axon activity from all axons during freezing (Fig. 2.4b and d, peach-shaded regions) to activity during running (unshaded regions; Fig. 2.4d, 3-4 s long freezing epochs shown; all epochs in Supplementary Fig. 2.3e). Pre-shocks in either context, NR-axons did not significantly modulate their activity between running and freezing epochs (Fig. 2.4d). However, post-shocks, we found that NR-axons significantly increased their activity during freezing epochs, compared to reduced activity during running epochs. This was observed during all post-shocks freezing epochs, in both the shocked and control contexts (Fig. 2.4c).

Third, we characterized the dynamics of NR-axon activity within each freezing epoch on pre-shock day 0 and compared to retrieval day 1. To do so we aligned NR-axons by dividing each freezing or running epoch into 5 even bins, each containing a mean normalized $\Delta f/f$ of NR-axon peaks, then took the within-bin mean across all epochs pre and post-shocks. This enabled us to effectively 'stretch' or 'shrink' all epoch lengths to a uniform standard. Using this method, we found that mean axon activity ramped up rapidly in the beginning of a freezing epoch, plateaued, then fell right before freezing transitioned to running (Fig. 2.4e). Such temporal dynamics were absent during the freezing epochs pre-shocks (Fig. 2.4e).

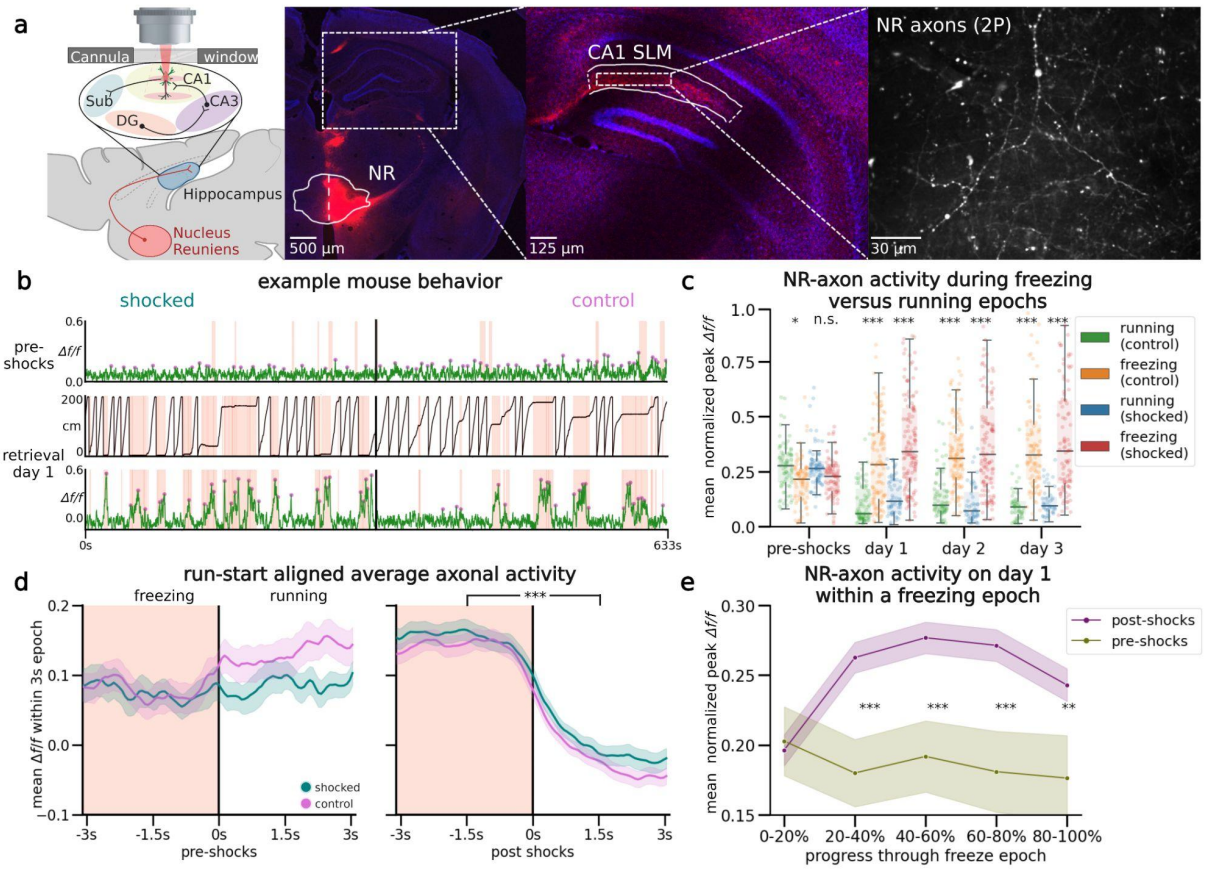

Figure 2.4. Nucleus reuniens axons in CA1 tune to fearful freezing epochs following CFC

a, Left: Schematic representation of NR axonal imaging. Mice were trained as in Fig. 2.1b. Middle left: NR mRuby expression in nucleus reuniens under confocal imaging. Middle right: Axonal expression in the hippocampus limited to the SLM layer in subiculum and CA1. Right: Example average FOV of NR axons in SLM through 2-photon during mouse behavior. **b,** Example mouse NR axonal activity pre-shocks (Top) and on day 1 (Bottom) in the shocked (left) and control (right) contexts. Red shading indicates freezing epochs. Middle trace is the mouse position on retrieval day 1. **c,** Normalized mean *Δf/f* of axonal peaks per freezing epoch plotted as dots, boxplot indicates median, 25-75th interquartile range, whiskers include all data points not determined to be outliers. In both contexts, mean normalized axonal activity increases post-shocks compared to pre-shocks, and remains elevated during post-shocks freezing epochs as compared to running epochs. Correspondingly, activity in running epochs decreased post-shocks from pre-shocks (Student's T, P control, shocked = pre-shocks: 0.09, 0.04, day 1: 9.42e-15, 6.27e-13, day 2: 6.64e-18, 7.19e-13, day 3: 1.22e-12, 1.47e-13). **d,** Peach shading (left) indicates freezing epochs. Line indicates mean, shading indicates 95% CI. Pre-shocks, NR-CA1 pathway axonal activity is comparable during freezing and running epochs. Post-shocks, activity is significantly elevated during freezing when compared to running epochs (Wilcoxon Rank Sum left: P = 0.14, right: P = 1.37e-52). Freezing epochs displayed are 3-4 s long, additional epoch windows are shown in Supplementary. Fig 2.2e). **e,** Normalized mean *Δf/f* of axonal peaks were binned into 5 categories based on percent progress throughout the freezing epoch. For the majority of the pause (20-100% freeze progress) in both contexts (purple), axonal activity was

significantly increased than activity before shocks (green) (Student's T, P = 0-20%: 0.81, 20-40%: 1.71e-3, 40-60%: 9.75e-4, 60-80%: 1.46e-3, 80-100%: 8.70e-3).

We were additionally able to track the same axon across multiple days in a subset of NR-axons (N = 4; example axon shown in Fig. 2.5a) and independently quantified their activity to observe how individual axon dynamics change throughout CFC, retrieval, and extinction (Fig. 2.5, analysis of all multi-day tracked axons averaged together). We again found similar levels of baseline activity between the freezing and running contexts pre-shocks (Fig. 2.5b, far-left). On day 1 post-shocks, these axons significantly increased their activity during freezing, and decreased their activity during running (Fig 2.5b, middle-left). This tuning to freezing behavior remained true on day 2 (Fig. 2.5b, middle-right), but fell back towards baseline levels on day 3, albeit with a small but significant difference between freezing and running remaining (Fig. 2.5b far-right). This effect is also reflected in the mean normalized *Δf/f* of peaks, with differences in peak *Δf/f* between freezing and running smaller on day 3 (Fig. 2.5c). The general shape of activity during a freezing epoch remained 'bow' shaped (Fig. 2.5d). Overall in the multi-day tracked axonal subset, we found very similar results to our overall NR-axon data (Fig. 2.4), with the notable decrease of axonal activity on retrieval day 3. This decrease corresponds with behavioral extinction of fear responses (Supplementary Fig. 2.1h), potentially indicating that individual axons may decrease their activity during freezing as the fear-suppressive effect is no longer needed. These results collectively show that NR-axons projecting to CA1 strongly tune their activity to fearful freezing epochs during CFMR, and this post-CFC activity is context-independent.

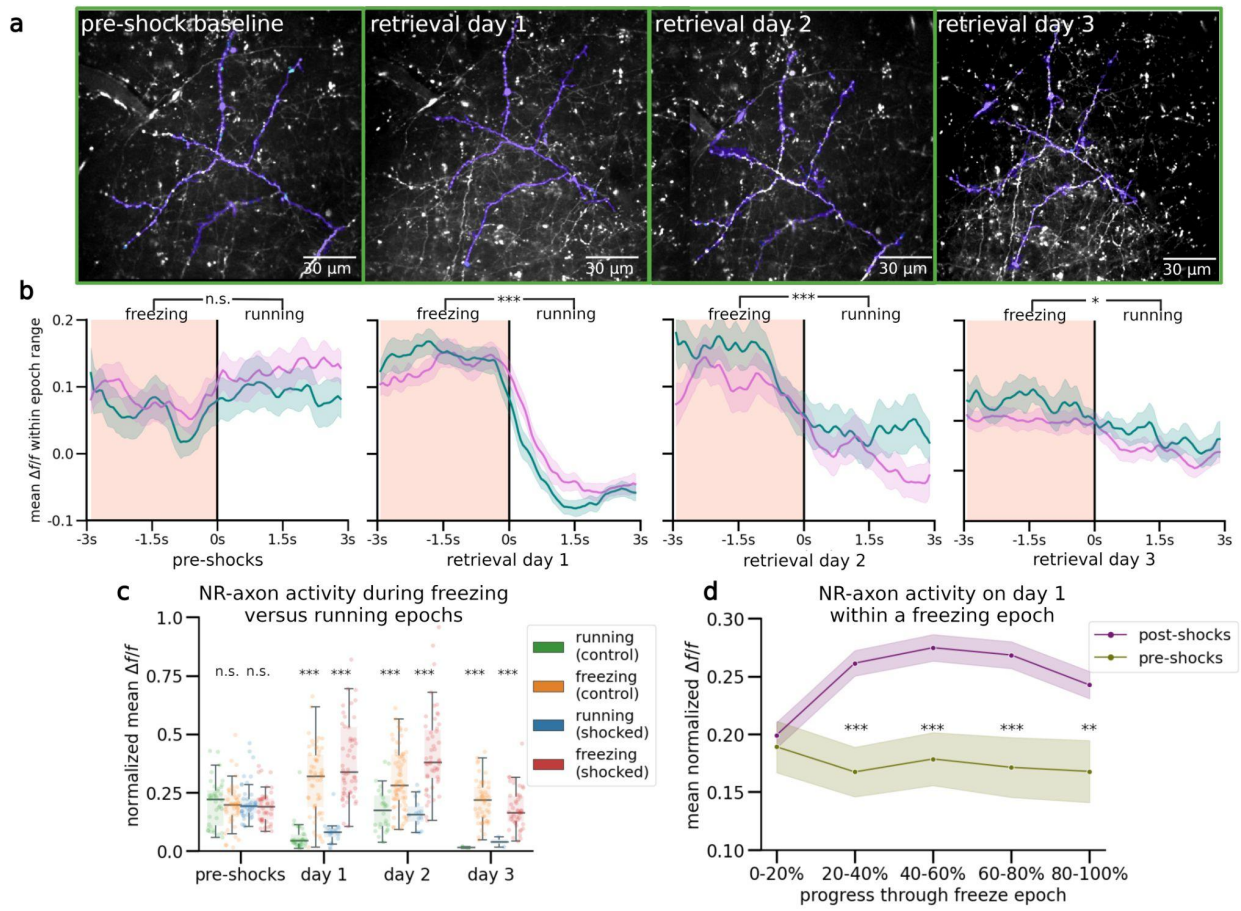

Figure 2.5. Multi-day Tracking of the Same NR Axons Reveals Fearful Freezing Tuning that Decays with Extinction

a, Example field-of-view (FOV) tracked across 4 days. FOVs are directly outputted from the Suite2p mean image with no color correction modifications applied, with Suite2p ROIs that comprised the final combined trace highlighted in purple on all four days, demonstrating our capacity to track the same NR-CA1 axonal structure over multiple days (N = 4/10 imaged mice). **b,** Average axon activity of the 4 multi-day tracked axons. Peach shading (left side of each panel) indicates freezing epochs. Colored pink line indicates mean in control context and teal line indicates mean in shocked context, shading indicates 95% CI. Pre-shocks (far left panel), NR-CA1 axonal activity is comparable during freezing and running epochs in both contexts. On post-shock retrieval days 1-3, activity is significantly elevated during freezing compared to running epochs in both the shocked and control contexts (Wilcoxon Rank Sum, P = pre-shocks: 0.33, day 1: 3.75e-8, day 2: 2.17e-5, day 3:0.02). **c,** Analyses same as Fig. 2.4c, but on the subset of multi-day imaged axons (Students' T-test: shocked P = pre-shocks: 0.21, day 1: 4.10e-3, day 2: 3.76e-8, day 3: 1.34e-6, control P = pre-shocks: 0.59, retrieval day 1: 2.19e-13, retrieval day 2: 1.26e-5, retrieval day 3: 4.2e-4). **d,** Analyses same as Fig. 2.2e, but on the subset of multi-day imaged axons (Students' T-test P = 0-20%: 0.67, 20-40%: 8.60e-6, 40-60%: 1.24e-4, 60-80%: 1.46e-4, 80-100%: 3.94e-3).

Encoding Model Predicts NR-CA1 Axonal Activity, But Only Following CFC

To further quantify the relationship between behavior and NR-axon activity, and to rule out other potential causes driving activity dynamics other than freezing, such as pupil diameter fluctuations or contributions from other behavioral variables, we developed a quantitative encoding boosted trees decision model to predict axonal activity from behavioral and pupil variables. We trained the model using XGBoost[302] to use behavioral information about freezing epochs, running epochs, velocity, location on the track, and pupil diameter to predict NR-axon activity. We separately trained on 80% of traversals and tested on the remaining 20% of traversals in each mouse, on each day, and in each context (Fig. 2.6; Supplementary Fig. 2.4; Methods: Boosted Trees Model). Model prediction most heavily relied on behavioral parameters pertaining to whether the mouse was freezing or running, its velocity, and duration passed or remaining within a freezing or running epoch (Fig. 2.6h, Supplementary Fig. 2.4b). Overall, the model predicted NR-axon activity well in both the shocked and control contexts on retrieval days post-shocks (with a context/day-combined 0.43 r^2 goodness of fit; Fig. 2.6g), but predicted axonal activity poorly in both contexts pre-shocks (with a context-combined 0.01 r^2; Fig. 2.6g). In the example mouse shown in Fig. 2.6c-f, the maximum model accuracy pre-shocks was r^2 of 0.06 (Fig. 2.6c) compared to a much higher r^2 of 0.86, 0.88, and 0.78 on retrieval days 1, 2, and 3, respectively (Fig. 2.6d-f).

Because there was variability in the fluorescence signal recorded from the axons, we checked whether model accuracy was related to the signal-to-noise. Indeed, model accuracy was correlated with axon activity - the greater the change in the normalized fluorescence signal from baseline, the better the model performed (Fig. 2.6b). The model performed significantly above chance in predicting NR-axon signal in 8/10 mice, on retrieval days 1-3. In 2/10 mice, model

prediction was poor on retrieval days, due to lower signal-to-noise ratio (SNR) of the fluorescence signal. However, changes in SNR did not account for the poor model performance pre-shock, as model accuracy was still low in animals with higher axon activity. Although overall activity was higher in post-shock days, pre-shock activity in longitudinally-tracked axons reached similar peak heights as in post-shock days (Supplementary Fig. 2.3a), and all mice included in analysis had at least 2 peaks reaching a minimum of 0.1 $\Delta f/f$ in the recording session, ensuring that poor model performance was not simply due to a lack of signal to predict. In summary, using an encoding model, we demonstrated that NR-axon activity recorded in hippocampal CA1 can be predicted from freezing behavior during CFMR, but not before the animal is fear-conditioned, revealing the development of predictable structure in NR-axon activity tuned to CFMR.

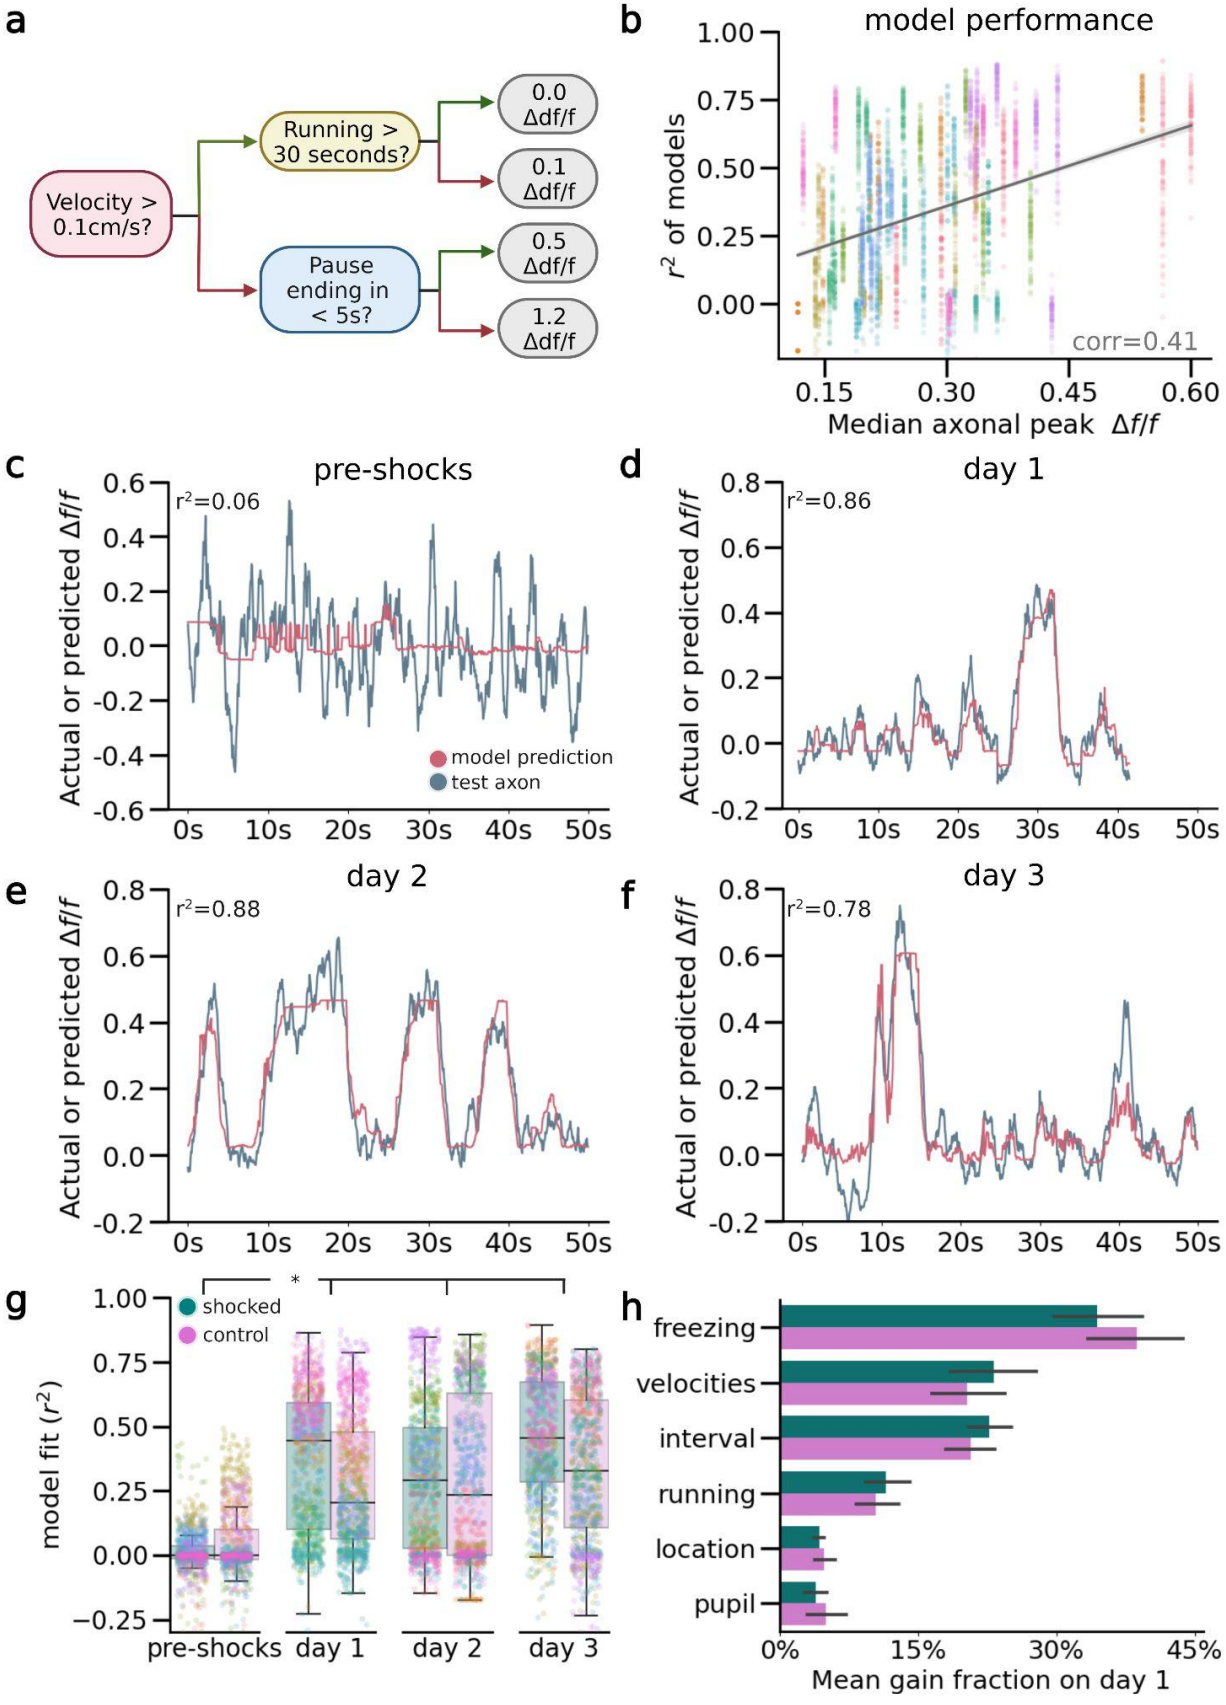
a
Velocity > 0.1cm/s?
Running > 30 seconds?
Pause ending in < 5s?
0.0 Δdf/f
0.1 Δdf/f
0.5 Δdf/f
1.2 Δdf/f
b
model performance
r² of models
corr=0.41
Median axonal peak Δf/f
c
pre-shocks
r²=0.06
Actual or predicted Δf/f
model prediction
test axon
d
day 1
r²=0.86
e
day 2
r²=0.88
f
day 3
r²=0.78
g
shocked
control
model fit (r²)
pre-shocks
day 1
day 2
day 3
h
freezing
velocities
interval
running
location
pupil
Mean gain fraction on day 1

Figure 2.6. Nucleus reuniens-CA1 axon activity is accurately predicted by a computational model following CFC

a, Simplified schematic of decision tree prediction. **b,** Model performed better on axons with higher fluorescence signals. Goodness of model fit r^2 was calculated for all model runs and plotted against median unnormalized axonal peak height as a proxy for data quality. Graded color dots are coded per mouse to visualize within-mouse stability (N = 8000 total runs; 10 mice, run 100 times per mouse, day, and context). **c-f,** Examples of model prediction for the same axon in the same mouse tracked across days in the shocked context. Matched control context model examples are shown in Supplementary Fig. 2.4d). **g,** Points indicate goodness of fit r^2 for each model run, color coded by mouse, boxplot indicates median r^2, 25-75th interquartile range, whiskers include all data points not determined to be outliers. Median model performance improved for both the shocked (pink) and control (teal) contexts across all days post-shocks, compared to pre-shock, in 8/10 imaged mice (Wilcoxon Rank Sum, P = Shocked: day 1: 0.03, day 2: 0.02, day 3: 0.02, Control: day 1: 0.04, day 2: 0.04, day 3: 0.03). **h,** Mean gain fraction plotted per category, error bar indicates SEM. Model parameters pertaining to information about pausing, velocity, and duration of time paused or remaining in either a pausing or running interval ('interval') were used more than model parameters pertaining to running information, location on the track, or pupil information. Full gain fractions are shown in Supplementary Fig. 2.4b.

Discussion

Our findings expand on a previous canon of work that indicates both the mPFC-NR projection and NR itself are required for extinction of contextual fear and prevention of fear over-generalization[97,117,118,279]. Our results suggest that in addition to these roles, NR reduces time spent freezing following CFC by suppressing CFMR as it is occurring during freezing epochs. We found that the NR-CA1 pathway is a key component of the circuit responsible for mediating the fear suppressive function of NR. This is supported by our observation that NR axons in CA1 become selectively tuned to freezing epochs, 'ramping' up activity after the initiation of freezing epochs following CFC, and inhibiting the NR-CA1 pathway lengthens freezing epochs. The function of the NR-CA1 pathway in CFMR suppression is not restricted to the context in which shocks were presented, but extends to similar contexts where shocks never occurred. This seems to limit over-generalization as shown by NR-CA1 inhibition reducing context discrimination.

Lastly, the process of suppressing ongoing CFMR by the NR-CA1 pathway also has longer term effects, as shown by reduced extinction on day 2 following NR-CA1 inhibition on day 1. In summary, our observations support a framework in which the NR-CA1 pathway actively suppresses fear responses by disrupting ongoing hippocampal-dependent CFMR to promote non-fearful behavior, and this process also limits fear overgeneralization and promotes fear extinction. Supplementary Fig. 2.5 brings all our main findings together in a cartoon model.

Interestingly, we did not observe a significant difference in NR-axon activity during CFMR between contexts on retrieval days in our NR-axon recordings, despite NR-CA1 inactivation reducing discrimination between these contexts. It could be that while the NR-axons in CA1 are not contextually modulated, their activity induces postsynaptic dynamics in CA1 that encode differences in context. This is supported by previous work showing that CA1 is specifically necessary for the context-dependence of fear extinction[303]. We also found the difference between NR-axonal activity between freezing and running epochs following CFC (activity tuned to freezing epochs) did not decrease over days, even as mice decreased their time spent freezing. However, we did observe a decrease in axon activity tuned to freezing on retrieval day 3 in our multi-day tracked axon dataset, indicating that individual CA1-projecting NR neurons may become untuned to freezing epochs as extinction occurs, while population responses may not. Previous work shows that extinction does not erase previously-learned contextual fear memories, as reactivation of hippocampal fear memories rapidly reinduces fear behavior[204,207]. This suggests that fear memories are retained but are dormant after extinction. Continued differential activity of NR-axons between freezing and running epochs in CA1, even after extinction, may be necessary to prevent the maladaptive retrieval of dormant fear memories, therefore enabling successful extinction learning.

A recent study showed that during freezing epochs in remote post-conditioning CFMR, optogenetic activation of NR significantly shortened freezing epochs, while inactivation lengthened freezing epochs[120] - in agreement with our NR-CA1 inhibition results. These authors revealed a transient increase in NR activity before the termination of freezing epochs, and showed a similar signal in the NR-BLA (basolateral amygdala) pathway. The profile of the NR and NR-BLA activity during freezing epochs they report differs from the profile we report, as they ramped up at the end of a freezing epoch and remained high during running. What could be causing this discrepancy? One key difference is the time period in which the NR and NR-BLA signals occur, compared to our reported NR-CA1 signal. We recorded 1 day following CFC, whereas NR and NR-BLA signals were measured 30 days following CFC, a time period in which memories are considered remote and no longer dependent on the hippocampus[304]. In addition, the authors showed that the NR-BLA pathway was not necessary to facilitate extinction one day following shocks, instead only being activated for remote memory retrieval. In our results, we show that the NR-CA1 pathway facilitates extinction across the timescale of a day. This suggests that CA1 and BLA projecting neurons may come from separate populations of NR neurons with different activity dynamics, a hypothesis supported by evidence that NR partially segregates into distinct subpopulations by projection region[305].

Another possibility for divergent results is that our axonal measurements are recording localized axonal spiking activity decoupled from somatic spiking activity, which has been observed in other circuits[306,307]. In this case, local synaptic modulation in CA1 may be exerting excitatory or inhibitory impacts on NR-axon activity, decoupling its activity profile in CA1 from somatic activity in NR. Future work using closed-loop optogenetic stimulation of the NR-CA1

pathway during freezing epochs, examining the existence or incidence of separate NR neurons projecting to BLA and CA1, investigating NR-CA1 activity at remote time points, and examining the activity of specifically NR-CA1 projecting somata is needed to directly test these hypotheses.

Interestingly, we found that after the NR-CA1 signal ramps up upon the start of freezing epochs it starts to ramp back down before the end of freezing epochs. We hypothesize that the ramping down at the end of freezing epochs could indicate a neural preparation or decision-making phase during which fear memory retrieval in the hippocampus has been suppressed but the decision to move has yet to occur. It also takes time for the brain to prime the motor circuits and coordinate the necessary movements before the actual initiation of movement. The CA1 is not a motor region, so the suppression of a fear memory in this region is unlikely to immediately cause movement. The ramping down of the NR-CA1 signal may reflect the completion of fear memory suppression, allowing the animal to respond quickly when the decision to move is made, but the decision may take time. Animals may engage in an internal evaluation or computation of relevant information, such as assessing potential risks or benefits associated with the impending movement. The decision-making processes and the subsequent activation of motor commands may be occurring as the NR-CA1 signal ramps down having exerted memory suppression.

The input driving the NR-CA1 pathway is most likely from the mPFC, encompassing both the prelimbic (PL) and infralimbic (IL) regions. While PL is needed for fear acquisition and retrieval, IL is necessary for the opposing task of fear suppression and preventing over-generalization[71,308–310]. The likely opposing influences of IL and PL on NR during CFMR illustrates the importance of understanding NR output pathways. Our results indicate that a fear

suppression signal circuit may be transmitted from IL, through NR, and into CA1 during CFMR. Of note, a small population of NR neurons that project both to CA1 and either PL or IL may have a key role in facilitating cross-regional theta and beta synchrony associated with hippocampal-dependent memory processing[47,103,114]. We cannot rule out that some of our recorded NR-axons collaterally project to mPFC, however, since this population makes up a small subset of all NR neurons (~3-9%[103]), we expect the majority of our recordings to be from non-dual projecting neurons. It additionally remains to be seen if the NR-CA1 exclusively projecting versus the NR-CA1 dual projecting populations have distinct dynamics during CFMR.

The amygdala, specifically the BLA, is a critical region that associates contextual information with fear[311,312]. Ultimately, short term CFMR suppression and long-term extinction likely involves changing the contextual information sent out of the hippocampus to the BLA, either directly or indirectly, to an output not associated with fear in BLA. A key question that arises from our work is how the NR-CA1 pathway potentially disrupts CFMR-associated neural dynamics in CA1. NR exclusively projects to the SLM within CA1, where the distal dendritic tuft of pyramidal neurons receive targeted synaptic input from both medial and lateral EC and local inhibitory interneurons[54,313–315]. Whether NR directly synapses on these dendrites is under contention, with contradictory anatomical and electrophysiological reports supporting evidence for and against these direct synapses[97,98,100,101]. Electrophysiological stimulation of NR projections to CA1 in rodent slice work has largely supported that NR projections depolarize, but do not directly drive firing in pyramidal neurons [96,99,100,301], with one notable early exception[94].

Interestingly, NR and EC has been proposed to both project to the same dendritic compartments, and dual activation of NR and EC projections in slice amplifies nonlinear dendritic spiking, implying that NR/EC interactions may be important *in vivo* for synaptic

plasticity[99,305]. Such dendritic-spike-induced plasticity has been associated with the formation of new place fields in novel environments[147] could provide a mechanism through which NR both disrupts CFMR and promotes extinction learning. Additionally, either NR or EC projections to SLM, when coincident with CA3 inputs through schaffer collaterals, induce burst firing in CA1 pyramidal cells[316–319]. CA1 pyramidal cell bursts are also capable of inducing new place fields in CA1 through behavioral timescale synaptic plasticity (BTSP)[316,318–320]. If our newly-reported NR input to CA1 pyramidal cell apical tuft dendrites during fearful freezing epochs coincides with CA3 inputs, their combined activity could induce burst firing and initiate BTSP. The bursts themselves could disrupt population dynamics to "jolt " the network out of CFMR, enabling the behavioral transition from freezing to running, while also inducing new place cell representations to form (remapping) through BTSP to support extinction learning[321].

This framework could explain why inhibiting the NR-CA1 pathway on retrieval day one reduced fear extinction on retrieval day 2. In effect, we may have prevented BTSP from inducing remapping and thus prevented extinction learning. Alternatively, reduced fear extinction on day 2 could be caused by the increased freezing behavior on day 1 during NR-CA1 inhibition, i.e., more freezing on day 1 leads to more freezing on day 2, although this effect was not observed in the control context. Whether delayed extinction on day 2 results from disruption to circuit level processes on day 1 or from increases in freezing behavior on day 1 remains to be investigated. However, regardless of specific cause, our data support the conclusion that the NR-CA1 pathway is necessary for normal fear extinction.

NR could also disrupt CA1 dynamics through inhibition. It is well established that NR induces strong firing in various hippocampal interneuron populations with dendritic processes in SLM[94,99,100,301]. Which specific inhibitory populations are directly stimulated by NR is an open

question, and one that has wildly divergent implications for the overall impact of NR on CA1 activity. In the case where NR-axons in SLM exclusively target inhibitory interneuron postsynaptic partners, the overall impact of NR-CA1 pathway activation on CA1 pyramidal population activity could still be net inhibitory, net excitatory, or selectively mixed. NR-axons could activate inhibitory micro-circuits that disrupt awake replay of location-specific activity sequences of the shocked context during freezing[322], or silence temporally-restricted reactivation of engram cells[204] to induce fear memory suppression and enable extinction learning. Further research on the impact of NR on CA1 dendritic and somatic population dynamics is needed to unravel how the NR-CA1 pathway mechanistically induces suppression of ongoing CFMR.

Methods

Experimental Model and Subject Details

All experimental and surgical procedures were in accordance with the University of Chicago Animal Care and Use Committee guidelines. We used 10-20 week old male C57BL/6J wildtype (WT) mice (23-33 g). Male mice were used over female mice due to the size and weight of the headplates (9.1 mm x 31.7 mm, ~2 g) which were difficult to firmly attach on smaller female skulls, and low weights reached under water restriction in female mice making the additional weight of the tailcoat potentially burdensome and interfere with experimental results. Mice were individually housed in a reverse 12 hour light/dark cycle and behavioral experiments were conducted during the animal's dark cycle. We are unaware of any influence of strain or sex on the parameters analyzed in this study. A total of 79 mice were used, 46 of which were used for data analysis shown in this paper. 33 mice were excluded for not meeting running behavior criteria (See Methods: Behavior). Of these, 20 never reached the 4 traversals/minute cutoff after 14+ days of training, and 13 did not meet movement criterion after removal of water reward and addition of the tailcoat. Of the 46 mice that did meet behavioral criteria, 20 mice were used in the NR-CA1 intact group, 10 of which were imaged. Of the remaining 10, although they had imaging windows implanted, 6 did not have sufficient signal-to-noise on a pre-experimental day to image through CFC, and 4 imaging recordings were eliminated *post-hoc* for z-motion drift. In the remaining 26 mice, 9 mice were used for NR-CA1 DREADD inhibition and 17 mice were used as controls: 6 mice were used for no-shock control, 4 mice were used for mCherry DREADD control, and 7 mice were used for saline DREADD control.

Mouse surgery and viral injections

Mice were anesthetized (~1-2% isoflurane) and injected with 0.5 ml of saline (intraperitoneal IP injection) and 0.5 ml of Meloxicam (1-2 mg/kg, subcutaneous injection) before being weighed and mounted onto a stereotaxic surgical station (David Kopf Instruments). A small craniotomy (1-1.5 mm diameter) was made over the hippocampus (± 1.7 mm lateral, -2.3 mm caudal of Bregma) or nucleus reuniens (0.0 lateral, -0.6 caudal of Bregma). For NR imaging experiments, 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 4.1 mm below the surface of the dura) using a beveled glass micropipette leading to GCaMP6s and mRuby expression in a population of NR neurons. For DREADD experiments, first AAVrg-hSyn-DIO-hM4D(Gi)-mCherry (pAAV-hSyn-DIO-hM4D(Gi)-mCherry was a gift from Bryan Roth Addgene viral prep # 44362-AAVrg; RRID:Addgene_44362) was injected into bilateral hippocampal CA1 SLM (~50 nL per side at a depth of -1.5 mm below the surface of the dura). In the same surgical procedure, AAV9-hSyn-Cre (pENN.AAV.hSyn.Cre.WPRE.hGH was a gift from James M. Wilson Addgene viral prep # 105553-AAV9 ; http://n2t.net/addgene:105553 ; RRID:Addgene_105553) was injected into bilateral NR (~100 nL at a depth of -4.1 mm). For NR-DREADD Controls, AAVrg-hSyn-DIO-mCherry (pAAV-hSyn-DIO-mCherry was a gift from Bryan Roth Addgene viral prep # 50459-AAVrg; http://n2t.net/addgene:50459 ; RRID:Addgene_50459) was substituted for AAVrg-hSyn-DIO-hM4D(Gi)-mCherry. Afterwards, the site was covered using dental cement (Metabond, Parkell Corporation) and a metal head-plate (9.1 mm x 31.7 mm, Atlas Tool and Die Works) was also attached to the skull with the cement. Mice were separated into individual cages

and water restriction began the following day (0.8-1.0 ml per day). At least 7 days following injection surgery, and approximately 7 days prior to the beginning of mouse training, mice underwent another surgery to implant a hippocampal window as previously described[323]. Following implantation, the head-plate was reattached with the addition of a head-ring cemented on top of the head-plate which was used to house the microscope objective and block out ambient light. Post-surgery mice were given 1-2 ml of water/day for 3 days to enhance recovery before returning to the reduced water schedule (0.8-1.0 ml/day). Expression of axon-GCaMP6s reached a steady state ~50 days after the virus was injected, as monitored through 2p imaging. Expression of hM4D(Gi)-mCherry was validated using *post-hoc* confocal imaging.

Behavior

Our virtual reality (VR) and treadmill setup was designed similarly to previously described setups[297]. The virtual environments that the mice navigated through were created using VIRMEn[155]. Mice were head restrained with their limbs comfortably resting on a freely rotating styrofoam wheel ('treadmill'). Movement of the wheel caused movement in VR by using a rotary encoder to detect treadmill rotations and feed this information into our VR computer, as in (Heys et al., 2014; Sheffield et al., 2017). During training, mice received a water reward (4 µl) through a waterspout upon completing each traversal of the track (a lap), which was then associated with a clicking sound from the solenoid. Upon receiving the water reward, a short VR pause of 1.5 s was implemented to allow for water consumption and to help distinguish traversals from one another rather than them being continuous. Mice were then virtually teleported back to the beginning of the track and could begin a new traversal. Mice were also teleported to the beginning of a new contextual exposure.

Four distinct VR contexts were used in this experiment, and all contexts can be found on GitHub at https://github.com/hmacomber/NR-Analysis. The training VR has local cues of a black track with gray squares at regular intervals, and gray walls with black ovals. The end of the track has a gray wall with neon green dots. Additional local cues include a green and white checkered overhead 'bridge' and a black flag with a white star towards the end of the track. Distal cues include black and gray rectangles, and a cylindrical gray tube with lighter checkered pattern on an overhead oval. The first of the two VRs used as shocked or control contexts has a light gray track with four white diamonds for the first half of the track and a black track with two white lines for the second, black walls with white lines and gray dots at regular intervals, and bright solid green end wall with a black line in the middle. Additional distal cues are provided by a gray and black chevron pattern sky, neon green walls with dot patterns in white, black and neon blue cylinders in the later half of the track, and a black conical with white stripes at the end of the track. The second of the two VRs has a black track, surrounded by white lines and small gray dots. Green and yellow conical structures with white stripes dot the landscape. The far end of the track has a striped wall.

Mouse behaviors (running velocity, track position) were collected using a PicoScope Oscilloscope (PICO4824, Pico Technology). Pupil tracking was done through the imaging software (Scanbox, Neurolabware) at 15.49 Hz, using Allied Vision Mako U-130b camera with a 25 mm lens and a 750 nm longpass IR filter. IR illumination from the objective was used to illuminate the pupil for tracking. Behavioral training to navigate the virtual environment began ~7 days after window implantation (~30 minutes per day) and continued until mice reached a speed of greater than 4 traversals per minute, which took 10-14 days (although some mice never reached this level). This high level of training was necessary to ensure mice continued to traverse

the track similarly after reward was removed. Initial experiments showed that mice that failed to reach this criterion typically would not traverse the track as consistently without reward[297] a potential confound for post-shocks freezing data (data not shown). Mice that did not reach this criterion were not used for these experiments (28 mice removed across all conditions).

Contextual Fear Conditioning Paradigm

For mice that reached criteria in the training environment (>4 traversals per minute), mice were first exposed to two novel environments without water reward for 322 s (~5 minutes) each, with the addition of a custom-made tailcoat made of conductive fabric (Adafruit). Only mice that continued to maintain a speed of 4 traversals > minute without water rewards and with the tailcoat on were allowed to continue the experiment. Subselecting for mice with this consistent running behavior helped us to ensure that freezing responses recorded later were not due to the presence of the tailcoat or any discomfort from head-fixation or removal of reward. Here onwards, the tailcoat was kept on the mouse during the experimental sessions on all subsequent experimental days. Each contextual exposure was for a duration of ~ 5 minutes. Prior to experimental day 0, mA level of shock delivery was confirmed using an oscilloscope. On day 0, mice were exposed to both novel contexts, then shocked in one of the two contexts, administering 6 X 0.6 mA shocks delivered at an interval of 20-26 seconds each, (Coulbourn Instruments Precision Animal Shocker). Mice displayed rapid sprinting behavior when they received the tail shock, allowing us to confirm the delivery of shocks in real-time (Fig. 2.1e). On subsequent days, mice were exposed to both the shocked and non-shocked (control) contexts pseudorandomly, for 3 days.

DREADD Experimental Protocol

To activate the hM4D(Gi) receptor and silence a subset of NR glutamatergic neurons that project to CA1, we used Deschloroclozapine dihydrochloride (DCZ, MedChemExpress). DCZ was chosen over CNO as an h4MDi agonist as DCZ has a significantly increased potency, therefore enabling a 100-fold dosage reduction, heightened selectivity for h4MDi receptors over endogenous receptors, and a significantly more rapid onset than CNO. In addition, the DCZ metabolites C21 and DCZ- N-oxide are reported at negligible concentrations (< 0.2 nM), compared to the known tendency of CNO to metabolize to Clozapine at higher concentrations, a compound with significant off-target receptor binding in the mammalian brain[300]. Because of these factors, we chose to use DCZ for inactivation, as in our past work[297].

Once mice met training criteria, they were habituated to the injection process. They were exposed to the rewarded training environment for ~10 min. Afterwards, they were removed from the VR set up, placed in the holding room, and injected with ~150 μL of a 12% DMSO/Saline solution. After ~30-45 min, they were placed back in the VR setup and exposed to the rewarded training environment again for an additional 10 min. This was repeated for 3 days to acclimate mice to the injection procedure. Mice additionally received ~150 μL of a 12% DMSO/Saline solution on Day -1 of the experiment 30 minutes prior to first exposure to both neutral contexts to mimic conditions on Day 0.

For animals receiving DCZ injections, i.e. both the experimental NR-CA1 inhibited AAVrg-hSyn-DIO-hM4D(Gi)-mCherry group and the control NR-CA1 intact AAVrg-hSyn-DIO-mCherry group, DCZ was dissolved in DMSO at at .02 mg/mL concentration and stored at -80 °C on day 0. On retrieval day 1, DCZ solutions were thawed to room temperature and diluted to 0.01 mg/mL with DMSO/Saline. ~30 minutes prior to context

exposure, mice were brought to a holding room and IP injected with 0.1 mg/kg DCZ of a .02 mg/mL solution. A separate control group with hM4Di expression intact received DMSO/saline instead of DCZ on retrieval day 1. These mice were injected with a weight matched quantity (~100-150 μL) of saline in place of 0.01 mg/mL DCZ. In all groups, a quantity of DMSO/Saline solution identical to IP injection amount on Day 1 (~100-150 μL) was injected on all other experimental days, ~30 minutes prior to imaging, to control for the impact of any potential IP injection-induced stress. Imaging protocol for all DREADD NR-CA1 inhibited experimental mice and intact controls was kept identical to VR-CFC NR-axon imaged mice, with the addition of a 'dark' imaging session after context exposures of the same duration, where no context was displayed on screens, for ~5 minutes, to check for any impact of DCZ on movement (Fig. 2.3f).

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 16x/0.8 NA/3 mm WD water immersion objective (MRP07220, Nikon). axon-GCaMP6s was excited at 920 nm and mRuby was excited at 1040 nm with a femtosecond-pulsed two photon 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. A single imaging field of view (FOV) was positioned between 350-500 μm below the putative surface and 400-700 μm equally in the *x/y* direction to collect data from as many NR axonal segments as possible. Time-series images were collected through Scanbox (Neurolabware) and the PicoScope Oscilloscope was used to synchronize frame acquisition timing with behavior. When possible, the same axonal field was returned to across days (Fig. 2.5, N = 4/10 imaged mice).

Immunohistochemistry and Confocal Imaging

Expression of either hm4D(Gi)-mCherry or GCaMP6s-mRuby in glutamatergic neurons in NR were checked *post hoc*. Mice were anesthetized with isoflurane and perfused with ~10 ml phosphate-buffered saline (PBS) followed by ~20 ml 4% paraformaldehyde in PBS. Brains were removed and immersed in 30% sucrose solution overnight before being sectioned at 30 µm-thickness on a cryostat. Brain slices were collected into well plates containing PBS. Slices were washed 5 times with PBS for 5 min then were blocked in 1% Bovine Serum Albumin, 10% Normal goat serum, 0.1% Triton X-100 for 2hrs. Brain slices were then incubated with either 1:500 rabbit-α-mCherry (ab167453, Abcam) or 1:500 goat-α-mRuby (STJ140251, St John's Laboratory) in a blocking solution at 4°C. After 48 hrs, the slices were incubated with either 1:1000 goat-α-rabbit Alexa Fluor 488 secondary antibody (A32731, ThermoFisher) or 1:1000 rabbit-a-goat Alexa Fluor 488 secondary antibody (A27012, ThermoFisher) respectively, for 2 hrs. Brain slices were then collected on glass slides and mounted with a mounting media with DAPI (SouthernBiotech DAPI-Fluoromount-G Clear Mounting Media, 010020). Whole-brain slices were imaged under ×10 and x40 with a Caliber I.D. RS-G4 Large Format Laser Scanning Confocal microscope from the Integrated Light Microscopy Core at the University of Chicago.

Image Processing and ROI Selection

Time-series images were preprocessed using Suite2p (Pachitariu et al., 2017). Movement artifacts were removed using rigid and non-rigid transformations and assessed to ensure the absence of drifts in the *z*-direction. Datasets with visible *z*-drift were discarded (N = 4). All datasets collected during shock administration on Day 0 were discarded, due to the high velocity

post-shocks sprinting behavior of mice making FOVs too unstable for reliable analysis. Regions of interest (ROIs) were also defined using Suite2p (Fig. 2.5a) and manually inspected for accuracy. Baseline corrected *Δf/f* traces across time were then generated for each ROI.

In addition, to control for in-experiment motion artifacts for small axonal segments, a red mRuby channel was recorded simultaneously to GCaMP6s channel recordings. Per ROI, a savitzky-golay filter was applied to both channels to smooth the signal. Then, the demeaned red channel was 'subtracted' from the demeaned green channel, by orthogonalizing their vectors in variance space. That is, we took the projection of the red channel onto the green channel as $\frac{Cov(Green, Red)}{Cov(Red, Red)} \cdot Red$, and then subtracted that vector from the green channel. This results in a new vector which is guaranteed to have zero covariance with the red channel, thus removing any linear effects of the background fluorescence on the trace. All ROIs were analyzed for covariance, and any ROIs exceeding the 99th percentile of a shuffle distribution were combined using PCA and the first PC taken, in a method similar to Kaufman et al. 2020[324]. To ensure traces had sufficient activity for analysis, all mice used were required to have one axon per FOV with activity that exceeded 10% *Δf/f* twice on each experimental day (N = 10 mice). The activity of each axon was then internally rescaled per day to the 99th percentile of max activity to account for inter-axonal differences in calcium brightness. Peaks were calculated using the scipy.signal.find_peaks package with a required minimum height of 10% *Δf/f,* distance of 0.5 s, and prominence of 0.1. One axon per mouse was selected in each FOV for numerous reasons. In our dataset, NR axons in CA1 are highly branched. In most FOVs, a single axon branched visibly across the majority of the FOV, and were determined to be visually connected as well as having highly correlated activity among ROI segments. In the majority of cases, ROI segments with detectable activity determined by Suite2p were correlated with one another above an r^2 of

>0.2. Because of this relatively high correlation among identified segments, we could not rule out that all segments in each FOV originated from the same neuron in NR, even when they visually did not connect, and therefore we were conservative with our axonal selection to visually connected segments above our ROI correlation threshold, and only these ROIs were considered to be from the same axon. Further, only ROIs in a FOV exceeding the 99th percentile of a shuffle distribution were combined. We did the combining using PCA, and took the first PC, the remaining ROIs were discarded.

Pupil measures

To obtain images with dark pupils and high contrast around the borders of the pupils, pupil images were inverted, and their brightness/contrast was adjusted in ImageJ. Pupil area, pupil center of mass (COM), Pupil x and y positions, and blinking area were obtained using FaceMap (Stringer et al. 2019). Pupil data during blinking periods (frames where blinking area < mean – twice the standard deviation of the blinking area) was removed and the pupil data was interpolated to match the 2-photon imaging frame rate (15.49 Hz). Pupil area and x and y position data were smoothed with a savitzky-golay filter.

Boosted Trees Model

The encoding model used is the python implementation of the open-source gradient boosted trees algorithm XGBoost[302]. Behavioral model parameters (described below) were used to predict axon trace values. For reproducibility, the seed was set to 42. Data were then split into laps, and split using an 80/20 train/test regime. Model was run either per mouse (Fig. 2.6) or across mice (Supplementary. Fig. 2.4), per, day, and context paradigm, for a total of 8,000 runs (N = 10 mice, 4 days, 2 contexts, 100 draws). Chance performance was determined by shuffling neural activity by traversal compared to behavioral readout per mouse, across contexts and days. Model

hyperparameters were set to: gamma = 1, learning_rate = 0.01, n_estimators = 1000, base_score = 1, early_stopping_rounds = 5. The coefficient of determination r^2 is defined as $1 - \frac{u}{v}$ where u is the residual sum of squares $\sum (y_{true} - y_{pred})^2$ and v is the total sum of squares $\sum (y_{true} - y_{mean})^2$. The best possible r^2 score is 1.0, and the r^2 score can be negative because the model can be arbitrarily worse than chance. For ease of interpretability in Fig. 2.6h, the following groupings of related behavioral variables were made and the mean contribution found within each group: freezing = ('freeze', 'is freezing', 'freeze remaining', 'is postfreeze', 'freeze progress', 'freeze elapsed'), velocities = ('recorded velocity', 'velocity back 15 frames', 'velocity back 8 frames', 'velocity forward 8 frames', 'velocity forward 15 frames'), running = ('is running', 'running progress','running remaining', 'is backtracking','running_elapsed', interval = ('interval_elapsed', 'interval_remaining','interval_progress', location = 'location', pupil = ('pupil area', 'pupil x position', 'pupil y position'). We used the importance type 'gain' parameter to determine the importance of each figure to the model's overall performance. 'Gain' is how much an individual feature contributed to model accuracy (i.e. the distance between predicted and actual r^2 values) on each branch. For each feature's use in the model, that value is summed, then the median was taken across all models by context. Full gain fractions for each parameter are shown in Supplementary Fig. 2.4b.

Behavioral Parameters

All parameters described below were calculated per mouse, day, and context, and used in model training, with the exceptions of total displacement and shocks.

Time to complete a traversal: This was calculated as the total time (in seconds) taken by the animal to run from 0 to 200 cm. Frames recorded within the teleportation window were dropped from analysis.

Total Displacement: Total displacement was calculated as the distance traversed per mouse, per context, per day.

Freezing: Freezing epochs were determined as uninterrupted epochs where mouse velocity fell below 0.001 cm/s for at least 12 consecutive frames (~0.75 s). All epochs of velocity below 0.001 cm/s but not reaching 12 consecutive frames were not considered freezing or running, and were discarded from future analysis. Freezing epochs were then counted up, and each not in a freezing epoch assigned a '0', while each frame in a freezing epoch given a numeric value corresponding to the number of epochs in that recording (i.e. all frames that contained the 4th freeze of the recording would be assigned the integer '4'). Subsequent freeze features were then calculated, including the binary variable 'is freezing' which assigns a 1 to frames considered freezing, and 0 to frames not considered freezing, two sawtooth functions 'freeze remaining', and 'freeze elapsed, which counts the frames from the beginning of a freeze up or down until the end of a freeze, respectively, and 'freeze progress' which tracks the progress of a freeze as a fraction from 0 to 1.

Running: Running was determined as any epoch where forward progress velocity was sustained over 0.001 cm/s for 2 consecutive frames. The variables 'is running', 'running remaining', 'running elapsed', and 'running progress' are calculated using the running epoch data in the same fashion as their freezing counterparts.

Backward movement: Some mice demonstrated backward movement behavior in the virtual environment post-shocks, where they made backwards movement through the context. This

behavior was analyzed separately from running or pausing in Supplementary Fig. 2.1e. The binary variable 'is backtracking' assigns a 1 to frames considered backtracking, and 0 to frames not considered backtracking.

Shocks: Shock delivery was recorded through the Picoscope. Shock location on track and stereotyped post-shocks sprinting behaviors are quantified in Fig. 2.1h and Fig. 2.1c.

Velocity: Velocity was both directly measured through the picoscope encoder, and recalculated from position, to assess for accuracy. Recorded velocity was used for all velocity calculations and model training. Values were converted into cm/s for presentation.

Velocity offsets: Future and past velocity at ~1 s and ~0.5 s were calculated by offsetting the velocity to frames. The resulting non-existent 8 or 15 velocity frames at the beginning or end of the trace were extrapolated from the prior 15 frames.

Acceleration: Acceleration was calculated as the first derivative of recorded velocity.

Intervals: Three variables, 'interval elapsed', 'interval progress', and 'interval remaining combine pausing and running information into one datastream. Interval elapsed takes the component parts 'freeze elapsed' and 'running elapsed', and counts the time elapsed in either a pausing or running interval, before resetting at a switch point. 'Interval remaining' and interval progress do the same calculation, but using freeze remaining'/'running remaining' and 'freeze progress'/'running progress'

Location: Animal's position on virtual track was determined for each frame, and binned in 1 cm bins along the virtual track.

Pupil Area: Pupil area was calculated by FaceMap as previously described, then filtered with a savitzky-golay filter for smoothing.

Pupil horizontal (x) movement: Pupil x movement was calculated by FaceMap as previously described then filtered with a savitzky-golay filter for smoothing.

Pupil vertical (y) movement: Pupil y movement was calculated by FaceMap as previously described, then filtered with a savitzky-golay filter for smoothing.

Statistics

For all data distributions, a Shapiro-Wilk test was performed to verify the data was normally distributed before undergoing further statistical tests. For behavioral calculations of percent time spent freezing between contexts (shocked, control), conditions (NR-intact, NR-inhibited, unshocked), and days (pre-shock, retrieval day 1, retrieval day 2, retrieval day 3) a three-way repeated measures ANOVA in R was conducted using the afex package. A similar ANOVA structure was created for the behavior in the subset of mice that were also imaged, and a two-way repeated measures ANOVA was constructed for the behavior of mice in the dark context, which could not be incorporated into the larger behavioral ANOVA due to the unbalanced experimental design, i.e., mice were not exposed to the dark context on the pre-shock day. Pairwise comparisons of estimated marginal means were then conducted on model outputs in R using the emmeans package, with a bonferroni multiple comparisons adjustment applied *post-hoc* for a total of 51 tests. Any p-value corrected above 1 was reported as 1. Mann-Whitney U tests were used to compare distributions displayed as histograms. For axonal data, a paired Wilcoxon signed rank test, unpaired Student's T test, or an unpaired Mann-Whitney U test was used. For samples with five data points or less, only a non-parametric test was used. Box and whisker plots were used to display data distributions where applicable. The box in the box and whisker plots represent the first quartile (25^{th} percentile) to the third quartile (75^{th} percentile) of the distribution, showing the interquartile range (IQR) of the distribution. The black line across the

box is the median (50th percentile) of the data distribution. The whiskers extend to 1.5*IQR on either side of the box. A data point was considered an outlier if it was outside the whiskers or 1.5*IQR. Significance tests were performed with and without outliers. Data distributions were considered statistically significant only if they passed significance ($p < 0.05$) both with and without outliers. Significance numbers reported are without outliers. To model the probability distribution in the datasets and get an accurate idea of the data shape, a kernel density estimate was fitted to the data distribution and is shown alongside histograms. Cumulative probability distribution functions were compared using a Kolmogrov-Smirnov test. Correlations were performed using Pearson's correlation coefficient. $p < 0.05$ was chosen to indicate statistical significance and p-values presented in figures are as follows: *, $p < 0.05$, **, $p < 0.01$, ***, $p < 0.001$, N.S. not significant. Darker lines in the center of line plots are the mean, and shading is the 95% confidence interval, unless stated otherwise in text or figure legends. All regression analysis was conducted using the statsmodels Robust Linear Model package, which estimates a robust linear model via iteratively reweighted least squares, given a robust criterion estimator. The M-estimator minimizes the function $Q(e_i, \rho) = \sum_i \rho(\frac{e_i}{s})$ where ρ is a symmetric function of the residuals and *s* is an estimate of scale. We used standardized median absolute deviation for *s* and Huber's loss function, as it is less sensitive to outliers. Shading on regressions indicate 95% CI. (see https://www.statsmodels.org/dev/examples/index.html#robust-regression for additional details). Some data preprocessing was done with MATLAB (Mathworks, Version R2018a). All other data and statistical analyses were conducted in Python 3.7.4, with primary data accrued in Pandas DataFrames, and data figures were made in Python 3.7.4 using the Seaborn and Matplotlib packages (https://www.python.org/). Schematic figures (Fig. 2.1a, Fig. 2.1b. Fig. 2.4a,

Fig. 2.6a), some figure text, and figure layouts were made with BioRender (https://biorender.com/).

Data availability

Raw imaging data are extremely large and not feasible for upload to an online repository but are available upon request. Processed data are available on GitHub at https://github.com/hmacomber/NR-Analysis.

Code availability

Original code used to create figures from pre-processed data are available on GitHub at https://github.com/hmacomber/NR-Analysis

Acknowledgements

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, a New Innovator grant from the National Institutes of Health (1DP2NS111657-01) awarded to M.S., and a T32 training grant (T32DA043469) from National Institute on Drug Abuse awarded to S.K. We thank the University of Chicago imaging core for assistance with confocal imaging, and the University of Chicago animal care staff for ensuring the well-being of experimental animals. We thank Chad Heer for early help with imaging protocols. We thank Valerie Barreto and Cherry Wang for assistance training animals, *post-hoc* perfusion, immunohistochemistry, and confocal imaging. We thank Timothy Ratigan for assistance with data analysis. We thank Antoine Madar for assistance with statistical methodology. We thank Rossten Rad for discussions on experimental design. Finally, we thank Douglas Goodsmith, Jim Heys, Timothy Ratigan, and Rossten Rad for their invaluable comments on previous versions of the manuscript.

Contributions

S.K. and M.S. conceived of, designed, and tested the VR-CFC protocol. H.R. modified the VR-CFC protocol in collaboration with S.K. H.R. and M.S. conceived of and designed the experiments. H.R. performed surgeries. H.R. and S.S. collected all in-vivo behavioral and imaging data. H.R. and S.S. collected post-hoc data. H.R. wrote the analysis code and analyzed all data. H.R and M.S. interpreted the data and wrote the manuscript, with significant contributions from S.K.

Ethics Declarations

The authors declare no competing interests.

Key Resources Table

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
pENN.AAV9.axon.GCaMP6.P2A.mRuby3	Broussard et al. 2018[325]	Addgene #112005-AAV9
pENN.AAVrg.hSyn.DIO.hM4D(Gi).mCherry	Krashes et al. 2011[326]	Addgene # 44362-AAVrg
pENN.AAV.hSyn.Cre.WPRE.hGH	Wilson Lab Plasmids (unpublished)	Addgene # 105553-AAV9
pENN.pAAV.hSyn.DIO.mCherry	Roth lab DREADDs (unpublished)	Addgene # 50459-AAVrg;
Experimental models: Organisms/strains		
Mouse: C57BL/6J	Jackson Laboratories	JAX 000664 - C57BL/6J
Antibodies		
rabbit-α-mCherry	Abcam	ab167453
goat-α-mRuby	St. John's Laboratory	STJ140251
goat-α-rabbit Alexa Fluor 488	ThermoFisher	A32731, ThermoFisher
rabbit-a-goat Alexa Fluor 488	ThermoFisher	A27012, ThermoFisher
Software and Algorithms		
Fiji	Schindelin et al., 2012[327]	https://imagej.net/software/fiji/; RRID:SCR_002285
Suite2p	Pachitariu et al., 2016[328]	https://github.com/MouseLand/suite2p
MATLAB	MATLAB. (2018). *9.7.0.1190202 (R2018a)*.	https://www.mathworks.com/help/matlab/release-notes-R2018a.html

Python	Python 3.10.8	https://www.python.org/downloads/release/python-3108/
Pandas	Pandas 1.1.4	https://pandas.pydata.org/
XGBoost	XGBoost 1.5.0	https://xgboost.readthedocs.io/en/stable/python/python_api.html
SciPy	SciPy 1.9.3	https://scipy.org/install/
Seaborn	0.12.0	https://seaborn.pydata.org/installing.html
R	R-4.3.1	https://www.r-project.org/

Supplementary Figures

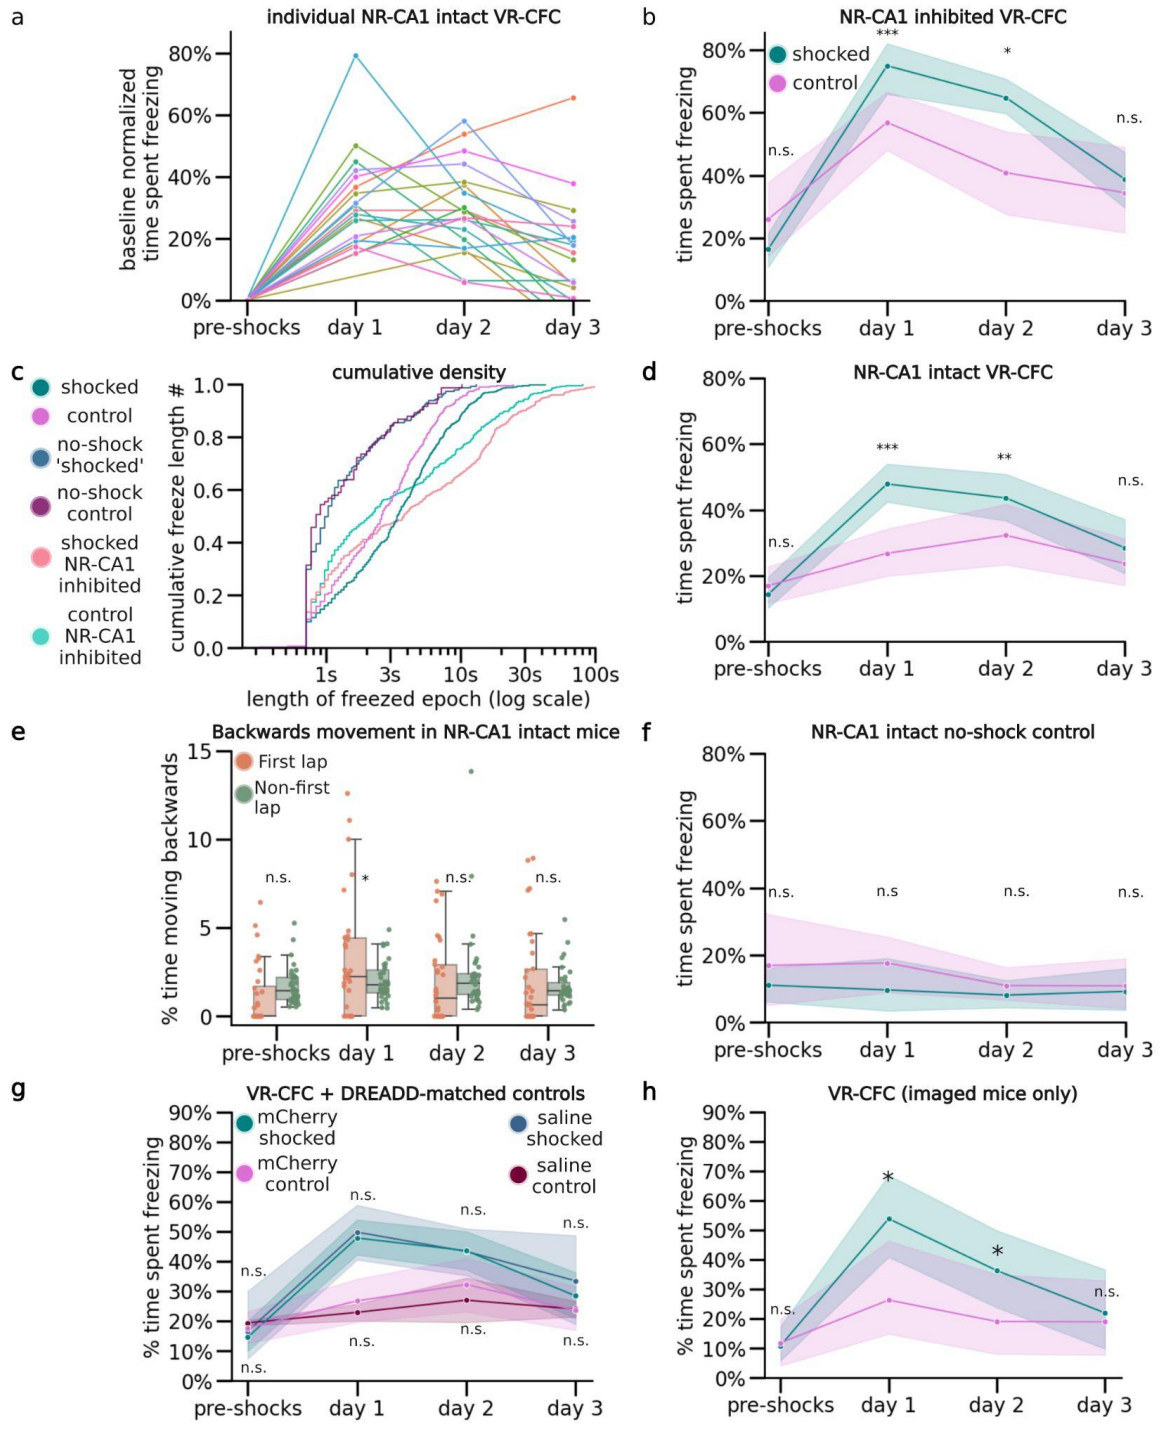

Supplementary Figure 2.1: Additional VR-CFC behavioral responses, controls and un-normalized analyses

a, Percent time freezing normalized to baseline for each mouse plotted individually in the NR-CA1 inhibited condition plotted across days in the shocked context. **b,** Same dataset as Fig.

2.2e, but plotted without normalization. Statistical comparison is between the shocked and control context on each day in the NR-CA1 inhibited condition, not compared to baseline (EMM P = pre-shocks: 1, day 1: 7.73e-5, day 2: 0.01, day 3: 1). **c,** Cumulative density plot of freeze lengths on retrieval day 1 in both contexts and in all three conditions (shocked condition, shocked with NR-CA1 inhibited condition, never shocked condition)**. d,** Same dataset as Fig. 2.2e, but plotted without normalization. Statistical comparison is between the shocked and control context on each day in the NR-CA1 intact condition, not comparing to baseline (Wilcoxon Rank Sum, P = pre-shocks: 1, day 1: 8.49e-4, day 2: 4.71e-3, day 3: 1). **e,** Backward movement behavior in NR-CA1 intact mice in the shocked context. Dots indicate average percent time spent moving backwards in each freezing epoch, either in the first traversal (green), or all other traversals (orange), boxplot indicates median, 25-75th interquartile range, whiskers include all data points not determined to be outliers. We observed instances of backwards movement behavior, where mice attempted to move 'backwards' on the track. This was significantly more common on the first traversal of the track on retrieval day 1 than all other traversals (Student's T, P = pre-shocks: 0.51, day 1: 2.30e-3, day 2: 0.30, day 3: 0.10). We interpret this backwards movement as an attempt by the mouse to escape the context by backing out of it, and classify it as a fearful behavior. **f,** Same dataset as Fig. 2.1f, but plotted without normalization. Statistical comparison is between the shocked and control context on each day (no shocks actually delivered), not compared to baseline (EMM P = 1 per day). **g,** DREADD construct does not alter behavior. Statistical comparison is between contexts for the NR-CA1 intact condition and the DREADD mCherry control condition (see Methods: DREADD experimental protocol. EMM P = 1). **h,** The freezing behavior of a subset of mice that were also imaged (N = 10). Statistical comparison is between the shocked and control context on each day, not compared to baseline (Wilcoxon Rank Sum, P = pre-shocks: 1, day 1: 0.01, day 2: 0.04, day 3: 1).

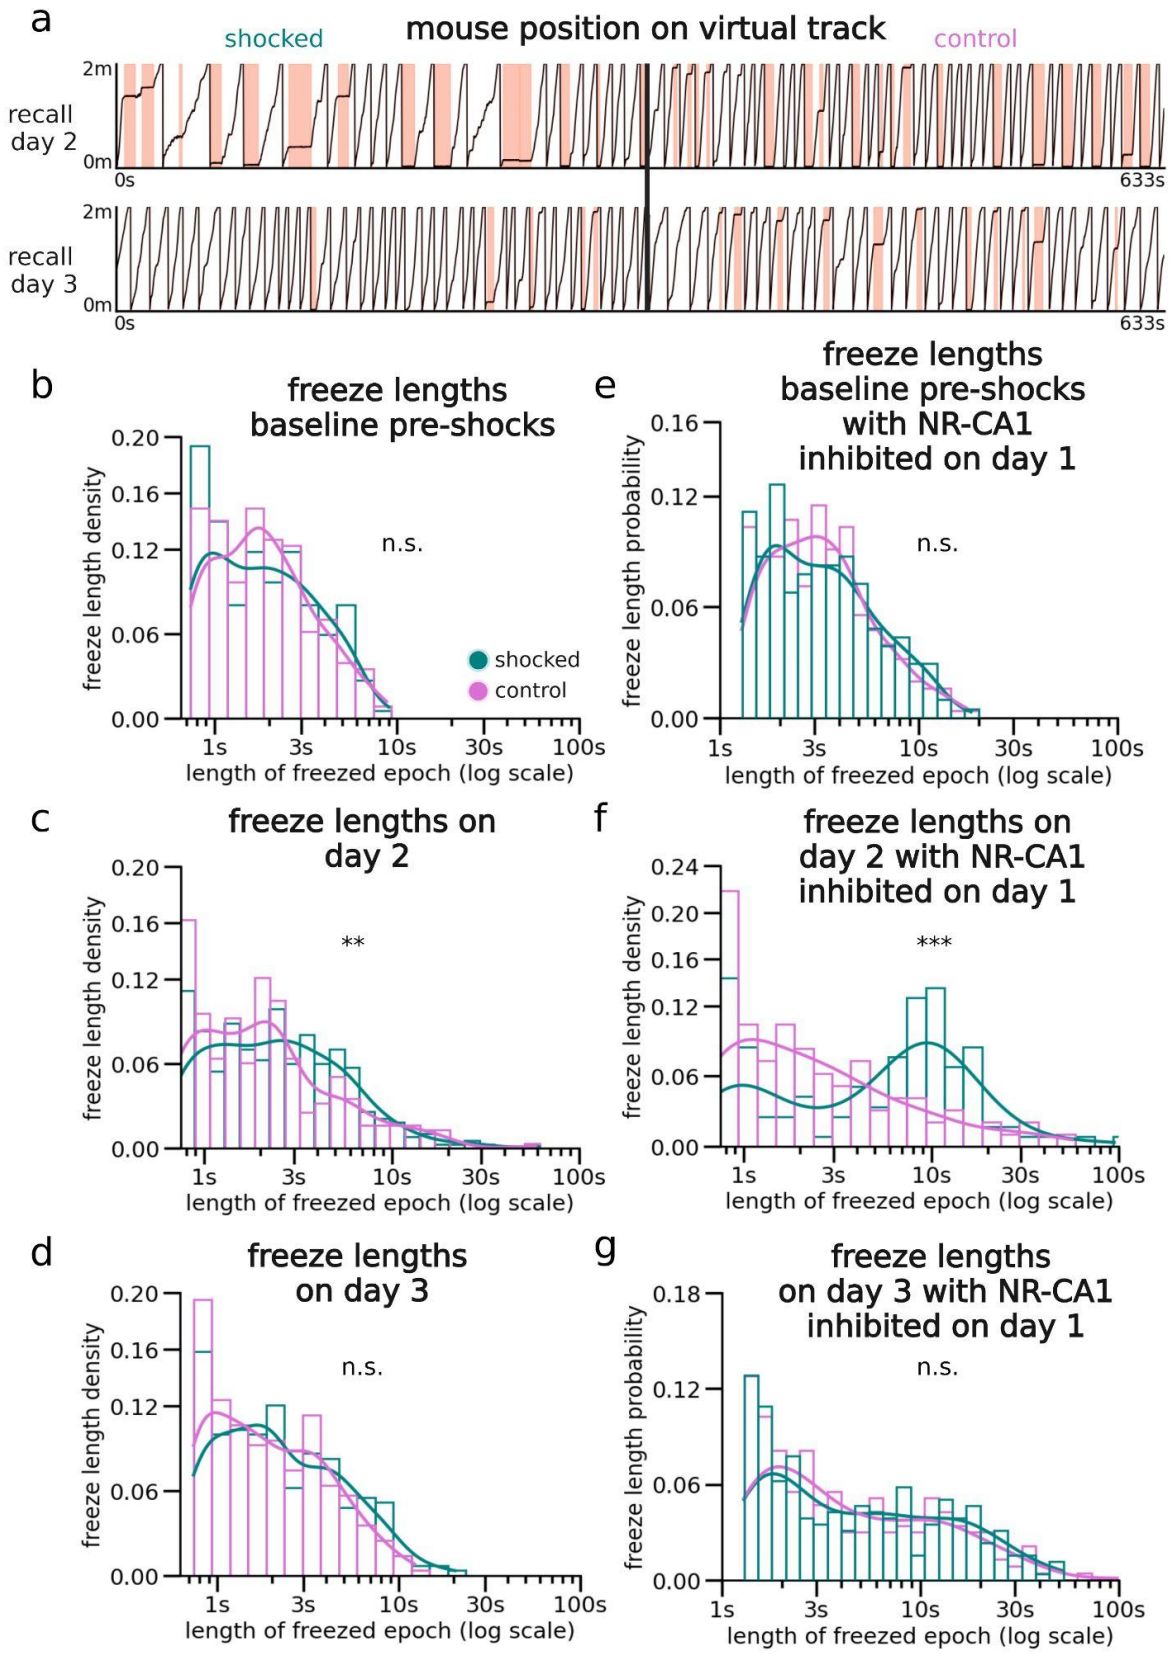
a
shocked
mouse position on virtual track
control
recall day 2
recall day 3
2m
0m
0s
633s
b
freeze lengths baseline pre-shocks
n.s.
shocked
control
freeze length density
length of freezed epoch (log scale)
1s
3s
10s
30s
100s
e
freeze lengths baseline pre-shocks with NR-CA1 inhibited on day 1
freeze length probability
c
freeze lengths on day 2
**
f
freeze lengths on day 2 with NR-CA1 inhibited on day 1
***
d
freeze lengths on day 3
g
freeze lengths on day 3 with NR-CA1 inhibited on day 1

Supplementary Figure 2.2: Additional analyses of freeze lengths with and without NR-CA1 pathway inhibition

a, Example behavioral traces from retrieval day 2 (top) and retrieval day 3 (bottom) from the same example NR-CA1 intact mouse and displayed in the same fashion as Fig. 2.1d. **b-g,** Freeze lengths of individual freeze epochs calculated as in Fig. 2.1g, day and condition indicated on graph. All significance stars were determined using Mann-Whitney U test (P = b: 0.47 c: 7.92e-3, d: 0.24, e: 0.31 f: 6.82e-3, g: 0.86).

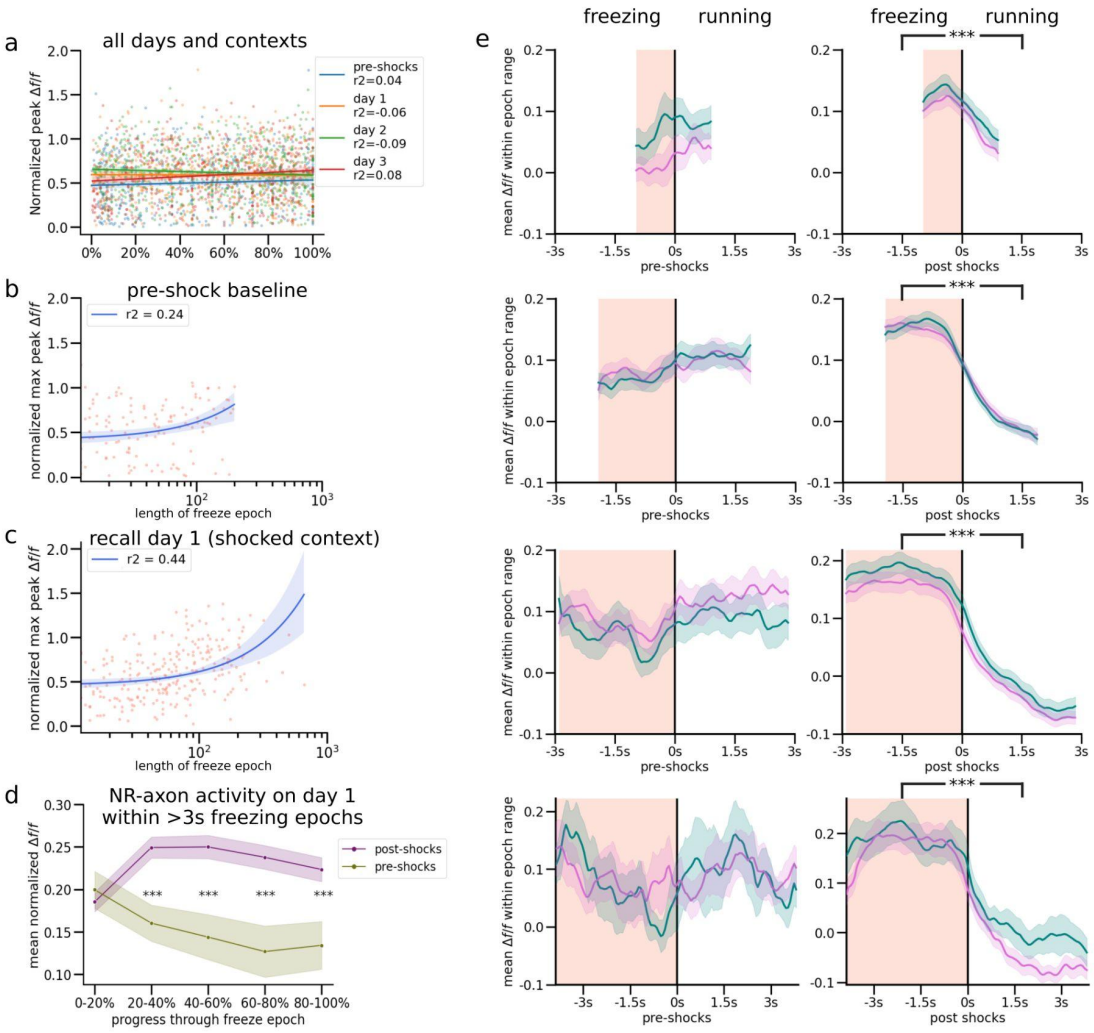

Supplementary Figure 2.3. Additional Comparisons of NR-CA1 axonal activity before and after CFC

a, The max activity of each axon per freeze epoch on each day in both contexts (dots, color-coded by experimental day), plotted against when it occurred within the freeze epoch (0% is the start of a freeze epoch, 100% is the end of a freeze epoch). A robust linear regression was then fitted to each day. This analysis indicates that the highest point of axonal activity within a freeze epoch can occur anywhere temporally within a freezing epoch. **b,** To test the impact of freeze length on max NR-CA1 axonal activity, we plotted the peak within each freeze epoch against the length of the freeze epoch in the NR-CA1 intact condition in the shocked context before shocks (behavior plotted in Fig. 2.1d) then fit the data with a robust quadratic regression (see Method Details: Statistics). The regression suggests a slight tendency for increased maximum amplitude in longer freezing epochs **c,** Analysis is the same as in 3b, but on retrieval day 1 in the shocked context. While the r^2 of the robust regression increased in the shocked context post-shocks compared to pre-shocks, since average freeze lengths also increased (as mice freeze for longer epochs post-shocks), it is difficult to assert that fearful freezing is inducing

changes in normalized max axonal peak as a function of time spent freezing from these analyses. **d,** Same analysis conducted as Fig. 2.2e, but on only a subset of longer pauses (3 s + in length) showing a similar average shape of activity throughout these freeze epochs to Fig. 2.2e. **e,** Analyses conducted the same as Fig. 2.2d, except the freeze and post-freeze running epoch window is restricted to different window lengths of 1-2 s, 2-3 s, 4-5 s, and 5-6 s, from top to bottom. All freezing epochs remained significantly elevated compared to running epochs post-shocks (Wilcoxon Rank Sum), with the same general underlying shape evident irrespective of the window chosen.

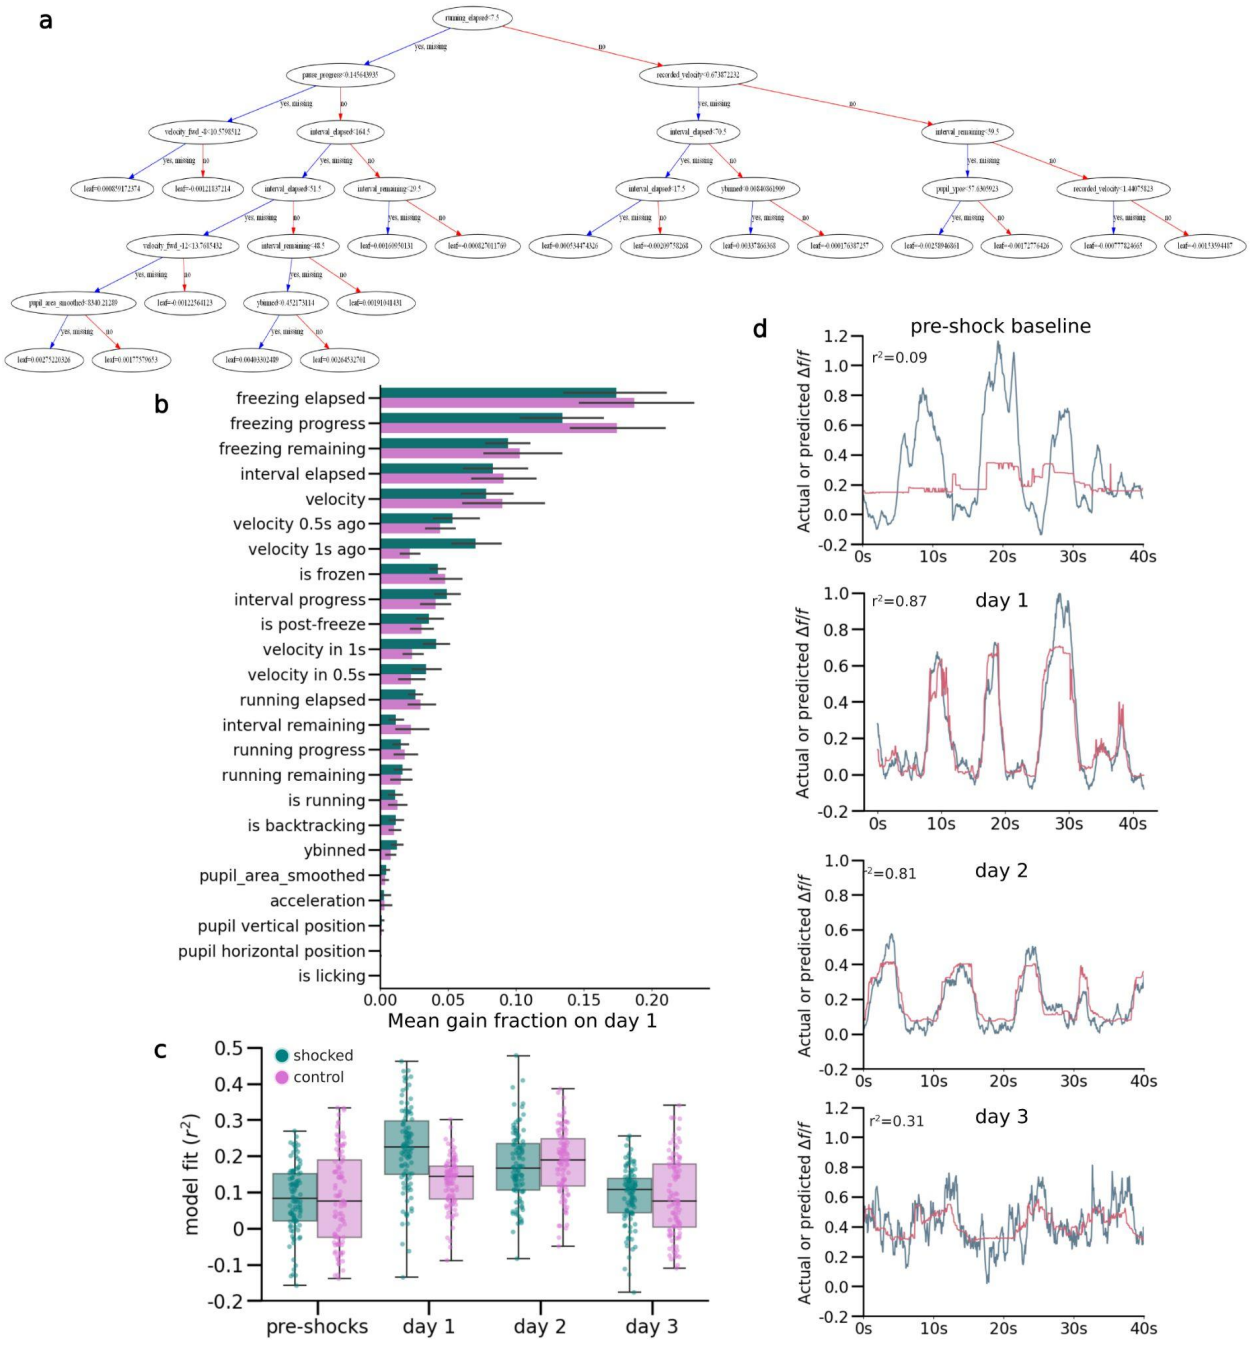

Supplementary Figure 2.4. Computational Model Details and Additional Examples

a, Real example of a model tree on retrieval day 1 in the shocked context from the same example mouse shown in Fig. 2.6d. **b,** Full ungrouped version of Fig. 2.6h. **c,** Same analyses as Fig 2.6g, except from a model that ran without differentiating and building models per individual mouse, but instead across all mice and differentiating only on day and context, demonstrating that high inter-mouse variability in axonal activity and amplitude led to decreased median r^2 goodness-of-fit model performance compared to inter-mouse modeling shown in Fig. 2.6g. **d,**

Analyses same as Fig. 2.6c-f in the same mouse shown in Fig. 2.6c-f, but in the control instead of the shocked context.

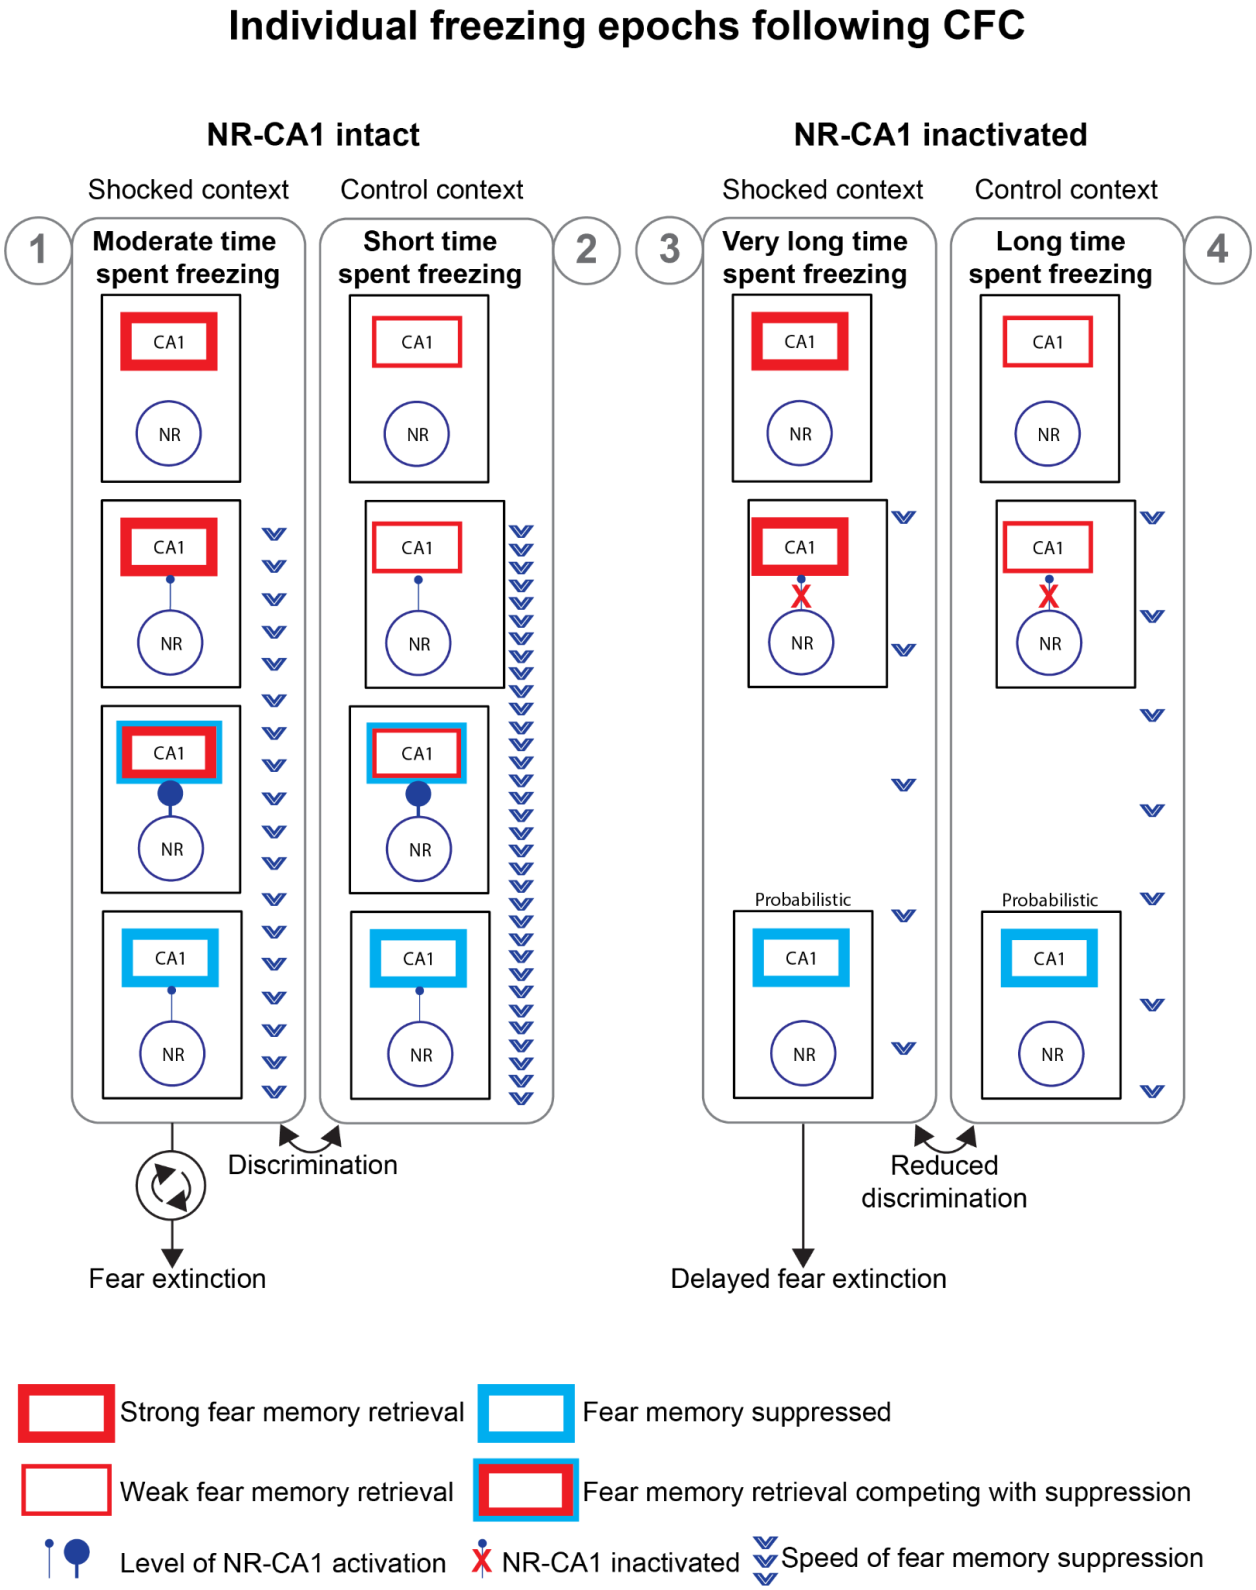

Supplementary Figure 2.5. Cartoon showing NR-CA1 speeds up contextual fear memory retrieval suppression and extinction, and amplifies context discrimination.

The cartoon brings together all the findings from the paper. It depicts the process of contextual fear memory retrieval suppression in CA1 during individual freezing epochs following CFC under 4 different conditions. Each condition is labeled with a number next to each vertical panel. From left to right the 4 conditions are: (1) a freezing epoch in the shocked context with NR-CA1 intact; (2) a freezing epoch in the control context with NR-CA1 intact; (3) a freezing epoch in the shocked context with NR-CA1 inactivated; (4) a freezing epoch in the control context with NR-CA1 inactivated. In each condition, the strength of contextual fear memory retrieval is dependent on the context the animal is in. In the "shocked " context, the contextual cues have been strongly associated with shocks, so they evoke strong fear memory retrieval in CA1, depicted by the thickness of the red box around CA1. The control context can also evoke freezing epochs because the context has some similarities with the shocked context. These weakly associated cues evoke weak fear memory retrieval in CA1, depicted by a thin red box. In the NR intact condition (1 and 2), the NR-CA1 pathway starts to ramp up as soon as the animal is in the contextual fear memory retrieval state. After some time, the NR-CA1 activity reaches maximum levels providing a disrupting signal that promotes a state change in CA1 that tries to jolt it out of fear memory retrieval. Eventually the NR-CA1 signal causes the state change in CA1 and the NR signal starts to ramp back down. The ramping down phase could indicate a neural preparation or decision-making phase during which fear memory retrieval in the hippocampus has been suppressed but the decision to move has yet to occur. The ramping down could be mechanistically caused by a feedback signal from CA1 to NR, which anatomically exists but has not been investigated. This process also occurs in the control context during freezing epochs (2), but it occurs much faster because the CA1 is in weak memory retrieval state, making it easier and quicker for the NR-CA1 signal to jolt it out of fear memory retrieval - successfully suppressing the memory. The blue V's depict the speed in which fear memory retrieval in CA1 is successfully suppressed (higher V frequencies indicate faster speeds). The process of NR-CA1-dependent fear memory retrieval suppression during individual freezing epochs also promotes long-term fear extinction. Repeated NR-CA1-induced suppression speeds up fear extinction through an unknown mechanism. Further, the difference in speed at which NR-CA1 suppresses strong versus weak memory retrieval, i.e., in the shocked versus control context, allows for context discrimination as freezing epochs are shorter in the control versus shocked context. In the NR-CA1 inactivation condition, the CA1 remains in fear memory retrieval state for much longer because the NR-CA1 disruption signal is absent. The CA1 will eventually jolt out of this state, but only in a probabilistic way. This leads to much longer freezing epochs. It also reduces context discrimination, because even in the control context when animals are weakly retrieving the fear memory, the CA1 remains in this state until it probabilistically jolts out of it. This causes similar freezing epochs in control versus shocked contexts, reducing behavioral discrimination. Lastly, because repeated NR-CA1-induced fear memory suppression does not take place during NR-CA1 inactivation, fear extinction is delayed.

CHAPTER 3

THALAMIC NUCLEUS REUNIENS MODULATES HIPPOCAMPAL CA1 POPULATION ACTIVITY THROUGHOUT FEAR CONDITIONING TO INDUCE CONTEXTUAL DISCRIMINATION AND FEAR MEMORY EXTINCTION

I was primary lead on all research conducted in this chapter, and will be listed as first author on any future research publications on this work. I was assisted in data collection and post-processing by Cherry Wang and Valerie Barreto.

Abstract

Successfully remembering aversive experiences by correctly situating them in their appropriate contexts is a critical function of the mammalian brain. This function is largely reliant on the hippocampus, a brain structure well-recognized for its role in contextual memory formation, retrieval, and modification. How the hippocampus forms, recalls, and alters aversive memories at a population level has been investigated through multiple avenues, as hippocampal somata have been shown to tune to specific locations through place fields, to freezing onset, and simultaneously preferentially activate in small groups to support memory recall. The selective input from the nucleus reuniens of the thalamus to the CA1 subregion of the hippocampus is critical for normal contextual fear memory retrieval and extinction after virtual reality contextual fear conditioning. How this pathway impacts dCA1 hippocampal population dynamics to exert its influence on contextual fear memory retrieval, discrimination, and extinction is unknown.

Here, we investigate the impact of context-dependent shocks on fearful freezing in male mice with either NR-dCA1 intact or with NR-dCA1 inhibited on post-shock fear memory retrieval day 1 (day 1). We first look into a variety of tunings that dCA1 somata can have to

salient features of a contextual fear memory on a per-somata basis: preferential firing in the shocked over the neutral context, tuning to location through place fields, tuning to shocks, and tuning to the onset and offset to freezing epochs. We next identified fast co-activations of coordinated groups of dozens of somata, termed co-activations (COAs), and how the size, frequency, and composition of COAs change throughout VR-CFC. We reveal that with the NR-dCA1 pathway intact, subsets of dCA1 somata preferentially fire in the shocked over the neutral context on day 1, massively remap by changing the location of and somata recruited to have place fields, tune to shocks, and tune to freezing onset and offset. We additionally reveal that COA frequency and size both increase in the shocked context on day 1.

In contrast, inhibiting the NR-dCA1 pathway modifies how subsets of dCA1 somata are active throughout VR-CFC. It increases non-selective firing in the neutral context, decreases remapping by increasing the stability of pre-shock place fields, and heightens the quality of both shock and freeze tuning. The frequency, size, and composition of COAs are likewise impacted, with NR-dCA1 removal selectively decreasing the frequency of COAs in the shocked context, decreasing context-dependent COA size selectivity on day 1, and increasing the proportion of shock and freeze-tuned somata recruited into COAs. Overall, our findings suggest that broad changes occur in the activity of dCA1 somata throughout VR-CFC to support contextual fear memory encoding, retrieval, discrimination, and extinction. They further indicate that the NR-dCA1 pathway is critical for regulating these changes by supporting neural discrimination, inducing place field remapping, and suppressing inappropriate synchronized memory retrieval of freeze and shock-tuned somata.

Introduction

The hippocampal formation is crucial for encoding contextual memories[138,140,329,330]. As animals go through hundreds of contextual experiences each day, most of which are mundane, the hippocampus needs to preferentially encode behaviorally salient memories, i.e., those memories that are particularly surprising, emotionally charged, or that encode information important for behavioral survival[4]. Indeed, such memories, especially those of aversive events that threaten the animal's survival, are more strongly encoded by the hippocampus than memories of threat-neutral events[331,332]. A key question in neuroscience has been understanding the mechanisms through which neural representations of fearful events are created, modified, and stored at a population level within the hippocampal formation. One such mechanism is through place cells, excitatory pyramidal neurons (ePNs) first discovered in dorsal hippocampal CA1 (dCA1) that are comprised of ~1/4-1/3 of all active ePNs and are preferentially active when an animal is located within a specific location in space[139,140,171]. The place fields formed by these neurons tile the entirety of a context and stay relatively stable on a day-to-day basis within a single context[145,146]. In a different context, a different subset of ePNs encode a unique map of the space.

The change in ePN state between these contexts is called 'remapping', and it is thought that this remapping is crucial for telling the difference between contexts[172,177,333]. Remapping can occur globally, such that the maps between contexts are orthogonal to one another, or locally, often termed 'partially', such that a subset of the place fields are stable while another subset remaps. Partial remapping can happen when individual ePNs change their place field location to a different place, aplace field disappears, or new ePNs previously without a place field form a

new one[139,333–335]. These partial remaps happen when smaller-scale environmental changes occur, such as changing a context's color, odor, or texture, modulating the context's physical shape slightly, or removing an object[150,336–339]. Partial remapping has also been observed when context salience is altered through adding novel cues, manipulating reward, or adding an aversive stimulus such as a shock[156,297,340,341].

Further evidence for the use of dCA1 place maps in memory are the existence of replays, sequences of ePNs identified as place cells that reactivate in either forward or backward sequences at ~20x speed of real-time navigation (~40-100ms) during sharp wave ripples (SWRs) that occur during sleep, restful awakeness, or pausing periods prior to decision-making[163,164,266,332,342]. Interrupting these replay events disrupts learning, decision-making, and long-term memory, implying that place maps are actively important components of hippocampal-dependent episodic memory[343]. Therefore, place maps can be thought of as the scaffolding on which context-dependent fear memories are built. It has thus been proposed that dCA1 hippocampal place maps are essential for forming episodic memories of aversive events[159,214,340,341].

In an almost entirely separate investigative vein, the hippocampal formation has been singled out as a hotspot for engrams, a sparse subset of ePNs that are activated, modified, and reactivated by a specific contextual experience. These small groups of ePNs - on the scale of dozens to hundreds - are identified throughout the hippocampal formation using immediate early gene technique and can be selectively activated or inactivated to impact animal behavior[188,194,195,198]. Contextual fear conditioning has been a focus of engram studies, revealing that the hippocampus encodes robust engrams of fear experiences and reactivates them during fear memory retrieval[204,344]. Pertinent examples of engram studies include erasing fear memories

and corresponding fear responses by inactivating ePNs active during CFC and artificially inducing a fear response in a neutral context by activating ePNs active in a feared context during neutral context exposure[204,205]. Fear extinction memory engrams can also be identified and modified the same way throughout the hippocampus[207].

As hippocampal place cells and replay are studied using wholly independent techniques from engrams, how these separately-identified populations may interact or overlap is still largely unknown, with few studies bridging the gap. Some of these studies assert that place cells largely make up engrams, while others assert that engrams and place cells largely do not overlap[119,209–211,345]. Additionally, some recent investigations have examined hippocampal co-activation, i.e., short epochs where a significant number of neurons are simultaneously synchronously active, and found that the number of somata participating in a co-activation increases during fear memory re-activation[213,340]. Unlike replay events, which are limited to identification during SWRs while the mouse is not moving, co-activation happens both within running and freezing epochs and can occur on longer time scales than would be anticipated from a SWR, indicating it is a unique phenomenon. These co-activation events may instead represent a form of memory recall that simultaneously re-activates both non-place engram cells, non-engram place cells, engramic place cells into one cohesive hippocampal representation of a fear memory[213].

To better understand population dynamics within hippocampal CA1 throughout contextual fear encoding, we recorded from thousands of dCA1 ePNs in male mice throughout our previously-established contextual fear conditioning paradigm with and without the input of thalamic nucleus reuniens to dCA1 pathway (NR-dCA1) on fear memory retrieval day 1, which is known to induce significant increases in fear responses[346]. We additionally examined activity

in a control 'no shock' paradigm to examine multi-day drift in somatic representations or intrinsic context-dependent biases independent of fear-induced behavioral changes. We further divided mouse behavior into three categories based on whether they exhibited shocked context-specific fear, overgeneralized fear, or no fear responses. We next investigated the dynamics of somatic neural population activity through context-dependent neural discrimination, place fields, and co-activation events throughout VR-CFC.

Our findings first reveal that somata demonstrate preferential activity in the shocked context on the first day of context re-exposure post-shocks when NR is intact. In contrast, mice with NR-dCA1 inhibited have more active somata that are not differentially tuned. We additionally reveal that increased neural discrimination is only correlated with increased freezing in the shocked context for mice that discriminate well between the contexts. We next demonstrate that place maps in the shocked context largely remap between the pre-shocks baseline and recall day 3, and are comprised of wider and more excitable place fields when the NR-dCA1 pathway is intact. On the other hand, when NR-dCA1 is inhibited on day 1, place maps are more stable overall and do not show differences in place field qualities. However, both of these maps are less stable than when mice were never shocked, indicating that increased remapping between pre and post-CFC is a feature of encoding the aversive experience.

Expanding beyond place tuning, we demonstrate that dCA1 possesses both shock-tuned and freeze onset/offset-tuned independent somatic populations. Mice that better discriminate between the shocked and neutral contexts on day 1 demonstrate higher quality shock tuning. Mice with NR intact also significantly only have freeze on/offset-tuned somata in the shocked context on recall day 1 be significantly better-tuned in the subpopulation that well-discriminate

between the contexts. Without NR-dCA1 online, mice that generalize well also see marked increases in the quality of freeze-tuned somata.

We next investigated how somata with tunings to context, place, shocks, and freezing, as well as somata without these tunings, are activated in synchronized co-active neural ensembles over the course of VR-CFC. We demonstrate that rapid synchronized co-activations (COAs) are a common feature in dCA1 throughout CFC. With NR intact, the frequency of COAs increase in both the shocked and neutral contexts on retrieval day 1 as compared to baseline pre-shocks, with only the shocked context retaining an increase on day 2, and neither context demonstrating an increase on day 3. With NR-dCA1 inhibited, the number of COAs in the shocked context remains slightly elevated on day 1, but not in the neutral context or in either context on subsequent days. We additionally show that with NR intact, the number of somata recruited to each COA increases selectively on retrieval day 1 in the shocked context and decreases selectivity in the neutral context. With NR-dCA1 inhibited, this differentiation is erased.

Lastly, we demonstrate that somata tuned to context, place, shocks, and freezing transition make up a significant fraction of COAs across all 4 days of VR-CFC. The composition of COAs in the shocked context is changed on retrieval day 1, with a higher fraction of the COA being comprised of freeze-transition tuned somata. Individual somata with tunings to shocks, freeze-transitions, and place either during pre-shocks baseline or post-shocks retrieval day 3 were more likely to participate in a higher number of COAs post-shocks. With the NR-dCA1 pathway inhibited, more highly freeze-transition tuned somata were recruited into more COAs, a selective form of memory recall that potentially underlies their increased freezing behavior. Overall, we conclude that VR-CFC induces significant changes to neural encoding made by dCA1 by increasing context-dependent activity, inducing significant remapping, and inducing

tunings to shocks and freezing transitions, with population-level synchronized co-activations of dCA1 somata being selectively modulated in both frequency and size in the shocked context post-shocks. We additionally conclude that the NR-dCA1 pathway is a critical modulator of post-CFC-induced changes to dCA1 population encodings, ensuring neutral contextual discrimination is maintained, inducing increased remapping, preventing over-representation of shock and freeze-tuned somata, and moderating how somata activity are synchronized.

Results

Virtual Reality Contextual Fear Conditioning Induces Context-Selective Freezing

We trained mice to run in VR-CFC as previously described (Fig 3.1A)[346]. In brief, we trained water-restricted mice to run in a water-rewarded VR context, then removed the reward and ensured mice would run sufficiently (over >3 laps/minute) in two novel VR contexts with the tail-coat connected, as necessary to deliver shocks. Mice were then shocked in one context ('shocked') but not the other ('neutral'). They were then re-exposed to the two contexts each day for 3 days (Fig. 3.1A, Methods: Contextual Fear Conditioning Paradigm). Mice were exposed to one of three conditions. The first was no manipulation of the nucleus reuniens to the dCA1 pathway (NR intact). The second was using a DREADD strategy as previously described to selectively inhibit the NR to dCA1 pathway on the first day of post-shock context re-exposure (day 1) in a subset of mice (NR inhibited)[346]. The third was a control group of mice that underwent all other aspects of VR-CFC, including not receiving reward throughout the paradigm and running with the tailcoat, but were not shocked on day 0 (unshocked). All three groups of mice received intraperitoneal saline on all days except retrieval day 1, when they received the DREADD agonist DCZ, enabling both the NR intact and unshocked groups to serve as positive

controls for any effect of DCZ on day 1 (Methods:DREADD experimental protocol). All mice that were shocked showed a characteristic reflexive sharp increase in velocity from ~10 cm/s to ~68 cm/s immediately during and for ~1s immediately following shocks (Fig. 3.1B) before returning to baseline ~3s following shock initiation, consistent with stereotyped escape behavior observed in previous experiments[213,346]. This response additionally served as a check to ensure all mice in shocked groups received all shocks appropriately. Percent time spent freezing was used as a readout of contextual fear memory retrieval (CFMR), calculated as previously described [213,346](Methods: Behavioral Parameters).

All statistical calculations reported in this section are relative paired t-tests corrected for multiple comparisons for N>5 and Mann-Whitney U tests corrected for multiple comparisons for N<5 unless otherwise specified (see Methods: Statistics). Overall, mice in the NR intact condition (N = 9) froze significantly more in the shocked context on retrieval day 1 as compared to pre-shock baseline (Fig 3.1A). They froze 18.3±12.5% more in the shocked context compared to the pre-shocks baseline (P = 0.02) and froze significantly more than the neutral condition (P = 0.02), where they did not freeze more than baseline at 7.3±7.5% (P = 0.20). Mice maintained an elevated freezing level on day 2 at 16.0±8.7% compared to the pre-shocks baseline in the shocked context (P = 7.0e-3) but no longer significantly different from the neutral context freezing at 11.4±10.7% (P = 0.31). By day 3, freezing in the shocked context returned to baseline at 4.0±9.2% (P = 0.31), while freezing in the neutral context was slightly elevated at 11.1±11.3%% (P = 0.046), with the difference between them remaining insignificant (P = 0.24). We show that mice successfully learn to freeze in a shocked context while discriminating the shocked from the neutral context and subsequently learn to extinguish this fear response, in agreement with previous work[346].

Mice in the NR inhibited condition (N = 9) froze more in the shocked context as compared to baseline baseline on day 1, freezing 47.8±14.7% of the time (P = 6.8e-5), and froze significantly more than in the neutral context where they froze 30.0±13.2% of the time (P = 0.02). They additionally froze significantly more than mice in the NR intact condition in the shocked (P = 0.04) but not the neutral (P = 0.34) context, although freezing in the neutral context was significantly elevated from baseline (P = 6.7e-3), indicating NR inhibited mice increased their fear response to the neutral context (Fig. 3.1E, 3.1G). On day 2, freezing remained elevated in the shocked context at 37.6±13.9% as compared to the pre-shock baseline (P = 1.8e-4), and freezing also remained higher than in the neutral context (P = 5.9e-3). However, freezing in the neutral context dropped to 14.2± 6.7%, and when compared to the pre-shock baseline, was no longer significant (P = 0.11). By day 3, freezing in the shocked context remained slightly elevated above baseline at 12.1±8.7% (P = 0.02) but showed no difference from the neutral context (P = 0.67), which remained at baseline 9.6±7.8% (P = 0.18), or the NR intact 'shocked' context (P = 0.20). Our results largely replicate previous work, indicating that mice with NR inhibited on day 1 are significantly more fearful in the shocked context, generalize more to the neutral context, and take longer to extinguish the fearful freezing response[346].

Mice in the unshocked condition (N = 6) unexpectedly froze significantly more in the neutral context on day 1, freezing 26.5±5.1% of the time as compared to pre-shocks (P = 3.4e-3), despite not being shocked in the 'shocked' context, but not in the 'shocked' context where they froze 7.9±5.9% of the time (P = 0.27, Fig. 3.1F), with the difference between the contexts being significant (P = 1.5e-3). This effect disappeared on day 2 with freezing at 11.4±3.9% in the neutral (P = 0.07) and 7.8±8.2% in the shocked (P = 0.16). Freezing remained similar to the pre-shock baseline on day 3, with freezing at 19.7±15.7% in the neutral (P = 0.13) and 1.8±9.2%

in the shocked (P = 0.40), with the difference between contexts becoming insignificant (P = 0.33, P = 0.051). The significantly increased freezing time in the neutral context on day 1 and slightly elevated, although insignificant, freezing time on day 3 may also be partially a result of the experimental design, which had mice exposed to the 'neutral' context second and the 'shocked' context first on these days. This condition additionally had fewer mice (N = 6) than either the NR intact (N = 9) or NR inhibited (N = 9) conditions, contributing to higher behavioral variability in this condition to match previous results[346]. We expect this result to disappear with the addition of further mice to match other conditions.

As expected, the main cause of increased freezing time was the elongation of freeze length times, with freeze lengths elongated post-shocks on day 1 in the shocked as compared to the neutral context in both NR intact (shocked = 10.7 seconds ± 3.6%, neutral = 8.9s ± 5.7%, Mann-Whitney U P =5.2e-5) and NR inhibited (shocked = 14.2s ± 6.0%, neutral = 9.5s ± 3.8%, Mann-Whitney U P = 8.0e-3) conditions, with the unshocked condition showing no change between contexts ('shocked' = 6.1s ± 2.3%, neutral = 7.9s ± 3.3%, Mann Whitney U P = 0.63 Fig. 3.1I). Overall, these behavioral results mostly replicate previous work, except for the unshocked group having an elevated freezing preference in the neutral context.

Mice That Better Discriminate Contexts Better Extinguish Contextual Fear

We noted that across conditions, mice would respond differently to the paradigm, and additionally wanted to determine if there existed within-condition behavioral categorizations that distinguish mouse behavior using non-normalized freezing percentages. We split mice across conditions into three categories based on freezing behavior on day 1, then measured their freezing behavior on days 1, 2, and 3. Mice that learned and responded to the paradigm appropriately, i.e. mice that froze at least 5% more in the shocked context on day 1 than pre-shocks, and 5% more than the neutral context on day 1, were termed as 'discriminators' (Fig

3.1J). These mice made up 1/6 of the unshocked, 4/9 of the NR intact, and 5/9 of the NR inhibited conditions. The remaining mice could respond inappropriately in one of two ways. They could over-encode the fear by freezing at least 5% more in the shocked context on day 1 than pre-shocks, but additionally freezing within 5% of that level in the neutral context, termed 'generalizers' (Fig 3.1K). These mice made up 3/6 unshocked, 2/9 of the NR intact, and 3/9 of the NR inhibited groups. They could also learn the task incorrectly by not encoding the fear at all by freezing within 5% on day 1 as pre-shocks in the shocked context, termed 'runners' (Fig. 3.1L). These mice made up 2/6 of the unshocked, 3/9 of the NR intact, and 1/9 of the NR inhibited conditions.

These behavioral categories revealed that the generalizing mice froze at a significantly higher level compared to baseline of 58.9±12.8% in the shocked context (P = 8.2e-3) and 58.3±11.3% in the neutral context (P = 5.6e-3) for NR inhibited mice, and 29.9±14.9% in the shocked context (P = 0.04) and 14.7±3.2% (P = 0.01) in the neutral context for the NR intact mice, with no difference between the contexts (P = 0.48, P = 0.84). Generalizers in both the NR inhibited and NR intact conditions also did not return to baseline freezing in the shocked context by day 3 (P = 0.02, P = 0.04, Fig 3.1K). In contrast, as expected, the discriminator mice did show a signficiant difference between the contexts on day 1 in both the NR inhibited (shocked = 50.0±14.5% , neutral = 17.9±8.8%, P = 0.03) and NR intact (shocked = 34.0±5.6% , neutral = 0.0±8.8%, P = 8.7e-3) conditions, and by day 3 returned to baseline in the shocked context for both the NR inhibited (9.1±10.0% P = 0.24) and the NR intact (12.2±7.8%, P = 0.27) conditions. Since only one mouse was categorized as a discriminator in the unshocked condition, we did not conduct statistical analysis on him, but he froze 26.1% in the neutral and 16.8% in the shocked context on day 1, and 27.2% in the neutral and 1.0% in the shocked context on day 3,

discriminating on both days in the opposite direction (higher freezing in neutral than 'shocked') than the either of the shocked conditions. These results indicate that discriminator mice not only learned both more precise contextual fear responses but also more readily learned contextual fear extinction in the absence of additional shocks (Fig. 3.1J).

In mice deemed runners in the NR intact condition, mice exhibited no significant increase in freezing in the shocked (3.2±0.2, P = 0.97) or neutral (17.1±6.1%, P = 0.11) contexts on day 1, or on day 3 (shocked = 2.2±1.6%, P = 0.85, neutral = 3.7±8.5%, P = 0.29). In the unshocked condition, runner mice also exhibited no significant increase in freezing from baseline in the 'shocked' (3.2±0.2%, P = 0.36) or neutral (17.1±6.1%, P = 0.07) contexts on day 1, or on day 3 ('shocked' = 2.2±1.6%, P = 0.21, neutral = 3.7±8.5%, P = 0.39). Since only one mouse was identified as a 'runner' in the NR inhibited condition, we did not conduct statistical analyses on his behavior, but he did freeze at low levels in both the shocked (4.2%) and neutral (9.0%) contexts on day 1 and day 3 (shocked = 5.6%, neutral = 2.5%). These results suggest that runners across conditions did indeed fail to correctly associate the shocks with the shocked context.

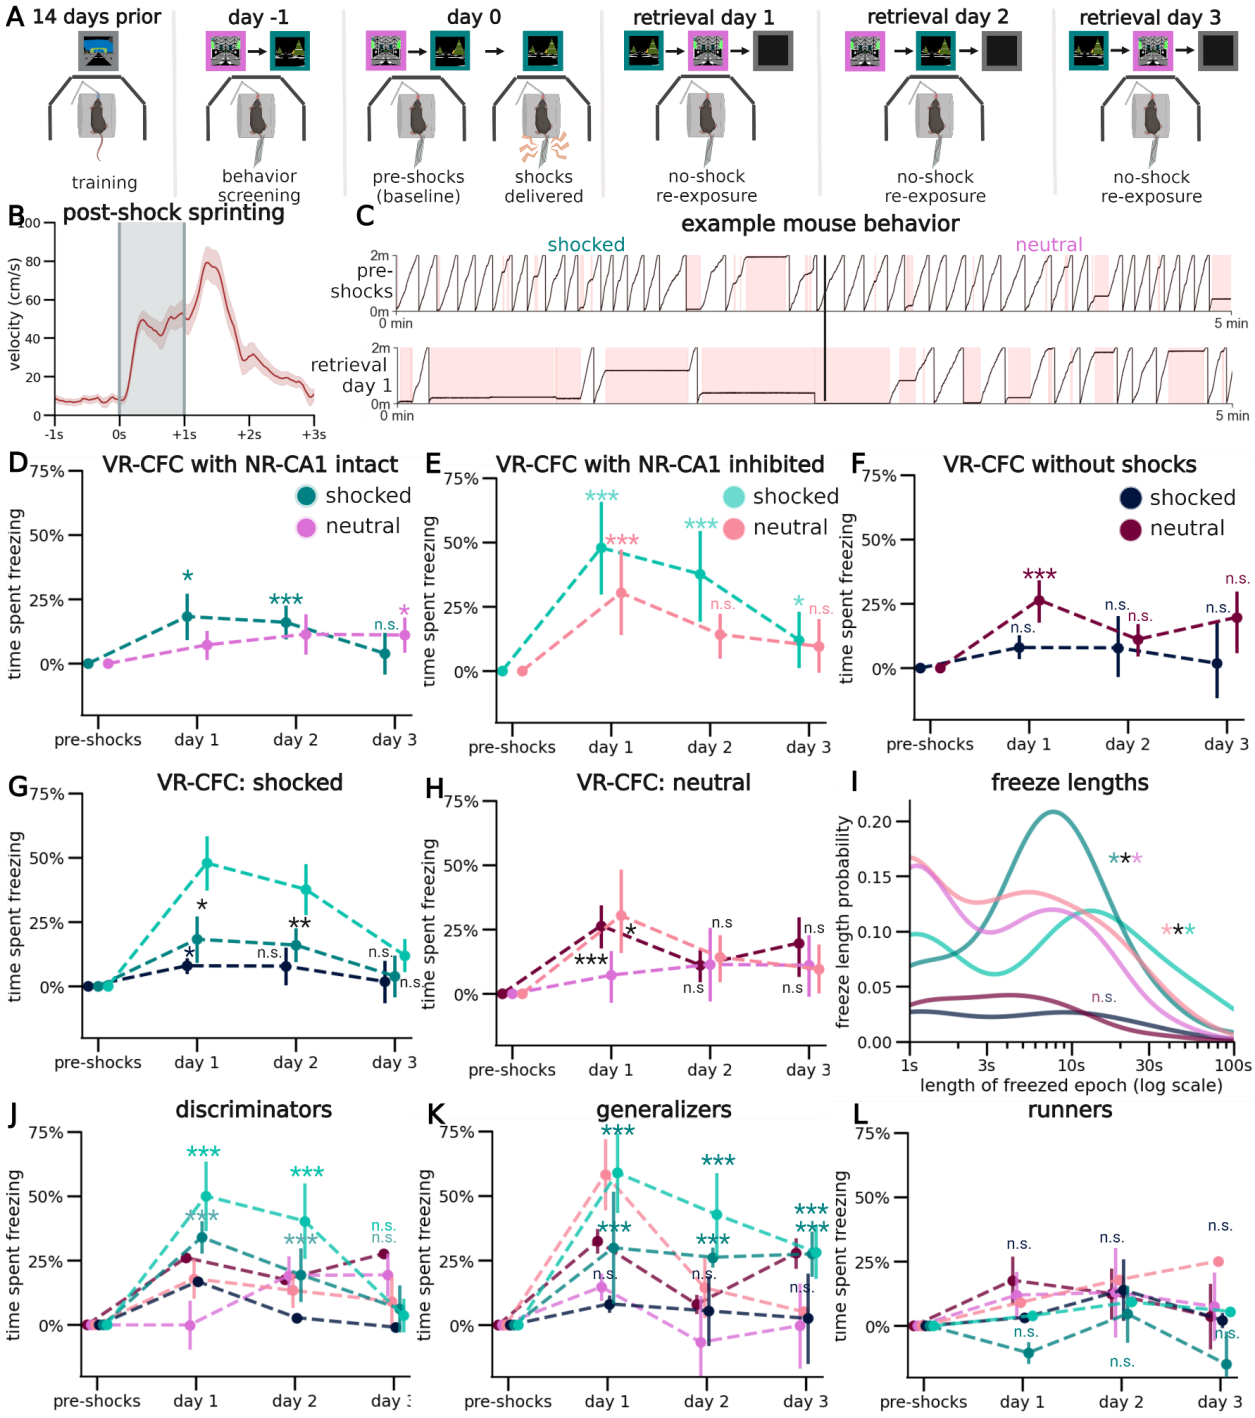

Figure 3.1: VR-CFC across contexts, conditions, and behavioral categories.
A) Schematic of experimental paradigm. Mice train for ~10-14 days in a familiar context with water, before being screened for suitability under VR-CFC by running with the tail-coat on in two novel, unrewarded VR contexts. Mice are then administered twelve 0.6mA tail shocks each 1s long 26s apart at pseudorandom locations along the track in one of the two contexts. Mice are next re-exposed to the two contexts, 'shocked' (teal) and 'neutral' (pink), as well as a dark context

with no VR display, for 3 subsequent days with no further shock administration. All mice received body-weight appropriate IP saline on all days except day 1, when they received DCZ. Unshocked mice did not receive shocks on day 0. **B)** All mice demonstrated during and immediately post-shock sprinting. Velocities were trial-aligned and shock velocity calculated. Dark gray lines indicate bounds of the shock, light gray indicates active shocking. Red line indicates average velocity 1s prior to, 1s during, and 2s after shock administration. Red shading indicates ±SEM. **C)** Example mouse behavior from an NR-inhibited mouse. Top: position on virtual track on day 0 before shocks in the to-be shocked (left) and neutral (right) contexts. Bottom: position on the virtual track on day 1 during NR-dCA1 inhibition in both the shocked (left) and neutral (right) contexts. Red shading indicates freezing epochs. **D)** Per mouse, we tested the difference in percent freezing within each context from pre-shocks baseline in the NR intact condition. N = 9 mice, error bars indicate ±SEM, significance stars indicate difference per-context from baseline. Teal indicates shocked context, orchid indicates neutral context. Mice froze significantly more in the shocked context post-shocks than pre-shocks on retrieval days 1 and 2, but not 3 (Student's T-test with holm-sidak correction, P = day 1: 0.02, day 2: 7.0e-3 , day 3: 0.31). Shocked mice did not freeze more in the neutral context post-shocks, except on day 3 (P = day 1:0.20 , day 2:0.23 , day 3:0.046). The difference between contexts was significant only on day 1 (significance stars not shown, P = day 1: 0.02, day 2: 0.31, day 3: 0.24). **E)** Same as D, but for NR-inhibited mice (N = 9), turquoise indicates shocked context, salmon indicates neutral context. Mice froze significantly more in the shocked context post-shocks than pre-shocks on retrieval days 1-3 (P = day 1: 6.8e-5, day 2: 1.8e-4, day 3: 0.02), and generlized to the neutral context on day 1 (P = day 1: 6.7e-3, day 2: 0.11 , day 3: 0.18). The difference between contexts was significant on days 1 and 2, but not day 3 (significance stars not shown, P = day 1: 0.02, day 2: 5.9e-3, day 3: 0.67). **F)** Same as D but for unshocked mice (N = 6), navy indicates 'shocked' context (i.e. the context mice would have been shocked in, were they in the NR-intact or NR-inhibited conditions), maroon indicates 'neutral' context. Mice did not freeze more in the 'shocked' context than pre-shocks baseline across days (P = day 1: 0.27, day 2: 0.13 , day 3:0.79). Mice did freeze significantly more in the neutral context over pre-shock baseline on day 1, but not days 2 or 3 (P = day 1: 3.4e-3, day 2: 0.07 , day 3: 0.16). The difference between contexts was significant on day1, but not days 2 or 3 (significance stars not shown, P = day 1: 1.5e-3, day 2: 0.33, day 3: 0.051). **G)** Same data as D-F, but plotting only the shocked/'shocked' context against each other. Color-coding remains constant, with significance stars indicating difference from NR inhibited to NR intact, or NR intact to unshocked. Freezing in the NR inhibited shocked context was significantly increased from the NR intact shocked context on days 1 and 2, but not day 3 (P = day 1: 0.04, day 2: 5.9e-3, day 3: 0.20). Freezing in the NR intact shocked context was significantly increased from the unshocked 'shocked' context on day 1, but not days 2 or 3 (P = day 1: 0.02, day 2: 0.19, day 3: 0.42). **H)** Same data as D-F, but plotting only the neutral context against each other. Color-coding remains constant, with significance stars indicating difference from NR inhibited to NR intact, or NR intact to unshocked. Freezing in the NR inhibited neutral context was significantly increased from the NR intact neutral context on day 1, but not days 2 or 3 (P = day 1: 0.02, day 2: 0.37, day 3: 0.57). Freezing in the NR intact neutral context was significantly decreased from the unshocked condition on day 1, but not days 2 or 3 (P = day 1: 8.5e-3, day 2: 0.49, day 3: 0.86). **I)** Freeze lengths were calculated as previously described. Significance stars indicate comparison between shocked and neutral for each condition (P = NR intact: 5.2e-5, NR inhibited: 8.0e-3, unshocked: 0.63). **J)** Percent time spent freezing color-coded by context and condition, for mice categorized as discriminators (N=10).

Significance comparisons are to pre-shocks baseline. Only shocked/'shocked' context comparisons are shown for visual clarity. Shocked NR inhibited (P = day 1:5.6e-3 , day 2: 0.01, day 3: 0.02), shocked NR intact (P = day 1: 4.5e-4 , day 2: 0.046 , day 3: 0.24). Neutral NR inhibited (P = day 1: 0.02 , day 2: 0.03, day 3: 0.17), neutral NR intact (P = day 1: 0.50, day 2: 0.15, day 3: 0.23). **K)** Same as J, except for generalizers. Shocked NR inhibited (P = day 1: 8.1e-3, day 2: 0.03, day 3: 0.02), shocked NR intact (P = day 1: 0.01 , day 2: 7.2e-3, day 3: 0.04), 'shocked' unshocked (P = day 1: 0.02, day 2: 0.34 , day 3: 044). Neutral NR inhibited (P = day 1: 5.7e-3 , day 2: 0.12 , day 3: 0.01), neutral NR intact (P = day 1: 0.04, day 2: 0.63, day 3: 0.50), neutral unshocked (P = day 1: 6.9e-4, day 2: 0.03, day 3: 3.2e-2). **L)** Same as J, except for runners. Shocked NR intact (P = day 1: 0.98, day 2: 0.33, day 3: 0.85), 'shocked' unshocked (P = day 1: 0.99, day 2: 0.84, day 3: 0.79). Neutral NR intact (P = day 1: 0.88, day 2: 0.76, day 3: 0.70), neutral unshocked (P = day 1: 0.08, day 2: 0.15, day 3: 0.39).

Population neural dynamics reveal link between better behavioral discrimination and more preferential activity in the shocked context

Inspired by Kinsky et al., we next set out to determine if these behavioral categories were reflected in dorsal CA1 neural dynamics across our three conditions[214]. To do so, we recorded changes in fluorescence from stable fields of view in dorsal CA1 across all conditions, across all 4 days of VR-CFC, by using CaMKII-GCAMP8f (Fig. 3.2A). We recorded near-exclusively from excitatory neurons, verified both visually through careful examination of morphology of identified ROIs in each field of view, and through co-expression of CaMKII (Fig. 3.2B: red). Only mice for which we were confident we stably recorded the same neurons over the entirety of the experiment were included in any analysis, neural or behavioral (see Methods: Image Processing and ROI selection). We ensured the quality of recordings were comparable across days with similar proportions of cells active on each day out of the total number of ROIs identified across all days, as compared to the first pre-shocks day of recording (P = 0.64, P = 0.59, P = 0.93), identifying transients in ~95% of somata active each day, and by not observing a visual difference in transient quality during analysis (Fig. 3.C-D). We next examined differences in neural activity between contexts by calculating a neural discrimination index (nDI), where the

more negative the metric, the more differentially active the somata were in the shocked context versus the neutral context (Methods: Neural Discrimination Index)[214].

We found that in the NR intact condition, somata trended towards being preferentially active in the shocked context on day 1 with a negative mean nDI of -0.08±0.04, while that trend reversed on day 2 to 0.06±0.03, and returned to baseline similar activity on day 3 of 0.01±0.03, although none of these trends reached significance. In the NR inhibited condition, the nDI trended positive at 0.04±0.05, which was significantly more positive as compared to the nDI in the NR intact condition ($P = 0.049$) while activity on subsequent, non-NR inhibited days was similar to that of the NR intact condition (Fig 3.2E, $P = 0.68$, $P = 0.55$). A more positive nDI on day 1 could either be caused by more activity in the neutral context, or less activity in the shocked context. We examined the root cause of this difference in nDI, and found that an increase of activity in the neutral, from a mean 15.2 peaks/s in NR intact mice to a mean 18.3 peaks/s, and a slight increase in activity in the shocked, from a mean 17.1 peaks/s to 18.4 peaks/s, was driving the more comparable nDI. This result indicates that inhibiting the NR-dCA1 pathway increases excitability in dCA1 somata.

We next examined the behavioral discrimination index (bDI), calculated as previously described by finding the difference in freezing activity between the contexts, where the more negative the metric the more differential freezing occurred in the shocked context (Methods:Behavioral Parameters - Behavioral Discrimination Index)[97,346]. Here, as expected based on previous work, we found a significantly more negative bDI between the NR inhibited and unshocked groups on day 1 ($P = 0.01$), and between the NR inhibited and NR intact on day 2 ($P = 0.03$, Fig. 3.1D-F). These results indicate that NR inhibited mice have a strong, contextually modulated freezing preference on day 1, and that they better discriminate between

contexts on day 2 than NR intact mice - which is aligned with the still-significantly heightened freezing on day 2 in the shocked context, and with the neutral context's day 2 reduction in freezing, as reported earlier (Fig 1E). We hypothesized that the increase between days 1 and 2 in the nDI - i.e. the flip from a negative bDI preferring the shocked context to a positive bDI preferring the neutral context - was driven by mice that discriminated well between the contexts. As expected, the discriminators showed the most negative bDI, significant difference from both the runners (P = 1.8e-3) and generalizers (P = 1.9e-4), as well as a significant decrease from zero (P = 3.9e-4, Fig. 3.2H).

We lastly compared the nDI with the bDI across first conditions, then across behavioral categories (Fig 3.2G, 3.2J). We found that across conditions, the more preferentially active the somata were in the shocked context -i.e. a more negative nDI - the more preferentially mice froze in the shocked context - i.e. a more negative bDI. On retrieval day 1, this was true in both the shocked conditions NR intact and NR-dCA1 inhibited, although the majority of the nDIs for NR intact mice were negative, while the majority of nDIs for NR-dCA1 inhibited mice were positive (Fig. 3.2G). As expected, we did not observe this trend in unshocked mice. Additionally, on retrieval day 1, when we split mice by behavioral category, we found that discriminators were largely driving this association, indicating that mice with better neural discrimination were also better at associating fear with the appropriate context (Fig. 3.2J). Overall, our results indicate that inhibition of the NR-dCA1 pathway disrupts the selective increase in firing rates in the shocked context for mice that discriminate between contexts by increasing overall firing in both contexts.

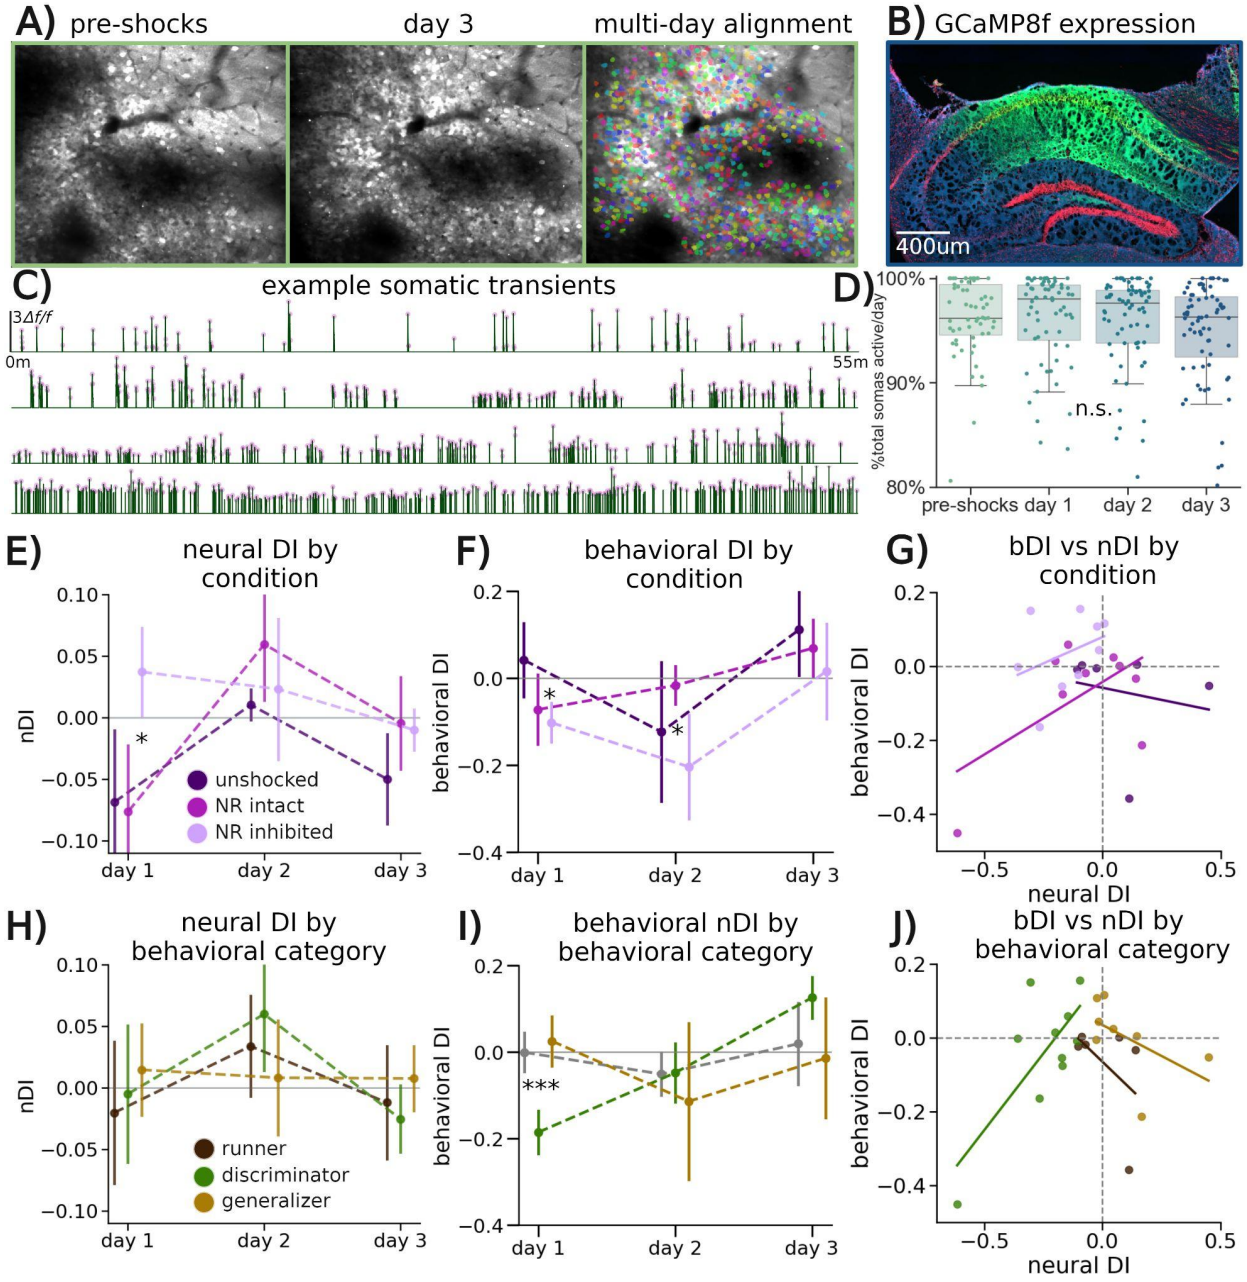

Figure 3.2: Multi-day somatic imaging reveals increased neural selectivity for the shocked context is diminished during NR-dCA1 inhibition

A) Population of somata across days for an example mouse in the NR intact condition. Left; field-of view (FOV) pre-shocks. Middle: FOV post shocks on day 3. Right: FOV for all 4 days with somata identified in randomized multi-color (N = 756 somata). **B)** Example confocal image of GCaMP8f expression in dCA1 (green). CaMKII stain is observed in red, and DAPI in blue. All cells expressing CGaMP8f also expressed CaMKII. Note tissue removed above dCA1 for cannula. **C)** Four example somatic traces across the entire paradigm are shown from 4 different example mice/somata, demonstrating representative activity across all days, and similar peak *Δf/f* levels across days.Top, middle-top: transient examples from NR intact mouse. Middle bottom:

transient example from an NR inhibited mouse. Bottom: transient example from an unshocked mouse. **D)** Boxplot indicating fraction of somata active per experimental day across all conditions, with no significant differences from pre-shock baseline mean percent active of 95.9%. (percent active day 1: 95.6%, day 2: 95.7%, day 3: 95.1%), P = day 1: 0.64, day 2: 0.59, day 3: 0.93). No differences were observed between days per condition or behavioral category (data not shown). **E)** Neural discrimination index (nDI) was calculated as number of transient peaks per soma in the shocked/'shocked' context minus transient peaks per soma in the neutral context divided by the sum of transients in both, then averaged across all somata per mouse, with the dot representing the cross-mouse average per each condition. Error bars indicate ±SEM, significance star indicates difference between NR inhibited and NR intact on day 1. (P = day 1: 0.049, day 2: 0.68 , day 3: 0.55). Remaining cross-condition comparisons were not significant. **F)** Behavioral discrimination index (bDI) was calculated as fraction time spent freezing in the shocked/'shocked' context minus fraction time spent freezing in the neutral context, divided by the sum of both fractions spent freezing. Significance stars represent NR inhibited to unshocks on day 1 (P = 0.01) and NR inhibited to NR intact on day 2 (P = 0.03). Remaining cross-condition comparisons were not significant. **G)** Linear regressions were computed between the nDI and bDI of each condition for day 1. NR intact (slope = 0.39, intercept = -0.04, $r^2 = 0.35$, p = 0.09), NR inhibited (slope = 0.29, intercept = 0.08, $r^2 = 0.12$, P= 0.34), unshocked (slope = -0.13, intercept = -0.058, $r^2 = 0.04$, P = 0.71). **H)** nDI calculated as in E, except across behavioral categories. No comparisons were significant. **I)** bDI calculated as in F, except across behavioral categories. Significance star is between discriminators and baseline (P = 3.9e- 4). No other comparisons were significant. **J)** Linear regressions were computed between the nDI and bDI of each behavioral category across both shocked conditions for day 1. Discriminators (slope = 0.83, intercept = 0.17, $r^2 = 0.55$, P = 0.01), generalizers (slope = -0.34, intercept = 0.04, $r^2 = 0.27$, P = 0.18), runners (slope = -0.62, intercept = -0.07, $r^2 = 0.23$, P = 0.34).

Place fields globally remap following VR-CFC in NR intact, but not NR inhibited or unshocked mice

We next examined the spatial tuning of the somatic population before and after VR-CFC. We used a modified peak identification method that identified an average of 38 place fields (PFs) per mouse in the shocked context and 53 PFs in the neutral context pre-shocks (day 0), 48 PFs in the shocked context and 61 PFs in the neutral context on retrieval day 3 (Methods: Place Field Identification). The specifics of our method were calibrated to minimize identification of place fields in the dark context, which would necessarily be false positives, while retaining a majority of place fields in the shocked and neutral contexts, resulting in a slightly decreased fraction of PFs than reported previously (pre-shocks shocked context = 10%, pre-shocks neutral context =

13.9%, day 3 shocked context = 12.0%, day 3 neutral context = 15.5%, Fig 3.3B). We did not find any differences in the number of place fields identified between days or contexts (Fig 3.3B). We did not examine place-tuning on days 1 and 2 due to a low level of running induced by VR-CFC preventing reliable place field identification.

We found that in our unshocked mice, place maps in both contexts were still relatively unstable by day 3, with 25.2% of somata in the 'shocked' and 21.5% of PFs in the neutral remaining stable, as identified by retaining a place field in the same location. Of the PFs that destabilized, the vast majority vanished, with a small proportion of 4.6% in the 'shocked' and 10.0% in the neutral remapping to a new location. No statistically significant differences were found between contexts in any of these three place field fates. This level of instability has been reported previously on this timescale, and could be explained by the removal of reward triggering partial remapping over 4 days[145,297,347]. However, the unshocked mice still retained a coherent map of the environment, primarily by recruiting >65% new somata to take over with new place fields (Fig 3.3A). The new map had similar qualities between the first and last day of recording across three place field metrics, with the exception of a decrease in average lap emergence in the 'neutral' context ($P = 0.04$, Fig 3.3D-F).

In contrast, mice that underwent VR-CFC experienced heightened remapping in the shocked context, with 3/8 mice presenting completely orthogonal place maps between pre-shocks and retrieval day 3 (Fig 3.3C). Overall, mice in the NR intact condition retained only 7.2% stable place fields in the shocked context after extinction, and 7.7% stable place fields in the neutral context. Similar portions of the non-stable population remapped in the NR intact condition as compared to the unshocked condition: in both the shocked and neutral contexts, 5.7% and 5.4% respectively, with the remaining vanishing. The new map post-extinction,

however, did show distinct dissimilarities when compared to the unshocked condition map in the shocked context on experimental day 3, largely in being a more flexible and less stably precise representation (Fig. 3.3D-F). Comparing every context to the assumption that there is no change between pre-shocks and day 3, the place map is made of wider PFs in the shocked context (P = 8.0e-3), has a higher percentage of out-of-field firing (P = 8.9e-3), and unchanged lap emergence as compared to the pre-shocks map (P = 0.87). These effects were heightened in the shocked as compared to the neutral context in width (P = 0.04) and lap emergence (P = 0.03), although not in out-of-field firing (P = 0.56). These results indicate that post VR-CFC extinction learning drives an increase in more distinct contextual representations, potentially to support context-dependent memories.

Lastly, we examined place map instability in NR-inhibited mice. We found that they remapped at comparable levels to unshocked mice, retaining 18.1% in the shocked and 23.5% in the neutral context of their place fields, with the remapping level at 8.7% in the shocked and 10.6% in the neutral, and the remaining vanishing. This represented significantly less remapping than the NR intact group, especially in the number of stable somata in the shocked context (P = 4.0e-3), and was comparable to levels in the unshocked group (P = 0.36, Fig 3.3 C). The place field metrics of the post-CFC map reflected this higher stability, as change in PF width and change in PF out of field firing were statistically similar to those in the unshocked condition for both contexts. Lap emergence was the only metric that differed, with place fields emerging on average 2.2 laps earlier in the shocked (P = 8.8e-3) and 1.0 laps earlier in the neutral context (P = 0.02.

Overall, we found remapping occurred on an experiential gradient across our conditions. In unshocked mice, remapping is already significant over the four day span, likely due to high

representational drift from the lack of water reward or goal-motivated behaviors, but the overall quality of the map remains stable with time. There is significantly more remapping in mice that have undergone CFC, with some NR intact mice showing global remapping, and the post-CFC map is comprise of wider somata that fire outside their place fields more often, indicating the incorporation of fear and extinction learning impacts how dCA1 is encoding place. Lastly the NR-dCA1 inhibited mice fall in the middle, remapping more than unshocked mice, but less than NR intact mice, despite having a larger behavioral increase in freezing. This result indicates that post-CFC remapping does not drive freezing behavior, but instead may represent extinction learning, a process that is interrupted in NR-dCA1 inhibited mice.

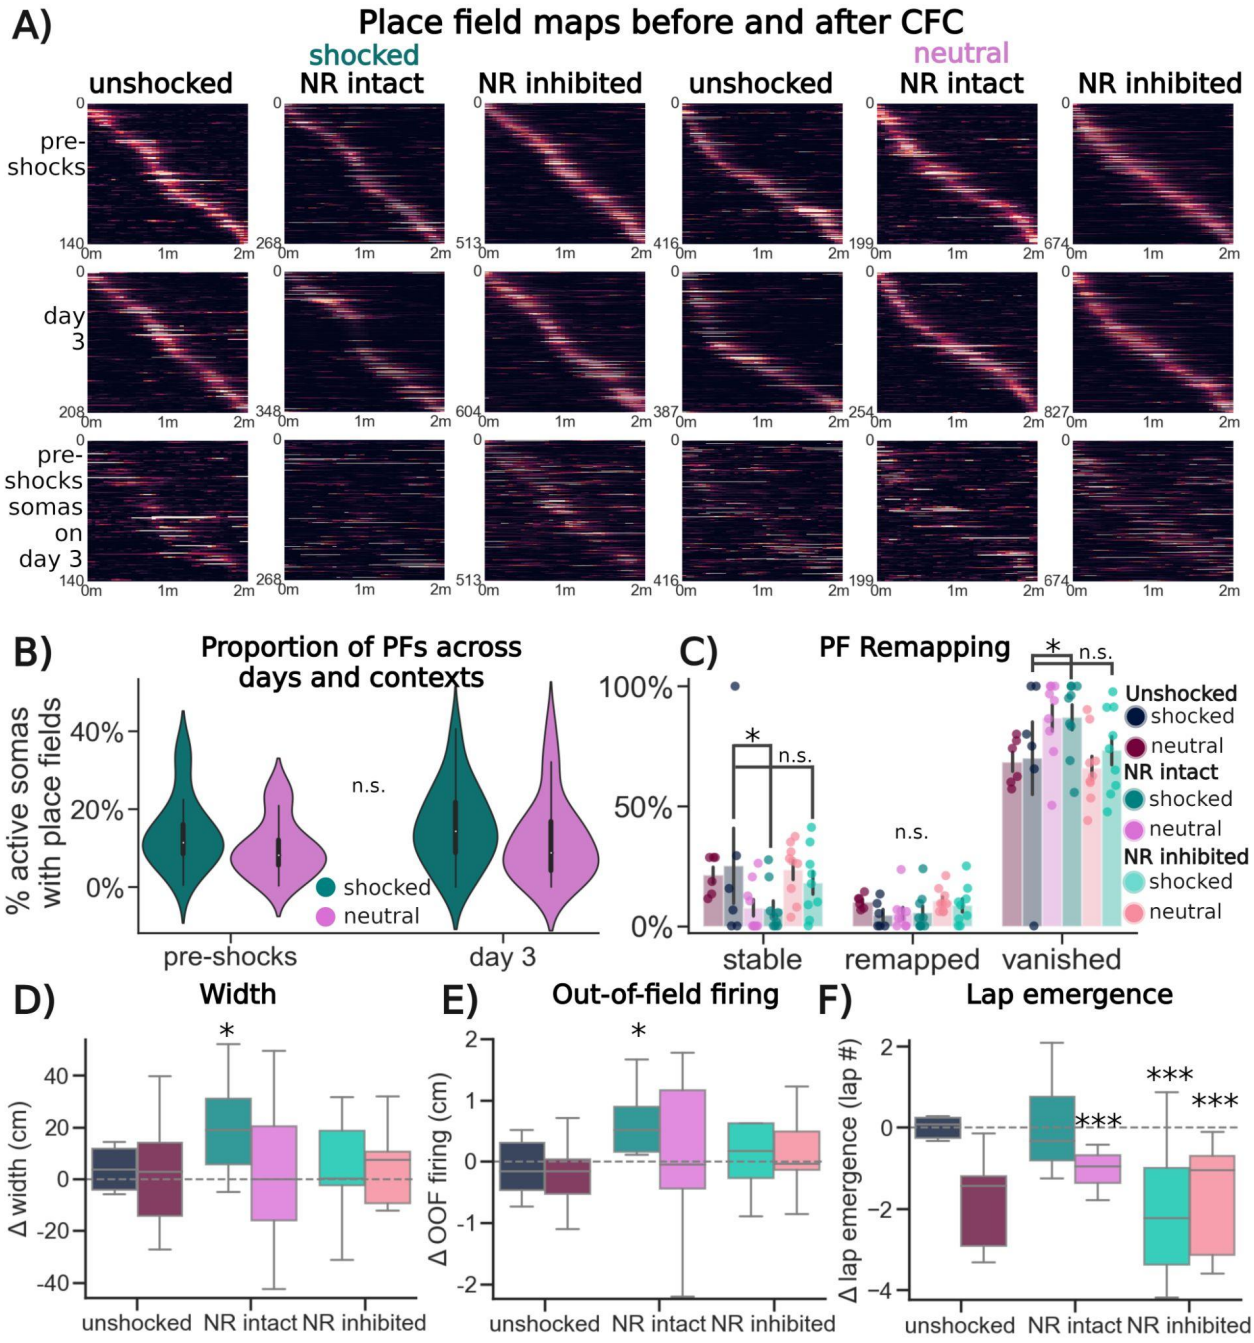

Figure 3.3: Place fields remap more in NR intact than NR-dCA1 inhibited mice, and more than unshocked mice

A) Place fields (PFs) across all mice where average COM location across all laps is plotted as a single row. Scale is black = no activity, white = 20*Δf/f*. PFs before CFC in pre-shocks (top row), after CFC on day 3 (middle row), and PFs from pre-shocks plotted after CFC, i.e. remapping (bottom row) are plotted across contexts (shocked: left three columns, neutral: right three columns) and conditions (left to right, unshocked, NR intact, NR inhibited). **B)** Percent of active somas, defined as somas with at least 1 transient above 0.1*Δf/f* in the context/day pairing, with PFs, are shown as violinplots. No significant differences between contexts (P pre-shocks = 0.06,

day 3 = 0.25) or within-context before and after CFC (P shocked =0.23, P neutral = 0.28), were found. **C)** PF fate between pre-shocks and day 3 were categorized as stable, i.e. a soma retaining a PF on day 3 within a context with any overlap to its PF on pre-shocks, remapped, i.e. a soma with a PF on pre-shocks moving to a new location but retaining its identity as a PF, or vanished, i.e. a soma with a PF on pre-shocks no longer retaining a PF anywhere on day 3. Significance stars represent comparison between the unshocked 'shocked' context to shocked context per PF fate of the NR intact (P stable: 0.03, remapped: 0.61, vanished: 0.01) or NR inhibited (P stable: 0.08, remapped: 0.17, vanished: 0.41) condition. Neutral comparisons are omitted for visual clarity and reported as follows between the unshocked neutral context to the neutral context per PF fate of the NR intact (P stable: 0.08, remapped: 0.09, vanished: 0.01) or NR inhibited (P stable: 0.36, remapped: 0.41, vanished: 0.64) condition. **D)** PF width was calculated as the difference in centimeters between the first and last bin with in-field activity. Graph shows the change in PF size between pre-shocks and day 3, with negative values indicating smaller PFs on day 3, and positive values indicating larger PFs on day 3. Statistical comparisons are per context to baseline(P shocked NR intact = 8.0e-3, all other comparisons not significant). **E)** Out-of-field firing was calculated as the mean $\Delta f/f$ in bins within the PF divided by the $\Delta f/f$in bins outside the PF. Graph shows the change in out-of-field firing between pre-shocks and day 3, with negative values indicating less out-of-field firing and positive values indicating more out-of-field firing on day 3. Statistical comparisons are per context to baseline (P shocked NR intact =8.9e-3, all other comparisons not significant). **F)** Lap emergence was calculated as the first lap with activity within the PF within a context, per day. Graph shows change in lap emergence between pre-shocks and day 3, with negative values indicating earlier lap emergence and positive values indicating later lap emergence on day 3. Statistical comparisons are per context to baseline (P neutral NR intact = 5.7e-4, P shocked NR inhibited = 8.8e-3, P neutral NR inhibited = 0.02).

dCA1 tunes strongly to shocks and freezing epochs

Due to sharp increases in running post-shocks inducing motion instability in imaging (Fig 3.1B), only 5/18 shocked mice had stable enough FOVs to analyze the shock window. All shock-tuning analyses are limited to these mice. Despite this reduced sample size, we still found a robust pool of excitatory pyramidal neurons tuned to shocks, with 540 total out of 2,617 active somata, i.e. ~20%, correlated to the the periods in which shocks were administered, above a shuffled threshold of 99% (Fig 2.4A-C, Methods: Shock tuning). We next looked to see if mice with better or more shock-tuned somata discriminated the contexts better or worse. We found that behavioral categories had significantly different shock-tuning correlations, with discriminators being better-tuned than runners (P = 0.03) or generalizers (P 4.6e-6), and runners

being better tuned than generalizers (P = 1.3e-5), indicating that better shock tuning in somata on day 1 may be predictive of better fearful contextual discrimination post-shocks. (Fig 3.4C).

We next looked for tuning to freezing epochs, both looking for an increase in activity within ±1s of a freezing epoch start, or separately, within ±1s of a freezing epoch end (Fig 3.4D-I). We chose these timings due to other groups reporting freeze-tunings in this range[159,214,340,348]. We found a small portion of somata tuned to each, with a subset tuned to both that we will focus on, due to large freezing epochs making the realistic distinction between the groupings difficult to identify (Fig 3.4D-E). Unexpectedly, we found no relationship between somata being shock-tuned during shocks and later showing significant freeze on/offset tuning, indicating that while freeze-tuned cells do predict later mouse behavior, they do not do so through later tuning to freeze-transitions (Fig 3.4F). However, the quality of freeze tuning in freeze on/offset tuned somata was heightened in the shocked as compared to the neutral context on post-shock day 1 in discriminators in the NR intact condition (P = 0.03), but not on any other experimental day or behavioral category indicating that freeze-tuning may play a separate role in post-shock context discrimination (Fig 3.4G,I). In the NR inhibited condition on day 1, freeze-tuning was significantly heightened in the shocked context as compared to the neutral context in both discriminators (P = 0.03) and generalizers (P = 4.1e-4).

We additionally found that the quality of freeze tuning in the NR inhibited condition was significantly heightened as compared to the NR intact group on day 1 across behavioral conditions in the shocked (P = 0.01) but not the neutral (P = 0.46) context, which disappeared when NR-dCA1 was no longer actively inhibited on days 2 (P shocked = 0.17, P neutral = 0.41) or 3 (P shocked = 0.16, P neutral = 0.09). These results indicate that the NR-dCA1 pathway may play an important role in suppressing the activity of freeze-tuned somata in dCA1.

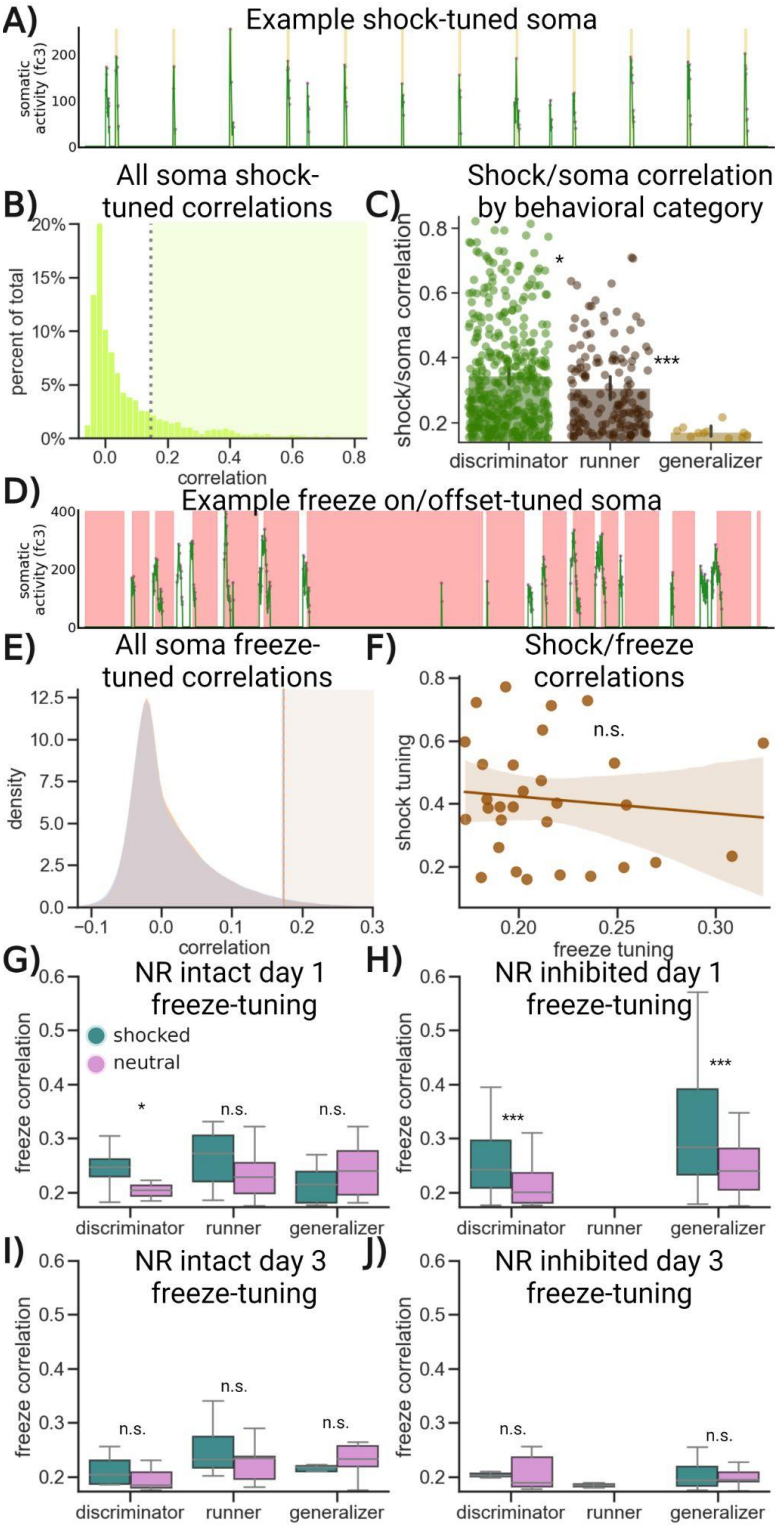

Figure 3.4: Somata tune to shocks and predict post-shock contextual discrimination, while freeze-tuning is suppressed by the NR-dCA1 pathway

A) Example *Δf/f* of a strongly shock-tuned soma in green trace. Pink dots indicate identified peaks. Lime-green shading indicates shock periods. This soma was 0.81 correlated with shock periods. **B)** Histogram of the correlations of all active somata measured (N = 2,617 against shock periods). The 99th percentile of a shuffle distribution is shown in a dotted grey line at 0.16. Yellow shading indicates somata identified as significantly shock-tuned (N = 540 somata). **C)** Barplot of the shock/soma correlations separated by behavioral category. Stars indicate differences between discriminators and runners (Mann-Whitney U P = 0.03) and runners and generalizers (P = 1.3e-5). Discriminators and generalizers were also highly significantly different (P = 4.6e-6). **D)** Example *Δf/f* of a strongly freeze on/offset tuned soma in green trace of a NR-inhibited mouse in the shocked context on day 1. Pink dots indicate identified peaks. Red shading indicates freezing epochs. This example soma has 0.61 correlation to freezing onset and 0.49 correlation to freezing offset. **E)** Kernel density estimate histograms of freeze-onset (cornflower blue) and freeze offset (rust orange) correlations across all somata (N = 10,034) show similar distributions and tuning cutoffs (orange/blue line) at 0.17 for significance. Brown shaded region is significant somata (N= 889

somata tuned to both freeze onset and offset). **F)** A linear regression was run on the subset of somata with both significant shock tuning and significant post-shock freeze onset and offset tuning in the shocked context on day 1 (N = 29), showing no relationship between somata that were previously shock tuned and being freeze-tuned in the shocked context post-shocks (slope = -0.76, intercept = 0.5 r2 = 0.03, P = 0.37). **G)** For the subset of freeze onset and offset tuned somata (N = 889), we computed the on/offset mean correlation, then plotted this correlation by context across behavioral categories in the NR intact condition. Only discriminator mice showed a significant difference between the contexts on day 1 (P = 0.03). **H)** As in G, except for the NR inhibited condition. Both discriminators (P = 0.01) and generalizers (P = 4.1e-4) showed significant increases in freeze-tuned correlations in the shocked as compared to the neutral context. No dual-tuned somata were identified in the mouse identified as a runner in the NR-inhibited condition. **I)** As in G, except on day 3. No significant context-dependent tuning was observed in any behavioral category. **J)** As in G, except in the NR inhibited condition and on day 3. No significant context-dependent tuning was observed in any behavioral category.

Synchronous co-activations are both prevalent and behaviorally relevant in dCA1

We next wanted to understand how coordinated ensembles of activity, potentially related to hippocampal replays or engrams, may be identified and studied using calcium imaging techniques. We developed a new technique, in conversation with previous work, to identify ensembles of co-active somata above a shuffled threshold baseline[213,348]. In brief, we found frames requiring that at least two somata had 50% or greater of their peak prominence located within that frame, identifying them as putative components of synchronized co-activations (COAs), although in practice no retained group of somata had fewer than ten somata. We then used a peak finding algorithm to identify series across those frames that had at least 50% of the peak number of somata active, which formed the temporal bounds of the co-activation (Fig 3.5A-B). We lastly used a threshold of the 99th percentile in shuffled data separately determined for each mouse, day, and context, and disregarded any putative COAs identified below this threshold.

We identified 10,117 COAs total across all mice and days of recordings, for a median of 81 COAs per mouse/day. COA lengths temporally ranged from 0.07s to 4.7s, with a median length of 0.52s (Fig 3.5C). COAs recruited anywhere from 10 to 331 participating somata, with

an overall median of 22 somata participating per COA (Fig 3.5D). We next examined how frequently any particular somata was recruited into a COA, and found that on average, each somata participated in a median 4 COAs, with somata participating in a range from 1 to 64 total COAs (Fig 3.5E). Somewhat unexpectedly, many COAs are ruled out from being considered putative replay events, as 75.9% of COAs occur during running epochs, despite mice running epochs only representing 59.4% of all recorded velocities. This indicates that COAs preferentially occur during running epochs.

We were next interested in characterizing if the number of COAs changed by condition throughout VR-CFC in the shocked or neutral context, as compared to the number of COAs in each context by condition pre-shocks (Fig 3.5 F-H). We found both no difference between the shocked and neutral context or any difference from pre-shock baseline levels in the unshocked mice, indicating that despite the lack of reward, preferential freezing in the control context, and significant place cell remapping, COA rates remain steady. COA rates remain constant across days (Fig. 3.5F). In sharp contrast, the number of COAs in post-shock NR intact mice increased by a median 100% in the shocked and 50% in the neutral context as compared to the number of COAs pre-shocks (Fig. 3.5G). This increase from baseline persists in the shocked, but not the neutral context, on day 2 (P shocked = 0.04, P neutral =0.1), with both returning to baseline on day 3 (P shocked = 0.06, P neutral = 0.06). Given the general preference of COAs to happen during running epochs, and the significant increase in freezing epochs after shocks, we do not believe this increase is simply due to longer pausing times but is instead linked to the behavioral salience of the shocks. In contrast, COA counts in the NR inhibited group were only significantly heightened from baseline in the shocked context on day 1, being 53% more likely to occur (P = 0.01) and failed to reach significance in the neutral context on all post-shock days or shocked

context on subsequent days, indicating that an active NR-dCA1 pathway during the first re-exposure to the shocked context increases the number of these coordinated somatic events (Fig. 3.5H).

We next examined how many somata were participating in each COA. Once again, in the unshocked condition we did not observe any differences between contexts or from baseline in this metric, further indicating the stability of COAs across days when the context carries no behavioral salience (Fig. 3.5I). Once again in contrast, the post-shocked NR intact mice increased the number of somata recruited to the somata by 23.9% ($P = 0.01$), but selectively in the shocked as compared to the neutral context, which selectively decreased 11.5% ($P = 0.02$), both only on day 1 (Fig. 3.5J). In the NR inhibited condition, the increase in number of somata recruited in the shocked context was similarly significant to the NR intact condition at 23.5%, but the selectivity for the shocked context disappeared, with the the neutral context also elevating at 21.4%, with the difference between the neutral contexts being significant $P = 3.1e\text{-}3$, Fig. 3.5K). Intriguingly, on retrieval day 2 the COAs in both contexts of the NR inhibited condition recruited more somata, with a 58.1% increase in the shocked and 40.1% increase in the neutral, indicating a 'rebound' effect, which returned back to day 1 levels by day 3 at a 20.9% increase in the shocked and 23.3% increase in the neutral context (Fig. 3.5K). Overall, these results indicate that an intact NR-dCA1 pathway is important to regulate the selectivity between the number of somata recruited in the shocked and neutral contexts by suppressing inappropriate neutral context somatic recruitment on post-shock retrieval day 1.

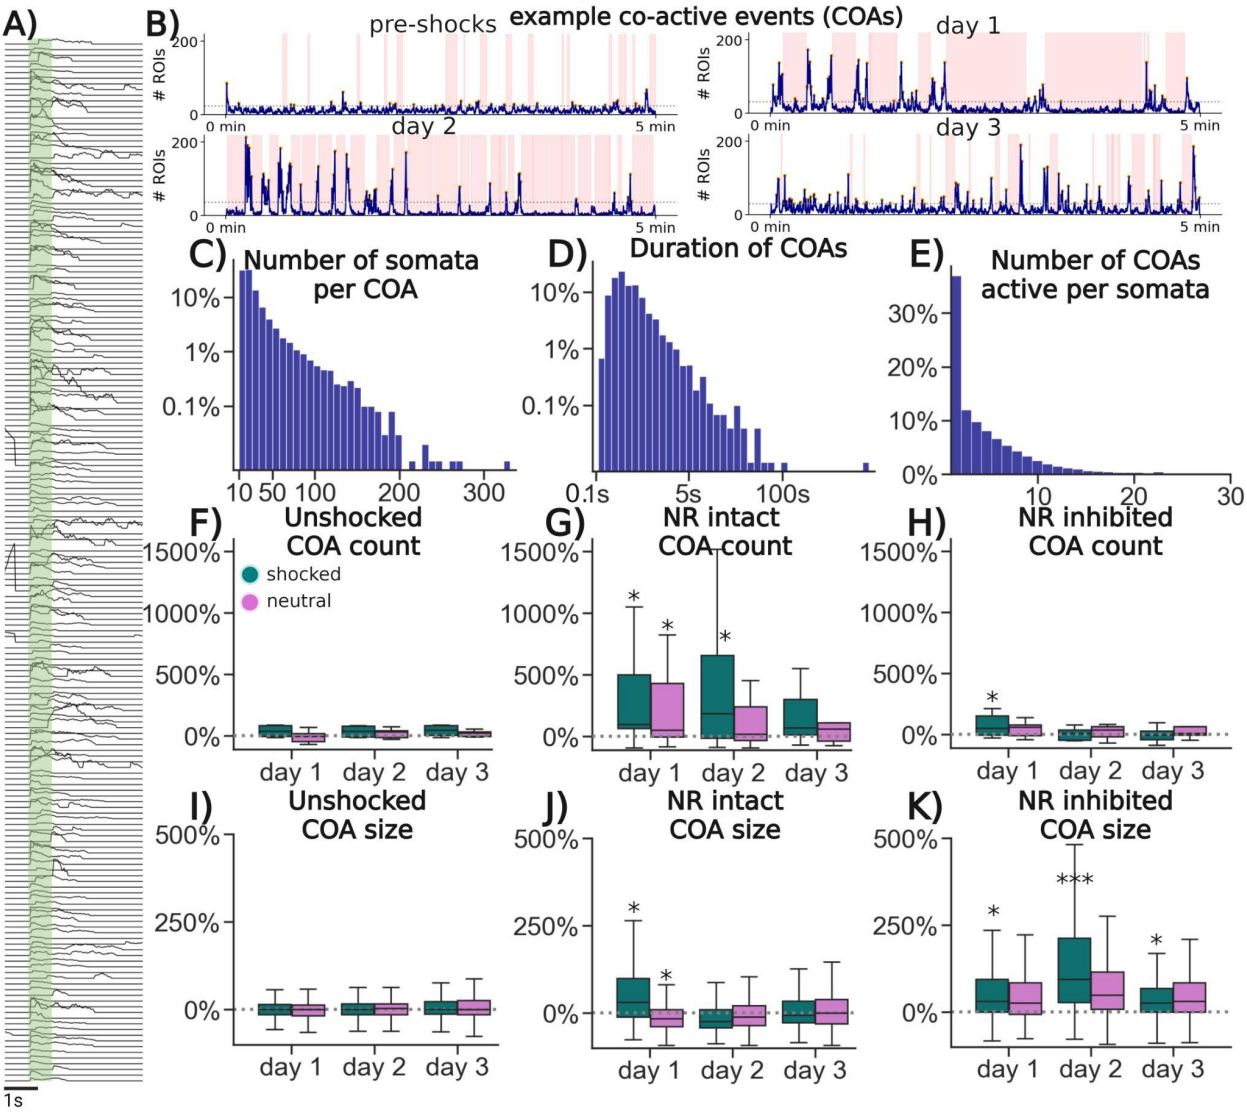

Figure 3.5: Synchronized co-activations are induced by the NR-dCA1 pathway post-shocks
A) Example co-activation (COA) from an NR inhibited mouse in the shocked context on day 1 involving 212 somata. Each black line is one somatic *Δf/f* transient. Temporal duration of co-activation is highlighted in green (0.78s). **B)** Example COA summaries from each day of CFC in an NR inhibited mouse.The blue line tracks average simultaneous activity occurring at per-frame resolution. Grey dotted line indicates threshold for significant COAS. Orange dots indicate identified COA peak. Red shading indicates freezing epochs. **C)** Percent-wise histogram of the number of somata involved in each COA across all mice, days, and conditions. **D)** Percent-wise histogram of the duration of COAs across all mice, days, and conditions. **E)** Number of COAs each soma that is active in at least one COA participates in across all mice, days, and conditions **F)** Percent increase in number of COAs in the unshocked condition in the 'shocked' (teal) and neutral (orchid) contexts, calculated per day. Dotted line indicates no change from pre-shock baseline. No significant change from pre-shock baseline or between contexts was observed in any day or condition in unshocked mice. **G)** As in F, but for the NR intact condition. Significance stars indicate significant change from baseline in one-sample students' t-test with

holm-sidak correction (shocked P = day 1: 0.03, day 2 = 0.04, day 3 = 0.06, neutral P = day 1: 0.03, day 2: 0.1, day 3 = 0.06). **H)** As in F, but for the NR inhibited condition (shocked P = day 1: 0.01, day 2 = 0.28, day 3 = 0.29, neutral P = day 1: 0.12, day 2: 0.13, day 3 = 0.14). **I)** Percent increase in mean number of somas involved per COA in the unshocked condition in the 'shocked' (teal) and neutral (orchid) contexts. Dotted grey line indicates no change from pre-shock baseline. No comparisons were significant. **J)** As in I, but for the NR intact condition. The shocked context on day 1 had significantly more somata recruited to its coas than the pre-shock baseline (P = 0.01), while the neutral context had significantly fewer somata recruited to its coas than the pre-shock baseline (P = 0.02). No comparisons were significant on subsequent days. **K)** As in I, but for the NR inhibited condition (Shocked P = day 1: 0.01, day 2: 1.4e-3, day 3: 0.01, neutral P = day 1: 0.06, day 2: 0.06, day 3: 0.07).

Co-activations in relation to other somatic tunings

We next wanted to investigate the makeup of these co-activation somata in relation to other tunings to place, shock, or freezing as previously discussed, with a focus on the shocked context in both NR intact and NR inhibited conditions (Fig 3.6A). Omitted in this figure are the unshocked condition and somata without tunings investigated previously, labeled here as 'untuned' somata, although they potentially have responses to phenomena external or internal we did not investigate. These somata make up the majority of COAs in all cases, although their number is altered by condition. In the unshocked condition, the vast majority of COA somata are untuned, at 80.8% across all 4 days. In the NR intact condition, 67.3% of COA somata are untuned across all 4 days, a significant decrease (P = 4.8e-57). That fraction drops significantly again in the NR inhibited condition as compared to the NR intact condition, to 53.7% (P = 1.4e-60). This result indicates that shocks induce an increase in COA recruitment of relevantly tuned somata, and that the NR-dCA1 pathway moderates that increase.

We next investigated how somata with tunings to context, place, shocks, and freeze-transitions participate in COAs (Fig 3.6A). These tuning values do not necessarily add to 100%, as some somata could fall into multiple tunings, i.e a soma that is both shock-tuned when shocks are administered and later has a place field on day 3. We found that on day 1, shock

transition-tuned soma made up an average of 4.5% of somas within COAs, while freeze-transition tuned made up an average of 4.4%, somata that had place field in the to-be shocked context pre-shocks made up an average of 14.7%, somata that would have place fields in the shocked context on day 3 made up an average of 15.0%, and somata with strong (i.e. less than -0.5) preferential activity in the shocked context made up an average of 14.0% (Fig 3.6A). Shock-transition tuned somata stayed consistently recruited across days in the NR-dCA1 inhibited mice, but not in the NR intact mice, which barely recruited shock-tuned somata on post-shock days. Freeze-tuned somata notably increased their participation from pre-shocks baseline (P = 1.5e-5), then decreased participation significantly again on subsequent days, an effect again largely driven by NR-dCA1 inhibited mice (Fig 3.6A, P = day 2: 3.8e-3, day 3: 54.5e-3). Expectedly, the number of place field-tuned somata were only increased significantly on the day and context pairings that matched their place fields as compared to day 1 (pre-shocks PF tuned P = 7.8e-7, day 3 PF tuned P = 1.1e-10), indicating that perhaps out-of-field firing in these place fields may be attributable to participation in COAs. Lastly, nDI tuned somata stayed stable throughout pre-shocks to day 2, but notably increased on day 3 (P = 1.2e-6).

We next more intricately investigated the relationship between nDI and COA participation (Fig 3.6B). We found no differences in condition or on number of COAs and nDI (data not shown), but did find in the subset of shocked mice that generalizers tended towards a relationship with more negative nDIs predicting higher somata per COA (P = 0.03, Fig 3.6B). We next investigated if somata with shock-tuning on day 1 had any relationship with COAs on subsequent days. We found a positive relationship between somata being highly shock-tuned and participating in a higher number of COAs in only the NR-dCA1 inhibited condition, and not the NR intact condition on day 1 (Fig. 3.6C). Despite making up only a small percentage of COAs

overall, shock-tuned somata seem to independently be recruited into many COAs in the absence of NR-dCA1. This finding indicates that the NR-dCA1 pathway suppresses preferential recruitment of shock-tuned somata into COAs.

Our next examination was of the freeze transition-tuned somata on day 1 (Fig 3.4E-J), hypothesizing that this activity was preferentially happening during COAs. We again examined only the subpopulation of somata that were dual freeze on/offset tuned (Fig 3.6D). Our findings reveal that in the shocked context on day 1, NR intact freezing transition tuned-somata participated in fewer COAs (12.6 on average), and the quality of the tuning did not predict COA participation ($P = 0.83$). In contrast, NR inhibited freezing-tuned somata participated in more COAs (15.6 on average, $P = 2.9e-3$), and the quality of freeze-tuning positively correlated with increased COA participation ($P = 8.7e-10$). This relationship disappeared on subsequent days (data not shown), indicating that an intact NR-dCA1 pathway is necessary to suppress the re-activation of shock-tuned somata in subsequent COAs.

We next wanted to investigate if having a place field in the shocked context on day 0 was predictive of the soma being recruited into more COAs in the shocked context on day 1. Indeed, we found that having a place field on day 0 in the to-be shocked context predicted higher COA participation in the shocked context on day 1 in both the NR intact ($P = 1.4e-15$) and NR inhibited ($P = 1.1e-5$), conditions with no difference between the conditions ($P = 0.18$), while the unshocked conditions showed no change in PF recruitment ($P = 0.06$, Fig 3.6D). This result indicates that while shocks increased the number of relevant place fields recruited into subsequent COAs, this recruitment is unaffected by inactivation of the NR-dCA1 pathway. We found a similar trend for somas with place fields on day 3. Somata that will have place fields on day 3 in the shocked context participate in more COAs in the shocked context on day 1 in both

the NR intact (P = 4.6e-21) and NR inhibited (P = 3.27e-8) conditions. However here, the PF on day 3 tuned somata in the NR intact condition did participate in more COAs than the NR inhibited condition (P = 0.02), while the unshocked conditions again showed no change (P = 0.11). This difference may arise from the less stable overall map in NR intact mice, as somata participating in higher COAs on day 1 may be more likely to be recruited as new place fields for the NR intact mice, while NR inhibited mice maps remain more stable.

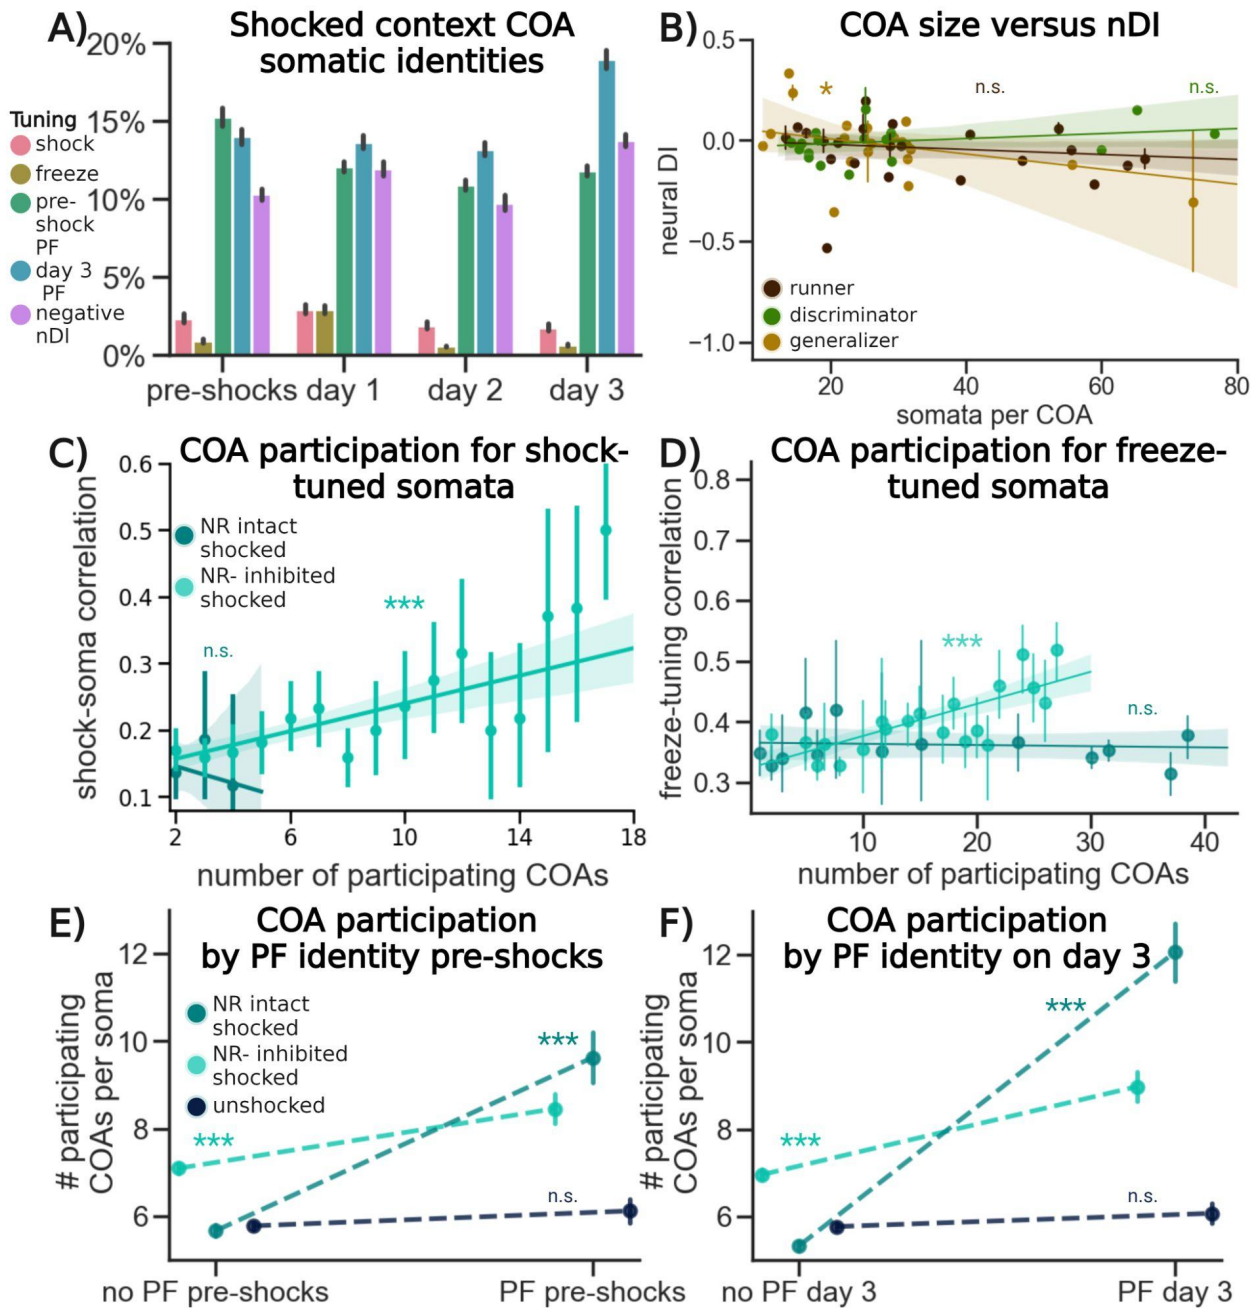

Figure 3.6: The NR-dCA1 pathway increases tuned somata participation in synchronized co-activations post-shocks

A) Mean percent change of makeup of co-activations in the shocked context in the unshocked condition, comprised of tuning as described elsewhere to shock on day 0 (salmon) freezing in the shocked context per-day (gold), having a place field pre-shocks in the shocked context (green), having a place field on day 3 in the shocked context (blue), or having a strongly negative (equal to or less than -0.5) neural discrimination index. Remainder of somata are untuned to these specified metrics (not shown), comprising on average ~59% of somata. **B)** COA size is unaffected by nDI on day 1, except for generalizers, which have a slight correlation between increased neural DI and increased mean number of somata per COA (discrimintors: slope =

0.001, intercept = - 0.04, r^2 = 0.02 , P = 0.49. Runners: slope = 0, intercept = 0, r^2 = 0, P = 0.89. Generalizers: slope = 0.003, intercept = -0.1, rvalue = 0.5, pvalue = 0.03). **C)** Shock-tuned somata plotted against the total number of somas participated in, in the shocked context on day 1 in both shocked conditions. Only mice in the NR-dCA1 inhibited group show a positive correlation between the number of COAs a soma participates in and its tuning to shocks on day 0 (NR inhibited: slope = 4.7, intercept = 0.37 r^2 = 0.23, P = 4.6e-09. NR intact: slope = -0.8, intercept = 0.28, r^2 = 0.01 P = 0.5). **D)** Somata in the NR intact condition had a higher lower average COA participation than NR inhibited somata and no relationship between freeze-tuning quality and COA participation (NR intact: slope = - 4.9, intercept = 14.4, r^2 = 0, P = 0.83), NR inhibited somata demonstrated a positive relationship between freeze-tuning correlation and the number of participating COAs and higher average COA participation quantity (slope = 39.5, intercept = -0.67, r^2 = 0.20, P= 8.7e-10. **E)** Somata with place fields on day 0 in the to-be shocked context participate in more COAs in the shocked context on day 1 in both the NR intact (P = 1.4e-15) and NR inhibited (P = 1.1e-5) conditions with no difference between the increases in the two conditions (P = 0.18), while the unshocked conditions showed no change in place field recruitment (0.06). **F)** Somata that will have place fields on day 3 in the shocked context participate in more COAs in the shocked context on day 1 in both the NR intact (P = 4.6e-21) and NR inhibited (P = 3.27e-8) conditions with the NR intact condition participating in more (P = 0.02), while the unshocked conditions showed no change (P = 0.11).

Discussion

Our findings both behaviorally validate and experimentally expand upon our previous work showing the importance of the NR-dCA1 pathway in contextual fear conditioning [346]. We re-affirm that that inactivation of the NR-dCA1 induces increases in freezing and generalization in shocked mice, and further show that this increase is driven by both mice that well-discriminate and mice that generalize between contexts on day 1 (Fig 3.1D-E, J-K). We further demonstrate that the same excitatory pyramidal somata can be tracked across days of a VR-CFC paradigm (Fig 3.2A-D), that inhibiting the NR-CA1 pathway prevents relative nDI increases on retrieval day 1 (Fig 3.2E), and that the relationship between more negative nDI and bDI is driven by discriminating mice (Fig 3.2J). We also show that inactivating NR-dCA1 has a stabilizing impact on place fields, with post-shock extinction place maps being more comparable to pre-shock maps in the shocked context and more comparable to unshocked maps after 4 days of reward-free context exposure (Fig. 3.3).

We next demonstrate that a significant portion of somata tune to shocks, that increased quality of shock-tuning predicts behavioral discrimination (Fig. 3.4A-C), and that a significant portion of somata tune to freezing immediately post-CFC as compared to pre-shock freeze tuning in both NR intact and NR inhibited mice, with a higher fraction tuning to freezing in NR inhibited generalizers (Fig. 3.4G-J). We further describe a novel identification method for somatic synchrony in the form of relatively fast co-activation events that occur during running and freezing, the quality of which is impacted by NR-dCA1 inactivation by decreasing the overall number of events while increasing the overall number of somata involved in each event (Fig 3.5). Lastly, we demonstrate that somata with tuning to shocks, to freezing, and to location both pre- and post- CFC all participate in co-activations, but that a correlation between well-tuned somata and COA recruitment only exists in shock and freeze-transition tuned somata when the NR-CA1 pathway is offline (Fig 3.6). Overall, we demonstrate that the NR-dCA1 pathway helps prevents inappropriate memory retrieval of somata tuned to salient features of contextual fear experiences.

Our findings suggest a complex role for NR in dCA1. On one hand, it seems to have some excitatory impacts on the circuit by increasing instability and remapping in place fields, and by increasing the overall number of COAs. The increase in excitatory NR input activity during freezing to the apical dendritic tuft could be largely responsible for these effects, especially if NR input, when coincident with MEC, LEC, or CA3 input, is capable of inducing strong burst firing[99]. The increased width of place fields and out-of-field firing observed in the place maps post-CFC can both be explained by this increase in excitation to the apical tuft, inducing a less-static, more flexible place map able to potentially incorporate extinction information more readily. The increase in number of COAs overall in both contexts post-CFC

could be explained by this decrease as well, if NR axons synapse on multiple dCA1 pyramidal neurons and simultaneously drive depolarization during freezing epochs, lower the firing threshold for these somata and driving increased COA participation.

However, there are a number of findings within this work that do not support a simple excitation to the apical tuft model for nucleus reuniens's role in modifying somatic activity. An intact NR helps drive a decrease in in differential peak firing, inducing differential mean peaks/soma of 17.2 in the shocked and and 15.2 peaks/soma in the neutral context, compared to an overall more equally excitable 18.3 and 18.5 peaks/soma when NR is taken offline, respectively. How this selective decrease is impacted by the non-selective NR activity increase in both contexts during freezing epochs is unclear. Partially, because mice freeze less in the neutral context, they may be receiving overall less excitatory NR input, driving higher differential activity. Partially, the overall slight decrease in firing under NR intact conditions may be driven by the impacts NR has on inhibition in the circuit.

It is well-documented that activation of the NR axons projecting to dCA1 *in vitro* synapse upon and induce robust GABAergic firing in interneurons in dCA1[95,99–101,301]. If some of these interneurons synapse directly onto dCA1 excitatory pyramidal neurons, an increase in NR activity could be driving a decrease in overall activity gated by inhibitory populations. Similarly, NR inactivation drives an increase in re-activation of freeze-tuned somata in generalizing mice. In the engram literature, re-activation of a small group of somata can induce contextual freezing[204,205,207]. It is possible that these freeze-tuned somata may be part of this subset that can induce freezing, and NR-dCA1 typically suppresses their activity to enable extinction learning, a role well-characterized for NR generally in the contextual fear circuit[97,117,118]. While inhibition of the NR-dCA1 pathway itself does not impact the number of somata recruited to each COA, it

does induce an increase in an inappropriate somatic recruitment in the neutral context on retrieval day 1, potentially contributing to freezing generalization in the neutral context.

The reinstatement of the NR-dCA1 pathway on retrieval day 2 interestingly induces a massive 'rebound' increase in activity in the shocked context. If this increase in COA somata recruitment is important for memory of fearful experiences, it would potentially explain why NR-dCA1 mice continue to freeze at high levels on retrieval day 2, despite the lack of NR inhibition. Indeed, the somata most commonly recruited to these inappropriate COA events are freeze and shock-tuned, and specifically shock-tuned somata are more likely to be recruited to multiple COAs, further providing evidence in agreement with their potential overlap with immediate early gene identified engrams.

The nature of COAs themselves deserves some additional discussion. Initial hypotheses for these coordinated synchronized events were as putative replays, which are difficult to identify using calcium imaging due to its slow speed compared to electrophysiology, despite a few recent attempts[349]. Interestingly, albeit using different methods similar co-active ensembles have been identified during CFC and described as putative sharp-wave-ripple events, despite a described preference for activation during running epochs that we additionally observe[340]. We have also observed these events in both hippocampal dCA1 and dCA3 after CFC in other preliminary work[213]. Altogether, while a subset of COAs do meet the criterion as putative replays both temporally and behaviorally, the majority do not. What role, if any, do the rest of these COAs serve in hippocampal encoding and modification of contextual fear memory? Recent work indicates that the type of activity that induces strong cFOS activation is not, as previously hypothesized, singular intense firing, but instead regular coordinated synchronous activations[350]. Our COAs would be temporally scaled to induce exactly this type of strong cFOS activation,

strengthening our hypothesis that some subset of somata involved in COAs are likely to overlap with IEG-identified engrams.

Future Directions

How the freezing-tuned signal that NR sends to dCA1 during contextual fear extinction impacts the variety of somatic population-level effects seen during and after NR-dCA1 inactivation is an open question[346]. As discussed, NR can impact CA1 somata either directly through synapsing on apical tuft dendrites where it can interact with other excitatory apical tuft inputs from lateral or medial entorhinal cortex, or through activating a variety of inhibitory interneurons which can then either directly decrease, or through inhibiting other interneurons, increase excitatory somatic activity [95,98,99,301].

Future work is needed to understand both how NR axons interact directly with distal apical tuft dendrites of dCA1 pyramidal neurons to understand the circuit mechanisms through which NR impacts dCA1 population encoding. Additional future work is sorely needed to understand first which inhibitory interneuron populations NR synapses onto in dCA1, and then how its activity triggers activation in interneuron populations *in vivo* and throughout CFC, how the activity of NR-triggered interneurons interacts with excitatory populations, and furthermore how perturbation of the NR-dCA1 pathway impacts these interactions. The realm of NR interneuron interactions *in vivo* is, to our knowledge, entirely unexplored at present.

Further work is additionally needed to link high context discrimination, freeze, shock, and place-tuned somata preferentially reactivated during COAs with the IEG-based literature. Current studies linking activity of engram-identified somata to activity during calcium imaging in the hippocampus are limited and conflicting[210,211,345,351]. Dual IEG monitoring and calcium imaging during VR-CFC and NR-dCA1 manipulation would help unravel the possible

connections between COAs and engrams. In the present study, we have only explored how a subset of obvious tunings to contextual rate firing increases shock, freezing, and place. A multitude of other, more complex tunings likely exist in dCA1 somata, as well as potential changes to low-dimensional latent embeddings identified in CA1. These might be captured by combining sophisticated behavioral recordings and pose-estimation techniques such as DeepLabCut with modern dimensionality reduction tools such as CEBRA [352,353].

Lastly, much remains to be explored about other aspects of how NR interacts with the dCA1 circuit. *In vitro,* it has been shown that NR inputs to distal dCA1 pyramidal dendrites interact superlinearly with inputs from MEC and LEC, which have been demonstrated to carry behaviorally relevant information *in vivo* [99,295,354]. If EC/NR inputs occur simultaneously naturally *in vivo*, what information they carry, and how simultaneous activation from these inputs impacts somatic activity in dCA1 remains to be seen. NR inputs additionally project to subiculum and to entorhinal cortex from different sub-populations of NR somata[305], which potentially impact and modulate the overall circuit. It is also worth noting that NR projects up to 10 times more strongly to ventral CA1 than to dorsal CA1, and what information NR is sending to vCA1, and how it impacts vCA1 somatic populations, remains completely unclear[101,106]. NR has also been implicated in a wide range of other complex behaviors requiring medial prefrontal cortical synchrony with the hippocampus, and how NR axons are selectively active in these tasks remains unknown[111,355–358].

We have overall described a critical role for nucleus reuniens projections to the dorsal CA1 subregion of the hippocampus during contextual fear memory retrieval and extinction in inducing extinction-related place field remapping, and suppressing inappropriate re-activation of shock and freeze-tuned somata. NR is likely impacting the somatic circuit through both direct

inputs to distal dCA1 excitatory dendrites, and through robust activation of regulatory interneurons. Future research will hopefully help disentangle the impact of the NR-dCA1 pathway on hippocampal population dynamics throughout contextual fear conditioning.

Methods

Experimental Model and Subject Details

All experimental and surgical procedures were in accordance with the University of Chicago Animal Care and Use Committee guidelines. We used 10-20 week old male C57BL/6J wildtype (WT) mice (24-35 g). Male mice were used over female mice due to the size and weight of the headplates (9.1 mm x 31.7 mm, ~2 g) which were difficult to firmly attach on smaller female skulls. Mice were individually housed in a reverse 12 hour light/dark cycle and behavioral experiments were conducted during the animal's dark cycle. We are unaware of any influence of strain or sex on the parameters analyzed in this study. A total of 61 mice were used in this study, 34 of which were used in the final dataset. 4 mice were excluded for cannula or headplate instability. 13 mice were excluded for not reaching the running behavior criteria (See Methods: Behavior). 10 mice were excluded for not having sufficiently bright identified GCaMP8f expression in the somatic layer. The remaining 4 mice were excluded for field-of-view instability over the 4 days of imaging.

Mouse surgery and viral injections

Mice were anesthetized (~1-2% isoflurane) and injected with 0.5 ml of saline (intraperitoneal IP injection) and 0.5 ml of Meloxicam (1-2 mg/kg, subcutaneous injection) before being weighed and mounted onto a stereotaxic surgical station (David Kopf Instruments). In all mice, a small craniotomy (1-1.5 mm diameter) was made over the hippocampus (± 1.7 mm lateral, -2.3 mm caudal of Bregma). Either CamKIIa-jGCaMP8f-WPRE (AAV-CamKIIa-jGCaMP8f-WPRE was a gift from GENIE Project; $\geq 7\times10^{12}$ vg/mL; viral prep # 176750-AAV1 ; http://n2t.net/addgene:176750) was diluted 1:1 with PBS or a 1:1 mixture of AAV-CaMKII0.4.Cre-SV40 (pENN.AAV.CamKII 0.4.Cre.SV40 was a gift from James M.

Wilson, ≥ 7×10^{12} vg/mL; viral prep #105558-AAV1 ; http://n2t.net/addgene:105558) and pGP-AAV-syn-FLEX-jGCaMP8f-WPRE (pGP-AAV-syn-FLEX-jGCaMP8f-WPRE was a gift from GENIE Project, ≥ 7×10^{12} vg/mL; viral prep #162379-AAV1; http://n2t.net/addgene:162379) was injected into left hippocampal CA1 at the somatic layer (~100 nL at a depth of -1.25mm below the surface of the dura) to achieve the end effect of a CaMKII-based expression of GCaMP8f, limiting indicator expression largely to excitatory neurons. For DREADD experiments, AAVrg-hSyn-DIO-hM4D(Gi)-mCherry (pAAV-hSyn-DIO-hM4D(Gi)-mCherry was a gift from Bryan Roth Addgene viral prep # 44362-AAVrg; RRID:Addgene_44362) was additionally injected into bilateral hippocampal CA1 SLM (~50 nL per side at a depth of -1.5 mm below the surface of the dura). In the same surgical procedure, AAV9-hSyn-Cre (pENN.AAV.hSyn.Cre.WPRE.hGH was a gift from James M. Wilson Addgene viral prep # 105553-AAV9 ; http://n2t.net/addgene:105553 ; RRID:Addgene_105553) was injected into bilateral NR (~100 nL at a depth of -4.1 mm, at coordinates 0.0 lateral, -0.6 caudal of Bregma). Afterwards, the site was covered using dental cement (Metabond, Parkell Corporation) and a metal head-plate (9.1 mm x 31.7 mm, Atlas Tool and Die Works) was also attached to the skull with the cement. Mice were separated into individual cages and water restriction began the following day (0.8-1.0 ml per day). At least 7 days following injection surgery, and approximately 7 days prior to the beginning of mouse training, mice underwent another surgery to implant a hippocampal window as previously described[323]. Following implantation, the head-plate was reattached with the addition of a head-ring cemented on top of the head-plate which was used to house the microscope objective and block out ambient light. Post-surgery mice were given 1-2 ml of water/day for 3 days to enhance recovery before returning to the reduced water schedule (0.8-1.0 ml/day). Mice weights

were monitored daily to ensure healthy weight maintenance post-water deprivation. Expression of GCaMP8f reached a steady state ~14 days after the virus was injected, as monitored through 2-photon imaging.

Behavior

Our virtual reality (VR) and treadmill setup was designed similarly to previously described setups[297]. The virtual environments that the mice navigated through were created using VIRMEn[155]. Mice were head restrained with their limbs comfortably resting on a freely rotating styrofoam wheel ('treadmill'). Movement of the wheel caused movement in VR by using a rotary encoder to detect treadmill rotations and feed this information into our VR computer, as in (Heys et al., 2014; Sheffield et al., 2017). During training, mice received a water reward (4 µl) through a waterspout upon completing each traversal of the track (a lap), which was then associated with a clicking sound from the solenoid. Upon receiving the water reward, a short VR pause of 1.5 s was implemented to allow for water consumption and to help distinguish traversals from one another rather than them being continuous. Mice were then virtually teleported back to the beginning of the track and could begin a new traversal. Mice were also teleported to the beginning of a new contextual exposure. Seven distinct VR contexts were used in this experiment. The 'familiar' context was the context that mice trained in prior to imaging. Two new contexts were used for contextual fear conditioning, and mice spent at minimum 30 minutes in each context prior to imaging on pre-shocks. The dark context was completely dark with no visual cues. Mouse behaviors (running velocity, track position) were collected using a PicoScope Oscilloscope (PICO4824, Pico Technology). Behavioral training to navigate the virtual environment began ~7 days after window implantation (~30 minutes per day) and continued until mice reached a speed of greater than 4 traversals per minute, which took 10-14 days. This high level of training was necessary to ensure mice continued to traverse the track similarly after

reward was removed. Initial experiments showed that mice that failed to reach this criterion typically would not traverse the track as consistently without reward[297]. Mice that did not reach this criterion were not used for these experiments.

Contextual Fear Conditioning Paradigm

For mice that reached criteria in the training environment (>4 traversals per minute), mice were first exposed to two novel environments without water reward for 322 s (~5 minutes) each, with the addition of a custom-made tailcoat made of conductive fabric (Adafruit). Only mice that continued to maintain a speed of 4 traversals > minute without water rewards and with the tailcoat on were allowed to continue the experiment. Subselecting for mice with this consistent running behavior helped us to ensure that freezing responses recorded later were not due to the presence of the tailcoat or any discomfort from head-fixation or removal of reward. Here onwards, the tailcoat was kept on the mouse during the experimental sessions on all subsequent experimental days. Each contextual exposure was for a duration of ~5 minutes. Prior to experimental day 0, mA level of shock delivery was confirmed using an oscilloscope. On day 0, mice were exposed to both novel contexts, then shocked in one of the two contexts, administering 12 X 0.6 mA shocks delivered at an interval of 20-26 seconds each, (Coulbourn Instruments Precision Animal Shocker). Mice displayed rapid sprinting behavior when they received the tail shock, allowing us to confirm the delivery of shocks in real-time and post-hoc. On subsequent days, mice were exposed to both the shocked and non-shocked (neutral) contexts pseudorandomly, for 3 days.

DREADD Experimental Protocol

To activate the hM4D(Gi) receptor and silence a subset of NR glutamatergic neurons that project to CA1, we used Deschloroclozapine dihydrochloride (DCZ, MedChemExpress). DCZ was

chosen over CNO as an h4MDi agonist as DCZ has a significantly increased potency, therefore enabling a 100-fold dosage reduction, heightened selectivity for h4MDi receptors over endogenous receptors, and a significantly more rapid onset than CNO. In addition, the DCZ metabolites C21 and DCZ- N-oxide are reported at negligible concentrations (< 0.2 nM), compared to the known tendency of CNO to metabolize to Clozapine at higher concentrations, a compound with significant off-target receptor binding in the mammalian brain[300]. Because of these factors, we chose to use DCZ for inactivation, as in our past work[297,346].

Once mice met training criteria, they were habituated to the injection process. They were exposed to the rewarded training environment for ~10 min. Afterwards, they were removed from the VR set up, placed in the holding room, and injected with ~150 μL of a 12% DMSO/Saline solution. After ~30-45 min, they were placed back in the VR setup and exposed to the rewarded training environment again for an additional 10 min. This was repeated for 3 days to acclimate mice to the injection procedure. Mice additionally received ~150 μL of a 12% DMSO/Saline solution on Day -1 of the experiment 30 minutes prior to first exposure to both neutral contexts to mimic conditions on Day 0.

All animals received DCZ injections, regardless of group, to control for the impact of DCZ. We previously demonstrated no impact of DCZ, additional surgical procedure, or of the viral constructs necessary to achieve NR-CA1 specific DREADD expression on mouse behavior[346]. DCZ was dissolved in DMSO at at .02 mg/mL concentration and stored at -80 °C on day 0. On retrieval day 1, DCZ solutions were thawed to room temperature and diluted to 0.01 mg/mL with DMSO/Saline. ~30 minutes prior to context exposure, mice were brought to a holding room and IP injected with 0.1 mg/kg DCZ of a .02 mg/mL solution. A quantity of DMSO/Saline solution identical to IP injection amount on Day 1 (~100-150 μL) was injected on all other experimental

days, ~30 minutes prior to imaging, to control for the impact of any potential IP injection-induced stress. Imaging protocol for all DREADD NR-CA1 inhibited experimental mice and intact mice was kept identical. All mice experienced a 'dark' imaging session after context exposures of the same duration, where no context was displayed on screens, for ~5 minutes, to check for any impact of DCZ on movement (Fig. 2.3f).

Two-photon imaging

Imaging was conducted 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 16x/0.8 NA/3 mm WD water immersion objective (MRP07220, Nikon). GCaMP8f was excited at 920 nm with a femtosecond-pulsed two photon laser (Insight DS+Dual, Spectra-Physics) and emitted fluorescence was collected using a GaAsP PMT (H11706, Hamamatsu). The average power of the laser measured after the objective ranged between 9-38mW. Each field of view (FOV) was positioned ~150 μm below the putative surface to target the somatic layer, and 400-700 μm equally in the *x/y* direction. Planes were selected to be over the highest number of visible active CA1 pyramidal cells. Each plane was imaged at the same magnification per mouse and day. Time-series images were collected through Scanbox (Neurolabware) and the PicoScope Oscilloscope (Pico Technology) was used to synchronize frame acquisition timing with behavior.

Image Processing and ROI Selection

Time-series images were preprocessed using Fiji to manually align fields-of-view that had xy rotations over days. Aligned and concatenated tifs were then processed through Suite2p (Pachitariu et al., 2017). Movement artifacts were removed using rigid and non-rigid transformations, and manually assessed to ensure absence of drifts in the *z*-direction and good

cross-day alignment. Some datasets collected during shock administration were discarded, due to the high velocity post-shocks sprinting behavior of mice rendering those FOVs too unstable for reliable analysis. Regions of interest (ROIs) were also defined using Suite2p and manually inspected for accuracy. Baseline corrected *Δf/f* traces across time were then generated for each ROI. Per ROI, a savitzky-golay filter was applied to both channels to smooth the signal. From this trace, the Fc3 was calculated as previously described and used for further analysis [147]. Only ROIs that exceeded 10% *Δf/f* at least once per experimental day were retained. Peaks were calculated using the scipy.signal.find_peaks package with a required minimum height of 10% *Δf/f,* distance of 0.3s, and prominence of 0.1.

Neural Discrimination Index

Neural discrimination index (nDI) was calculated as described in [214]. Briefly, for each ROI, the number of transient peaks were counted across a single context, in a single day. The nDI was then calculated across context or day, as described in the text, by calculating a comparison number of transient peaks for the comparison context, subtracting the first context from the second, then dividing by the sum of transients in both.

Place Field Identification

Place fields were identified using a modification of the 'peak' method described in Grijseels et al, 2021. [359]. Briefly, the Fc3 of ROI, after processing as described above, was binned into 40 bins along the 2 meter virtual track. Activity in each bin was identified for putative place field behavior based on being active on at least 4 laps in that bin, activity post-emergence being present in 30% of subsequent laps, and being at least 2 bins wide and no more than 17 bins wide. These putative fields were then shuffled 1000 times against ROI data and tested for significance, requiring firing in at least one location along the track to be equal to or greater than two standard

deviations above the shuffled activity. This method allowed for multiple place fields per ROI, which were each treated separately for subsequent analysis with the other IDed place field's activity disregarded for properties such as out-of-field-firing and precision.

Place Field Metrics

All metrics described below were conducted per mouse, day, plane, context, and post-processed ROI, excepting remapping which requires computation across context. Most comparisons were between days pre-shocks and day 3, as days 1 and 2 had too low of laps due to fear expression to appropriately calculate place fields.

Active fraction: Each ROI with activity in at least one of the contexts on a per-day basis was counted, and the fraction was computed as the number of ROIs that contained at least one place field within that context divided by the number of active ROIs.

Emergence: Starting from the first lap the mouse traversed in the context, we searched lap-by-lap for a lap with a significant calcium transient present within the boundaries of the future place field, as calculated from all the laps in the session. This lap was identified as the place field emergence lap.

Reliability: From the emergence lap to the last lap the mouse traversed in the context, we searched lap by lap for significant calcium transients. This metric is the number of laps with significant in-field calcium transients divided by the total number of laps traversed post-place field emergence.

Out/in field ratio: The mean $\Delta f/f$ in bins within the place field was divided by the $\Delta f/f$in bins outside of the place field.

Center of Mass: We first measured $\Delta f/f$ in each bin. We then used the following equation to calculate the COM for each traversal n (COMn):

$$COM_n = \frac{\sum_i DF_i \cdot x_i}{\sum_i DF_i}$$

Where DFi is the somatic *Δf/f* in bin i and xiis the distance of bin i from the start of the track. We then calculated the weighted average COM (COMw) from all traversals n (COMn from each traversal was weighted by the peak transient *Δf/f* on that traversal (An)):

$$COM_w = \frac{\sum_n A_n \cdot COM_n}{\sum_n A_n}$$

Precision: Spatial precision (SP) was calculated as the inverse of the COM standard deviation:

$$SP = \frac{1}{\sqrt{\frac{\sum_n A_n (COM_n - COM_w)^2}{\sum_n A_n}}}$$

Width: The difference between the first and last bin with in-field activity, expressed as number of bins.

Slope: Slope represents a linear regression as calculated using the scipy.stats linregress module from the lap identified to have peak activity and the center of mass, thus yielding a calculation of place field shifting over time. A positive metric would indicate forward shifting, while a negative metric would indicate backwards shifting.

r^2: Goodness-of-fit of the slope as calculated previously.

Remapping: To measure spatial correlation across contexts in place fields, we found somata that had place fields in the first shocked and neutral contexts pre-shocked and then calculated the Pearson's correlation coefficient between the mean activity along the track (in 40 bins) for all laps between that context and post-shock activity on day 3.

Co-Activation Detection

We determined clusters of synchronized co-activations of neurons by each animal, day, and context using the following method: First, we filtered for ROIs with at least one peak within the given window. We next computed the onset, offset, peak, and duration of each transient using the scipy peaks method as previously described[346]. Second, we computed a shuffle matrix 100 times, interval shuffling the peak regions within-each ROI (i.e. randomly reordering intervals between events for each ROI), to create a significance cutoff at three sigma for the number of cells that needed to be simultaneously co-active at a time point to be above chance (i.e. bigger than the 99th percentile). We next compared the number of transient peaks active for each time point for both the test and shuffle data distributions by running find_peaks once again on a linear plot of the number of simultaneously co-active cells per frame. We considered the local peaks of co-activity to be the center time point of the co-activation, and the bounds of the co-activation to require at least 50% of IDed somata to still be active on both the left (onset) and right (offset) sides of the center timepoint, with the span between them considered the duration. The identity of individual neurons of the ensemble for each event, likelihood of individual neurons to participate in an event, and event likelihood within behavioral and categorical parameters, were additionally examined.

Freeze-Transition Tuning

We considered 3 separate possibilities for tuning to freezing: onset-tuned, offset-tuned, and duration-tuned. Onset-tuned was defined as activity (as determined by the presence of a previously-described peak - see Methods: Image Processing and ROI Selection) within 0.5s of the onset of a freezing epoch. Offset-tuned was defined as activity within 0.5s of the conclusion

of a freezing epoch. Duration-tuned was defined as ROIs where the discrimination index, (see Methods: neural discrimination index) using freezing and non-freezing epochs as the two comparisons, that was significantly negative - indicating more differential activity during freezing epochs as compared to running epochs. All significance levels were determined using 100 per-ROI shuffle distributions and taking the 99th percentile above the shuffle distribution.

Shock Tuning

For mice in which we possessed imaging data during shocks, we identified shock-tuned neurons by performing a cross-correlation between the Fc3 of each ROI and a boolean determining if the mouse was currently receiving a shock or had received a shock within the last second. ROI correlations were compared against 100 shuffle distributions of shock timings maintaining the structure and number of shocks, and only ROIs above the 99th percentile of the shuffle distribution were considered shock-tuned.

Behavioral Parameters

All parameters described below were calculated per mouse, day, plane, and context.

Freezing: Freezing epochs were determined as uninterrupted epochs where mouse velocity fell below 0.001 cm/s for at least 12 consecutive frames (~0.75 s). All epochs of velocity below 0.001 cm/s but not reaching 12 consecutive frames were not considered freezing or running, and were discarded from future analysis. Freezing epochs were then counted up, and each not in a freezing epoch assigned a '0', while each frame in a freezing epoch given a numeric value corresponding to the number of epochs in that recording (i.e. all frames that contained the 4th freeze of the recording would be assigned the integer '4'). Subsequent freeze features were then calculated, including the binary variable 'is freezing' which assigns a 1 to frames considered freezing, and 0 to frames not considered freezing, two sawtooth functions 'freeze remaining', and

'freeze elapsed, which counts the frames from the beginning of a freeze up or down until the end of a freeze, respectively, and 'freeze progress' which tracks the progress of a freeze as a fraction from 0 to 1.

Running: Running was determined as any epoch where forward progress velocity was sustained over 0.001 cm/s for 2 consecutive frames. The variables 'is running', 'running remaining', 'running elapsed', and 'running progress' are calculated using the running epoch data in the same fashion as their freezing counterparts.

Backward movement: Some mice demonstrated backward movement behavior in the virtual environment post-shocks, where they made backwards movement through the context. This behavior was analyzed separately from running or pausing. The binary variable 'is backtracking' assigns a 1 to frames considered backtracking, and 0 to frames not considered backtracking.

Shocks: Shock delivery was recorded through the Picoscope.

Velocity: Velocity was both directly measured through the picoscope encoder, and recalculated from position, to assess for accuracy. Recorded velocity was used for all velocity calculations and model training. Values were converted into cm/s for presentation.

Acceleration: Acceleration was calculated as the first derivative of recorded velocity.

Location: Animal's position on virtual track was determined for each frame.

Behavioral Discrimination Index: The behavioral discrimination index (bDI) was calculated per day across conditions and across behavioral groupings as: (% time spent freezing in shocked context - % time spent freezing in control context)/total % time spent freezing in both contexts.

Statistics

For all data distributions, a Shapiro-Wilk test was performed to verify the data was normally distributed before undergoing further statistical tests. For dendritic data with over five data points, a paired Wilcoxon signed rank test, unpaired Student's T test, or an unpaired Mann-Whitney U test was used as noted. For samples with five data points or less, only a non-parametric test was used. Box and whisker plots were used to display data distributions where applicable. The box in the box and whisker plots represent the first quartile (25^{th} percentile) to the third quartile (75^{th} percentile) of the distribution, showing the interquartile range (IQR) of the distribution. The black line across the box is the median (50^{th} percentile) of the data distribution. The whiskers extend to 1.5*IQR on either side of the box. A data point was considered an outlier if it was outside the whiskers or 1.5*IQR. Horizontal dotted lines indicate group median, minimum possible value, or zero, as indicated in figure legend. Significance tests were performed with and without outliers. Data distributions were considered statistically significant only if they passed significance ($p < 0.05$) without outliers. Significance numbers reported are without outliers. Correlations were performed using Pearson's correlation coefficient. $p < 0.05$ was chosen to indicate statistical significance and p-values presented in figures are as follows: *, $p < 0.05$, **, $p < 0.01$, ***, $p < 0.001$, N.S. not significant. Darker lines in the center of line plots are the mean, and shading is the 95% confidence interval, unless stated otherwise in text or figure legends. All regression analysis was conducted using the statsmodels. Some data preprocessing was conducted with MATLAB (Mathworks, Version R2022a). All other data and statistical analyses were conducted in Python 3.7.4, with primary data accrued in Pandas DataFrames, and data figures were made in Python 3.7.4 using the Seaborn and Matplotlib packages (https://www.python.org/). Schematic figures, some figure text, and figure layouts were made with BioRender (https://biorender.com/).

Data availability

Raw imaging data are extremely large and not feasible for upload to an online repository but are available upon request. Processed data will be available on GitHub after publication of this work in an academic, peer-reviewed journal.

Code availability

Original code used to create figures from pre-processed data will be available on GitHub after publication of this work in an academic, peer-reviewed journal.

Acknowledgements

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, a New Innovator grant from the National Institutes of Health (1DP2NS111657-01) awarded to M.S., and a BRAIN Initiative grant from the National Institutes of Health (RF1NS127123). We thank the University of Chicago imaging core for assistance with confocal imaging, and the University of Chicago animal care staff for ensuring the well-being of experimental animals. We thank Timothy Ratigan for assistance with data analysis. We thank Antoine Madar for assistance with statistical methodology. We thank Rossten Rad for discussions on experimental design.

Contributions

Heather C. Ratigan (H.R.), Valerie Barreto (V.B.), Chery Cherian (C.C.), and Mark E.J. Sheffield (M.S.) contributed to this work. H.R. and M.S. conceived of and designed the experiments. H.R., V.B., and C.C. performed surgeries, animal training, collected all *in-vivo* behavioral and imaging data, and collected all *post-hoc* data. H.R. wrote the analysis code and analyzed all data. H.R and M.S. interpreted the data. H.R. created all figures and wrote the manuscript.

Ethics Declarations

The authors declare no competing interests.

Key Resources Table

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and viral strains		
pENN.AAV-CamKIIa-jGCaMP8f-WPRE	Zhang et al. 2023[360]	Addgene #176750-AAV1
pENN.AAV.CamKII 0.4.Cre.SV40	AAV Plasmids[361]	Addgene #105558-AAV1
pGP-AAV-syn-FLEX-jGCaMP8f-WPRE	Zhang et al. 2023[360]	Addgene #162379-AAV1
Experimental models:	**Organisms/strains**	
Mouse: C57BL/6J	Jackson Laboratories	JAX 000664 - C57BL/6J
Software and Algorithms		
Fiji	Schindelin et al., 2012[327]	https://imagej.net/software/fiji/; RRID:SCR_002285
Suite2p	Pachitariu et al., 2016[328]	https://github.com/MouseLand/suite2p
MATLAB	MATLAB. (2018). *9.7.0.1190202 (R2018a)*.	https://www.mathworks.com/help/matlab/release-notes-R2018a.html
Python	Python 3.10.8	https://www.python.org/downloads/release/python-3108/
Pandas	Pandas 1.1.4	https://pandas.pydata.org/
SciPy	SciPy 1.9.3	https://scipy.org/install/
Seaborn	0.12.0	https://seaborn.pydata.org/installing.html

CHAPTER 4

HIPPOCAMPAL CA1 BASAL AND APICAL PYRAMIDAL DENDRITES DIFFERENTIALLY ENCODE SPATIAL AND NON-SPATIAL ASPECTS OF EXPERIENCE

I was primary lead on all research conducted in this chapter, and will be listed as first author on any future research publications on this work. I was assisted in data collection and post-processing by Cherry Wang and Valerie Barreto.

Abstract

The hippocampal formation is crucial for the formation, modification, and storage of episodic memories, which combines spatial and non-spatial contextual information into a coherent representation. In dorsal hippocampal CA1, excitatory pyramidal neurons encode spatial information through place cells, which fire in response to a specific location. However, these neurons also encode non-spatial information about novelty and reward, while flexibly changing their representations over time - known as representational drift. Pyramidal somata retain a unique dendritic arbor, segregating its glutamatergic inputs to the basal and proximal apical arbor from CA3 from its glutamatergic inputs to the distal apical tuft arbor from the entorhinal cortex and thalamus. These dendrites are well-situated to receive segregated spatial and non-spatial information, leading to integrated CA1 contextual representation encoded by somatic firing. However, whether spatial and non-spatial information are differentially encoded in apical tuft dendrites versus basal dendrites is unknown. Here, we investigated the activity of basal and apical dendritic activity using the genetically encoded calcium indicator GCaMP8f in male mice navigating virtual reality space through familiar, novel, and unrewarded contexts. We

reveal a functional segregation between predominantly spatial representations in the basal dendrites, and predominantly non-spatial reward and novelty representations in the apical dendrites. These findings shed light into how CA1 integrates diverse contextual information to encode stable yet flexible episodic memories.

Introduction

The hippocampus plays a critical role in contextual memory formation, with its pyramidal neurons of specific interest for their critical role in 'mapping' both spatial and non-spatial properties of contextual memories[140,153,155,156,362]. Hippocampal dorsal CA1 (dCA1) is of particular importance as the primary output region of the hippocampus, and the region that encodes both spatial and non-spatial aspects of contextual memories[134]. In mammals, the role of dCA1 in encoding spatial contexts has been well-studied through the lens of 'place fields', activity in excitatory pyramidal neurons that fire regularly at a specific location in space, combining to tile the entirety of a given context[139,140]. The role of dendritic spikes (dSpikes) in shaping neural activity to encode memory has been long-established across the brain. dSpikes provide the depolarization necessary for synaptic plasticity via Hebbian potentiation, which in dCA1 synapses forms the backbone of memory formation and storage[192,243,244]. dSpikes across the dendritic arbor at Behavioral Time Scale Plasticity (BSTP) have been shown repeatedly to be critical for forming new and modifying existing somatic place fields[241,254–256,363].

The unique architecture of dCA1 lends support to the importance of dendrites in dCA1 memory encoding. dCA1 is the final synapse of the 'trisynaptic loop', the canonical glutamatergic circuit of information flow from medial entorhinal cortex (MEC), through hippocampal dentate

gyrus (DG), hippocampal CA3, and finally hippocampal CA1[134]. Within dCA1, the pyramidal neurons are densely packed and aligned, such that their basal dendritic arbor is dorsal of the somatic layer by ~20-80 microns, and their long trunks extend ventrally before spreading into the distal apical tuft ~150 microns below the somatic layer in the stratum lacunosum moleculare (SLM), forming a bi-conal structure[231]. Basal and apical dendrites have a host of differing intrensic properties [364]. While basal dendritic spikes are more likely to induce somatic spikes, distal apical dendritic spikes are critical for coincident detection of inputs and can, when aligned with basal and proximal apical inputs, induce massive somatic bursting, known as complex-spikes, over a multi-second time scale, that leads to neuron-wide synaptic remodeling known as BTSP[251–253].

Glutamatergic inputs in dCA1 are strongly spatially segregated[134]. dCA1 receives the majority of its glutamatergic inputs from hippocampal CA3 through Schaffer collaterals to its basal dendrites, inputs strong enough to induce firing in dCA1 excitatory pyramidal neurons. While DG and CA3 receive glutamatergic input from layer II of the entorhinal cortex (EC), dCA1 alone additionally receives direct glutamatergic inputs to its apical tuft dendrites through direct inputs from layer III of the medial (MEC) and lateral (LEC) portions of EC, and from thalamic nucleus reuniens (NR)[54,313]. All three of these direct pathways are of similar strength and selectively target the CA1 SLM - the location of the distal apical dendrites - and avoid other layers or hippocampal regions[99].

The role of basal dendrites *in vivo* has been investigated in recent reports[147,227,342,365,366]. Basal dendrites have been shown to have strong spatial information encoding, predicting the formation of future somatic place fields, and glutamatergic input to the basal arbor demonstrates spatial tuning [147,365,366]. Apical dendritic calcium spikes, in particular, are critical for complex

somatic bursting, a feature that is important for learning. Coincident dendritic and somatic activity during sharp-wave ripples thus leads to learning-related coupling[227,342]. The activity of CA1 apical tuft dendrites *in vivo* during behavior is presently unknown, as the small diameter and significant distance from the surface of CA1 (~300-400 microns) rendered optical access difficult *in vivo*.

However, some clues can be found by investigating the activity of glutamatergic inputs from EC to the SLM of dCA1. This approach has stumbling blocks, as inactivating EC itself has the confound of disrupting the trisynaptic loop, while specifically inactivating either the MEC-CA1 or LEC-CA1 direct pathway is difficult due to nearby EC inputs to CA3 and DG[367–369]. One recent study has imaged the activity of MEC and LEC axons in dCA1 SLM during a range of tasks in virtual reality[354]. While axons from both regions were strongly spatially tuned, axons from LEC showed robust reward-centric spatial encoding even when the spatial context changed[354]. What is currently known about the activity of the NR-CA1 input *in vivo* is elucidated in Chapter 2 of this dissertation, which imaged NR axons in CA1 for the first time during spatial navigation and throughout a contextual fear conditioning paradigm[346]. However, the overall effect of these inputs on dendritic activity is hard to predict given the various types of receptors, their locations along the dendritic arbor that can interact with one another, and various postsynaptic mechanisms for synaptic integration that are involved. Direct measurements from dendrites can help determine what the overall integrative effect is of these inputs, and what information is then sent to the soma to impact its firing.

Hints as to the roles of basal and apical dendrites during behavior may additionally be found in differences between CA3 and CA1 somatic populations, as CA3 is the largest contributor to basal dendritic and proximal dendritic activity, and the major architectural

difference between the regions is the infusion of the glutamatergic cortical and thalamic inputs to CA1. Such differences can be found in both encoding of novelty and reward manipulation[223,297,370]. While CA1 more rapidly forms place fields in a novel context, these fields have a higher tendency to shift throughout contextual exposure and over days, and exhibit higher levels of representational drift when compared to CA3 place fields, which take longer to form but retain their spatial information more stably[223]. Additionally, CA1 place fields degrade and partially remap when reward is removed from an environment, while CA3 place fields seem more robust to removal of reward-signaling regions like VTA[297,370]. Therefore, manipulating novelty exposure and reward expectation may be behaviors that differentially involve basal and distal apical dendritic involvement.

In summary, present knowledge of dCA1 basal dendritic activity *in vivo* is limited, while knowledge of apical dendritic activity *in vivo* is non-existent. To better understand the role of these distinct dendritic regions, we designed a spatial contextual paradigm for manipulating novelty exposure and reward expectation while imaging either basal or apical dendrites. Our results reveal a differentiated role for dCA1 pyramidal dendrites, with basal dendrites primarily responsible for encoding spatial information and apical dendrites primarily responsible for encoding non-spatial information about reward expectation and contextual change.

Results

Experimental Design and Mouse Behavior

To examine the differential role of basal and apical dendrites in behaviors known to engage CA1 pyramidal neurons, we first designed a two-day paradigm that enabled optical

access to the basal and apical tuft during novelty and reward manipulation (Fig. 1A; Methods: Novelty and Reward Manipulation Paradigm). We did not observe any differences in mouse behavior, as measured through velocity and licking between apical or basal fields of view, regardless of which was imaged first in the counterbalanced design (Supp. Fig. 1A-B). A total of 15 male mice were successfully imaged through this paradigm in both their apical and basal arbors, with apical or basal fields of view (FOVs) remaining stable throughout imaging on a day-to-day basis (Fig. 1B). We found that while we overall recorded a higher number of transient peaks from a higher number of identified regions of interest (ROIs) in the apical tuft than the basal tuft, (N basal = 914 ROIs, 8,663 peaks, apical = 1,357 ROIs, 14,240 peaks), the overall distribution of dendritic peak amplitudes was similar between the two populations (Fig. 1C), with the vast majority of peaks remaining under 1 *Δf/f* and 95% of peaks in both planes not exceeding 4 *Δf/f.* Similarly, there was no difference between basal and apical tuft transient widths or amplitudes across days or contexts (Supp. Fig. 2A-B). Additionally, dendritic populations across planes had similar profiles of activity across experimental paradigms (Fig. 1D-E). Similarities in these measures provide support that differences found between these populations during behavior were less likely to be due to planar differences in recording quality.

Characterization of Basal and Apical Dendritic Activity

We then examined differences in dendritic peak activity during behavior, differentiating between when the mouse was running or pausing, and licking or not licking (Supp. Fig. 4.1C-F). While most metrics yielded no differences between the planes, we did observe a trend towards higher activity in the apical dendrites during licking compared to the basal dendrites (Supp. Fig. 4.1D). However, due to the low lick rate across planes and experiments (0.13 licks/s; Supp. Fig.

4.1B) and relatively low firing rate of both dendritic types (basal 0.011 peaks/s; apical 0.012 peaks/s), the absolute number of peaks during licking (25 total) remains too small to make a statistical determination of difference. We then examined the differences between basal and apical dendrites during the beginning and end of context exposure, hypothesizing that apical tuft dendrites may be more sensitive to novelty due to inputs from LEC and, therefore, more responsive at the beginning of novel contextual exposure. Indeed, we found that both peak width and peak count overall were significantly higher at the beginning of the novel context when compared to either apical activity at the end of the context or basal activity at the beginning of the context (Supp. Fig. 4.2D-E). Basal activity showed no preference for the beginning versus the end of the context (Supp. Fig. 4.2C,E,F). Interestingly this increase was true even in familiar environments, although the increase in peak width especially was highest after reward removal and reinstatement in the reward manipulation experiment in the apical tuft. These results indicate that CA1 apical tuft pyramidal dendrites are preferentially responsive at the beginning of contextual exposure, particularly when that context has manipulated reward exposure.

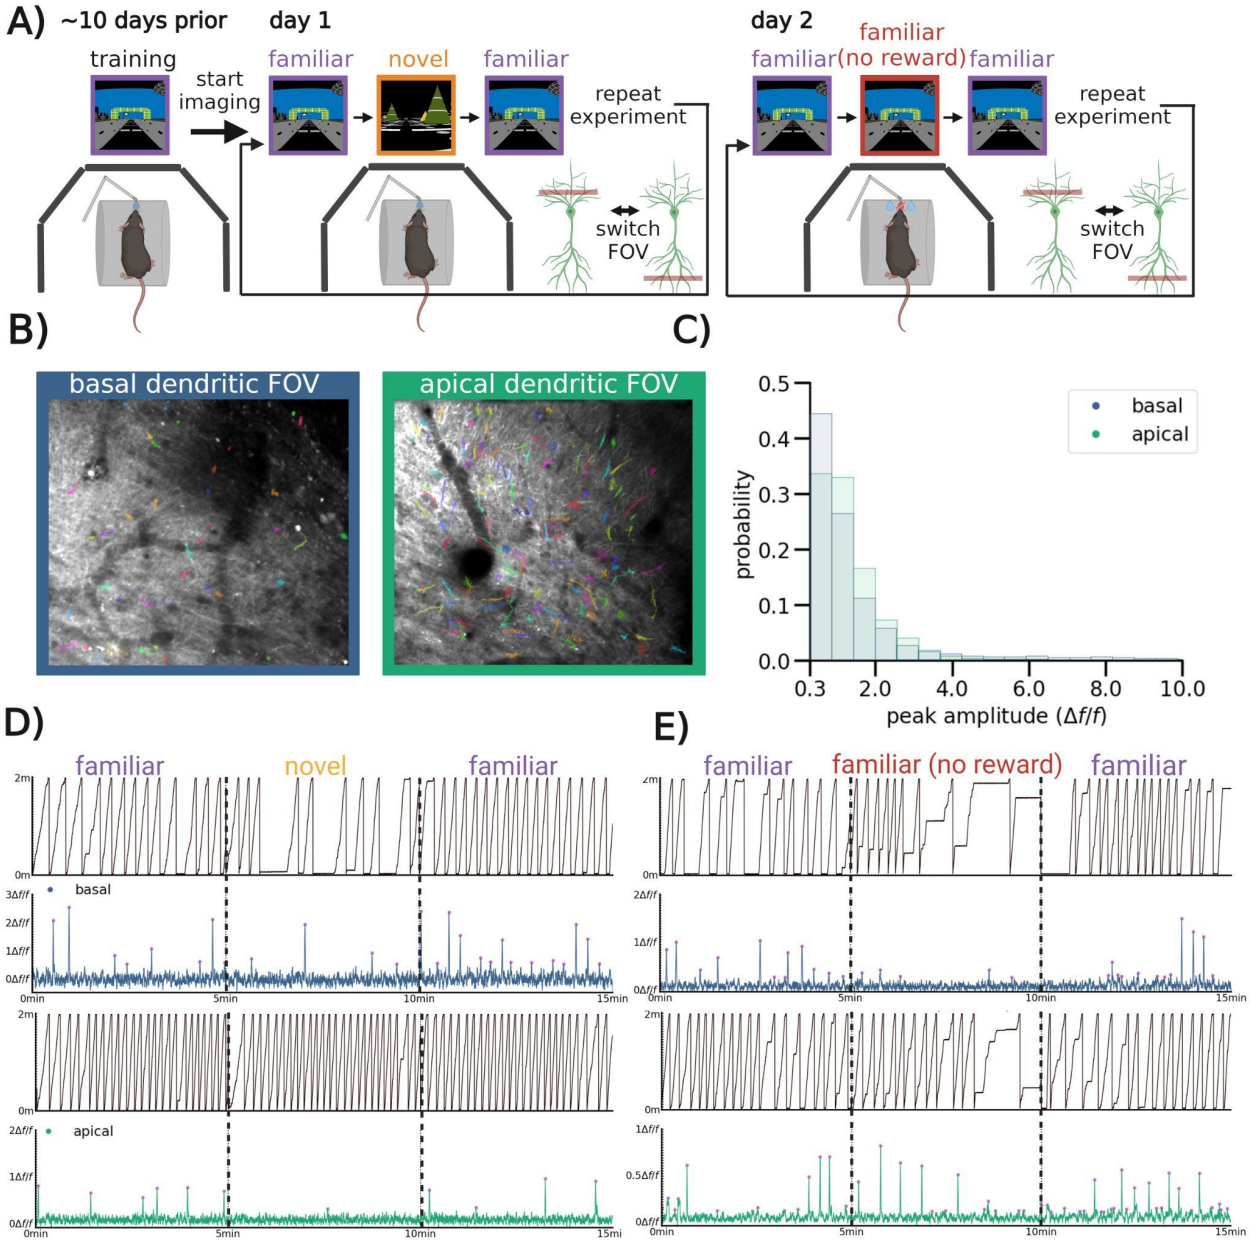

Figure 4.1. Basal and apical dendrites are active during contextual novelty and reward manipulation.

A) Behavioral paradigm. After training in a familiar context and having a basal and apical field of view identified, mice were exposed to a familiar, novel, and familiar environment (novelty manipulation). The field of view (FOV) was switched from apical to basal or from basal to apical, and mice were exposed to the familiar, a new novel, and familiar context again. On the second day, mice were exposed to a familiar context, familiar context with reward removed, and familiar context again with water reinstated (reward manipulation). The FOV was switched again, and mice underwent the same experiment. **B)** Two example FOVS for basal (left) and apical (right) dendrites with identified ROIs highlighted on top in an arbitrary color scheme. **C)**

Probability distribution histogram of peak amplitude for all peaks identified in basal and apical dendrites, demonstrating similar (i.e., not significantly different) distributions of amplitudes of calcium transients across planes. **D)** Representative examples of basal (top) and apical (bottom) dendritic traces after baseline normalization and savitzky-golay smoothing during the novelty manipulation experiment. Peaks are identified in pink dots. **E)** Representative examples of basal (top) and apical (bottom) dendritic traces after baseline normalization and savitzky-golay smoothing during the reward manipulation experiment. Peaks are identified in pink dots.

Place Fields in Basal and Apical Dendrites

As basal dendrites have been shown previously to contain place fields and be instrumental in the formation of somatic place fields in dCA1, we next investigated basal and apical place fields (Fig. 4.2). We found that a much smaller percentage of dendritic activity was place-related as compared to somatic place activity, with 6.1% of basal dendrites and 4.3% of apical dendrites forming place fields (PFs) in any given context (Fig. 2a). A two-way ANOVA revealed that plane, but not context or the interaction of context and plane, was significant (Fig. 4.2A, P = 0.035). The overall character of representative PFs are shown in Fig. 4.2B, with basal PFs (top) appearing more robustly stable and active in comparison to apical PFs (bottom). Indeed, when we measured the number of laps the dendrite was active within its place field after the emergence of that field (reliability), we observed that basal place fields in the familiar unaltered context were nearly twice as reliable as apical place fields (Supp. Fig. 4.3A, rank-sums, P = 3.41e-36). Interestingly, in both planes, the reliability of the field significantly decreased in both the novel and unrewarded contexts as compared to the familiar contexts on the same day in the same plane (Fig. 4.2C-D, Student's T-test, apical novel P = 0.01, apical unrewarded P = 9.21e-3, basal novel P = 0.035, basal unrewarded P = 0.026). This decreased reliability is also observed in novel and unrewarded contexts in place field somata [147,297]. As both basal and apical dendrites demonstrated decreased reliability during novelty or reward manipulation, we hypothesize that the source of reduced dendritic reliability results from

backpropagation of action potentials from the soma through the dendritic arbor rather than dendritic inputs.

Since apical dendrites fired less reliably, we postulated that this unreliability could emerge from apical place fields appearing earlier temporally in a new context, and therefore having more laps in which it could fire. To investigate this possibility, we examined the lap number where the place fields emerged in both the basal and apical tufts. Indeed, we found that apical place fields were significantly more likely to emerge earlier in the track, on average 2.6 laps into the context, while basal place fields emerged significantly later, on average 4.4 laps into the context (Supp. Fig. 4.3B, P= 4.8e-11). We then examined within each context to see if a specific context was driving this difference, and found that in all contexts except the novel context, basal place fields emerged significantly slower than apical place fields (Fig. 4.2E-F). The increased similarity in lap emergence between the basal and apical dendrites in the novel context was driven by a significantly decreased lap emergence, i.e. earlier emergence of place fields, in the basal dendrites in the novel context as compared to the first familiar context, as apical lap emergence was invariant across contexts (Fig. 4.2E, apical P= 9.0e-3: Fig. 4.2F, basal P n.s.). This observation may reflect that the basal dendritic transients are more heavily influenced by back-propagating action potentials from the soma, induced by lower dendritic inhibition in the basal than the apical tuft. In addition to being significantly more reliable and slower to emerge, basal place fields are slightly wider than apical place fields, on average spanning 3.8 cm compared to 3.0 cm for apical place fields (P = 4.3e-10, Supp. Fig. 4.3C).

We next observed the center-of-mass of the place fields across the track. While in aggregate both basal and apical dendrites mapped across the length of the 2 meter track in all contexts, we found that as a population in the familiar context the basal dendrites reliably skew

towards over-representing the beginning of the track, while apical dendrites reliably skew towards over-representing the end of the track (P = 7.1e-3, Supp. Fig. 4.3D). This skew appears to be driven by a cluster of apical place fields near the end of the track, which notably disappear when reward is removed (Supp. Fig. 4.4F, bottom). Overall, basal place fields are significantly more reliable, wider, slower to emerge, and skew towards the beginning of the track, while apical place fields are significantly less reliable, faster to emerge, and skew towards the end of the track unless reward has been removed.

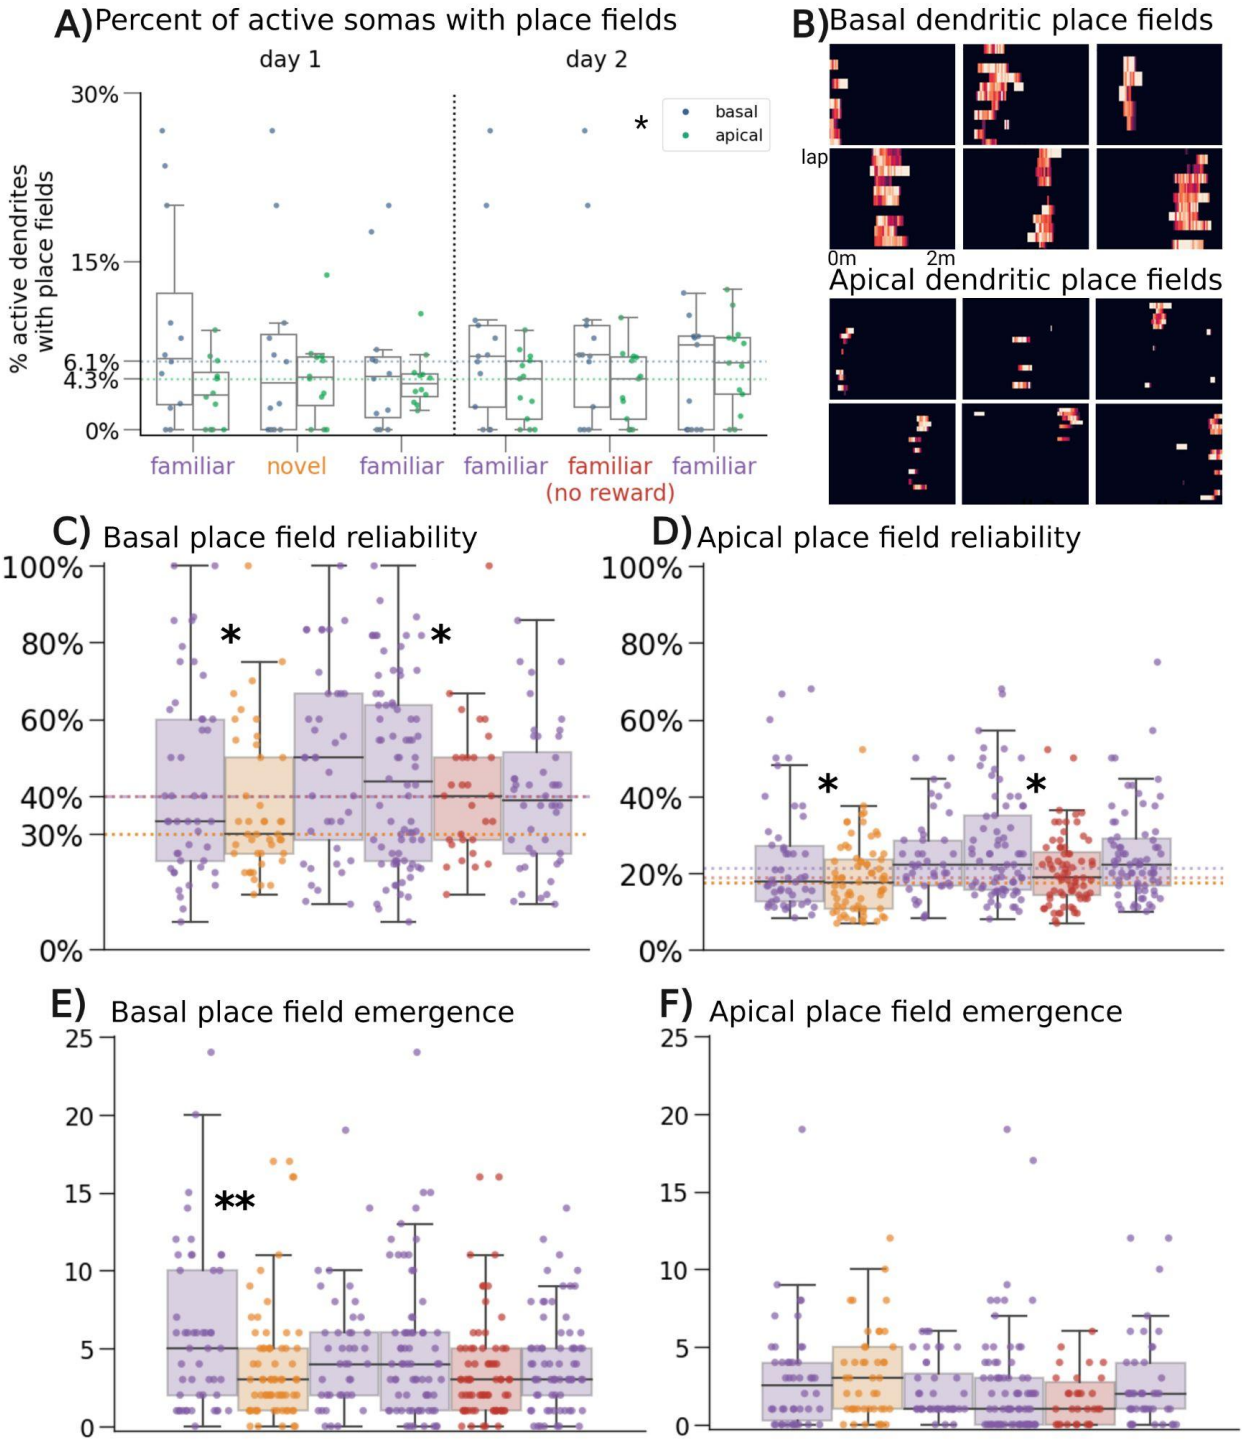
A) Percent of active somas with place fields
day 1
day 2
30%
15%
6.1%
4.3%
0%
% active dendrites with place fields
basal
apical
*
familiar
novel
familiar
familiar
familiar (no reward)
familiar
B) Basal dendritic place fields
lap
0m
2m
Apical dendritic place fields
C) Basal place field reliability
100%
80%
60%
40%
30%
0%
*
*
D) Apical place field reliability
100%
80%
60%
40%
20%
0%
*
*
E) Basal place field emergence
25
20
15
10
5
0
**
F) Apical place field emergence
25
20
15
10
5
0

Figure 4.2. Basal place fields are more common, more reliable, and emerge later than apical place fields.
A) Percent of active somata with place fields. Basal place fields are overall more prevalent than apical place fields, while remaining significantly suppressed as compared to somatic place fields under similar conditions Two-way ANOVA, P = 0.024) **B)** Top; representative example basal dendritic place fields. The x-axis represents position along a 2m track binned into 40 bins, while the y axis represents individual laps (different for each example dendrite). Each example place field was taken from a different dendrite in a different context, from top left to bottom right: familiar, novel, familiar, familiar, familiar (no reward), familiar. Bottom; representative example apical dendritic place fields, plotted according to the same standards as basal dendrites. **C)** Basal place field reliability, as defined by the percent of laps the place dendrite fired after the emergence of the place field within a context. Dotted lines indicate median across the familiar contexts, novel, or familiar (no reward) contexts. After two-way ANOVA indicated significance in both context and plane but not the interaction (P context = 5.0e-4, P plane 1.1e-39), pairwise t-tests were performed with holm-sidak multiple correction comparisons comparing the first familiar context to subsequent contexts per day (P first familiar to novel = 0.035, P first familiar to familiar (no reward) = 0.026). **D)** Same as C but with apical place fields (P first familiar to novel =0.01, P first familiar to familiar (no reward) = 9.2e-3). Comparison between basal and apical independent and not independent of context can be found in Supp. Fig.s 34.3A and 4.4A, respectively. **E)** Basal place field emergence, as defined by the first lap within a context in which the place field emerged. After two-way ANOVA indicated significance in both context and plane but not the interaction (P context = 1.5e-2, P plane = 6.24e-12), pairwise t-tests were performed with holm-sidak multiple correction comparisons comparing the first familiar to subsequent contexts per day(P first familiar to novel = 9.0e-3). **F)** Same as E but with apical place fields (P= n.s.).

Remapping In Dendritic Place Maps

We next investigated population-wide shifts in activity in the basal and apical tufts throughout our experiment. We first calculated the neural discrimination index (nDI) for each dendritic ROI, then averaged across all dendrites for each mouse per FOV, per day [214]. In brief, the nDI takes the number of peaks in the first familiar context per day, subtracts the number of peaks in the comparison context, and divides it by the total number of peaks across both contexts. A positive nDI would equate to higher firing in the first familiar context, a negative nDI would equate to higher firing in the comparison context, and a zero nDI would equate to similar firing in both contexts. While the nDI for basal dendrites did not show a significant skew away from zero in any comparison, the nDI of apical dendrites significantly skewed negative in

the novel and familiar-no-reward contexts, indicating higher firing in the apical dendrites under these two conditions (Fig. 4.3A, novel P = 0.01, familiar (no reward) P = 0.03). Additionally, this decrease in nDI was significantly different from the basal dendritic nDI in the novel and familiar (no reward) contexts (Fig. 4.3A, novel P = 2.45e-3, familiar (no reward) P = 9.9e-3). These results suggest that apical tuft dendrites selectively increase their activity under novel conditions.

We next investigated the rates of remapping in basal and apical dendrites by comparing the correlation between the mean activity of place fields on the first familiar context per plane and experimental day to the activity of those same dendritic ROIs in subsequent contexts. We found that overall, basal dendrites were more stable across contexts than apical dendrites with average basal place field stability at 21.6% and apical place field correlation at 15.9%, although even the heightened stability of basal place fields were still far less stable than somatic place fields recorded under similar experimental conditions [223,297]. The remainder of place fields in the familiar context that were not place fields subsequently vanished, as we recorded no dendritic place fields that remapped to another place field location between the familiar context and any other context, either familiar, novel, or unrewarded. Interestingly, a small subset of place fields in both basal and apical dendrites retained their place selectivity even in the novel context, potentially due to context-invariant spatial information contained of direct EC inputs to the apical tuft or CA3 inputs to the basal tuft (Fig. 4.3B). We then visualized the overall place maps across mice in both basal and apical dendrites. In the basal dendrites, we observed between 29-68 place fields per context, each independently tiling the track (Fig. 4.3C; top). In both the novel and unrewarded contexts, we noticed a decrease in the mean activity as compared to the first familiar context on each day, which was largely attributable to the decrease in percent laps active (Fig. 4.2C; left). In the apical tuft we observe between 41-67 place fields per context, also

tiling the track (Fig. 4.3D; top). In all contexts except the unrewarded context, apical place fields over-represented the end of the track (place field COMs above the white dotted line), aligned with our COM skew findings (Fig. 4.3D; top). Aligned with the overall increase in remapping observed in summary data, comparing the first familiar to subsequent contexts per day yielded highly degraded maps. Interestingly, on the novelty manipulation maps the place fields that encoded for the end of the track were more likely to retain encoding for the end of the track, while on the reward manipulation maps the place fields encoding for the end of the track virtually all disappeared with the removal of reward and did not reappear with the reinduction of reward.

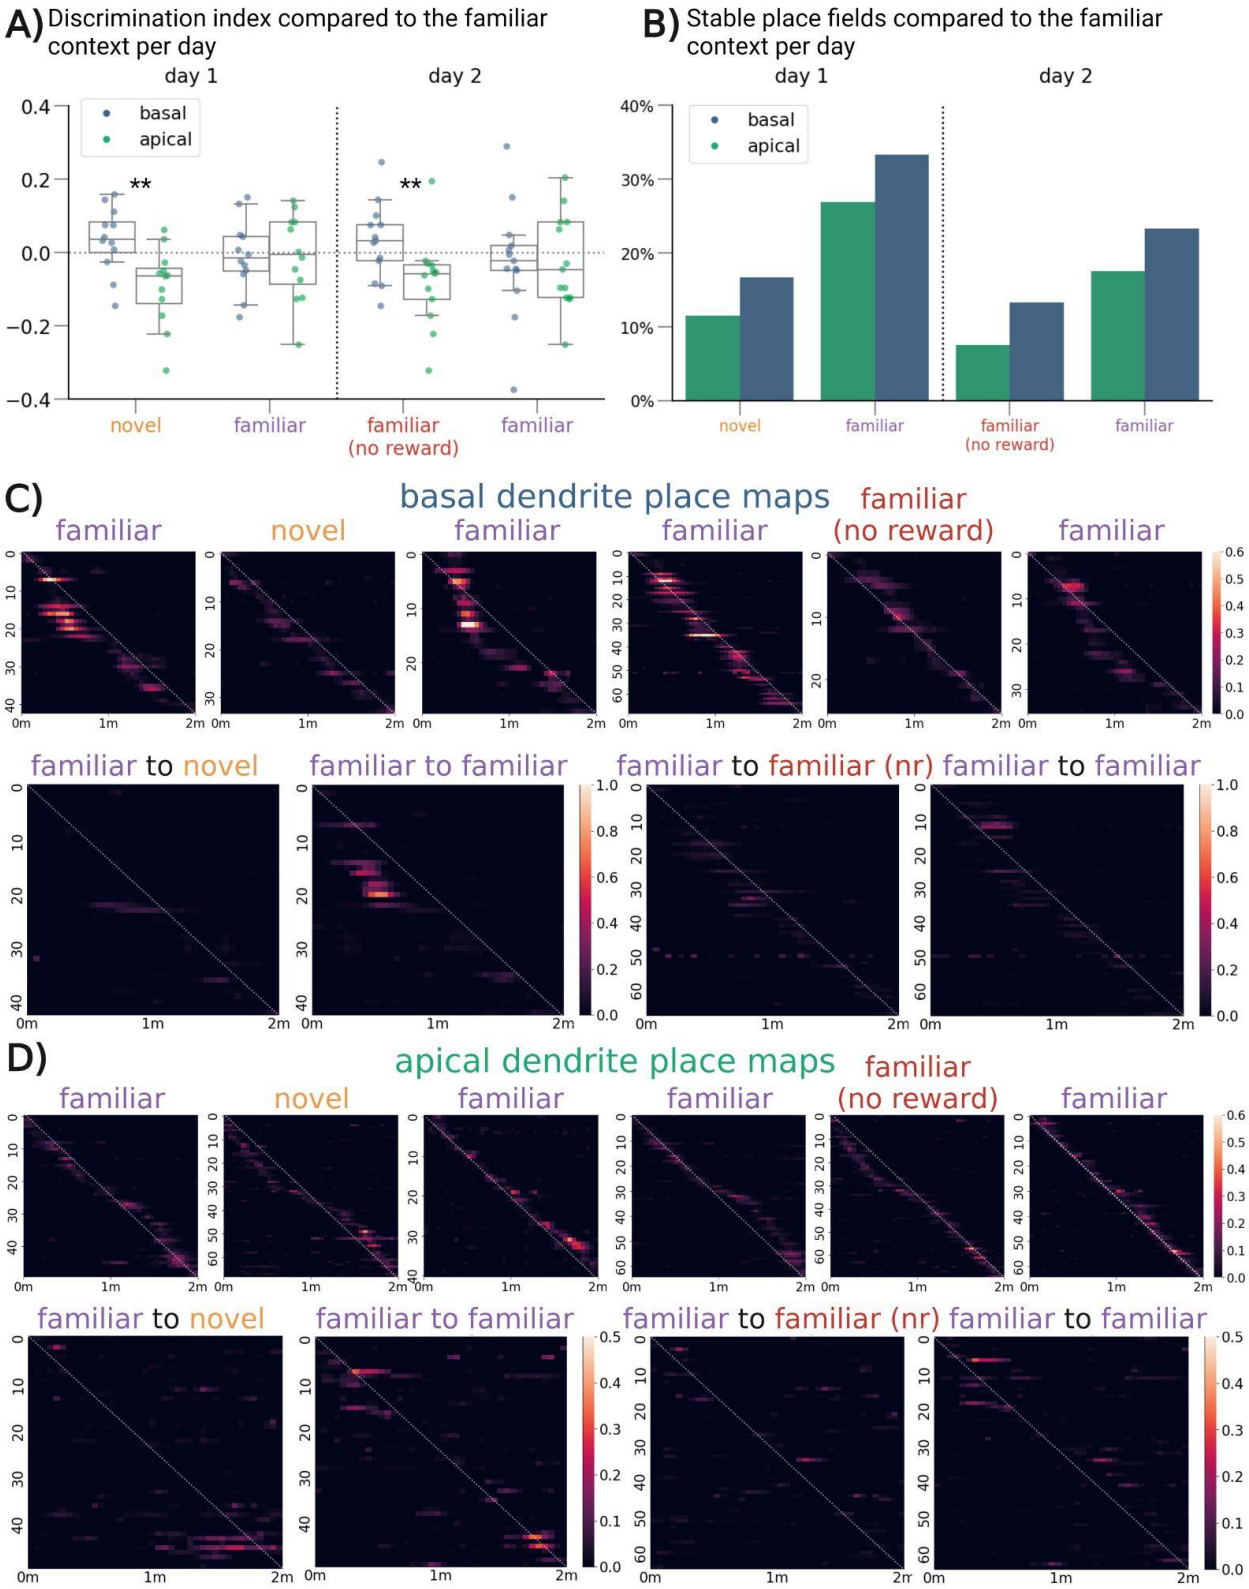
A) Discrimination index compared to the familiar context per day
day 1
day 2
basal
apical
**
**
0.4
0.2
0.0
−0.2
−0.4
novel
familiar
familiar (no reward)
familiar
B) Stable place fields compared to the familiar context per day
day 1
day 2
basal
apical
40%
30%
20%
10%
0%
novel
familiar
familiar (no reward)
familiar
C)
basal dendrite place maps
familiar
novel
familiar
familiar
familiar (no reward)
familiar
familiar to novel
familiar to familiar
familiar to familiar (nr)
familiar to familiar
D)
apical dendrite place maps
familiar
novel
familiar
familiar
familiar (no reward)
familiar
familiar to novel
familiar to familiar
familiar to familiar (nr)
familiar to familiar
0m
1m
2m

Figure 4.3. Basal place maps are more stable, while apical place maps over-represent reward.
A) Neural discrimination index per-day calculated per ROI as comparison between firing rate in the first familiar and in subsequent contexts per plane. Pairwise t-test with holm-sidak correction indicated that DI was both lower for the apical dendrites than expected if there was no change between contexts and lower than basal dendrites (P novel from baseline = 6.2e-3 P familiar (no reward) from baseline = 0.02, P novel from basal = 2.45e-3, P familiar (no reward) from basal =9.9e-3) **B)** Percent stable place fields, calculated per ROI with a place field in the first familiar context as having overlapping place field bin locations first familiar and subsequent contexts per plane. In every context, basal dendrites were more stable than apical dendrites. **C)** Top; basal place maps across mice, per context. White dotted line across the diagonal indicates hypothetical unbiased coverage of the track. Bottom; basal remapping comparisons, plotting the activity of somata identified in the first familiar context per day against their activity in subsequent contexts. **D)** Same plots as basal, but for apical dendrites.

Discussion

Using 2-photon GCaMP8f imaging, we recorded from basal and apical dendritic populations in mice during novelty exposure and reward manipulation. We were able to robustly identify individual dendritic peaks of similar amplitude and width in both populations, with no differences between populations in mouse behavior. We were then able to characterize place field activity, identifying place fields in both basal and apical dendrites. We identified more place fields in basal than apical dendrites, and these basal dendrites had multiple differences in characteristics from apical dendrites. Namely, basal dendrites were more reliable, emerged later, had wider fields, skewed towards the beginning of the track, and remapped less over days. In comparison, apical dendrites were less reliable, emerged earlier, had narrower fields, and skewed towards the end of the track, vanishing more over days with the exception of place fields which represented reward at the end of the track. Lastly, apical dendrites were preferentially active in surprising contexts as compared to basal dendrites. To our knowledge, this investigation provides the first *in vivo* imaging of apical tuft dendrites during behavior, and the first investigation of reward manipulation in dCA1 pyramidal basal dendrites. It likewise provides the first

investigation directly comparing the characteristics of activity between basal and apical dendritic dCA1 pyramidal populations *in vivo*.

Interpreting these findings relies first on a consideration of exactly what we have measured by measuring a genetically encoded calcium indicator. Biophysically speaking, dendritic spikes are generated when the membrane potential increases superlinearly, either triggered by depolarization at the synapse from ligand-gated ion channels, or from voltage-gated ion channels opening. Calcium (Ca) spikes can fall into either category, creating transient increases in Ca that can be measured using GCaMPs[259,371]. The first is through neurotransmitter binding opening calcium channels, with the main source of neurotransmitter-induced Ca influx in CA1 pyramidal neurons coming from glutamatergic NMDA receptors (NMDAr). NMDAr are required for large plateau potential Ca events that correspond to spontaneous complex burst firing, and in the basal arbor can induce modification or formation of new place cells[254,255,372]. The second is through the dendrite depolarizing enough to trigger the opening of voltage-gated Ca channels along the dendritic membrane. This depolarization threshold can be hit either by strong enough Ca or non-Ca dependent, i.e. from sodium channels, spatiotemporal input to a dendritic branch, or by somatically generated back-propagated action potentials (bAPs). Thus, our recordings may be coming from any one, or a mixture, of these sources, and subsequent characterizations must discuss these separate possibilities.

Our findings indicate that the basal dendritic place fields are generally more reliable and stable across days than those of apical dendrites, encoding on average wider place fields that form temporally later. These characteristics all align with more CA3 than 'CA1-like' traits, as CA3 pyramidal neurons also encode place fields that are more reliable and stable across days, wider, and later-forming than CA1 neurons recorded simultaneously[223]. We hypothesize that the

majority of the place fields we are recording from in the basal tuft are being driven by Schaffer collaterals from CA3, as the basal dendritic arbor in CA1 is so largely driven by these massive inputs. This theory is partially supported by a lack in remapping across contexts in individual dendrites, which entirely vanished rather than change place field position, indicating that these dendrites may be receiving stable spatially segregated input (Fig 4.3B). Alternatively, these basal dendrites could be reflecting back-propagation from CA1 somatic firing, where we unintentionally record from a subset of CA1 somata that disappear instead of remapping in after changing contexts. Inactivating CA3 inputs to CA1 and recording from the basal tuft dendrites would reveal if these inputs are largely being driven by CA3 or by CA1 backpropagation.

Apical tuft activity imaging revealed a decreased spatial preference, with generally unreliable and unstable place fields rendered from a smaller population of active dendrites. This spatial information, albeit weakened, present in the apical tuft could represent either spatially salient information coming in from MEC or LEC axons, or a significantly lower fraction of back-propagation from the soma reaching the apical tuft. These apical do, however, show strong preferential encoding of changing, unexpected, and behaviorally salient information. This encoding is expressed through increased apical dendritic activity in the first few seconds of a new context, increased activity in novel and unrewarded contexts, and increased activity and place field over-representation near the rewarded location at the end of the track that disappeared with reward removal. As these skewed responses were not observed in the basal tuft, and any back-propagation strong enough to reach the apical tuft would almost certainly also reach the basal tuft, we are reasonably confident that the source of this information comes from inputs to the dendritic arbor. The reward preference specifically mirrors similar reward preferences

observed in LEC axons to the apical tuft, pointing to the LEC-CA1 pathway as their likely source[354].

Some of our findings are similar across the basal and apical arbor, such as a decrease from familiar context reliability in the novel and unrewarded contexts (Fig 4.2C-D). We hypothesize that these types of cross-arbor similarities are largely being driven by back-propagation from the somata, although it could be due to similarity in inputs from EC being preserved throughout the trisynaptic loop and to the apical tuft. This hypothesis is less likely, largely as inputs to the trisynaptic loop come from a different layer than the direct inputs to CA1[134]. It could also arise from interneuronal silencing or neuromodulation impacting calcium spiking in both dendritic regions.

Future directions

A major limitation of this work is the lack of simultaneous *in vivo* basal, apical, and somatic recordings. While basal and somatic recordings can be conducted reasonably easily together near-simultaneously by rapidly alternating the field of view, the distance from the apical tuft to the somata stretches the limits of this type of 'jumping' via piezoelectric mirror control, creating optical distortions that overtake the majority of the field of view. Even if this limitation could be overcome, the frame rate used to capture three fields of view simultaneously would reduce each to 10 frames per second, which would miss a large fraction of dendritic events as the peak duration of many events we recorded fell below 0.3s (Supp. Fig 4.2A). Indeed, we are likely still missing a significant fraction of events by being limited to 30 frames per second, a limitation difficult to overcome in multi-photon microscopy. While scanless multiphoton imaging techniques are under active development, they at present require significant trade-offs in

either resolution or field of view size [373,374]. Future technique development will ideally enable whole-neuron imaging *in vivo* that can simultaneously assess the role of the dendritic arbor in CA1 pyramidal activity[375].

Together, these results suggest that the role of basal dendrites may be to stably inherit spatial information from CA3, while the role of apical dendrites may be to contribute non-spatial information about novelty and reward from cortical and thalamic inputs, with the ultimate goal of CA1 somata flexibly encoding behaviorally salient contextual representations.

Methods

Experimental Model and Subject Details

All experimental and surgical procedures were in accordance with the University of Chicago Animal Care and Use Committee guidelines. We used 10-20 week old male C57BL/6J wildtype (WT) mice (24-35 g). Male mice were used over female mice due to the size and weight of the headplates (9.1 mm x 31.7 mm, ~2 g) which were difficult to firmly attach on smaller female skulls. Mice were individually housed in a reverse 12 hour light/dark cycle and behavioral experiments were conducted during the animal's dark cycle. We are unaware of any influence of strain or sex on the parameters analyzed in this study. A total of 39 mice were used in this study, 21 of which were used in the final dataset. 4 mice were excluded for cannula or headplate instability. 3 mice were excluded for not reaching the running behavior criteria (See Methods: Behavior). The remaining 11 mice were excluded for not having sufficiently bright identified GCaMP8f expression in both the basal and apical tuft.

Mouse surgery and viral injections

Mice were anesthetized (~1-2% isoflurane) and injected with 0.5 ml of saline (intraperitoneal IP injection) and 0.5 ml of Meloxicam (1-2 mg/kg, subcutaneous injection) before being weighed and mounted onto a stereotaxic surgical station (David Kopf Instruments). A small craniotomy (1-1.5 mm diameter) was made over the hippocampus (± 1.7 mm lateral, -2.3 mm caudal of Bregma). Either CamKIIa-jGCaMP8f-WPRE (AAV-CamKIIa-jGCaMP8f-WPRE was a gift from GENIE Project; $\geq 7\times10^{12}$ vg/mL; viral prep # 176750-AAV1 ; http://n2t.net/addgene:176750) was diluted 1:1 with PBS or a 1:1 mixture of AAV-CaMKII0.4.Cre-SV40 (pENN.AAV.CamKII 0.4.Cre.SV40 was a gift from James M. Wilson, $\geq 7\times10^{12}$ vg/mL; viral prep #105558-AAV1 ;

http://n2t.net/addgene:105558) and pGP-AAV-syn-FLEX-jGCaMP8f-WPRE (pGP-AAV-syn-FLEX-jGCaMP8f-WPRE was a gift from GENIE Project, $\geq 7\times10^{12}$ vg/mL; viral prep #162379-AAV1; http://n2t.net/addgene:162379) was injected into left hippocampal CA1 at the somatic layer (~100 nL at a depth of -1.25mm below the surface of the dura). Afterwards, the site was covered using dental cement (Metabond, Parkell Corporation) and a metal head-plate (9.1 mm x 31.7 mm, Atlas Tool and Die Works) was also attached to the skull with the cement. Mice were separated into individual cages and water restriction began the following day (0.8-1.0 ml per day). At least 7 days following injection surgery, and approximately 7 days prior to the beginning of mouse training, mice underwent another surgery to implant a hippocampal window as previously described[323]. Following implantation, the head-plate was reattached with the addition of a head-ring cemented on top of the head-plate which was used to house the microscope objective and block out ambient light. Post-surgery mice were given 1-2 ml of water/day for 3 days to enhance recovery before returning to the reduced water schedule (0.8-1.0 ml/day). Mice weights were monitored daily to ensure healthy weight maintenance post-water deprivation. Expression of GCaMP8f reached a steady state ~14 days after the virus was injected, as monitored through 2p imaging.

Behavior

Our virtual reality (VR) and treadmill setup was designed similarly to previously described setups[297]. The virtual environments that the mice navigated through were created using VIRMEn[155]. Mice were head restrained with their limbs comfortably resting on a freely rotating styrofoam wheel ('treadmill'). Movement of the wheel caused movement in VR by using a rotary encoder to detect treadmill rotations and feed this information into our VR computer, as in (Heys

et al., 2014; Sheffield et al., 2017). During training, mice received a water reward (4 µl) through a waterspout upon completing each traversal of the track (a lap), which was then associated with a clicking sound from the solenoid. Upon receiving the water reward, a short VR pause of 1.5 s was implemented to allow for water consumption and to help distinguish traversals from one another rather than them being continuous. Mice were then virtually teleported back to the beginning of the track and could begin a new traversal. Mice were also teleported to the beginning of a new contextual exposure. Seven distinct VR contexts were used in this experiment. The 'familiar' context was the context that mice trained in prior to imaging. Two novel contexts were used for novelty manipulation. A variant of the familiar context with the water reward removed at the end of the track was used for reward manipulation. Mouse behaviors (running velocity, track position) were collected using a PicoScope Oscilloscope (PICO4824, Pico Technology). Behavioral training to navigate the virtual environment began ~7 days after window implantation (~30 minutes per day) and continued until mice reached a speed of greater than 4 traversals per minute, which took 10-14 days. This high level of training was necessary to ensure mice continued to traverse the track similarly after reward was removed. Initial experiments showed that mice that failed to reach this criterion typically would not traverse the track as consistently without reward[297]. Mice that did not reach this criterion before the start of experiments did not continue in the paradigm.

Novelty and Reward Manipulation Paradigm

Mice were imaged in either their apical or basal dendritic tuft in dCA1 for ~5 minutes (322s) per context. Which plane mice had imaged first (apical or basal) was counterbalanced, with 8 mice experiencing apical and then basal tuft imaging on both days, and 7 mice experiencing basal and then apical tuft imaging on both days. First, in a single plane field-of-view (FOV), they navigated the familiar context with water rewards. They were then switched into a novel context

they had never experienced before, and navigated the novel context with water rewards. They were next switched back to the familiar context. The plane was then changed (from apical to basal N=8, from basal to apical N=7), and experiment repeated with a new novel context. The following day, mice were placed back in the familiar context with water rewards at the end of the track, then the familiar context with water removed, then the familiar context with water reinstated. The plane was then changed (from apical to basal N=8, from basal to apical N=7), and the experiment was repeated.

Two-photon imaging

Imaging was conducted 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 16x/0.8 NA/3 mm WD water immersion objective (MRP07220, Nikon). GCaMP8f was excited at 920 nm with a femtosecond-pulsed two photon laser (Insight DS+Dual, Spectra-Physics) and emitted fluorescence was collected using a GaAsP PMT (H11706, Hamamatsu). The average power of the laser measured after the objective ranged between 20-42mW for basal dendritic imaging and 53-85 mW for apical dendritic imaging. Each field of view (FOV) was positioned between -50-100 μm below the putative surface for basal dendritic imaging and -300-400 μm below the putative surface always retaining at least -150 μm distance from the center of the somatic layer for apical tuft dendritic imaging in the *z* direction, and 400-700 μm equally in the *x/y* direction. Planes were selected to be over the highest number of visible active CA1 pyramidal cells with visible activity in both the basal and apical FOVs prior to imaging. Each basal and apical plane was imaged at the same magnification per mouse and day. Time-series images were collected through Scanbox (Neurolabware) and the PicoScope Oscilloscope (Pico Technology) was used to synchronize frame acquisition timing with behavior.

Image Processing and ROI Selection

Time-series images were preprocessed using Suite2p (Pachitariu et al., 2017). Movement artifacts were removed using rigid and non-rigid transformations and assessed to ensure absence of drifts in the *z*-direction. All datasets collected during shock administration were discarded, due to the high velocity post-shocks sprinting behavior of mice rendering FOVs too unstable for reliable analysis. Regions of interest (ROIs) were also defined using Suite2p and manually inspected for accuracy. Baseline corrected *Δf/f* traces across time were then generated for each ROI. Per ROI, a savitzky-golay filter was applied to both channels to smooth the signal. Then, ROIs were analyzed for covariance, and any ROIs exceeding the 99th percentile of a shuffle distribution were combined using PCA and the first PC taken, in a method similar to Kaufman et al. 2020[324]. From this trace, the fc3 was calculated as previously described and used for further analysis, with the modification that four sigma was required instead of two sigma to account for noisier baselines in dendritic data and reduce false positives[147]. Only ROIs that exceeded 30% *Δf/f* at least once per experimental day were retained. Peaks were calculated using the scipy.signal.find_peaks package with a required minimum height of 30% *Δf/f,* distance of 0.3s, and prominence of 0.1.

Neural Discrimination Index

Neural discrimination index (nDI) was calculated as described in [214]. Briefly, for each ROI, the number of transient peaks were counted across a single context, in a single day. The nDI was then calculated across context or day, as described in the text, by calculating a comparison number of transient peaks for the comparison context, subtracting the first context from the second, then dividing by the sum of transients in both.

Place Field Identification

Place fields were identified using a modification of the 'peak' method described in Grijseels et al, 2021. [359]. Briefly, the fc3 of ROI, after processing as described above, was binned into 40 bins along the 2 meter virtual track. Activity in each bin was identified for putative place field behavior based on being active on at least 2 laps in that bin, activity post-emergence being present in 30% of subsequent laps, and being at least 1 bin wide and no more than 17 bins wide. These putative fields were then shuffled 1000 times against ROI data and tested for significance, requiring firing in at least one location along the track to be equal to or greater than two standard deviations above the shuffled activity. This method allowed for multiple place fields per ROI, which were each treated separately for subsequent analysis with the other IDed place field's activity disregarded for properties such as out-of-field-firing and precision. Only one ROI with multiple place fields was identified.

Place Field Metrics

All metrics described below were conducted per mouse, day, plane, context, and post-processed ROI, excepting remapping which requires computation across context.

Active fraction: Each ROI with activity in at least one of the contexts on a per-day basis was counted, and the fraction was computed as the number of ROIs that contained at least one place field within that context divided by the number of active ROIs.

Emergence: Starting from the first lap the mouse traversed in the context, we searched lap-by-lap for a lap with a significant calcium transient present within the boundaries of the future place field, as calculated from all the laps in the session. This lap was identified as the place field emergence lap.

Reliability: From the emergence lap to the last lap the mouse traversed in the context, we searched lap by lap for significant calcium transients. This metric is the number of laps with significant in-field calcium transients divided by the total number of laps traversed post-place field emergence.

Out/in field ratio: The mean *Δf/f* in bins within the place field was divided by the *Δf/f*in bins outside of the place field.

Center of Mass: We first measured *Δf/f* in each bin. We then used the following equation to calculate the COM for each traversal n (COMn):

$$COM_n = \frac{\sum_i DF_i \cdot x_i}{\sum_i DF_i}$$

Where DFi is the somatic *Δf/f* in bin i and xiis the distance of bin i from the start of the track. We then calculated the weighted average COM (COMw) from all traversals n (COMn from each traversal was weighted by the peak transient *Δf/f* on that traversal (An)):

$$COM_w = \frac{\sum_n A_n \cdot COM_n}{\sum_n A_n}$$

Precision: Spatial precision (SP) was calculated as the inverse of the COM standard deviation:

$$SP = \frac{1}{\sqrt{\frac{\sum_n A_n (COM_n - COM_w)^2}{\sum_n A_n}}}$$

Width: The difference between the first and last bin with in-field activity, expressed as number of bins.

Slope: Slope represents a linear regression as calculated using the scipy.stats linregress module from the lap identified to have peak activity and the center of mass, thus yielding a calculation of place field shifting over time. A positive metric would indicate forward shifting, while a negative metric would indicate backwards shifting.

r^2: Goodness-of-fit of the slope as calculated previously.

Remapping: To measure spatial correlation across contexts in place fields, we found dendrites that had place fields in the first familiar context, and then calculated the Pearson's correlation coefficient between the mean activity along the track (in 40 bins) for all laps between that context and subsequent context exposures.

Behavioral Parameters

All parameters described below were calculated per mouse, day, plane, and context.

Freezing: Freezing epochs were determined as uninterrupted epochs where mouse velocity fell below 0.001 cm/s for at least 12 consecutive frames (~0.75 s). All epochs of velocity below 0.001 cm/s but not reaching 12 consecutive frames were not considered freezing or running, and were discarded from future analysis. Freezing epochs were then counted up, and each not in a freezing epoch assigned a '0', while each frame in a freezing epoch given a numeric value corresponding to the number of epochs in that recording (i.e. all frames that contained the 4th freeze of the recording would be assigned the integer '4'). Subsequent freeze features were then calculated, including the binary variable 'is freezing' which assigns a 1 to frames considered freezing, and 0 to frames not considered freezing, two sawtooth functions 'freeze remaining', and 'freeze elapsed, which counts the frames from the beginning of a freeze up or down until the end of a freeze, respectively, and 'freeze progress' which tracks the progress of a freeze as a fraction from 0 to 1.

Running: Running was determined as any epoch where forward progress velocity was sustained over 0.001 cm/s for 2 consecutive frames. The variables 'is running', 'running remaining', 'running elapsed', and 'running progress' are calculated using the running epoch data in the same fashion as their freezing counterparts.

Backward movement: Some mice demonstrated backward movement behavior in the virtual, where they attempted to 'backtrack' through the context. This behavior was analyzed separately from running or pausing. The binary variable 'is backtracking' assigns a 1 to frames considered backtracking, and 0 to frames not considered backtracking.

Velocity: Velocity was both directly measured through the picoscope encoder, and recalculated from position, to assess for accuracy. Recorded velocity was used for all velocity calculations and model training. Values were converted into cm/s for presentation.

Acceleration: Acceleration was calculated as the first derivative of recorded velocity.

Location: Animal's position on virtual track was determined for each frame, and binned in 1 cm bins along the virtual track.

Statistics

For all data distributions, a Shapiro-Wilk test was performed to verify the data was normally distributed before undergoing further statistical tests. Two-way ANOVAs were performed per metric to identify if context, plane, or context|planes were significant factors. Post-hoc pairwise T tests were then performed with a holm-sidak multiple comparisons correction. Box and whisker plots were used to display data distributions where applicable. The box in the box and whisker plots represent the first quartile (25^{th} percentile) to the third quartile (75^{th} percentile) of the distribution, showing the interquartile range (IQR) of the distribution. The black line across

the box is the median (50th percentile) of the data distribution. The whiskers extend to 1.5*IQR on either side of the box. A data point was considered an outlier if it was outside the whiskers or 1.5*IQR. Horizontal dotted lines indicate group median, minimum possible value, or zero, as indicated in figure legend. Significance tests were performed with and without outliers. Data distributions were considered statistically significant only if they passed significance ($p < 0.05$) without outliers. Significance numbers reported are without outliers. Correlations were performed using Pearson's correlation coefficient. $p < 0.05$ was chosen to indicate statistical significance and p-values presented in figures are as follows: *, $p < 0.05$, **, $p < 0.01$, ***, $p < 0.001$, N.S. not significant. All regression analysis was conducted using the statsmodels package. Some data preprocessing was conducted with MATLAB (Mathworks, Version R2022a). All other data and statistical analyses were conducted in Python 3.7.4, with primary data accrued in Pandas DataFrames, and data figures were made in Python 3.7.4 using the Seaborn and Matplotlib packages (https://www.python.org/). Schematic figures, some figure text, and figure layouts were made with BioRender (https://biorender.com/).

Data availability

Raw imaging data are extremely large and not feasible for upload to an online repository but are available upon request. Processed data will be available on GitHub after publication of this work in an academic, peer-reviewed journal.

Code availability

Original code used to create figures from pre-processed data will be available on GitHub after publication of this work in an academic, peer-reviewed journal.

Acknowledgements

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, a New Innovator grant from the National Institutes of Health (1DP2NS111657-01) awarded to M.S., and a BRAIN Initiative grant from the National Institutes of Health (RF1NS127123). We thank the University of Chicago imaging core for assistance with confocal imaging, and the University of Chicago animal care staff for ensuring the well-being of experimental animals. We thank Timothy Ratigan for assistance with data analysis. We thank Antoine Madar for assistance with statistical methodology. We thank Rossten Rad for discussions on experimental design.

Contributions

Heather C. Ratigan (H.R.), Valerie Barreto (V.B.), Chery Cherian (C.C.), and Mark E.J. Sheffield (M.S.) contributed to this work. H.R. and M.S. conceived of and designed the experiments. H.R., V.B., and C.C. performed surgeries, animal training, collected all *in-vivo* behavioral and imaging data, and collected all *post-hoc* data. H.R. wrote the analysis code and analyzed all data. H.R and M.S. interpreted the data. H.R. created all figures and wrote the manuscript.

Ethics Declarations

The authors declare no competing interests.

Key Resources Table

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Bacterial and virus strains		
pENN.AAV-CamKIIa-jGCaMP8f-WPRE	Zhang et al. 2023[360]	Addgene #176750-AAV1
pENN.AAV.CamKII 0.4.Cre.SV40	AAV Plasmids[361]	Addgene #105558-AAV1

pGP-AAV-syn-FLEX-jGCaMP8f-WPRE	Zhang et al. 2023[360]	Addgene #162379-AAV1
Experimental models:	**Organisms/strains**	
Mouse: C57BL/6J	Jackson Laboratories	JAX 000664 - C57BL/6J
Software and Algorithms		
Fiji	Schindelin et al., 2012[327]	https://imagej.net/software/fiji/; RRID:SCR_002285
Suite2p	Pachitariu et al., 2016[328]	https://github.com/MouseLand/suite2p
MATLAB	MATLAB. (2018). *9.7.0.1190202 (R2018a)*.	https://www.mathworks.com/help/matlab/release-notes-R2018a.html
Python	Python 3.10.8	https://www.python.org/downloads/release/python-3108/
Pandas	Pandas 1.1.4	https://pandas.pydata.org/
SciPy	SciPy 1.9.3	https://scipy.org/install/
Seaborn	0.12.0	https://seaborn.pydata.org/installing.html

Supplementary Figures

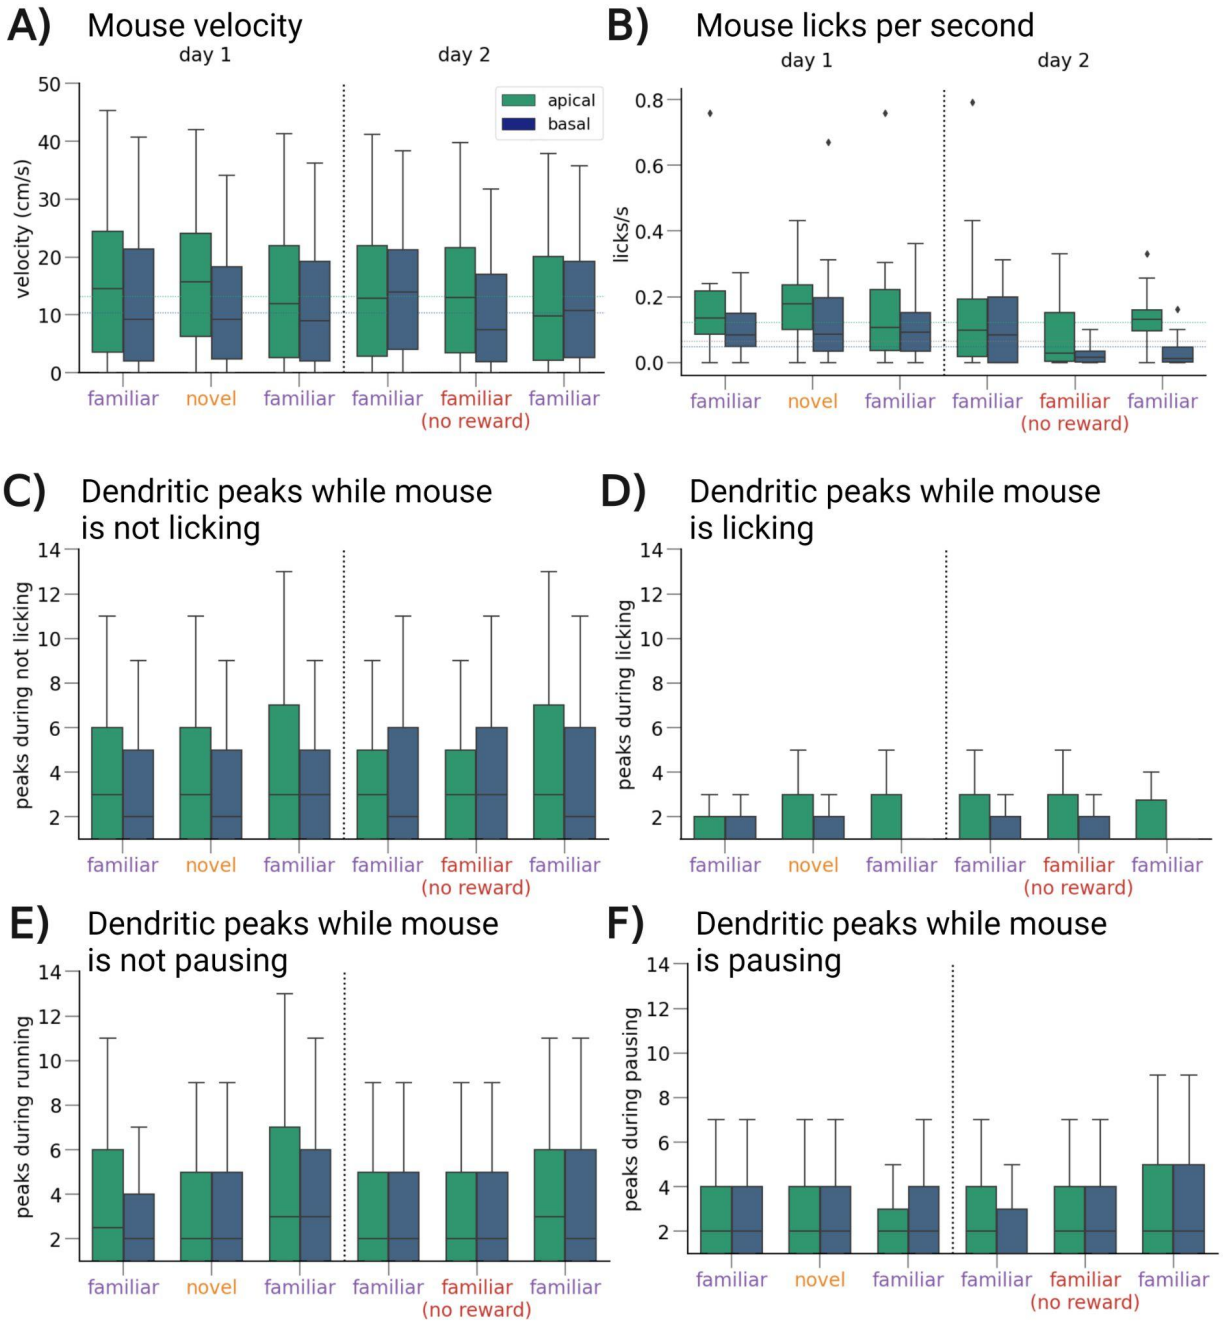

Supplementary Figure 4.1. Mouse behavior and correlation with peak activity.
A) Mouse velocity in centimeters per second in each context in all apical (green, left bars) vs. basal (blue, right bars) mice are similar across contexts and imaging planes. **B)** Mouse lick rates in apical and basal context as calculated by licks per second are similar across contexts and imaging planes **C)** Average dendritic peaks (in peaks per second) across all dendritic segments in the apical vs. basal planes while mice are not licking are similar across contexts and imaging planes. **D)** Average dendritic peaks (in peaks per second) across all dendritic segments in the

apical vs. basal planes while mice are licking are similar across contexts but different across imaging planes. However, low sample size due to low mean frequency of licking (~0.1 licks/s) results in low confidence. **E)** Average dendritic peaks (in peaks per second) across all dendritic segments in the apical vs. basal planes while mice are not pausing, i.e. are running, are similar across contexts and imaging planes. **F)** Same as E, but while mice are pausing.

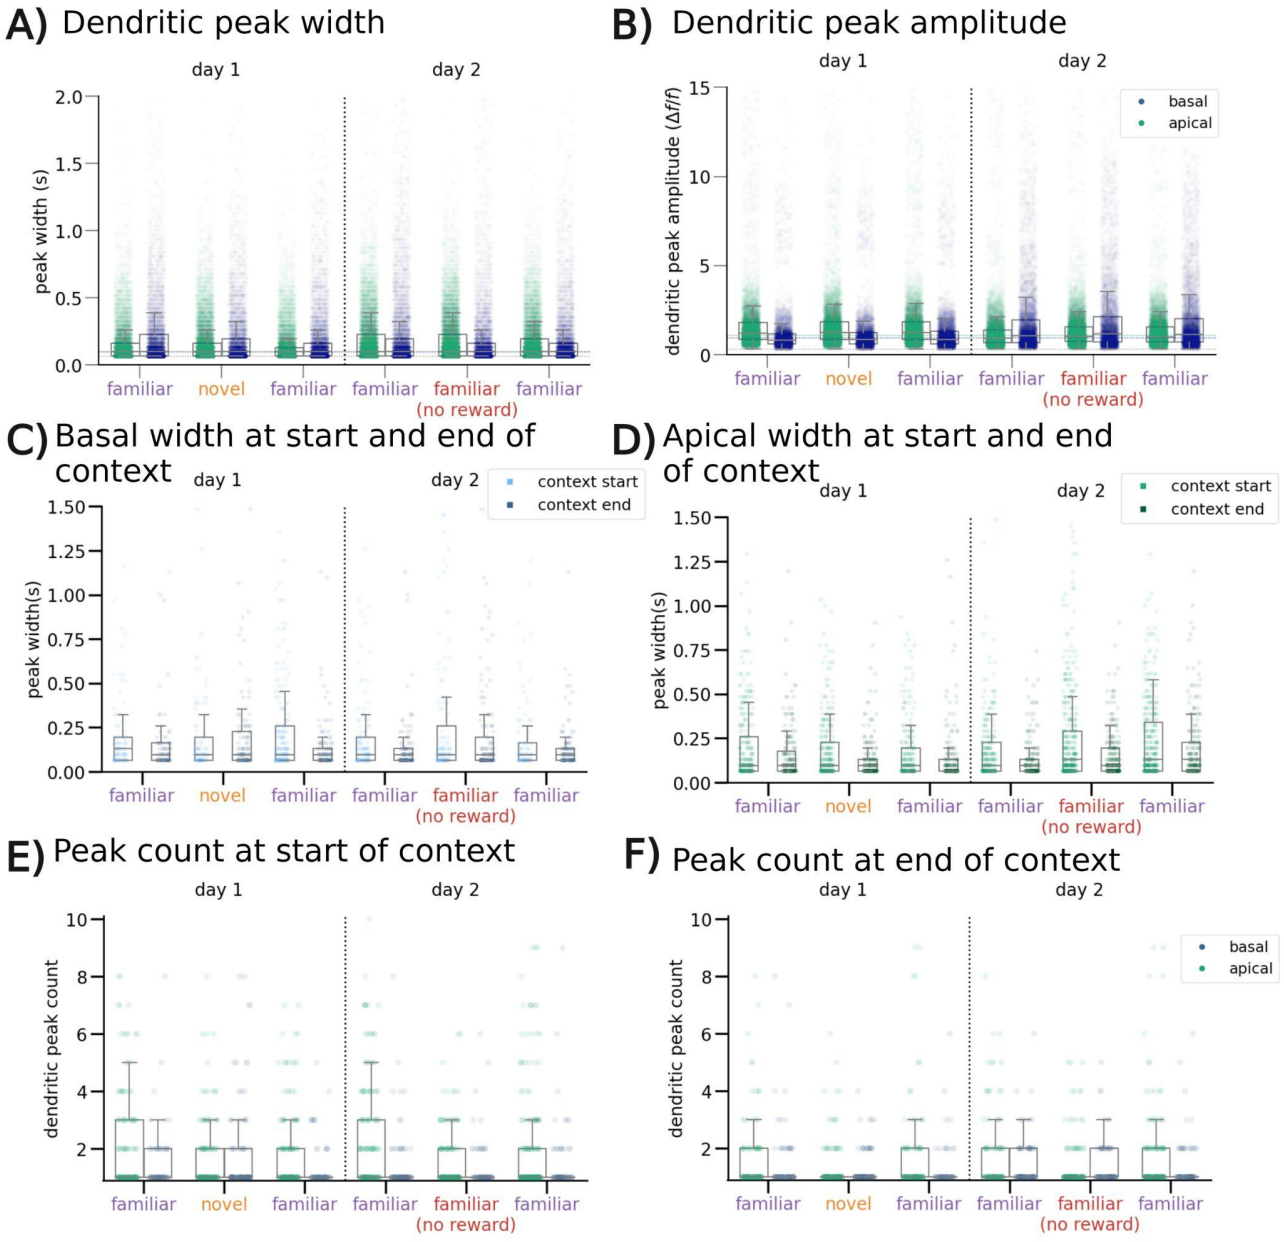

Supplementary Figure 4.2. Additional dendritic peak metrics.
A) Peak widths of each dendritic peak, segmented across plane and context, showing similar average peak width in both apical and basal dendritic planes in all contexts. **B)** Peak amplitudes of each dendritic peak, segmented across plane and context, showing similar average peak amplitude and range in both apical and basal dendritic planes in all contexts. **C)** Peak widths in the basal context within the first 200 frames (6.45s) or last 200 frames of the context showing no difference in peak width between beginning and end of the context exposure. **D)** Peak widths in

the apical context within the first 200 frames (6.45s) or last 200 frames of the context showing an increase in peak width at the beginning as compared to the end of the context exposure across all contexts. **E)** Peak counts in the apical versus the basal plane across contexts within the first 200 frames (6.45s) of each context showing increase in peak count in the apical as compared to the basal plane. **F)** Peak counts in the apical versus the basal plane across contexts within the first 200 frames (6.45s) of each context showing similarity in peak count in the apical as compared to the basal plane.

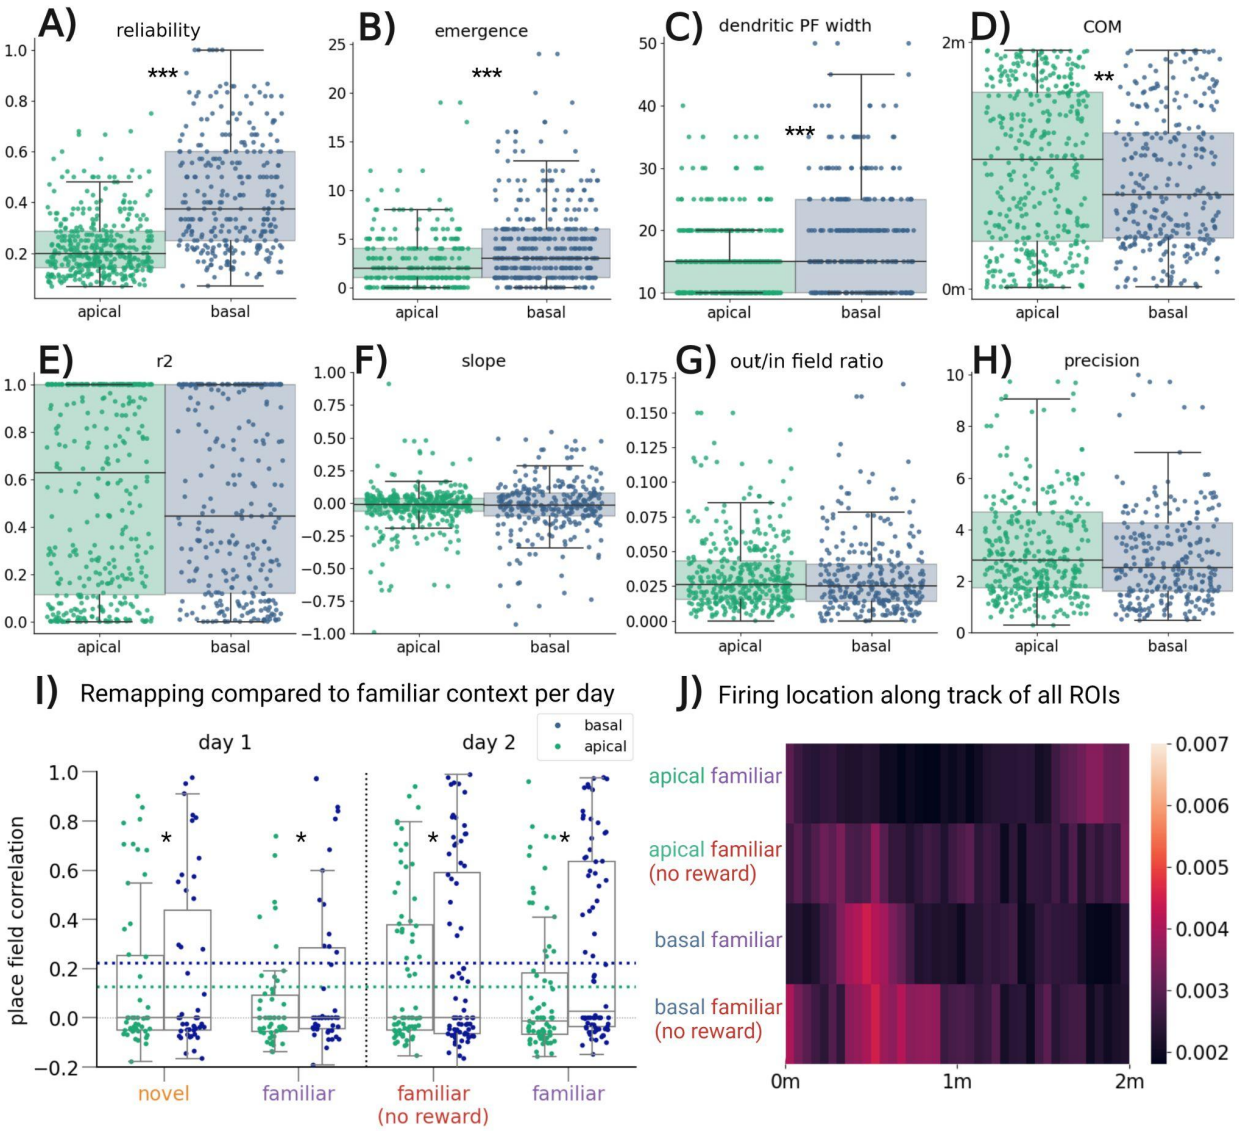

Supplementary Figure 4.3. Additional dendritic place field metrics by plane

A) Reliability is significantly higher in the basal than apical dendrites across contexts. **B)** Basal dendrites emerge on significantly later laps than apical dendrites. **C)** Basal dendrites have significantly wider place fields than apical dendrites. **D)** Basal dendrites have significantly forward-shifted average centers of mass compared to apical dendrites, which have significantly end-shifted average centers of mass. **E)** r2 is a metric of goodness-of-fit from testing for backwards or forward shift, which is similar between planes. **F)** Slope is the recorded slope of backwards (negative) or forwards (positive) shifting, which is not different from zero or from each other between planes. **G)** The firing outside of a place field's identified field over the firing inside the field was similar between planes. **H)** The precision of the place fields was similar between planes, although there was a slight trend towards lower precision (higher value) in apical place fields. **I)** Remapping as defined by place field correlation between each individual dendritic ROI with a place field present on the first familiar context and activity on subsequent context exposures. Basal dendrites had reliably higher correlations, indicating less remapping. **J)** Heatmap of all peaks in within-specified context(s). Top two rows are apical planes, bottom two

rows are basal planes. Top-most row is apical peaks across all familiar contexts, showing strong preference for the end of the track. Middle-top row is apical peaks in the unrewarded context, showing lack of preference for the end of the track. Middle-bottom row is basal peaks across all familiar contexts, showing preference for the beginning of the track. Bottom row is basal peaks in the unrewarded context, showing lack of change from the familiar rewarded context. Heatmap represents peaks/s.

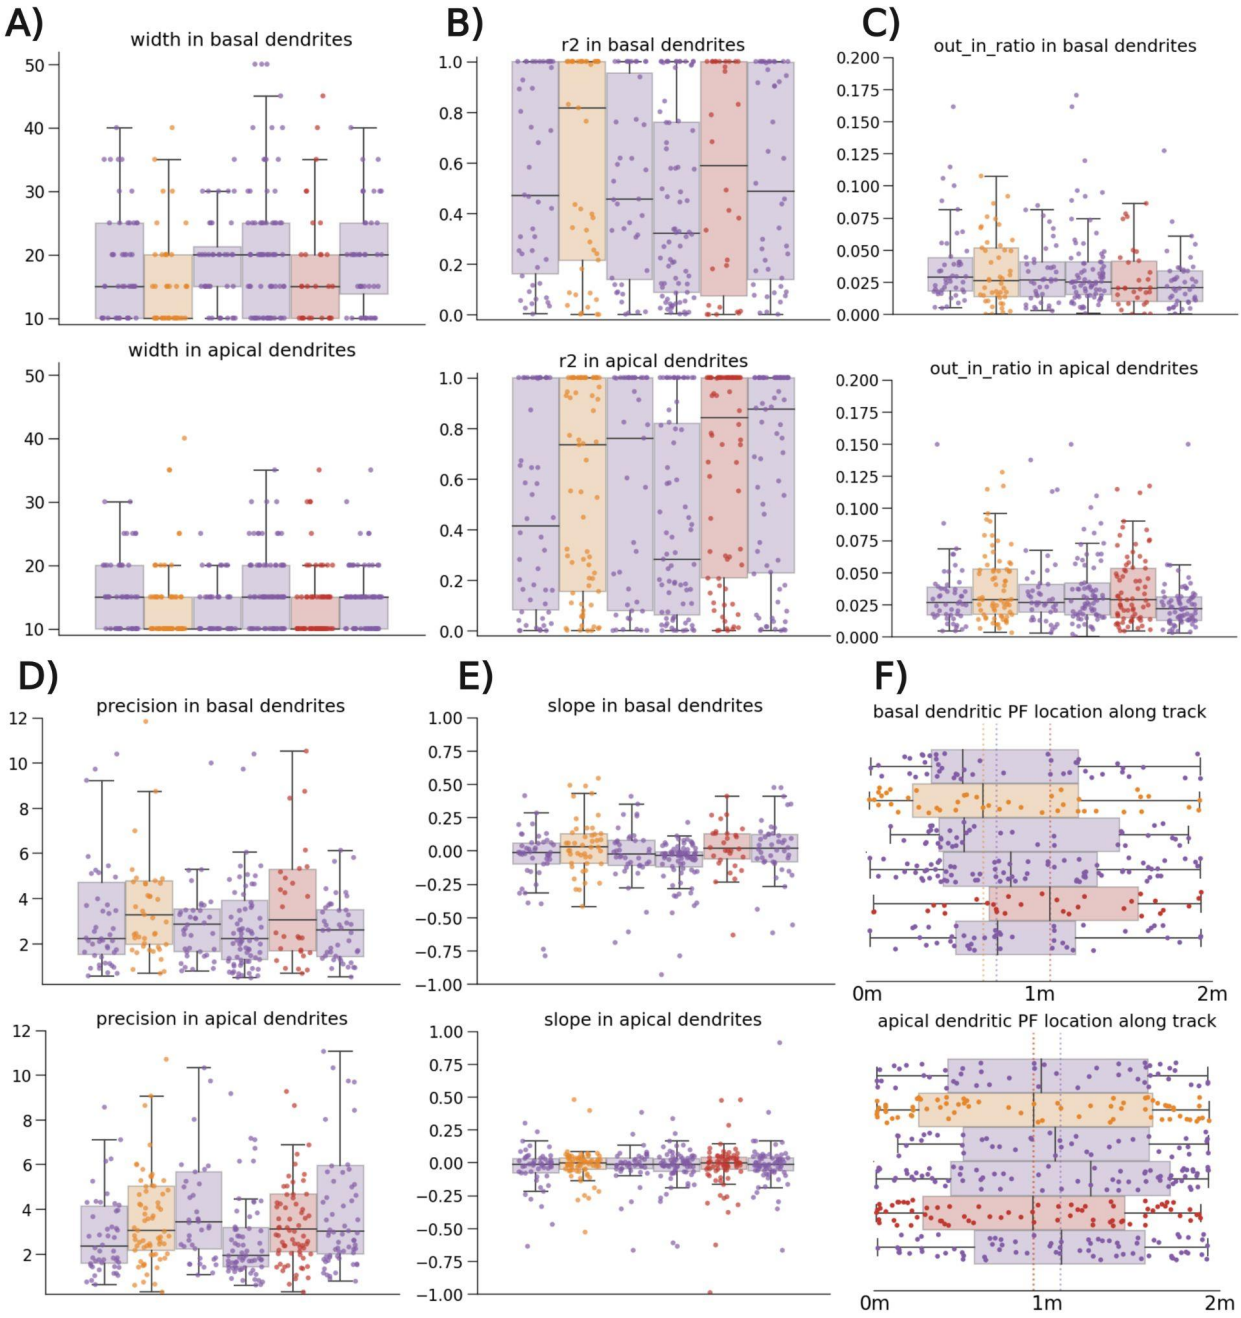

Supplementary Figure 4.4. Additional dendritic place field metrics by context
All metrics are for all place fields identified in that context. No metrics showed context-dependent effects in two-way ANOVAs. **A)** Width in basal (top) and apical (bottom) dendrites by context. **B)** r2 in basal (top) and apical (bottom) dendrites by context. **C)** out/in field firing ratio in basal (top) and apical (bottom) dendrites by context. **D)** Precision in basal (top) and apical (bottom) dendrites by context. **E)** Slope in basal (top) and apical (bottom) dendrites by context. **F)** Center-of-mass in basal (top) and apical (bottom) dendrites by context.

SYNTHESIS

Summary of Findings

This dissertation has explored the role of hippocampal dorsal CA1 (dCA1) in the encoding of aversive experiences through three different lenses, with a special emphasis on the role of thalamic nucleus reuniens (NR). Presented in Chapter 2, we used a novel virtual-reality contextual fear conditioning paradigm (VR-CFC) in parallel with chemogenetics and two-photon imaging to inactivate, record from, and model NR axons in dCA1 throughout VR-CFC. In Chapter 3 we built on these techniques to image the same population of dCA1 excitatory pyramidal somata across days throughout VR-CFC with and without NR-dCA1 inactivation. We analyzed these neural recordings for somatic contextual discrimination, place maps, shock and freeze-tuning, and co-activation events. Finally, in Chapter 4 investigated differential activity across the dendritic arbor during novelty and reward manipulation, including the first *in vivo* recordings of distal apical tuft dendrites in the hippocampus. Overall, this work represents a significant leap forward in our understanding of subcellular and population-level dynamics in the hippocampus during aversive experiences and the role of the NR-dCA1 pathway, including a range of discoveries and technique development that can be built upon in future investigations. Below, I will summarize the major findings from each chapter.

In Chapter 2, we first established that VR-CFC can induce context-dependent fearful freezing responses in mice, to our knowledge the first head-fixed fear conditioning paradigm that used shocks as the unconditioned stimulus to be more well-aligned with the freely-moving contextual fear conditioning literature, and the first head-fixed paradigm to elicit characteristic freezing responses. We next used a retro-Cre strategy to selectively inactivate NR somata that project to dCA1 on the first day of post-shock re-exposure to the shocked and neutral contexts.

This strategy elicited a significant increase in fearful freezing behavior in both the shocked and neutral contexts by increasing freezing-epoch durations, an increase which persisted across non-NR-inhibited retrieval days 2 and 3. We next imaged NR axons *in vivo* during behavior, tracking some of the same axons over days, and demonstrated that they selectively increased their activity during fearful freezing in both the shocked and neutral contexts. Modeling the activity of NR axons using a boosted-trees strategy, revealed that freezing dynamics and relatedly mouse velocity were the primary dynamics driving NR-axon activity in dCA1. This investigation led to the development of new techniques, including the first *in vivo* imaging of difficult-to-access NR axons in dCA1, demonstrated the vital importance of the NR-dCA1 pathway in contextual fear suppression, discrimination, and extinction through freeze-tuned activity, and laid a solid foundation for future investigation of the NR-dCA1 pathway.

In Chapter 3, we investigated how the NR-dCA1 pathway influenced population dynamics of excitatory pyramidal neurons to exert the effects reported in chapter 2. We first re-affirmed the critical role of NR-CA1 axons in suppressing inappropriate fear responses and enabling extinction. We then used behavioral analysis to categorize our mice into three groups: those that generalize their fear responses, those that fail to produce a fear response, and those that correctly discriminate between contexts. We found that discriminators are better at extinguishing the fear response than generalizers. We next observed that mice that discriminated well between the contexts behaviorally also expressed greater discrimination in their neural activity between contexts, and that inhibition of the NR-dCA1 pathway on retrieval day 1 prevented neural and decreased behavioral discrimination between contexts. We then showed that our VR-CFC paradigm induced near-global remapping in NR-dCA1 intact mice to more unstable, flexible maps post-extinction, while unshocked mice and NR-dCA1 inhibited mice remapped partially,

retaining more stable maps post-extinction. We additionally identified shock-tuned somata, to our knowledge the first identification of shock-tuned activity *in vivo* in dCA1. We re-affirmed the existence of freeze-tuned somata post-CFC, with NR-dCA1 inhibited generalizing mice preferentially re-activating freeze-tuned somata in the shocked context on retrieval day 1. Finally, we developed a novel method for identifying synchronized co-activations, i.e. groups of somata activating simultaneously, and demonstrated that these co-activations are modified by shocks to increase in size and frequency specifically in the shocked context on retrieval day 1. NR-dCA1 inhibition decreased the frequency of co-activations on retrieval day 1. On retrieval day 2, post-inhibition co-activation size increased in a 'rebound' effect. We showed that the composition of these co-activations was significantly overrepresented by somata with tunings to contextual discrimination, shock, freezing, and place both pre- and post-CFC. These results represent a meaningful contribution to our understanding of how dorsal CA1 encodes contextual fear and contextual fear extinction over time. These results additionally demonstrate the myriad ways in which the NR-dCA1 pathway impacts population dynamics throughout CFC. Furthermore, in the process we developed multiple new analysis techniques for future investigations of somatic dynamics beyond place maps.

In Chapter 4, we created a novel technique for imaging basal and distal apical tuft dendrites of excitatory pyramidal somata *in vivo* in dCA1, to our knowledge the first *in-vivo* measurement of this dendritic population. We demonstrated that basal dendritic place fields are more common, more reliable, and emerge later than apical tuft dendritic place fields, which we also showed are more differentially responsive to novelty and reward, are more biased toward rewarded locations, and more readily remap their place fields after reward removal. These results demonstrate that basal dendrites have more 'CA3-like' activity, potentially reflecting their

excitatory inputs via the Schaffer collateral from CA3. Meanwhile, apical tuft dendrites are fundamentally less spatially modulated, and more responsive to surprising or aversive experiences, like change, novelty, and reward removal. This potentially reflects synthesis of the less spatially-tuned excitatory inputs from lateral entorhinal cortex, as well as thalamic nucleus reuniens. Overall, this chapter offers novel insights into the role dendritic segmentation *in vivo* during reward and novelty manipulation, with basal dendrites demonstrating strong spatial-tuning and apical dendrites demonstrating reward preferences. Considering that the NR pathway synapses directly on dCA1 apical tuft dendrites, these experiments provide the foundation for future work exploring their role during contextual fear conditioning.

Remaining Questions

Despite the significant progress this dissertation represents, many questions remain. While a fuller exploration of future directions can be found in individual chapter discussions, here I highlight a few of the most pressing questions remaining across chapters as guidance for future investigations.

One crucial question is the dendritic mechanisms underlying VR-CFC in dorsal CA1. In this dissertation, I demonstrated differential processing between the basal and dendritic tuft, with significantly higher non-spatial and saliency tuning represented in the apical tuft. An investigation of the dendritic arbor throughout VR-CFC would help disentangle which aspects of somatic encoding are from inputs unique to the CA1 apical tuft, and which are inherited from CA3. A limitation of the present work to note is its primary basis in calcium imaging. Incorporating a wider suite of methodological approaches such as sodium-based voltage

imaging, *in vivo* glutamate detector imaging, or simultaneous local field potential recordings could imbue more nuance into NR-induced changes to excitatory dendrites.

Another important consideration that is left unexplored in this dissertation is the role of inhibitory interneurons in the hippocampal circuit. The most foundational missing experiments are *in vitro* anatomic and slice-based explorations to determine which inhibitory dCA1 subpopulations are innervated by NR axons. The results of this *in vitro* work would pinpoint relevant *in vivo* inhibitory interneuron populations for subsequent calcium imaging during VR-CFC and NR-dCA1 manipulation. Ideally, recording from and manipulating these interneuron populations will help reveal the mechanisms underlying how NR regulates dCA1 activity throughout VR-CFC.

On a population scale, incorporating a broader suite of techniques could be of use to detangle the role of synchronized co-activations (COAs) and related NR axonal activity in contextual fear conditioning. To investigate their relationship to hippocampal replay, the experiments outlined in chapters 2 and 3 would need to be repeated with simultaneous calcium imaging and electrophysiology. Relatedly, a better understanding of the relationship between COA activity and engrams would require simultaneous immediate early gene tagging/monitoring and calcium imaging. Lastly, our experiments were limited to a small suite of behaviors in a head-fixed preparation in exclusively male mice. Increasing the suite of behaviors, adding new behavioral recordings, expanding the NR-dCA1 pathway manipulation into activation and inactivation, and modifying the parameters of contextual fear conditioning all are important future directions. Ultimately, much work remains to fully understand role of the NR-dCA1 circuit in encoding, recalling, and adaptively modifying aversive experiences.

REFERENCES

1. Maren, S., Phan, K. L. & Liberzon, I. The contextual brain: implications for fear conditioning, extinction and psychopathology. *Nat. Rev. Neurosci.* **14**, 417–428 (2013).
2. Bouton, M. E. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychol. Bull.* **114**, 80–99 (1993).
3. Barrett, L. F. & Kensinger, E. A. Context is routinely encoded during emotion perception. *Psychol. Sci.* **21**, 595–599 (2010).
4. Stickgold, R. & Walker, M. P. Sleep-dependent memory triage: evolving generalization through selective processing. *Nat. Neurosci.* **16**, 139–145 (2013).
5. Cooper, R. A., Kensinger, E. A. & Ritchey, M. Memories Fade: The Relationship Between Memory Vividness and Remembered Visual Salience. *Psychol. Sci.* **30**, 657–668 (2019).
6. Fanselow, M. S. & Wassum, K. M. The Origins and Organization of Vertebrate Pavlovian Conditioning. *Cold Spring Harb. Perspect. Biol.* **8**, a021717 (2015).
7. Pavlov, P. I. Conditioned reflexes: An investigation of the physiological activity of the cerebral cortex. (1927).
8. Watson, J. B. & Rayner, R. Conditioned emotional reactions. *J. Exp. Psychol.* **3**, 1–14 (1920).
9. Mobbs, D. *et al.* Viewpoints: Approaches to defining and investigating fear. *Nat. Neurosci.* **22**, 1205–1216 (2019).
10. Adolphs, R. The biology of fear. *Curr. Biol.* **23**, R79–93 (2013).
11. Blanchard, R. J. & Blanchard, D. C. Crouching as an index of fear. *J. Comp. Physiol. Psychol.* **67**, 370–375 (1969).
12. Gibbbon, J. & Basalm, P. Spreading association in time. in *Autoshaping and Conditioning Theory* (ed. Herbert S. Terrace, J. G.) 219–253 (Academic Press, Inc, 1981).
13. Fanselow, M. S. Conditioned and unconditional components of post-shock freezing. *Pavlov. J. Biol. Sci.* **15**, 177–182 (1980).
14. Goosens, K. A. & Maren, S. Contextual and auditory fear conditioning are mediated by the lateral,

basal, and central amygdaloid nuclei in rats. *Learn. Mem.* **8**, 148–155 (2001).

15. Konorski, J. & Miller, S. On two types of conditioned reflex. *J. Gen. Psychol.* **16**, 264–272 (1937).
16. Konorski, J. *Conditioned Reflexes and Neuron Organization*. (University Press, 1948).
17. Konorski, J. Integrative activity of the brain. (1967).
18. Quirk, G. J. & Mueller, D. Neural mechanisms of extinction learning and retrieval. *Neuropsychopharmacology* **33**, 56–72 (2008).
19. Hull, C. L. *Principles of Behavior: An Introduction to Behavior Theory*. (D. Appleton-Century Company, incorporated, 1943).
20. Guttman, N. & Kalish, H. I. Discriminability and stimulus generalization. *J. Exp. Psychol.* **51**, 79–88 (1956).
21. Rohrbaugh, M. & Riccio, D. C. Stimulus generalization of learned fear in infant and adult rats. *J. Comp. Physiol. Psychol.* **66**, 530–533 (1968).
22. Beckers, T. *et al.* Understanding clinical fear and anxiety through the lens of human fear conditioning. *Nat Rev Psychol* **2**, 233–245 (2023).
23. Scientific Image and Illustration Software. https://www.biorender.com/.
24. Kimble, D. P. Hippocampus and internal inhibition. *Psychol. Bull.* **70**, 285–295 (1968).
25. LeDoux, J. E. Emotion circuits in the brain. *Annu. Rev. Neurosci.* **23**, 155–184 (2000).
26. Kim, J. J. & Fanselow, M. S. Modality-specific retrograde amnesia of fear. *Science* **256**, 675–677 (1992).
27. Maren, S. & Fanselow, M. S. Synaptic plasticity in the basolateral amygdala induced by hippocampal formation stimulation in vivo. *J. Neurosci.* **15**, 7548–7564 (1995).
28. Maren, S. & Fanselow, M. S. Electrolytic lesions of the fimbria/fornix, dorsal hippocampus, or entorhinal cortex produce anterograde deficits in contextual fear conditioning in rats. *Neurobiol. Learn. Mem.* **67**, 142–149 (1997).
29. Burwell, R. D., Bucci, D. J., Sanborn, M. R. & Jutras, M. J. Perirhinal and Postrhinal Contributions to Remote Memory for Context. *J. Neurosci.* **24**, 11023–11028 (2004).

30. Frankland, P. W., Bontempi, B., Talton, L. E., Kaczmarek, L. & Silva, A. J. The involvement of the anterior cingulate cortex in remote contextual fear memory. *Science* **304**, 881–883 (2004).
31. Beyeler, A. & Dabrowska, J. Neuronal diversity of the amygdala and the bed nucleus of the stria terminalis. *Handb Behav Neurosci* **26**, 63–100 (2020).
32. Klüver, H. & Bucy, P. C. 'Psychic blindness' and other symptoms following bilateral temporal lobectomy in Rhesus monkeys. *Am. J. Physiol.* **119**, 352–353 (1937).
33. Price, H., LaROSA, W. V. & Settle, E. B. Lipoidosis cutis et mucosae; report of a case. *Arch. Derm. Syphilol.* **55**, 42–51 (1947).
34. Weiskrantz, L. Behavioral changes associated with ablation of the amygdaloid complex in monkeys. *J. Comp. Physiol. Psychol.* **49**, 381–391 (1956).
35. Adolphs, R., Tranel, D., Damasio, H. & Damasio, A. Impaired recognition of emotion in facial expressions following bilateral damage to the human amygdala. *Nature* **372**, 669–672 (1994).
36. Adolphs, R., Tranel, D., Damasio, H. & Damasio, A. R. Fear and the human amygdala. *J. Neurosci.* **15**, 5879–5891 (1995).
37. Baxter, M. G. & Murray, E. A. The amygdala and reward. *Nat. Rev. Neurosci.* **3**, 563–573 (2002).
38. Kochli, D. E., Thompson, E. C., Fricke, E. A., Postle, A. F. & Quinn, J. J. The amygdala is critical for trace, delay, and contextual fear conditioning. *Learn. Mem.* **22**, 92–100 (2015).
39. Romanski, L. M. & LeDoux, J. E. Information cascade from primary auditory cortex to the amygdala: corticocortical and corticoamygdaloid projections of temporal cortex in the rat. *Cereb. Cortex* **3**, 515–532 (1993).
40. Hestrin, S. *et al.* Physiological properties of excitatory synaptic transmission in the central nervous system. *Cold Spring Harb. Symp. Quant. Biol.* **55**, 87–93 (1990).
41. Mahanty, N. K. & Sah, P. Excitatory synaptic inputs to pyramidal neurons of the lateral amygdala. *Eur. J. Neurosci.* **11**, 1217–1222 (1999).
42. Sun, Y., Gooch, H. & Sah, P. Fear conditioning and the basolateral amygdala. *F1000Res.* **9**, (2020).
43. Quirk, G. J., Repa, C. & LeDoux, J. E. Fear conditioning enhances short-latency auditory responses

of lateral amygdala neurons: parallel recordings in the freely behaving rat. *Neuron* **15**, 1029–1039 (1995).

44. Maren, S. & Holt, W. The hippocampus and contextual memory retrieval in Pavlovian conditioning. *Behav. Brain Res.* **110**, 97–108 (2000).
45. Onishi, B. K. A. & Xavier, G. F. Contextual, but not auditory, fear conditioning is disrupted by neurotoxic selective lesion of the basal nucleus of amygdala in rats. *Neurobiol. Learn. Mem.* **93**, 165–174 (2010).
46. Keifer, O. P., Jr, Hurt, R. C., Ressler, K. J. & Marvar, P. J. The Physiology of Fear: Reconceptualizing the Role of the Central Amygdala in Fear Learning. *Physiology* **30**, 389–401 (2015).
47. Dolleman-van der Weel, M. J. *et al.* The nucleus reuniens of the thalamus sits at the nexus of a hippocampus and medial prefrontal cortex circuit enabling memory and behavior. *Learn. Mem.* **26**, 191–205 (2019).
48. Lee, J. H., Lee, S. & Kim, J.-H. Amygdala Circuits for Fear Memory: A Key Role for Dopamine Regulation. *Neuroscientist* **23**, 542–553 (2017).
49. Anastasiades, P. G. & Carter, A. G. Circuit organization of the rodent medial prefrontal cortex. *Trends Neurosci.* **44**, 550–563 (2021).
50. Laubach, M., Amarante, L. M., Swanson, K. & White, S. R. What, If Anything, Is Rodent Prefrontal Cortex? *eNeuro* **5**, (2018).
51. Preuss, T. M. & Wise, S. P. Evolution of prefrontal cortex. *Neuropsychopharmacology* **47**, 3–19 (2022).
52. Kesner, R. P. & Churchwell, J. C. An analysis of rat prefrontal cortex in mediating executive function. *Neurobiol. Learn. Mem.* **96**, 417–431 (2011).
53. Le Merre, P., Ährlund-Richter, S. & Carlén, M. The mouse prefrontal cortex: Unity in diversity. *Neuron* **109**, 1925–1944 (2021).
54. Wouterlood, F. G., Saldana, E. & Witter, M. P. Projection from the nucleus reuniens thalami to the

hippocampal region: light and electron microscopic tracing study in the rat with the anterograde tracer Phaseolus vulgaris-leucoagglutinin. *J. Comp. Neurol.* **296**, 179–203 (1990).

55. Guillery, R. W. Anatomical evidence concerning the role of the thalamus in corticocortical communication: a brief review. *J. Anat.* **187 (Pt 3)**, 583–592 (1995).
56. Barbas, H. & Blatt, G. J. Topographically specific hippocampal projections target functionally distinct prefrontal areas in the rhesus monkey. *Hippocampus* **5**, 511–533 (1995).
57. Condé, F., Maire-Lepoivre, E., Audinat, E. & Crépel, F. Afferent connections of the medial frontal cortex of the rat. II. Cortical and subcortical afferents. *J. Comp. Neurol.* **352**, 567–593 (1995).
58. Laviolette, S. R., Lipski, W. J. & Grace, A. A. A subpopulation of neurons in the medial prefrontal cortex encodes emotional learning with burst and frequency codes through a dopamine D4 receptor-dependent basolateral amygdala input. *J. Neurosci.* **25**, 6066–6075 (2005).
59. Vertes, R. P. Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience* **142**, 1–20 (2006).
60. Hoover, W. B. & Vertes, R. P. Anatomical analysis of afferent projections to the medial prefrontal cortex in the rat. *Brain Struct. Funct.* **212**, 149–179 (2007).
61. Orsini, C. A., Kim, J. H., Knapska, E. & Maren, S. Hippocampal and prefrontal projections to the basal amygdala mediate contextual regulation of fear after extinction. *J. Neurosci.* **31**, 17269–17277 (2011).
62. Scheel, N., Wulff, P. & de Mooij-van Malsen, J. G. Afferent connections of the thalamic nucleus reuniens in the mouse. *J. Comp. Neurol.* **528**, 1189–1202 (2020).
63. Ji, G. & Neugebauer, V. Modulation of medial prefrontal cortical activity using in vivo recordings and optogenetics. *Mol. Brain* **5**, 36 (2012).
64. Little, J. P. & Carter, A. G. Subcellular synaptic connectivity of layer 2 pyramidal neurons in the medial prefrontal cortex. *J. Neurosci.* **32**, 12808–12819 (2012).
65. Little, J. P. & Carter, A. G. Synaptic mechanisms underlying strong reciprocal connectivity between the medial prefrontal cortex and basolateral amygdala. *J. Neurosci.* **33**, 15333–15342 (2013).

66. Holson, R. R. Mesial prefrontal cortical lesions and timidity in rats. I. Reactivity to aversive stimuli. *Physiol. Behav.* **37**, 221–230 (1986).

67. Morgan, M. A., Romanski, L. M. & LeDoux, J. E. Extinction of emotional learning: contribution of medial prefrontal cortex. *Neurosci. Lett.* **163**, 109–113 (1993).

68. Morgan, M. A. & LeDoux, J. E. Differential contribution of dorsal and ventral medial prefrontal cortex to the acquisition and extinction of conditioned fear in rats. *Behav. Neurosci.* **109**, 681–688 (1995).

69. Vidal-Gonzalez, I., Vidal-Gonzalez, B., Rauch, S. L. & Quirk, G. J. Microstimulation reveals opposing influences of prelimbic and infralimbic cortex on the expression of conditioned fear. *Learn. Mem.* **13**, 728–733 (2006).

70. Burgos-Robles, A., Vidal-Gonzalez, I. & Quirk, G. J. Sustained conditioned responses in prelimbic prefrontal neurons are correlated with fear expression and extinction failure. *J. Neurosci.* **29**, 8474–8482 (2009).

71. Giustino, T. F. & Maren, S. The Role of the Medial Prefrontal Cortex in the Conditioning and Extinction of Fear. *Front. Behav. Neurosci.* **9**, 298 (2015).

72. Kim, H.-S., Cho, H.-Y., Augustine, G. J. & Han, J.-H. Selective Control of Fear Expression by Optogenetic Manipulation of Infralimbic Cortex after Extinction. *Neuropsychopharmacology* **41**, 1261–1273 (2016).

73. Quirk, G. J., Russo, G. K., Barron, J. L. & Lebron, K. The role of ventromedial prefrontal cortex in the recovery of extinguished fear. *J. Neurosci.* **20**, 6225–6231 (2000).

74. Sierra-Mercado, D., Padilla-Coreano, N. & Quirk, G. J. Dissociable roles of prelimbic and infralimbic cortices, ventral hippocampus, and basolateral amygdala in the expression and extinction of conditioned fear. *Neuropsychopharmacology* **36**, 529–538 (2011).

75. Awad, W., Ferreira, G. & Maroun, M. Dissociation of the Role of Infralimbic Cortex in Learning and Consolidation of Extinction of Recent and Remote Aversion Memory. *Neuropsychopharmacology* **40**, 2566–2575 (2015).

76. Do-Monte, F. H., Manzano-Nieves, G., Quiñones-Laracuente, K., Ramos-Medina, L. & Quirk, G. J. Revisiting the role of infralimbic cortex in fear extinction with optogenetics. *J. Neurosci.* **35**, 3607–3615 (2015).

77. Kim, E. J. *et al.* Alterations of hippocampal place cells in foraging rats facing a 'predatory' threat. *Curr. Biol.* **25**, 1362–1367 (2015).

78. Bukalo, O. *et al.* Prefrontal inputs to the amygdala instruct fear extinction memory formation. *Sci Adv* **1**, (2015).

79. Milad, M. R. & Quirk, G. J. Neurons in medial prefrontal cortex signal memory for fear extinction. *Nature* **420**, 70–74 (2002).

80. Burgos-Robles, A., Vidal-Gonzalez, I., Santini, E. & Quirk, G. J. Consolidation of fear extinction requires NMDA receptor-dependent bursting in the ventromedial prefrontal cortex. *Neuron* **53**, 871–880 (2007).

81. Giustino, T. F., Fitzgerald, P. J. & Maren, S. Fear Expression Suppresses Medial Prefrontal Cortical Firing in Rats. *PLoS One* **11**, e0165256 (2016).

82. Eichenbaum, H. Prefrontal-hippocampal interactions in episodic memory. *Nat. Rev. Neurosci.* **18**, 547–558 (2017).

83. Mitchell, A. S. *et al.* Advances in understanding mechanisms of thalamic relays in cognition and behavior. *J. Neurosci.* **34**, 15340–15346 (2014).

84. Vertes, R. P. Major diencephalic inputs to the hippocampus: supramammillary nucleus and nucleus reuniens. Circuitry and function. *Prog. Brain Res.* **219**, 121–144 (2015).

85. Sherman, S. M. The thalamus is more than just a relay. *Curr. Opin. Neurobiol.* **17**, 417–422 (2007).

86. Halassa, M. M. & Sherman, S. M. Thalamocortical Circuit Motifs: A General Framework. *Neuron* **103**, 762–770 (2019).

87. Herkenham, M. The connections of the nucleus reuniens thalami: evidence for a direct thalamo-hippocampal pathway in the rat. *J. Comp. Neurol.* **177**, 589–610 (1978).

88. Bokor, H., Csáki, A., Kocsis, K. & Kiss, J. Cellular architecture of the nucleus reuniens thalami and

its putative aspartatergic/glutamatergic projection to the hippocampus and medial septum in the rat. *Eur. J. Neurosci.* **16**, 1227–1239 (2002).

89. Frassoni, C. *et al.* Postnatal development of calbindin and parvalbumin immunoreactivity in the thalamus of the rat. *Brain Res. Dev. Brain Res.* **58**, 243–249 (1991).
90. Drexel, M., Preidt, A. P., Kirchmair, E. & Sperk, G. Parvalbumin interneurons and calretinin fibers arising from the thalamic nucleus reuniens degenerate in the subiculum after kainic acid-induced seizures. *Neuroscience* **189**, 316–329 (2011).
91. Jones, E. G. A new view of specific and nonspecific thalamocortical connections. *Adv. Neurol.* **77**, 49–71; discussion 72–3 (1998).
92. Witter, M. P., Groenewegen, H. J., Lopes da Silva, F. H. & Lohman, A. H. Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. *Prog. Neurobiol.* **33**, 161–253 (1989).
93. Su, H. S. & Bentivoglio, M. Thalamic midline cell populations projecting to the nucleus accumbens, amygdala, and hippocampus in the rat. *J. Comp. Neurol.* **297**, 582–593 (1990).
94. Bertram, E. H. & Zhang, D. X. Thalamic excitation of hippocampal CA1 neurons: a comparison with the effects of CA3 stimulation. *Neuroscience* **92**, 15–26 (1999).
95. Dolleman-Van der Weel, M. J. & Witter, M. P. Nucleus reuniens thalami innervates gamma aminobutyric acid positive cells in hippocampal field CA1 of the rat. *Neurosci. Lett.* **278**, 145–148 (2000).
96. Morales, G. J., Ramcharan, E. J., Sundararaman, N., Morgera, S. D. & Vertes, R. P. Analysis of the actions of nucleus reuniens and the entorhinal cortex on EEG and evoked population behavior of the hippocampus. *Conf. Proc. IEEE Eng. Med. Biol. Soc.* **2007**, 2480–2484 (2007).
97. Xu, W. & Südhof, T. C. A neural circuit for memory specificity and generalization. *Science* **339**, 1290–1295 (2013).
98. Sun, Y. *et al.* Cell-type-specific circuit connectivity of hippocampal CA1 revealed through Cre-dependent rabies tracing. *Cell Rep.* **7**, 269–280 (2014).

99. Dolleman-van der Weel, M. J., Lopes da Silva, F. H. & Witter, M. P. Interaction of nucleus reuniens and entorhinal cortex projections in hippocampal field CA1 of the rat. *Brain Struct. Funct.* **222**, 2421–2438 (2017).

100. Goswamee, P., Leggett, E. & McQuiston, A. R. Nucleus Reuniens Afferents in Hippocampus Modulate CA1 Network Function via Monosynaptic Excitation and Polysynaptic Inhibition. *Front. Cell. Neurosci.* **15**, 660897 (2021).

101. Andrianova, L. *et al.* Hippocampal CA1 pyramidal cells do not receive monosynaptic input from thalamic nucleus reuniens. *bioRxiv* 2021.09.30.462517 (2021) doi:10.1101/2021.09.30.462517.

102. Vertes, R. P. Analysis of projections from the medial prefrontal cortex to the thalamus in the rat, with emphasis on nucleus reuniens. *J. Comp. Neurol.* **442**, 163–187 (2002).

103. Hoover, W. B. & Vertes, R. P. Collateral projections from nucleus reuniens of thalamus to hippocampus and medial prefrontal cortex in the rat: a single and double retrograde fluorescent labeling study. *Brain Struct. Funct.* **217**, 191–209 (2012).

104. Varela, C., Kumar, S., Yang, J. Y. & Wilson, M. A. Anatomical substrates for direct interactions between hippocampus, medial prefrontal cortex, and the thalamic nucleus reuniens. *Brain Struct. Funct.* **219**, 911–929 (2014).

105. Ohtake, T. & Yamada, H. Efferent connections of the nucleus reuniens and the rhomboid nucleus in the rat: an anterograde PHA-L tracing study. *Neurosci. Res.* **6**, 556–568 (1989).

106. Vertes, R. P., Hoover, W. B., Do Valle, A. C., Sherman, A. & Rodriguez, J. J. Efferent projections of reuniens and rhomboid nuclei of the thalamus in the rat. *J. Comp. Neurol.* **499**, 768–796 (2006).

107. McKenna, J. T. & Vertes, R. P. Afferent projections to nucleus reuniens of the thalamus. *J. Comp. Neurol.* **480**, 115–142 (2004).

108. Vertes, R. P., Linley, S. B. & Hoover, W. B. Limbic circuitry of the midline thalamus. *Neurosci. Biobehav. Rev.* **54**, 89–107 (2015).

109. Griffin, A. L. Role of the thalamic nucleus reuniens in mediating interactions between the hippocampus and medial prefrontal cortex during spatial working memory. *Front. Syst. Neurosci.* **9**,

29 (2015).

110. Cassel, J.-C. *et al.* The reuniens and rhomboid nuclei of the thalamus: A crossroads for cognition-relevant information processing? *Neurosci. Biobehav. Rev.* **126**, 338–360 (2021).

111. Ito, H. T., Zhang, S.-J., Witter, M. P., Moser, E. I. & Moser, M.-B. A prefrontal-thalamo-hippocampal circuit for goal-directed spatial navigation. *Nature* **522**, 50–55 (2015).

112. Hauer, B. E., Pagliardini, S. & Dickson, C. T. The Reuniens Nucleus of the Thalamus has an Essential Role in Coordinating Slow Wave Activity between Neocortex and Hippocampus. *eNeuro* (2019) doi:10.1523/ENEURO.0365-19.2019.

113. Ferraris, M. *et al.* The Nucleus Reuniens Controls Long-Range Hippocampo-Prefrontal Gamma Synchronization during Slow Oscillations. *J. Neurosci.* **38**, 3026–3038 (2018).

114. Jayachandran, M. *et al.* Nucleus reuniens transiently synchronizes memory networks at beta frequencies. *Nat. Commun.* **14**, 1–15 (2023).

115. Vetere, G. *et al.* Chemogenetic Interrogation of a Brain-wide Fear Memory Network in Mice. *Neuron* **94**, 363–374.e4 (2017).

116. Ramanathan, K. R., Jin, J., Giustino, T. F., Payne, M. R. & Maren, S. Prefrontal projections to the thalamic nucleus reuniens mediate fear extinction. *Nat. Commun.* **9**, 4527 (2018).

117. Ramanathan, K. R., Ressler, R. L., Jin, J. & Maren, S. Nucleus Reuniens Is Required for Encoding and Retrieving Precise, Hippocampal-Dependent Contextual Fear Memories in Rats. *J. Neurosci.* **38**, 9925–9933 (2018).

118. Ramanathan, K. R. & Maren, S. Nucleus reuniens mediates the extinction of contextual fear conditioning. *Behav. Brain Res.* **374**, 112114 (2019).

119. Ressler, R. L., Goode, T. D., Kim, S., Ramanathan, K. R. & Maren, S. Covert capture and attenuation of a hippocampus-dependent fear memory. *Nat. Neurosci.* 1–8 (2021).

120. Silva, B. A. *et al.* A thalamo-amygdalar circuit underlying the extinction of remote fear memories. *Nat. Neurosci.* **24**, 964–974 (2021).

121. Vasudevan, K., Ramanathan, K. R., Vierkant, V. & Maren, S. Nucleus reuniens inactivation does not impair consolidation or reconsolidation of fear extinction. *Learn. Mem.* **29**, 216–222 (2022).

122. Braak, H. & Braak, E. Neuropathological stageing of Alzheimer-related changes. *Acta Neuropathol.* **82**, 239–259 (1991).

123. Visser, C. P., Coene, L. N., Brand, R. & Tavy, D. L. The incidence of nerve injury in anterior dislocation of the shoulder and its influence on functional recovery. A prospective clinical and EMG study. *J. Bone Joint Surg. Br.* **81**, 679–685 (1999).

124. Lawrie, S. M. *et al.* Reduced frontotemporal functional connectivity in schizophrenia associated with auditory hallucinations. *Biol. Psychiatry* **51**, 1008–1011 (2002).

125. Meyer-Lindenberg, A. *et al.* Evidence for abnormal cortical functional connectivity during working memory in schizophrenia. *Am. J. Psychiatry* **158**, 1809–1817 (2001).

126. Venkataraman, A., Whitford, T. J., Westin, C.-F., Golland, P. & Kubicki, M. Whole brain resting state functional connectivity abnormalities in schizophrenia. *Schizophr. Res.* **139**, 7–12 (2012).

127. Hardenacke, K. *et al.* Stimulate or degenerate: deep brain stimulation of the nucleus basalis Meynert in Alzheimer dementia. *World Neurosurg.* **80**, S27.e35–43 (2013).

128. Argyelan, M. *et al.* Resting-state fMRI connectivity impairment in schizophrenia and bipolar disorder. *Schizophr. Bull.* **40**, 100–110 (2014).

129. Kafetzopoulos, V. *et al.* The nucleus reuniens: a key node in the neurocircuitry of stress and depression. *Mol. Psychiatry* **23**, 579–586 (2018).

130. Reeders, P. C., Rivera Núñez, M. V., Vertes, R. P., Mattfeld, A. T. & Allen, T. A. Identifying the midline thalamus in humans in vivo. *Brain Struct. Funct.* (2023) doi:10.1007/s00429-022-02607-6.

131. Interactive Atlas Viewer :: Atlas Viewer. http://atlas.brain-map.org/atlas?atlas=1.

132. Cajal, R. Y. *Histologie du Systeme Nerveux de L'Homme et des Vertebres*. vol. 2 (1911).

133. Allen, T. A. & Fortin, N. J. The evolution of episodic memory. *Proc. Natl. Acad. Sci. U. S. A.* **110 Suppl 2**, 10379–10386 (2013).

134. van Strien, N. M., Cappaert, N. L. M. & Witter, M. P. The anatomy of memory: an interactive

overview of the parahippocampal-hippocampal network. *Nat. Rev. Neurosci.* **10**, 272–282 (2009).

135. Dudek, S. M., Alexander, G. M. & Farris, S. Rediscovering area CA2: unique properties and functions. *Nat. Rev. Neurosci.* **17**, 89–102 (2016).

136. Fanselow, M. S. & Dong, H.-W. Are the dorsal and ventral hippocampus functionally distinct structures? *Neuron* **65**, 7–19 (2010).

137. Keller, D., Erö, C. & Markram, H. Cell Densities in the Mouse Brain: A Systematic Review. *Front. Neuroanat.* **12**, 83 (2018).

138. Scoville, W. B. & Milner, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).

139. O'Keefe, J. & Conway, D. H. Hippocampal place units in the freely moving rat: why they fire where they fire. *Exp. Brain Res.* **31**, 573–590 (1978).

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

141. The Nobel Prize in Physiology or Medicine 2014. *NobelPrize.org* https://www.nobelprize.org/prizes/medicine/2014/summary/.

142. Tolman, E. C. Cognitive maps in rats and men. *Psychol. Rev.* **55**, 189–208 (1948).

143. O'Keefe, J. & Nadel, L. The Hippocampus as a Cognitive Map. 570 (1978).

144. Moser, E. I. *et al.* Grid cells and cortical representation. *Nat. Rev. Neurosci.* **15**, 466–481 (2014).

145. Ziv, Y. *et al.* Long-term dynamics of CA1 hippocampal place codes. *Nat. Neurosci.* **16**, 264–266 (2013).

146. Wilson, M. A. & McNaughton, B. L. Dynamics of the hippocampal ensemble code for space. *Science* **261**, 1055–1058 (1993).

147. Sheffield, M. E. J., Adoff, M. D. & Dombeck, D. A. Increased Prevalence of Calcium Transients across the Dendritic Arbor during Place Field Formation. *Neuron* **96**, 490–504.e5 (2017).

148. Courellis, H. S. *et al.* Spatial encoding in primate hippocampus during free navigation. *PLoS Biol.* **17**, e3000546 (2019).

149. Sheintuch, L. *et al.* Multiple Maps of the Same Spatial Context Can Stably Coexist in the Mouse Hippocampus. *Curr. Biol.* (2020) doi:10.1016/j.cub.2020.02.018.

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

151. Dudchenko, P. A., Wood, E. R. & Eichenbaum, H. Non-Spatial Correlates of Hippocampal Activity. in *The Neural Basis of Navigation: Evidence from Single Cell Recording* (ed. Sharp, P. E.) 81–96 (Springer US, 2002).

152. Squire, L. R. Memory systems of the brain: a brief history and current perspective. *Neurobiol. Learn. Mem.* **82**, 171–177 (2004).

153. Eichenbaum, H. Time cells in the hippocampus: a new dimension for mapping memories. *Nat. Rev. Neurosci.* **15**, 732–744 (2014).

154. Tavares, R. M. *et al.* A Map for Social Navigation in the Human Brain. *Neuron* **87**, 231–243 (2015).

155. Aronov, D., Nevers, R. & Tank, D. W. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature* **543**, 719 (2017).

156. Gauthier, J. L. & Tank, D. W. A dedicated population for reward coding in the hippocampus. *Neuron* **99**, 179–193.e7 (2018).

157. Eichenbaum, H. What Versus Where: Non-spatial Aspects of Memory Representation by the Hippocampus. *Curr. Top. Behav. Neurosci.* **37**, 101–117 (2018).

158. Knudsen, E. B. & Wallis, J. D. Hippocampal neurons construct a map of an abstract value space. *Cell* **184**, 4640–4650.e10 (2021).

159. Blair, G. J. *et al.* Hippocampal place cell remapping occurs with memory storage of aversive experiences. *Elife* **12**, (2023).

160. Ólafsdóttir, H. F., Bush, D. & Barry, C. The Role of Hippocampal Replay in Memory and Planning. *Curr. Biol.* **28**, R37–R50 (2018).

161. Pavlides, C. & Winson, J. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* **9**, 2907–2918 (1989).

162. Wilson, M. A. & McNaughton, B. L. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679 (1994).

163. Foster, D. J. & Wilson, M. A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature* **440**, 680–683 (2006).

164. Huelin Gorriz, M., Takigawa, M. & Bendor, D. The Role of Experience in Prioritizing Hippocampal Replay. *bioRxiv* 2023.03.28.534589 (2023) doi:10.1101/2023.03.28.534589.

165. Davidson, T. J., Kloosterman, F. & Wilson, M. A. Hippocampal replay of extended experience. *Neuron* **63**, 497–507 (2009).

166. Dupret, D., O'Neill, J., Pleydell-Bouverie, B. & Csicsvari, J. The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nat. Neurosci.* **13**, 995–1002 (2010).

167. Pfeiffer, B. E. & Foster, D. J. PLACE CELLS. Autoassociative dynamics in the generation of sequences of hippocampal place cells. *Science* **349**, 180–183 (2015).

168. Squire, L. R., Genzel, L., Wixted, J. T. & Morris, R. G. Memory consolidation. *Cold Spring Harb. Perspect. Biol.* **7**, a021766 (2015).

169. Pfeiffer, B. E. The content of hippocampal 'replay'. *Hippocampus* **30**, 6–18 (2020).

170. Strange, B. A., Witter, M. P., Lein, E. S. & Moser, E. I. Functional organization of the hippocampal longitudinal axis. *Nat. Rev. Neurosci.* **15**, 655–669 (2014).

171. Harvey, C. D., Collman, F., Dombeck, D. A. & Tank, D. W. Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature* **461**, 941–946 (2009).

172. Phillips, R. G. & LeDoux, J. E. Differential contribution of amygdala and hippocampus to cued and contextual fear conditioning. *Behav. Neurosci.* **106**, 274–285 (1992).

173. Maren, S., Aharonov, G. & Fanselow, M. S. Neurotoxic lesions of the dorsal hippocampus and Pavlovian fear conditioning in rats. *Behav. Brain Res.* **88**, 261–274 (1997).

174. Corcoran, K. A. & Maren, S. Hippocampal inactivation disrupts contextual retrieval of fear memory after extinction. *J. Neurosci.* **21**, 1720–1726 (2001).

175. Halonen, J. D., Zoladz, P. R., Park, C. R. & Diamond, D. M. Behavioral and neurobiological

assessments of predator-based fear conditioning and extinction. *J. Behav. Brain Sci.* **06**, 337–356 (2016).

176. Anagnostaras, S. G., Gale, G. D. & Fanselow, M. S. Hippocampus and contextual fear conditioning: Recent controversies and advances. *Hippocampus* **11**, 8–17 (2001).

177. Frankland, P. W., Cestari, V., Filipkowski, R. K., McDonald, R. J. & Silva, A. J. The dorsal hippocampus is essential for context discrimination but not for contextual conditioning. *Behav. Neurosci.* **112**, 863–874 (1998).

178. Stiedl, O., Misane, I., Spiess, J. & Ogren, S. O. Involvement of the 5-HT1A receptors in classical fear conditioning in C57BL/6J mice. *J. Neurosci.* **20**, 8515–8527 (2000).

179. Bast, T., Zhang, W.-N. & Feldon, J. Dorsal hippocampus and classical fear conditioning to tone and context in rats: effects of local NMDA-receptor blockade and stimulation. *Hippocampus* **13**, 657–675 (2003).

180. Quinn, J. J., Loya, F., Ma, Q. D. & Fanselow, M. S. Dorsal hippocampus NMDA receptors differentially mediate trace and contextual fear conditioning. *Hippocampus* **15**, 665–674 (2005).

181. Czerniawski, J., Ree, F., Chia, C. & Otto, T. Dorsal versus ventral hippocampal contributions to trace and contextual conditioning: differential effects of regionally selective NMDA receptor antagonism on acquisition and expression. *Hippocampus* **22**, 1528–1539 (2012).

182. Wiltgen, B. J., Sanders, M. J., Anagnostaras, S. G., Sage, J. R. & Fanselow, M. S. Context fear learning in the absence of the hippocampus. *J. Neurosci.* **26**, 5484–5491 (2006).

183. Gidyk, D. C., McDonald, R. J. & Sutherland, R. J. Intact behavioural expression of contextual fear, context discrimination, and object discrimination memories acquired in the absence of the hippocampus. *J. Neurosci.* (2020) doi:10.1523/JNEUROSCI.0546-20.2020.

184. Rudy, J. W. & O'Reilly, R. C. Contextual fear conditioning, conjunctive representations, pattern completion, and the hippocampus. *Behav. Neurosci.* **113**, 867–880 (1999).

185. Corcoran, K. A., Desmond, T. J., Frey, K. A. & Maren, S. Hippocampal Inactivation Disrupts the Acquisition and Contextual Encoding of Fear Extinction. *J. Neurosci.* **25**, 8978–8987 (2005).

186. Ji, J. & Maren, S. Electrolytic lesions of the dorsal hippocampus disrupt renewal of conditional fear after extinction. *Learn. Mem.* **12**, 270–276 (2005).

187. Ji, J. & Maren, S. Hippocampal involvement in contextual modulation of fear extinction. *Hippocampus* **17**, 749–758 (2007).

188. Ryan, T. J., Roy, D. S., Pignatelli, M., Arons, A. & Tonegawa, S. Memory. Engram cells retain memory under retrograde amnesia. *Science* **348**, 1007–1013 (2015).

189. Semon, R. W. & Simon, L. *The Mneme*. (George Allen & Unwin, 1921, 1927).

190. Lashley, K. S. *Brain mechanisms and intelligence: A quantitative study of injuries to the brain*. (University of Chicago Press, 1929).

191. Hebb, D. O. The organization of behavior; a neuropsychological theory. *https://psycnet.apa.org › recordhttps://psycnet.apa.org › record* **335**, (1949).

192. Bliss, T. V. & Lomo, T. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *J. Physiol.* **232**, 331–356 (1973).

193. Malenka, R. C. & Bear, M. F. LTP and LTD: an embarrassment of riches. *Neuron* **44**, 5–21 (2004).

194. Curran, T. & Morgan, J. I. Superinduction of c-fos by nerve growth factor in the presence of peripherally active benzodiazepines. *Science* **229**, 1265–1268 (1985).

195. Morgan, J. I., Cohen, D. R., Hempstead, J. L. & Curran, T. Mapping patterns of c-fos expression in the central nervous system after seizure. *Science* **237**, 192–197 (1987).

196. Bullitt, E. Expression of c-fos-like protein as a marker for neuronal activity following noxious stimulation in the rat. *J. Comp. Neurol.* **296**, 517–530 (1990).

197. Radulovic, J., Kammermeier, J. & Spiess, J. Relationship between Fos Production and Classical Fear Conditioning: Effects of Novelty, Latent Inhibition, and Unconditioned Stimulus Preexposure. *J. Neurosci.* **18**, 7452–7461 (1998).

198. Guzowski, J. F., McNaughton, B. L., Barnes, C. A. & Worley, P. F. Environment-specific expression of the immediate-early gene Arc in hippocampal neuronal ensembles. *Nat. Neurosci.* **2**, 1120–1124 (1999).

199. Hall, J., Thomas, K. L. & Everitt, B. J. Cellular imaging of zif268 expression in the hippocampus and amygdala during contextual and cued fear memory retrieval: selective activation of hippocampal CA1 neurons during the recall of contextual memories. *J. Neurosci.* **21**, 2186–2193 (2001).

200. Reijmers, L. G., Perkins, B. L., Matsuo, N. & Mayford, M. Localization of a stable neural correlate of associative memory. *Science* **317**, 1230–1233 (2007).

201. Reijmers, L. & Mayford, M. Genetic control of active neural circuits. *Front. Mol. Neurosci.* **2**, 27 (2009).

202. Guenthner, C. J., Miyamichi, K., Yang, H. H., Heller, H. C. & Luo, L. Permanent genetic access to transiently active neurons via TRAP: targeted recombination in active populations. *Neuron* **78**, 773–784 (2013).

203. Sørensen, A. T. *et al.* A robust activity marking system for exploring active neuronal ensembles. *Elife* **5**, e13918 (2016).

204. Liu, X. *et al.* Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature* **484**, 381–385 (2012).

205. Ramirez, S. *et al.* Creating a false memory in the hippocampus. *Science* **341**, 387–391 (2013).

206. Josselyn, S. A., Köhler, S. & Frankland, P. W. Finding the engram. *Nat. Rev. Neurosci.* **16**, 521–534 (2015).

207. Lacagnina, A. F. *et al.* Distinct hippocampal engrams control extinction and relapse of fear memory. *Nat. Neurosci.* **22**, 753–761 (2019).

208. Jung, J. H. *et al.* Examining the engram encoding specificity hypothesis in mice. *Neuron* (2023) doi:10.1016/j.neuron.2023.03.007.

209. Goode, T. D., Tanaka, K. Z., Sahay, A. & McHugh, T. J. An Integrated Index: Engrams, Place Cells, and Hippocampal Memory. *Neuron* **107**, 805–820 (2020).

210. Tanaka, K. Z. *et al.* The hippocampal engram maps experience but not place. *Science* **361**, 392–397 (2018).

211. Ghandour, K. *et al.* Orchestrated ensemble activities constitute a hippocampal memory engram. *Nat.*

Commun. **10**, 2637 (2019).

212. Liu, Y.-Z., Wang, Y., Shen, W. & Wang, Z. Enhancement of synchronized activity between hippocampal CA1 neurons during initial storage of associative fear memory. *J. Physiol.* **595**, 5327–5340 (2017).

213. Dong, C. HIPPOCAMPAL CA1 AND CA3 NETWORK DYNAMICS UNDERLYING EPISODIC MEMORIES. (University of Chicago, 2022).

214. Kinsky, N. R., Orlin, D. J., Ruesch, E. A., Diba, K. & Ramirez, S. Erasable Hippocampal Neural Signatures Predict Memory Discrimination. *bioRxiv* 2023.02.02.526824 (2023) doi:10.1101/2023.02.02.526824.

215. Marr, D. Simple memory: a theory for archicortex. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **262**, 23–81 (1971).

216. Leutgeb, J. K., Leutgeb, S., Moser, M.-B. & Moser, E. I. Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science* **315**, 961–966 (2007).

217. Kesner, R. P. & Rolls, E. T. A computational theory of hippocampal function, and tests of the theory: new developments. *Neurosci. Biobehav. Rev.* **48**, 92–147 (2015).

218. Knierim, J. J. & Neunuebel, J. P. Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiol. Learn. Mem.* **129**, 38–49 (2016).

219. GoodSmith, D. *et al.* Spatial Representations of Granule Cells and Mossy Cells of the Dentate Gyrus. *Neuron* **93**, 677–690.e5 (2017).

220. Rolls, E. T. & Mills, P. The Generation of Time in the Hippocampal Memory System. *Cell Rep.* **28**, 1649–1658.e6 (2019).

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

222. Geiller, T., Fattahi, M., Choi, J.-S. & Royer, S. Place cells are more strongly tied to landmarks in deep than in superficial CA1. *Nat. Commun.* **8**, 14531 (2017).

223. Dong, C., Madar, A. D. & Sheffield, M. E. J. Distinct place cell dynamics in CA1 and CA3 encode

experience in new environments. *Nat. Commun.* **12**, 1–13 (2021).

224. Hainmueller, T. & Bartos, M. Dentate gyrus circuits for encoding, retrieval and discrimination of episodic memories. *Nat. Rev. Neurosci.* **21**, 153–168 (2020).

225. Tse, D. *et al.* Schemas and memory consolidation. *Science* **316**, 76–82 (2007).

226. Tao, S. *et al.* Whole-Brain Mapping the Direct Inputs of Dorsal and Ventral CA1 Projection Neurons. *Front. Neural Circuits* **15**, 643230 (2021).

227. Fischer, L. *et al.* Dendritic Mechanisms for In Vivo Neural Computations and Behavior. *J. Neurosci.* **42**, 8460–8467 (2022).

228. Schiller, J., Schiller, Y., Stuart, G. & Sakmann, B. Calcium action potentials restricted to distal apical dendrites of rat neocortical pyramidal neurons. *J. Physiol.* **505 (Pt 3)**, 605–616 (1997).

229. Golding, N. L. & Spruston, N. Dendritic sodium spikes are variable triggers of axonal action potentials in hippocampal CA1 pyramidal neurons. *Neuron* **21**, 1189–1200 (1998).

230. Schiller, J., Major, G., Koester, H. J. & Schiller, Y. NMDA spikes in basal dendrites of cortical pyramidal neurons. *Nature* **404**, 285–289 (2000).

231. Larkum, M. E., Wu, J., Duverdin, S. A. & Gidon, A. The Guide to Dendritic Spikes of the Mammalian Cortex In Vitro and In Vivo. *Neuroscience* **489**, 15–33 (2022).

232. Rall, W. Electrophysiology of a dendritic neuron model. *Biophys. J.* **2**, 145–167 (1962).

233. Larkum, M. E., Zhu, J. J. & Sakmann, B. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature* **398**, 338–341 (1999).

234. Stuart, G. J. & Häusser, M. Dendritic coincidence detection of EPSPs and action potentials. *Nat. Neurosci.* **4**, 63–71 (2001).

235. London, M. & Häusser, M. Dendritic computation. *Annu. Rev. Neurosci.* **28**, 503–532 (2005).

236. Larkum, M. E., Nevian, T., Sandler, M., Polsky, A. & Schiller, J. Synaptic integration in tuft dendrites of layer 5 pyramidal neurons: a new unifying principle. *Science* **325**, 756–760 (2009).

237. Major, G., Larkum, M. E. & Schiller, J. Active properties of neocortical pyramidal neuron dendrites. *Annu. Rev. Neurosci.* **36**, 1–24 (2013).

238. Palmer, L. M. Dendritic integration in pyramidal neurons during network activity and disease. *Brain Res. Bull.* **103**, 2–10 (2014).

239. O'Hare, J. K. *et al.* Compartment-specific tuning of dendritic feature selectivity by intracellular Ca^{2+} release. *Science* **375**, eabm1670 (2022).

240. Larkum, M. E., Senn, W. & Lüscher, H.-R. Top-down dendritic input increases the gain of layer 5 pyramidal neurons. *Cereb. Cortex* **14**, 1059–1070 (2004).

241. Losonczy, A., Makara, J. K. & Magee, J. C. Compartmentalized dendritic plasticity and input feature storage in neurons. *Nature* **452**, 436–441 (2008).

242. Letzkus, J. J., Wolff, S. B. E. & Lüthi, A. Disinhibition, a Circuit Mechanism for Associative Learning and Memory. *Neuron* **88**, 264–276 (2015).

243. Kilborn, K., Kubota, D., Lynch, G. & Granger, R. Parameters of LTP Induction Modulate Network Categorization Behavior. in *Computational Neuroscience: Trends in Research, 1997* (ed. Bower, J. M.) 353–358 (Springer US, 1997).

244. Bliss, T. V. P. & Collingridge, G. L. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* **361**, 31–39 (1993).

245. Nicoll, R. A. A Brief History of Long-Term Potentiation. *Neuron* **93**, 281–290 (2017).

246. Stuart, G. J. & Spruston, N. Dendritic integration: 60 years of progress. *Nat. Neurosci.* **18**, 1713–1721 (2015).

247. Palmer, L., Murayama, M. & Larkum, M. Inhibitory Regulation of Dendritic Activity in vivo. *Front. Neural Circuits* **6**, 26 (2012).

248. McIlwain, H. Extended roles in the brain for second-messenger systems. *Neuroscience* **2**, 357–372 (1977).

249. Spruston, N., Schiller, Y., Stuart, G. & Sakmann, B. Activity-dependent action potential invasion and calcium influx into hippocampal CA1 dendrites. *Science* **268**, 297–300 (1995).

250. Kamondi, A., Acsády, L., Wang, X. J. & Buzsáki, G. Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials.

Hippocampus **8**, 244–261 (1998).

251. Benardo, L. S., Masukawa, L. M. & Prince, D. A. Electrophysiology of isolated hippocampal pyramidal dendrites. *J. Neurosci.* **2**, 1614–1622 (1982).

252. Wong, R. K. & Stewart, M. Different firing patterns generated in dendrites and somata of CA1 pyramidal neurones in guinea-pig hippocampus. *J. Physiol.* **457**, 675–687 (1992).

253. Golding, N. L., Jung, H. Y., Mickus, T. & Spruston, N. Dendritic calcium spike initiation and repolarization are controlled by distinct potassium channel subtypes in CA1 pyramidal neurons. *J. Neurosci.* **19**, 8789–8798 (1999).

254. Bittner, K. C. *et al.* Conjunctive input processing drives feature selectivity in hippocampal CA1 neurons. *Nat. Neurosci.* **18**, 1133–1142 (2015).

255. Takahashi, H. & Magee, J. C. Pathway interactions and synaptic plasticity in the dendritic tuft regions of CA1 pyramidal neurons. *Neuron* **62**, 102–111 (2009).

256. Xu, N.-L. *et al.* Nonlinear dendritic integration of sensory and motor input during an active sensing task. *Nature* **492**, 247–251 (2012).

257. Hill, D. N., Varga, Z., Jia, H., Sakmann, B. & Konnerth, A. Multibranch activity in basal and tuft dendrites during firing of layer 5 cortical neurons in vivo. *Proc. Natl. Acad. Sci. U. S. A.* **110**, 13618–13623 (2013).

258. Wilson, D. E., Whitney, D. E., Scholl, B. & Fitzpatrick, D. Orientation selectivity and the functional clustering of synaptic inputs in primary visual cortex. *Nat. Neurosci.* **19**, 1003–1009 (2016).

259. Takahashi, N., Oertner, T. G., Hegemann, P. & Larkum, M. E. Active cortical dendrites modulate perception. *Science* **354**, 1587–1590 (2016).

260. Doron, G. *et al.* Perirhinal input to neocortical layer 1 controls learning. *Science* **370**, (2020).

261. Basak, R. & Narayanan, R. Robust emergence of sharply tuned place-cell responses in hippocampal neurons with structural and biophysical heterogeneities. *Brain Struct. Funct.* **225**, 567–590 (2020).

262. Likhtik, E., Pelletier, J. G., Paz, R. & Paré, D. Prefrontal control of the amygdala. *J. Neurosci.* **25**, 7429–7437 (2005).

263. Pendyam, S. *et al.* Fear signaling in the prelimbic-amygdala circuit: a computational modeling and recording study. *J. Neurophysiol.* **110**, 844–861 (2013).

264. Arruda-Carvalho, M. & Clem, R. L. Pathway-selective adjustment of prefrontal-amygdala transmission during fear encoding. *J. Neurosci.* **34**, 15601–15609 (2014).

265. Song, C., Ehlers, V. L. & Moyer, J. R., Jr. Trace Fear Conditioning Differentially Modulates Intrinsic Excitability of Medial Prefrontal Cortex-Basolateral Complex of Amygdala Projection Neurons in Infralimbic and Prelimbic Cortices. *J. Neurosci.* **35**, 13511–13524 (2015).

266. Pfeiffer, B. E. & Foster, D. J. Hippocampal place-cell sequences depict future paths to remembered goals. *Nature* **497**, 74–79 (2013).

267. Perusini, J. N. & Fanselow, M. S. Neurobehavioral perspectives on the distinction between fear and anxiety. *Learn. Mem.* **22**, 417–425 (2015).

268. Bloodgood, D. W., Sugam, J. A., Holmes, A. & Kash, T. L. Fear extinction requires infralimbic cortex projections to the basolateral amygdala. *Transl. Psychiatry* **8**, 60 (2018).

269. Senn, V. *et al.* Long-range connectivity defines behavioral specificity of amygdala neurons. *Neuron* **81**, 428–437 (2014).

270. Asede, D., Bosch, D., Lüthi, A., Ferraguti, F. & Ehrlich, I. Sensory inputs to intercalated cells provide fear-learning modulated inhibition to the basolateral amygdala. *Neuron* **86**, 541–554 (2015).

271. Berretta, S., Pantazopoulos, H., Caldera, M., Pantazopoulos, P. & Paré, D. Infralimbic cortex activation increases c-Fos expression in intercalated neurons of the amygdala. *Neuroscience* **132**, 943–953 (2005).

272. Marek, R., Strobel, C., Bredy, T. W. & Sah, P. The amygdala and medial prefrontal cortex: partners in the fear circuit. *J. Physiol.* **591**, 2381–2391 (2013).

273. Bouton, M. E. Context, ambiguity, and unlearning: sources of relapse after behavioral extinction. *Biol. Psychiatry* **52**, 976–986 (2002).

274. Sotres-Bayon, F., Sierra-Mercado, D., Pardilla-Delgado, E. & Quirk, G. J. Gating of fear in prelimbic cortex by hippocampal and amygdala inputs. *Neuron* **76**, 804–812 (2012).

275. Knapska, E. *et al.* Functional anatomy of neural circuits regulating fear and extinction. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 17093–17098 (2012).

276. Hübner, C., Bosch, D., Gall, A., Lüthi, A. & Ehrlich, I. Ex vivo dissection of optogenetically activated mPFC and hippocampal inputs to neurons in the basolateral amygdala: implications for fear and emotional memory. *Front. Behav. Neurosci.* **8**, 64 (2014).

277. Han, S.-L. & Xu, T.-L. Unraveling the Mechanisms of Memory Extinction. *Neurosci. Bull.* **34**, 385–388 (2018).

278. Moscarello, J. M. Prefrontal cortex projections to the nucleus reuniens suppress freezing following two-way signaled avoidance training. *Learn. Mem.* **27**, 119–123 (2020).

279. Totty, M. S., Ramanathan, K. R., Jin, J., Peters, S. E. & Maren, S. Thalamic nucleus reuniens coordinates prefrontal-hippocampal synchrony to suppress extinguished fear. *bioRxiv* 2022.11.11.516165 (2022) doi:10.1101/2022.11.11.516165.

280. Ludkiewicz, B. *et al.* Cholinergic innervation of parvalbumin- and calbindin-containing neurones in the hippocampus during postnatal development of the rat brain. *Folia Morphol.* **61**, 89–96 (2002).

281. Lissek, S. *et al.* Generalized anxiety disorder is associated with overgeneralization of classically conditioned fear. *Biol. Psychiatry* **75**, 909–915 (2014).

282. Garfinkel, S. N. *et al.* Impaired contextual modulation of memories in PTSD: an fMRI and psychophysiological study of extinction retention and fear renewal. *J. Neurosci.* **34**, 13435–13443 (2014).

283. Maren, S. Neurobiology of Pavlovian fear conditioning. *Annu. Rev. Neurosci.* **24**, 897–931 (2001).

284. Fanselow, M. S. & Tighe, T. J. Contextual conditioning with massed versus distributed unconditional stimuli in the absence of explicit conditional stimuli. *J. Exp. Psychol. Anim. Behav. Process.* **14**, 187–199 (1988).

285. Suzuki, A. *et al.* Memory reconsolidation and extinction have distinct temporal and biochemical signatures. *J. Neurosci.* **24**, 4787–4795 (2004).

286. An, B. *et al.* Amount of fear extinction changes its underlying mechanisms. *Elife* **6**, (2017).

287. Kida, S. Function and mechanisms of memory destabilization and reconsolidation after retrieval. *Proc. Jpn. Acad. Ser. B Phys. Biol. Sci.* **96**, 95–106 (2020).

288. Lee, C. *et al.* Hippocampal engram networks for fear memory recruit new synapses and modify pre-existing synapses in vivo. *Curr. Biol.* **33**, 507–516.e3 (2023).

289. Venkataraman, A. & Dias, B. G. Expanding the canon: An inclusive neurobiology of thalamic and subthalamic fear circuits. *Neuropharmacology* **226**, 109380 (2022).

290. Wood, E. R., Dudchenko, P. A., Robitsek, R. J. & Eichenbaum, H. Hippocampal neurons encode information about different types of memory episodes occurring in the same location. *Neuron* **27**, 623–633 (2000).

291. Miry, O., Li, J. & Chen, L. The Quest for the Hippocampal Memory Engram: From Theories to Experimental Evidence. *Front. Behav. Neurosci.* **14**, 632019 (2020).

292. Matsuo, N. Irreplaceability of Neuronal Ensembles after Memory Allocation. *Cell Rep.* **11**, 351–357 (2015).

293. Wolff, M., Alcaraz, F., Marchand, A. R. & Coutureau, E. Functional heterogeneity of the limbic thalamus: From hippocampal to cortical functions. *Neurosci. Biobehav. Rev.* **54**, 120–130 (2015).

294. Troyner, F., Bicca, M. A. & Bertoglio, L. J. Nucleus reuniens of the thalamus controls fear memory intensity, specificity and long-term maintenance during consolidation. *Hippocampus* **28**, 602–616 (2018).

295. Lovett-Barron, M. *et al.* Dendritic inhibition in the hippocampus supports fear learning. *Science* **343**, 857–863 (2014).

296. Rajasethupathy, P. *et al.* Projections from neocortex mediate top-down control of memory retrieval. *Nature* **526**, 653–659 (2015).

297. Krishnan, S., Heer, C., Cherian, C. & Sheffield, M. E. J. Reward expectation extinction restructures and degrades CA1 spatial maps through loss of a dopaminergic reward proximity signal. *Nat. Commun.* **13**, 6662 (2022).

298. Zhu, H. *et al.* Chemogenetic inactivation of ventral hippocampal glutamatergic neurons disrupts

consolidation of contextual fear memory. *Neuropsychopharmacology* **39**, 1880–1892 (2014).

299. Roth, B. L. DREADDs for Neuroscientists. *Neuron* **89**, 683–694 (2016).

300. Nagai, Y. *et al.* Deschloroclozapine, a potent and selective chemogenetic actuator enables rapid neuronal and behavioral modulations in mice and monkeys. *Nat. Neurosci.* **23**, 1157–1167 (2020).

301. Dolleman-Van der Weel, M. J., Lopes da Silva, F. H. & Witter, M. P. Nucleus reuniens thalami modulates activity in hippocampal field CA1 through excitatory and inhibitory mechanisms. *J. Neurosci.* **17**, 5640–5650 (1997).

302. Chen, T. & Guestrin, C. XGBoost: A Scalable Tree Boosting System. *arXiv [cs.LG]* (2016).

303. Ji, J. & Maren, S. Differential roles for hippocampal areas CA1 and CA3 in the contextual encoding and retrieval of extinguished fear. *Learn. Mem.* **15**, 244–251 (2008).

304. Goshen, I. *et al.* Dynamics of retrieval strategies for remote memories. *Cell* **147**, 678–689 (2011).

305. Dolleman-Van Der Weel, M. J. & Witter, M. P. Projections from the nucleus reuniens thalami to the entorhinal cortex, hippocampal field CA1, and the subiculum in the rat arise from different populations of neurons. *J. Comp. Neurol.* **364**, 637–650 (1996).

306. Sheffield, M. E. J., Best, T. K., Mensh, B. D., Kath, W. L. & Spruston, N. Slow integration leads to persistent action potential firing in distal axons of coupled interneurons. *Nat. Neurosci.* **14**, 200–207 (2011).

307. Liu, C. *et al.* An action potential initiation mechanism in distal axons for the control of dopamine release. *Science* **375**, 1378–1385 (2022).

308. Corcoran, K. A. & Quirk, G. J. Activity in prelimbic cortex is necessary for the expression of learned, but not innate, fears. *J. Neurosci.* **27**, 840–844 (2007).

309. Bayer, H. & Bertoglio, L. J. Infralimbic cortex controls fear memory generalization and susceptibility to extinction during consolidation. *Sci. Rep.* **10**, 1–13 (2020).

310. Thompson, B. M. *et al.* Activation of the infralimbic cortex in a fear context enhances extinction learning. *Learn. Mem.* **17**, 591–599 (2010).

311. Kim, W. B. & Cho, J.-H. Synaptic Targeting of Double-Projecting Ventral CA1 Hippocampal

Neurons to the Medial Prefrontal Cortex and Basal Amygdala. *J. Neurosci.* **37**, 4868–4882 (2017).

312. Kim, W. B. & Cho, J.-H. Encoding of contextual fear memory in hippocampal–amygdala circuit. *Nat. Commun.* **11**, 1–22 (2020).

313. Steward, O. & Scoville, S. A. Cells of origin of entorhinal cortical afferents to the hippocampus and fascia dentata of the rat. *The Journal of Comparative Neurology* vol. 169 347–370 Preprint at https://doi.org/10.1002/cne.901690306 (1976).

314. Kajiwara, R. *et al.* Convergence of entorhinal and CA3 inputs onto pyramidal neurons and interneurons in hippocampal area CA1--an anatomical study in the rat. *Hippocampus* **18**, 266–280 (2008).

315. Chittajallu, R. *et al.* Afferent specific role of NMDA receptors for the circuit integration of hippocampal neurogliaform cells. *Nat. Commun.* **8**, 152 (2017).

316. Sun, Q. *et al.* Frequency-Dependent Synaptic Dynamics Differentially Tune CA1 and CA2 Pyramidal Neuron Responses to Cortical Input. *J. Neurosci.* **41**, 8103–8110 (2021).

317. Milstein, A. D. *et al.* Bidirectional synaptic plasticity rapidly modifies hippocampal representations. *Elife* **10**, (2021).

318. Priestley, J. B., Bowler, J. C., Rolotti, S. V., Fusi, S. & Losonczy, A. Signatures of rapid plasticity in hippocampal CA1 representations during novel experiences. *Neuron* **110**, 1978–1992.e6 (2022).

319. Fan, L. Z. *et al.* All-optical physiology resolves a synaptic basis for behavioral timescale plasticity. *Cell* **186**, 543–559.e19 (2023).

320. Milstein, A. D. *et al.* Bidirectional synaptic plasticity rapidly modifies hippocampal representations independent of correlated activity. 2020.02.04.934182 (2020) doi:10.1101/2020.02.04.934182.

321. Wang, M. E., Yuan, R. K., Keinath, A. T., Ramos Álvarez, M. M. & Muzzio, I. A. Extinction of Learned Fear Induces Hippocampal Place Cell Remapping. *J. Neurosci.* **35**, 9122–9136 (2015).

322. Wu, C.-T., Haggerty, D., Kemere, C. & Ji, D. Hippocampal awake replay in fear memory retrieval. *Nat. Neurosci.* **20**, 571–580 (2017).

323. Dombeck, D. A., Harvey, C. D., Tian, L., Looger, L. L. & Tank, D. W. Functional imaging of

hippocampal place cells at cellular resolution during virtual navigation. *Nat. Neurosci.* **13**, 1433–1440 (2010).

324. Kaufman, A. M., Geiller, T. & Losonczy, A. A Role for the Locus Coeruleus in Hippocampal CA1 Place Cell Reorganization during Spatial Reward Learning. *Neuron* **105**, 1018–1026.e4 (2020).
325. Broussard, G. J. *et al.* In vivo measurement of afferent activity with axon-specific calcium imaging. *Nat. Neurosci.* **21**, 1272–1280 (2018).
326. Krashes, M. J. *et al.* Rapid, reversible activation of AgRP neurons drives feeding behavior in mice. *J. Clin. Invest.* **121**, 1424–1428 (2011).
327. Schindelin, J. *et al.* Fiji: an open-source platform for biological-image analysis. *Nat. Methods* **9**, 676–682 (2012).
328. Pachitariu, M. *et al.* Suite2p: beyond 10,000 neurons with standard two-photon microscopy. *bioRxiv* 061507 (2016) doi:10.1101/061507.
329. Whittington, J. C. R. *et al.* The Tolman-Eichenbaum Machine: Unifying space and relational memory through generalisation in the hippocampal formation. *bioRxiv* 770495 (2019) doi:10.1101/770495.
330. Dulas, M. R., Schwarb, H., Cannavale, C. N. & Cohen, N. J. Chapter One - Should context hold a special place in hippocampal memory? in *Psychology of Learning and Motivation* (eds. Federmeier, K. D. & Sahakyan, L.) vol. 75 1–35 (Academic Press, 2021).
331. Lipinska, G., Stuart, B., Thomas, K. G. F., Baldwin, D. S. & Bolinger, E. Preferential Consolidation of Emotional Memory During Sleep: A Meta-Analysis. *Front. Psychol.* **10**, 1014 (2019).
332. Igata, H., Ikegaya, Y. & Sasaki, T. Prioritized experience replays on a hippocampal predictive map for learning. *Proc. Natl. Acad. Sci. U. S. A.* **118**, (2021).
333. Latuske, P., Kornienko, O., Kohler, L. & Allen, K. Hippocampal Remapping and Its Entorhinal Origin. *Front. Behav. Neurosci.* **11**, 253 (2017).
334. Sharp, P. E., Kubie, J. L. & Muller, R. U. Firing properties of hippocampal neurons in a visually symmetrical environment: contributions of multiple sensory cues and mnemonic processes. *J.*

Neurosci. **10**, 3093–3105 (1990).

335. Fyhn, M., Hafting, T., Treves, A., Moser, M.-B. & Moser, E. I. Hippocampal remapping and grid realignment in entorhinal cortex. *Nature* **446**, 190–194 (2007).

336. Jeffery, K. J. & Anderson, M. I. Dissociation of the geometric and contextual influences on place cells. *Hippocampus* **13**, 868–872 (2003).

337. Leutgeb, S. *et al.* Independent codes for spatial and episodic memory in hippocampal neuronal ensembles. *Science* **309**, 619–623 (2005).

338. Fetterhoff, D., Sobolev, A. & Leibold, C. Graded remapping of hippocampal ensembles under sensory conflicts. *Cell Rep.* **36**, 109661 (2021).

339. Schonhaut, D. R., Aghajan, Z. M., Kahana, M. J. & Fried, I. A neural code for spatiotemporal context. *bioRxiv* 2022.05.10.491339 (2022) doi:10.1101/2022.05.10.491339.

340. Schuette, P. J. *et al.* Long-term characterization of hippocampal remapping during contextual fear acquisition and extinction. *J. Neurosci.* (2020) doi:10.1523/JNEUROSCI.1022-20.2020.

341. Moita, M. A. P., Rosis, S., Zhou, Y., LeDoux, J. E. & Blair, H. T. Putting fear in its place: remapping of hippocampal place cells during fear conditioning. *J. Neurosci.* **24**, 7015–7023 (2004).

342. Rolotti, S. V., Blockus, H., Sparks, F. T., Priestley, J. B. & Losonczy, A. Reorganization of CA1 dendritic dynamics by hippocampal sharp-wave ripples during learning. *Neuron* **110**, 977–991.e4 (2022).

343. Gridchyn, I., Schoenenberger, P., O'Neill, J. & Csicsvari, J. Assembly-specific disruption of hippocampal replay leads to selective memory deficit. *Neuron* **106**, 291–300.e6 (2020).

344. Han, J.-W. *et al.* Genome-wide association study in a Chinese Han population identifies nine new susceptibility loci for systemic lupus erythematosus. *Nat. Genet.* **41**, 1234–1237 (2009).

345. Pettit, N. L., Yap, E.-L., Greenberg, M. E. & Harvey, C. D. Fos ensembles encode and shape stable spatial maps in the hippocampus. *Nature* **609**, 327–334 (2022).

346. Ratigan, H. C., Krishnan, S., Smith, S. & Sheffield, M. E. J. Direct Thalamic Inputs to Hippocampal CA1 Transmit a Signal That Suppresses Ongoing Contextual Fear Memory Retrieval. *bioRxiv*

2023.03.27.534420 (2023) doi:10.1101/2023.03.27.534420.

347. Kinsky, N. R., Sullivan, D. W., Mau, W., Hasselmo, M. E. & Eichenbaum, H. B. Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. *Curr. Biol.* **28**, 3578–3588.e6 (2018).

348. Zaki, Y. *et al.* Aversive experience drives offline ensemble reactivation to link memories across days. *bioRxiv* 2023.03.13.532469 (2023) doi:10.1101/2023.03.13.532469.

349. Malvache, A., Reichinnek, S., Villette, V., Haimerl, C. & Cossart, R. Awake hippocampal reactivations project onto orthogonal neuronal assemblies. *Science* **353**, 1280–1283 (2016).

350. Anisimova, M. *et al.* Neuronal FOS reports synchronized activity of presynaptic neurons. *bioRxiv* 2023.09.04.556168 (2023) doi:10.1101/2023.09.04.556168.

351. Mahringer, D. *et al.* Expression of c-Fos and Arc in hippocampal region CA1 marks neurons that exhibit learning-related activity changes. *bioRxiv* 644526 (2019) doi:10.1101/644526.

352. Mathis, A. *et al.* DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat. Neurosci.* **21**, 1281–1289 (2018).

353. Schneider, S., Lee, J. H. & Mathis, M. W. Learnable latent embeddings for joint behavioural and neural analysis. *Nature* **617**, 360–368 (2023).

354. Bowler, J. C. & Losonczy, A. Direct Cortical Inputs to Hippocampal Area CA1 Transmit Complementary Signals for Goal-directed Navigation. *bioRxiv* 2022.11.10.516009 (2022) doi:10.1101/2022.11.10.516009.

355. Davoodi, F. G., Motamedi, F., Akbari, E., Ghanbarian, E. & Jila, B. Effect of reversible inactivation of reuniens nucleus on memory processing in passive avoidance task. *Behav. Brain Res.* **221**, 1–6 (2011).

356. Cassel, J.-C. & Pereira de Vasconcelos, A. Importance of the ventral midline thalamus in driving hippocampal functions. *Prog. Brain Res.* **219**, 145–161 (2015).

357. Layfield, D. M., Patel, M., Hallock, H. & Griffin, A. L. Inactivation of the nucleus reuniens/rhomboid causes a delay-dependent impairment of spatial working memory. *Neurobiol.*

Learn. Mem. **125**, 163–167 (2015).

358. Jung, D., Huh, Y. & Cho, J. The Ventral Midline Thalamus Mediates Hippocampal Spatial Information Processes upon Spatial Cue Changes. *J. Neurosci.* **39**, 2276–2290 (2019).

359. Grijseels, D. M., Shaw, K., Barry, C. & Hall, C. N. Choice of method of place cell classification determines the population of cells identified. *PLoS Comput. Biol.* **17**, e1008835 (2021).

360. Zhang, Y. *et al.* Fast and sensitive GCaMP calcium indicators for imaging neural populations. *Nature* **615**, 884–891 (2023).

361. Addgene: AAV plasmids. https://www.addgene.org/browse/article/28192475/.

362. Nieh, E. H. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80–84 (2021).

363. Bittner, K. C., Milstein, A. D., Grienberger, C., Romani, S. & Magee, J. C. Behavioral time scale synaptic plasticity underlies CA1 place fields. *Science* **357**, 1033–1036 (2017).

364. Häusser, M., Spruston, N. & Stuart, G. J. Diversity and dynamics of dendritic signaling. *Science* **290**, 739–744 (2000).

365. Sheffield, M. E. J. & Dombeck, D. A. Calcium transient prevalence across the dendritic arbour predicts place field properties. *Nature* **517**, 200–204 (2015).

366. Adoff, M. D. *et al.* The functional organization of excitatory synaptic input to place cells. *Nat. Commun.* **12**, 3558 (2021).

367. Wilson, D. I. G., Watanabe, S., Milner, H. & Ainge, J. A. Lateral entorhinal cortex is necessary for associative but not nonassociative recognition memory. *Hippocampus* **23**, 1280–1290 (2013).

368. Hales, J. B. *et al.* Medial entorhinal cortex lesions only partially disrupt hippocampal place cells and hippocampus-dependent place memory. *Cell Rep.* **9**, 893–901 (2014).

369. Brandon, M. P., Koenig, J., Leutgeb, J. K. & Leutgeb, S. New and distinct hippocampal place codes are generated in a new environment during septal inactivation. *Neuron* **82**, 789–796 (2014).

370. Martig, A. K. & Mizumori, S. J. Y. Ventral tegmental area disruption selectively affects CA1/CA2 but not CA3 place fields during a differential reward working memory task. *Hippocampus* **21**,

172–184 (2011).

371. Helmchen, F., Svoboda, K., Denk, W. & Tank, D. W. In vivo dendritic calcium dynamics in deep-layer cortical pyramidal neurons. *Nat. Neurosci.* **2**, 989–996 (1999).

372. Higley, M. J. & Sabatini, B. L. Calcium signaling in dendrites and spines: practical and functional considerations. *Neuron* **59**, 902–913 (2008).

373. Adam, Y. *et al.* Voltage imaging and optogenetics reveal behaviour-dependent changes in hippocampal dynamics. *Nature* (2019) doi:10.1038/s41586-019-1166-7.

374. Hsu, F.-C. *et al.* Light-field microscopy with temporal focusing multiphoton illumination for scanless volumetric bioimaging. *Biomed. Opt. Express* **13**, 6610–6620 (2022).

375. Weber, T. D., Moya, M. V., Kılıç, K., Mertz, J. & Economo, M. N. High-speed multiplane confocal microscopy for voltage imaging in densely labeled neuronal populations. *Nat. Neurosci.* **26**, 1642–1650 (2023).