

# UC San Diego

## UC San Diego Electronic Theses and Dissertations

### Title

On the Dynamics of Hippocampal CA1 Interneurons During Associative Memory Processing

### Permalink

<https://escholarship.org/uc/item/0fw817c2>

### Author

Riviere Ruiz, Pamela Damaris

### Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA SAN DIEGO

On the Dynamics of Hippocampal CA1 Interneurons During Associative Memory Processing

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

in

Cognitive Science

by

Pamela D. Rivière

Committee in charge:

Professor Lara Rangel, Chair  
Professor Mikio Aoi  
Professor Andrea Chiba  
Professor Eran Mukamel  
Professor Douglas Nitz  
Professor Bradley Voytek

2023

Copyright

Pamela D. Rivière, 2023

All rights reserved.

The Dissertation of Pamela D. Rivière is approved, and it is acceptable in quality and form for publication on microfilm and electronically.

University of California San Diego

2023

## DEDICATION

For Dorcas Ruiz Jiménez, Joel Rivière William, Ginette William Lahens, Jean Rivière Durocher, and Gilberto Ruiz Alicea. You all believed I would be scholarly, and here I am.

## TABLE OF CONTENTS

DISSERTATION APPROVAL PAGE.....	iii
DEDICATION.....	iv
TABLE OF CONTENTS.....	v
LIST OF FIGURES.....	vi
LIST OF SUPPLEMENTARY FIGURES.....	vii
EPIGRAPH.....	viii
ACKNOWLEDGEMENTS.....	ix
VITA.....	xii
ABSTRACT OF THE DISSERTATION.....	xiii
CHAPTER 1: INTRODUCTION.....	1
CHAPTER 2: DIFFERENTIAL RECRUITMENT OF INHIBITORY INTERNEURONS: FIRING RATES.....	19
CHAPTER 3: MODELING RELATIONSHIPS BETWEEN RHYTHMIC PROCESSES AND NEURONAL SPIKE TIMING.....	35
CHAPTER 4: BEHAVIORALLY DEPENDENT HIPPOCAMPAL INTERNEURON ENTRAINMENT TO MULTIPLE RHYTHMS.....	104
CHAPTER 5: DISCUSSION.....	154

## LIST OF FIGURES

<b>Figure 1.1:</b> Schematic of a basic microcircuit motif.....	4
<b>Figure 1.2:</b> Simplified schematic of interneuron diversity across hippocampal CA1 lamina.....	6
<b>Figure 2.1.</b> Spike-field coherence and firing profiles of a single interneuron in the CA1 region of the rat hippocampus.....	25
<b>Figure 2.2.</b> Information for odors and odor-positions evolves over the course of odor sampling.....	28
<b>Figure 3.1.</b> Data structure and model specification .....	41
<b>Figure 3.2.</b> Model results in a simulated dataset.....	47
<b>Figure 3.3.</b> Model comparisons for odor sampling CA1 inhibitory interneuron data, for theta- and low gamma-phase based models separately.....	50
<b>Figure 4.</b> PhaseSH models trained on data from distinct behavioral epochs capture shifts in CA1 interneuron spike-LFP relationships.....	56
<b>Figure 3.5.</b> Model comparisons over time reveal rapid shifts in spike-phase relationships to LFP rhythms during an odor sampling epoch.....	60
<b>Figure 3.6.</b> Model comparisons over time on held-out data.....	62
<b>Figure 4.1.</b> Behavioral variables and model fitting procedure.....	110
<b>Figure 4.2.</b> Interneuron spike trains cluster according to spike-phase relationships with the theta (4-12 Hz) CA1 oscillation.....	114
<b>Figure 4.3.</b> Theta, low gamma, and high gamma frequencies can entrain interneurons simultaneously and orthogonally.....	122
<b>Figure 4.4.</b> Flexibility and stability in CA1 interneuron oscillatory responses to olfactory and visuospatial inputs.....	126

## LIST OF SUPPLEMENTARY FIGURES

<b>Supplementary Figure S3.1.</b> Average firing rates across the CA1 interneuron population during odor sampling intervals.....	82
<b>Supplementary Figure S3.2.</b> Model results in a simulated dataset of spike trains with high average firing rates (40-60 Hz).....	83
<b>Supplemental Figure. S3.3.</b> Average Kullback-Leibler (KL) divergence between phaseSH and SH model probabilities separate rhythmic from non-rhythmic simulated spike trains.....	85
<b>Supplemental Figure S3.4.</b> Population log losses evaluated on training trials.....	87
<b>Supplementary Figure S3.5.</b> Model comparisons for approach epochs.....	89
<b>Supplementary Figure S3.6.</b> Model comparisons for CA1 pyramidal cell activity during odor sampling.....	91
<b>Supplemental Figure S3.7.</b> Model log loss permutation tests.....	93
<b>Supplemental Figure S3.8.</b> Model comparisons over time across the population across behavioral states.....	95
<b>Supplementary Figure S4.1.</b> Schematic illustration and equation describing the computation for the sum of the Kullback-Leibler divergences (kld).....	149
<b>Supplementary Figure S4.2.</b> Permutation test results for conditional and joint probabilities of rhythmic entrainment across frequency bands).....	151



## EPIGRAPH

When I was younger  
it was plain to me  
I must make something of myself.

Older now  
I walk back streets  
admiring the houses  
of the very poor:  
roof out of line with sides  
the yards cluttered  
with old chicken wire, ashes,  
furniture gone wrong;  
the fences and outhouses  
built of barrel-staves  
and parts of boxes, all,  
if I am fortunate,  
smeared a bluish green  
that properly weathered  
pleases me best  
of all colors.

No one  
will believe this  
of vast import to the nation.

William Carlos Williams  
“Pastoral”, *¡Al Que Quiere!*

## ACKNOWLEDGEMENTS

Electrophysiological studies accrue as much clutter, painstakingly wrought but eventually misaligned furniture, forgotten box parts—all adding up to the ghosts of aspirations—as those quietly captured in William Carlos Williams’ musings. I might replace some of these items with wagon wheel mazes, split and spiraled tungsten wires, screws that managed to snap at the most inopportune times. My temperament is not built for this work, I don’t think, which makes me admire all the more all the expert physiologists who helped make my training in this space possible, and—on what felt like rare occasions—absolutely and undeservedly magical.

I thank Lara for introducing me to neural oscillations. “All energy moves in waves and particles, but the ocean wave is the only wave the naked human eye can see. If time is the lesson willing to be learned, the wave is time’s expression. It comes again and again. It keeps saying, “Will you know me now?”” (David Milch, *Life's Work*, p. 190). Of the brain regions I've encountered—and in my short career to date, this is obviously not a very large number—it is the hippocampus that affords this deepest of experiences that only the ocean wave would otherwise offer. We are not entitled to witness beauty and structure in our world. We could all have been born to chaos, disorder, frenzy. When it is given to us to see structure, when we become privy to the happy coincidence of (in the brain) laminar organization, density of cell bodies whose dendritic processes just happen to lie in the same orientation like fronds along the ocean floor straining towards the light—we partake in the unfolding of a gift. This anatomical, physical structure, when captured with the right electrophysiological tools, allows us to bear witness to oscillatory structure, periodicity in waves of ions that ebb and flow into and out of neuronal membranes.

Unlike the ocean wave, access to the particular natural wonder of neural oscillations was wrested through human innovation rather than grace, and it comes at a cost, and the cost should be acknowledged, so I thank the many rats that—hesitantly—offered their lives and gave up some comfort to give us all insights into the wonderful world of microcircuits, neural oscillations, and most importantly, inhibitory interneurons.

This dissertation is first and foremost dedicated to the people who made it possible for me to resiliently forge ahead when I was prepared to jump ship. My parents reset my mood and my orientation to life with daily doses of exceptionally good humor, patient listening, and unplumbable depths of encouragement. Talking to them each and every night, with very few exceptions over the course of seven years, kept me grounded and feeling feather-light even on days when I felt terrible before their call.

James Priestley has been a steady source of inspiration, intellectually stimulating text threads, stunning shots of the Gunks taken from inadvisable heights, more recently of jaw-dropping pictures of Swiss mundanity—that is to say, the glory of the Alps in visible range—and, as always, impossibly clever puns. In graduate school, I have met people whose wonder I had no right to expect. There are too many to name—and what a privilege that is (!!!)—but I will say I have spent lovely days and evenings chatting and thinking, exploring vegan restaurants around San Diego, and enjoying work and fun times alike. There is nothing quite like the sudden realization that you are sitting among people who are similar to you in so many odd, funny ways, and somehow also differ in ways that you admire.

Most of all, I thank Sean Trott. With Sean I have learned discipline and focus along with exuberant doses of playfulness and joy, and I've learned to bound my days such that I am reminded that I am first a (beloved!) person who is wealthy in family and friends, and then a scientist. Although I am still working hard to become a more reasonable human—through Sean's example—with him I have at least become more reasonably human. Moreover, I've now been introduced into his wonderful family, who radiate warmth and care in small and large ways. Even if grad school hadn't worked out at all, it would all still be worth it for the deep delight that is life with Sean.

Now, to business:

Chapter 2, in full, is a reprint of the material as it appears in *Advances in the Mathematical Sciences: AWM Research Symposium, Los Angeles, CA, April 2017* (pp. 161-171). Springer International Publishing. Rivière, Pamela D; Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.

Chapter 3, in full, is a reprint of the material as it appears in the *Journal of Neurophysiology*. Rivière, Pamela D; Schamberg, Gabriel; Coleman, Todd P; Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.

Chapter 4, in full, is being prepared for submission to *Current Biology*. Rivière, Pamela D; Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.

## VITA

2015 Bachelor of Arts in Neurosciences, Boston University

2023 Doctor of Philosophy in Cognitive Science, University of California San Diego

ABSTRACT OF THE DISSERTATION

On the Dynamics of Hippocampal CA1 Interneurons During Associative Memory Processing

by

Pamela D. Rivière

Doctor of Philosophy in Cognitive Science

University of California San Diego, 2023

Professor Lara Rangel, Chair

Surviving and thriving requires minimizing energy expenditure while navigating (sometimes harrowing) environments in search of nourishment, mates, and shelter from predators and the elements. Associative learning is one of the tools available to organisms as they work to solve this problem, allowing them to efficiently use scents, sounds, and visuospatial features to predict the likelihood and location of rewards.

We combine *in vivo* electrophysiology, complex behavioral paradigms, and statistical modeling techniques to explore the activity of the CA1 subregion of the hippocampus, which integrates a range of converging inputs that are thought to provide the sensorimotor, visuospatial, and affective information required to associate properties of the environment with inherently valuable stimuli.

To organize these streams of information, the CA1 subregion must attenuate activity irrelevant to the organisms' immediate goals, while simultaneously allowing select neuronal populations to respond to the appropriate inputs. This thesis consequently characterizes the activity patterns of inhibitory interneurons, in an effort to determine how they work cooperatively with constellations of inputs to coordinate select subpopulations of pyramidal cells. In Chapter 2, I show that inhibitory interneurons selectively modulate their firing rate responses to odorants and spatial locations during associative memory processing, suggesting that interneurons can in fact be recruited by distinct, behaviorally relevant inputs. Chapter 3 documents the development of a statistical modeling strategy for identifying rapid shifts in interneuron spike *timing* relative to oscillatory synaptic currents. This tool facilitated the discovery of interneurons whose spike timing reliably changes over the course of associative memory processing. Lastly, Chapter 4 investigates interneuron spike timing in response to distinct combinations of olfactory and visuospatial inputs. We apply the modeling approach described in Chapter 3 to characterize spike-phase relationships with respect to theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) local field potential (LFP) oscillations. We show that entrainment into higher frequencies is predicted by theta phase relationships, and we find that interneuron spike-phase relationships can vary according to distinct combinations of inputs. These findings demonstrate that interneurons flexibly orchestrate network activity given a diverse array of inputs and changing behavioral demands.

## CHAPTER 1: INTRODUCTION

*“But when many uniform instances appear, and the same object is always followed by the same event; we then begin to **entertain the notion of cause and connexion**. We then feel a new sentiment or impression, to wit, **a customary connexion** in the thought or imagination **between one object and its usual attendant** ....”*

- *David Hume, An Enquiry Concerning Human Understanding, p. 50 (bold emphasis mine)*

### **1.1 Survival, Associative Memory, and Hippocampal Inhibition**

Wild and designed spaces alike pose challenges to organisms' survival. Desertscapes with meager and concealed water supplies, dense forests teeming with predators and tantalizing but toxic plants, dark city tunnels whose walls quake every so often with the roar of heavy machinery traveling at speed...these are just some environments where organisms have managed to set up shop and make homes for themselves. Environments such as these impose costly penalties on long, luxurious excursions in search of the essentials for survival, such as food, water, shelter, and mates. Energy spent searching aimlessly for water in a desert may spell dehydration and disaster, wandering into predators' territories in the jungle may ensure becoming someone else's meal, and venturing out of a nook as the train barrels through the darkness may result in fatal injuries. Dangerous worlds demand more than endurance—they call for heuristics in the form of associations.

Associative memory allows organisms to leverage features of the environment—its sights, scents, and sounds—to make inferences about what kinds of rewards or dangers to expect, and where to find or avoid them. In the desert, a novice might respond impartially to a spiky cactus plant. To those who have formed a productive association, however, the uninviting or banal features of the cactus might in fact presage the extreme relief that accompanies breaking its skin and consuming the juicy water within. This process involves linking what looks like an arbitrary sensory feature (the swollen and spiky visual profile of the plant) to the inherent value obtained only after interacting with the object in a particular way (that is, biting into it). By the same token, coming across the reek of a predator's dung or urine should motivate a wary animal to distance itself (quickly, if possible) from the marked region. Associations serve as cues



to action, and they can often only really be learned, as argued by David Hume nearly three or so centuries ago: organisms do not come into the world intuiting all the affordances of objects in their environments. Experience, and the ability to establish associations between an object's sensory properties and the object's "usual attendants" or consequences are necessary preconditions for successfully inferring cause and effect, and thriving in the world<sup>1</sup>.

The work in this dissertation is concerned with the neural mechanisms that might facilitate an organisms' ability to deploy learned associations as they seek rewards. I orient my efforts towards a region of the brain, the hippocampus, that has been implicated in the formation of associations through a long history of detailed human case studies<sup>2,3</sup> and elegant experimental work with a variety of animal models, the most enduring of which have been rodents<sup>4-7</sup> and non-human primates<sup>8</sup>. Hippocampal neuronal activity—as measured through the rate of action potential (spike) emissions—drastically reorganizes in tandem with macaques' accuracy as they perform a visuospatial associative learning task<sup>9,10</sup>, and evidence has accrued through work with rodents to suggest that a portion of hippocampal neurons exhibit exquisitely selective responses to combinations of visual and olfactory stimuli that predict reward<sup>10-12</sup>. Notably, the spatial selectivity of hippocampal principal cell activity—the reason they have been often referred to as "place cells"<sup>13,14</sup>—manifests contextually, and depends heavily on the configuration of egocentric and allocentric visual cues<sup>15</sup> as well as the affective value of occupied locations (and future destinations) in the environment<sup>16-21</sup>. In this way, principal neuron spatial activity patterns operate less like absolute Cartesian coordinates, and more like a combinatorial code that accounts for multiple dimensions of experience.

Excitatory principal cells, however, are not the only hippocampal constituents. In fact, the CA1 subregion of the hippocampus is home to various classes of inhibitory interneurons, neurons whose typical effect is to suppress the activity of their targets<sup>22-26</sup>. Classification schemes vary along several dimensions: morphology<sup>22,27,28</sup>, postsynaptic target compartment (e.g. dendritic, somatic, axonal)<sup>22,24</sup>, molecular and transcriptomic expression profiles<sup>22</sup>, membrane receptor and ion channel kinetics<sup>24,29,30</sup>, and spike timing relationships to ongoing network activity<sup>25,31</sup>. To be sure, many of these features are

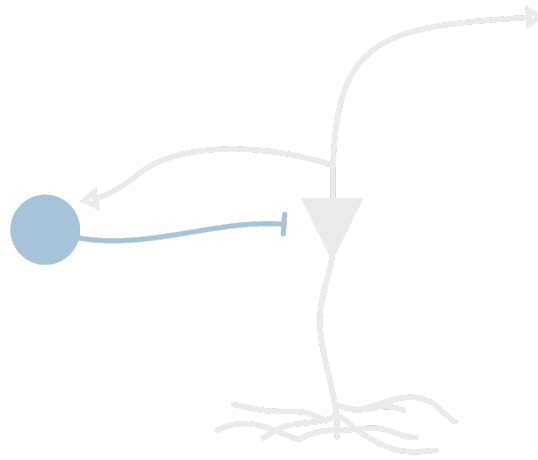
correlated, and one of the field's preoccupations from the late 1980s to the early 2000s involved mapping the relationships across these dimensions. The goal of this enterprise was to identify systematic functional roles for diverse interneuron classes, an outcome that would facilitate computational modeling of neural circuits, yield further insights into the machinery underlying cognition and behavior, and offer tractable—and increasingly precise—pharmacological targets for clinical interventions. Of particular interest for the work presented in this dissertation is the attempt to classify interneurons functionally during behavior, particularly with respect to the phase of oscillatory activity in the hippocampus.

Neural oscillations refer to systematic voltage fluctuations observable in the extracellularly recorded local field potential (LFP), a signal that emerges from the summation of all subthreshold ionic currents flowing in and out of cell membranes in the vicinity of the electrode<sup>32</sup>—with potential contributions from “volume-conducted” currents passively drifting from regions far from the recording electrode<sup>33</sup>. Oscillatory activity generated locally consequently reflects the collective and *temporally coordinated* activity of neurons in a network, as well as organization in the pattern of synaptic inputs provided by afferent fibers. Inhibition has been particularly implicated in the generation of sustained oscillations in the hippocampus, and the reliable spike timing patterns of these neurons relative to the hippocampal theta (4-12 Hz) oscillation observed in the 1970s earned them the name “theta cells”<sup>34</sup>.

Of the studies that then went on to perform inhibitory interneuron classifications on the basis of spike timing relative to theta rhythms (and in the 2010s, relative to faster gamma<sub>25-90 Hz</sub> oscillations) during behavior<sup>25,31</sup>, surprisingly little was or has been done to assess interneuron spike-phase relationships during the variety of associative memory behaviors that the hippocampus contributes to. The work here, motivated in detail through the remainder of this chapter, concerns itself with this apparent gap, and explores hippocampal interneuron dynamics—measured through variations in spike *rates* and spike *timing*—as they unfold during associative memory processing epochs. I will argue over the course of the work presented here that inhibitory interneurons are remarkably flexible in their responses to stimuli combinations associated with rewards, and that assessing these dynamics across the plurality of behaviors that organisms present is essential to unraveling the mysteries of neuronal machinery.

## 1.2 Inhibitory and Excitatory Neurons.

Neuronal tissues are diverse landscapes. At its broadest, variation across neurons manifests in the diversity of effects that a neuron exerts on its postsynaptic targets. *Excitation* (or the tendency to depolarize postsynaptic targets) and *inhibition* (precisely the opposite) are the main direct forms of influence observed in neural systems. Large brain regions comprise multiple smaller microcircuit motifs or templates. A common motif involves that of an inhibitory interneuron mutually interconnected with an excitatory principal cell (**Figure 1.1**). While the latter, with notable exceptions, serve as the primary outputs of the regions their cell bodies reside in, interneurons earned their moniker from their propensity to restrict their axonal arbors within a much more narrowly circumscribed volume of tissue (e.g. within the anatomically defined “boundaries” of a brain region).



**Figure 1.1.** Schematic of a basic microcircuit motif. A reduced inhibitory interneuron (blue) is mutually interconnected with an excitatory pyramidal cell (grey).

Inhibitory interneurons exert substantial influence on neural networks’ excitatory outputs. Although generally fewer in number relative to excitatory principal cells<sup>35–38</sup>, interneurons’ extremely dense axonal arbors target large groups of principal cells<sup>22,39,40</sup>. Interneurons also often exhibit ion channel and receptor distributions that result in relatively more depolarized resting membrane potentials<sup>41,42</sup>, a

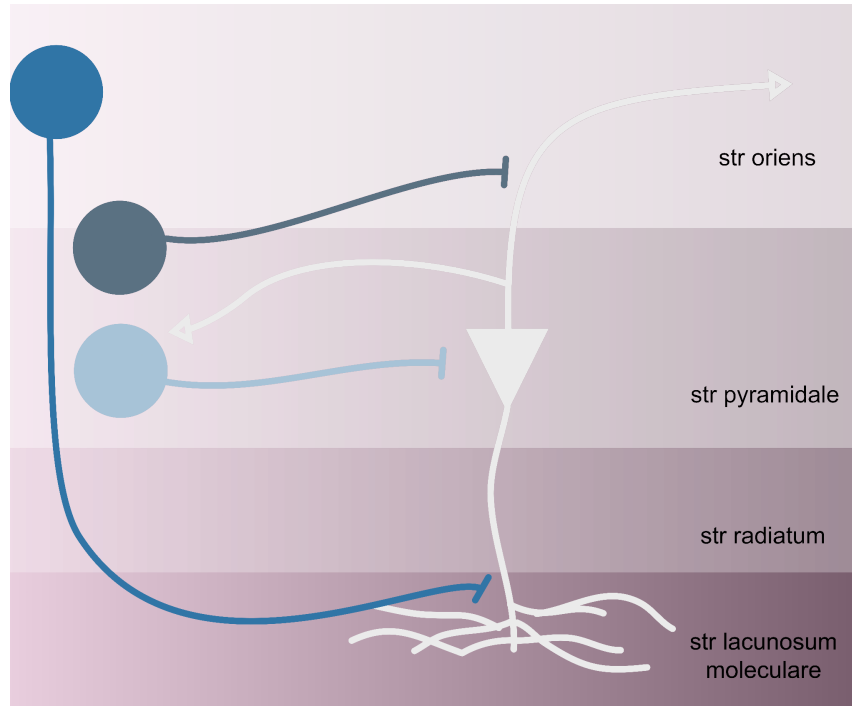
property that renders them more easily excitable on average than the principal neurons they innervate<sup>43</sup>. All these features coalesce to place inhibitory interneurons in an ideal position to organize the activity of principal cells.

### **1.3 Hippocampal CA1 and Network Architecture.**

Microcircuits like the (simplified) one illustrated in **Figure 1.1** abound in the hippocampal region, a network of structures—the dentate gyrus (DG), the CA3, and the CA1—that each possess unique microcircuit motifs and architectures and are connected serially and (mostly) unidirectionally, with the CA1 as the final node in the so-called trisynaptic pathway<sup>44,45</sup>. The CA1 subregion, the primary focus of this dissertation, boasts a wide variety of inhibitory cell types, each targeting distinct pyramidal cell dendritic, somatic, and axonal compartments<sup>22</sup>, ensuring that various inhibitory sources regulate inputs along all access points available to principal neurons (**Figure 1.2**).

The soma and dendrites of interneuron types (here defined according to the pyramidal cell compartments they preferentially target) occupy distinct layers of the CA1 subregion (**Figure 1.2**). Oriens lacunosum moleculare (OLM) cell soma occupy the stratum oriens CA1 layer, but their axons preferentially innervate the apical dendritic tufts of pyramidal cells, located in the CA1 stratum lacunosum moleculare and also the stratum radiatum<sup>22,46</sup>. Somatic-targeting “basket” and axonal-targeting axo-axonic interneuron soma largely occupy the pyramidal layer along with pyramidal neuron cell bodies<sup>22,46</sup>.

The laminar organization of different groups of interneuron cell bodies and dendrites determines which afferent inputs each class predominantly receives, and the location of interneuron axons in turn dictates which synaptic inputs they tend to regulate. This is a consequence of the exquisite laminar organization of input fibers traveling from the CA3 subregion of the hippocampus, from the entorhinal cortex, and from a profusion of neuromodulatory subcortical structures. In the following sections, I survey the variety of known afferents that recruit inhibitory and excitatory CA1 neurons, in the context of the kinds of sensory and affective information they are thought to provide.



**Figure 1.2.** Simplified schematic of interneuron diversity across hippocampal CA1 lamina. Interneurons vary in the pyramidal cell compartments their axons preferentially target.

#### 1.4 Olfactory Information: Entorhinal Cortex

Afferent fibers from entorhinal cortex layers 2/3 make up the Perforant Path, a name that arose from the way these fibers appear to “perforate” or impale hippocampal white matter tissues in early gross anatomical studies<sup>47</sup>. Synaptic terminals from the lateral entorhinal area (LEA) impinge at the level of CA1 stratum lacunosum moleculare, and consequently come into contact with the apical dendritic tufts of pyramidal cells<sup>48</sup>. As a consequence of this localization, LEA-driven excitatory postsynaptic currents in pyramidal cell dendrites are heavily regulated by OLM interneurons, whose axons lie in the vicinity of these contacts<sup>30,46</sup>. Any information conveyed through these fibers is consequently subject to inhibitory gating mechanisms.

But what is the nature of this “information”? In the rodent, the LEA receives direct inputs from the olfactory bulb<sup>49</sup>, a region of the brain whose activity is indirectly shaped—as its name implies—by odorants inhaled and capable of binding to olfactory receptor neurons located in the lining of the nasal

cavity, the olfactory epithelium<sup>50</sup>. Each olfactory receptor neuron is characterized by the olfactory receptors they express<sup>51</sup>, and odorants bind with different degrees of affinity to subsets of receptors, resulting in unique spatiotemporal patterns of neuronal spike emissions<sup>52</sup>. Spatial organization of olfactory receptor neurons in the olfactory epithelium is non-chemotopic; that is, the neurons do not cluster spatially according to the olfactory receptors they express. Instead, olfactory receptor neurons of a particular kind are scattered throughout the epithelium. Despite the (seeming) randomness, their projections are funneled into spatially organized bundles of fibers termed glomeruli, which comprise activity driven by a single olfactory receptor type. Olfactory bulb (OB) mitral cell dendrites constitute another major component of a given glomerulus. Mitral cells then provide inputs to, among other regions, the LEA<sup>53,54</sup>, which go on to express odor-selective activity patterns<sup>54,55</sup> as well as conjunctive responses for scented items in particular spatial locations<sup>56</sup>. This olfactory information may in turn influence CA1 microcircuits through the LEA's direct contacts with pyramidal cell apical dendrites.

Although inhibitory interneurons targeting pyramidal cell apical dendrites are poised to block or reduce LEA influence in the network, recent work demonstrates importantly nuanced LEA-driven inhibitory recruitment. Bilash et al. (2023) find that LEA inputs not only work to recruit pyramidal cells by eliciting dendritic spikes, they also simultaneously recruit a group of interneurons that *inhibits* inhibitory activity in the layer<sup>57</sup>. The ability of LEA fibers to recruit disinhibitory microcircuits to maximize their excitatory effects speaks to the diversity of hippocampal inhibitory dynamics, and the complex mechanisms involved—in this case—in processing olfactory information in the CA1 subregion.

## **1.5 Visuospatial Information: Entorhinal Cortex**

Although there exists a detailed anatomical case for an olfactory contribution, the rodent LEA also receives projections from the posterior parietal cortex and secondary visual area<sup>58</sup>. The LEA's activity patterns have consequently been explored extensively in the context of spatial tasks. In this domain, LEA neurons exhibit selectivity for objects, emitting action potentials at fixed distances away (or in the vicinity of) specific objects regardless of their location in allocentric space<sup>59</sup>, or responding only in the

locations that an object has previously occupied, often appearing to “follow” or “track” the object’s variously occupied positions in space<sup>60</sup>. Lesioning the LEA additionally provokes deficits in a novel object-context exploration task<sup>61</sup>, which presumably requires animals to have established familiarity with previously encountered objects. More recent work has shown LEA neuron activity patterns whose variance is explained by the rat’s egocentric orientation relative to a boundary or object, suggesting its ability to convey visuospatial information relative to egocentric spatial reference frames<sup>62</sup>.

The medial entorhinal area (MEA) receives what may be an even larger share of posterior parietal and visual associational innervation than the LEA<sup>63</sup>, and additionally projects Perforant Path fibers to CA1 pyramidal cell distal dendrites<sup>64</sup>. This region has been extensively studied in the context of spatial information processing. Its principal neurons often emit action potentials in spatially periodic patterns resulting in, so to speak, environment-specific tessellated firing patterns<sup>65</sup> when one observes the map of action potentials accumulated as a rat traversed all coordinates in a space. These firing patterns—notably also shaped by hippocampal feedback into deeper MEA layers<sup>66,67</sup>—can additionally be conjunctively modulated by the rat’s head direction and velocity, providing a substrate for combinatorial expressions of experience-dependent<sup>68,69</sup> interoceptive and spatial variables in MEA projection neuron dynamics<sup>70,71</sup>.

Visuospatial patterns of activity also emerge in hippocampal principal cells, though the direction of this information transfer has been called into question<sup>72</sup>, given that CA1 pyramidal cell place fields persist in the absence of MEA neurons’ typical grid tuning<sup>73</sup>. Whatever the case may be, both of these regions express spatially-anchored firing, with CA1 principal cells exhibiting multiple<sup>74,75</sup> and single tuning preferences for locations in space<sup>13,15</sup>.

Although CA1 inhibitory interneurons have traditionally been considered less spatially selective than their excitatory counterparts in the network, interneurons may in fact inherit a large amount of spatial information, perhaps partially through contacts with MEA *en passant* terminals<sup>76</sup> and pyramidal cell feedback projections<sup>76</sup>. Inhibitory interneurons are known for their very high average firing rates, which appear indiscriminate over space as rats traverse environments. As it turns out, however, these high rates of spike emissions vary systematically with space. Very early work first noted slight firing rate variations

across the different arms of an 8-arm radial maze<sup>77</sup>, and Kubie, Muller, & Bostock (1990) systematically explored interneuron spike rates across space and found a pronounced lack of uniformity in the spatial distribution of action potentials of interneurons, despite their high firing rates throughout the environment<sup>78</sup>. They can evince dramatic increases or drop-offs at narrowly circumscribed spatial locations in the environment, at the same spatial resolution, in fact, as hippocampal principal cells exhibit single place fields<sup>78–80</sup>.

### **1.9 Oscillatory Phase as a Coordinating Mechanism.**

The CA1 local field potential (LFP) exhibits a diverse oscillatory spectrum, whose expression is extremely behavior-dependent<sup>81</sup>. Originally termed “rhythmical slow activity,” the theta rhythm (4-12 Hz) dominates the CA1 LFP during locomotive behaviors but yields to large irregular waves during immobile states<sup>81,82</sup>. The LFP reorganizes yet again when considering associative memory processing epochs, during which fast gamma oscillations emerge more prominently relative to their presence during locomotion<sup>12,83–85</sup>.

Inhibitory interneuron activity in the CA1 is thought to contribute, along with various cortical and subcortical afferents, to the emergence of each of these oscillations. Early reports of inhibitory interneurons in the hippocampus documented systematic time-locking of these cells relative to theta oscillations, which suggested—though did not demonstrate—a relationship between GABAergic currents and theta oscillatory profiles<sup>34,86</sup>. *In vitro* bath application of a GABA<sub>A</sub> receptor antagonist abolished a locally generated CA1 theta oscillator<sup>87</sup>, and later work in the same preparation leveraged optogenetic silencing of parvalbumin-positive (PV+) interneurons to show this class of interneuron specifically contributes to sustaining locally generated theta oscillations<sup>88</sup>. Similarly, GABA<sub>A</sub> antagonists block hippocampal gamma oscillations centered around a ~40 Hz frequency<sup>89</sup>, with the frequency of the LFP oscillation determined by the duration of GABAergic inhibitory postsynaptic potentials in the targeted pyramidal cell population<sup>90</sup>.



Inhibition-based oscillatory coordination of pyramidal cell assemblies offers what appear to be advantageous sources of temporal coordination. First, enforcing periodic windows of inhibition and excitability in pyramidal cells may guarantee that only the most well timed afferents exert their influence on the pyramidal cell population, essentially establishing smooth interaction channels between certain inputs while blocking others<sup>91</sup>. Moreover, inhibition might serve to “sharpen” and ensure the integrity of excitatory outputs: without periodic inhibition, small temporal jitter in pyramidal cell assembly spiking would result in desynchronized assemblies. In the presence of regular inhibitory currents, pyramidal cell spiking rebounds in synchrony across the population; neurons that need to fire in tandem are then able to sustain this synchronicity over longer periods of time<sup>92,93</sup>.

### **1.10 Summary**

Interneuron activity profoundly shapes the output of the hippocampus, a brain region whose intact dynamics support associative memory and, consequently, organisms’ ability to navigate the demands of a complex world. To date, little work has characterized inhibitory interneuron activity during paradigms that specifically test associative memory processing. In the work that follows, I show that inhibitory interneurons flexibly vary their spike rates in response to distinct stimuli combinations during an associative memory task (Chapter 2). Given inhibitory interneuron involvement in sculpting oscillatory activity, I develop a method to quantify and predict spike-phase relationships to various rhythms over time. Applying this method across locomotive and associative memory processing epochs revealed groups of interneurons that shifted their spike-phase relationships as a function of behavioral state (Chapter 3). Finally, I identify the presence of systematic interneuron spike timing relationships to theta and gamma rhythmic circuit processes—again, a property that depended on the animal’s behavioral state, as well as the combination of olfactory and visuospatial information available during associative memory processing. Evidence from this work converges to demonstrate the dynamism inherent to hippocampal inhibition, and its potential contributions to associative memory and survival.

## 1.11 References

1. Hume, D. *An Inquiry Concerning Human Understanding*. (1861).
2. Scoville, W. B. & Milner, B. Loss of recent memory after bilateral hippocampal lesions. *J. Neurol. Neurosurg. Psychiatry* **20**, 11–21 (1957).
3. Zola-Morgan, S., Squire, L. R. & Amaral, D. G. Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *J. Neurosci.* **6**, 2950–2967 (1986).
4. Jarrard, L. E. Selective hippocampal lesions: differential effects on performance by rats of a spatial task with preoperative versus postoperative training. *J. Comp. Physiol. Psychol.* **92**, 1119–1127 (1978).
5. Morris, R. G., Garrud, P., Rawlins, J. N. & O'Keefe, J. Place navigation impaired in rats with hippocampal lesions. *Nature* **297**, 681–683 (1982).
6. Eichenbaum, H. Using olfaction to study memory. *Ann. N. Y. Acad. Sci.* **855**, 657–669 (1998).
7. Alvarez, P., Wendelken, L. & Eichenbaum, H. Hippocampal formation lesions impair performance in an odor-odor association task independently of spatial context. *Neurobiol. Learn. Mem.* **78**, 470–476 (2002).
8. Zola-Morgan, S., Squire, L. R. & Amaral, D. G. Lesions of the amygdala that spare adjacent cortical regions do not impair memory or exacerbate the impairment following lesions of the hippocampal formation. *J. Neurosci.* **9**, 1922–1936 (1989).
9. Suzuki, W. A. Making new memories: the role of the hippocampus in new associative learning. *Ann. N. Y. Acad. Sci.* **1097**, 1–11 (2007).
10. Eichenbaum, H. Hippocampus: cognitive processes and neural representations that underlie declarative memory. *Neuron* **44**, 109–120 (2004).
11. Komorowski, R. W., Manns, J. R. & Eichenbaum, H. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* **29**, 9918–9929 (2009).

12. Rangel, L. M. *et al.* Rhythmic coordination of hippocampal neurons during associative memory processing. *Elife* **5**, e09849 (2016).
13. 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).
14. 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).
15. Muller, R. U. & Kubie, J. L. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *J. Neurosci.* **7**, 1951–1968 (1987).
16. Lee, I., Griffin, A. L., Zilli, E. A., Eichenbaum, H. & Hasselmo, M. E. Gradual translocation of spatial correlates of neuronal firing in the hippocampus toward prospective reward locations. *Neuron* **51**, 639–650 (2006).
17. Ainge, J. A., Tamosiunaite, M., Woergoetter, F. & Dudchenko, P. A. Hippocampal CA1 place cells encode intended destination on a maze with multiple choice points. *J. Neurosci.* **27**, 9769–9779 (2007).
18. Hok, V. *et al.* Goal-related activity in hippocampal place cells. *J. Neurosci.* **27**, 472–482 (2007).
19. Lee, H., Ghim, J.-W., Kim, H., Lee, D. & Jung, M. Hippocampal neural correlates for values of experienced events. *J. Neurosci.* **32**, 15053–15065 (2012).
20. 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).
21. Hollup, S. A., Molden, S., Donnett, J. G., Moser, M. B. & Moser, E. I. Accumulation of hippocampal place fields at the goal location in an annular watermaze task. *J. Neurosci.* **21**, 1635–1644 (2001).
22. Freund, T. F. & Buzsáki, G. Interneurons of the hippocampus. *Hippocampus* **6**, 347–470 (1996).
23. Parra, P., Gulyás, A. I. & Miles, R. How many subtypes of inhibitory cells in the hippocampus? *Neuron* **20**, 983–993 (1998).
24. Vida, I., Halasy, K., Szinyei, C., Somogyi, P. & Buhl, E. H. Unitary IPSPs evoked by interneurons at the stratum radiatum-stratum lacunosum-moleculare border in the CA1 area of the rat hippocampus

- in vitro. *J. Physiol.* **506 ( Pt 3)**, 755–773 (1998).
25. Klausberger, T. *et al.* Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* **421**, 844–848 (2003).
  26. Harris, K. D. *et al.* Classes and continua of hippocampal CA1 inhibitory neurons revealed by single-cell transcriptomics. *PLoS Biol.* **16**, e2006387 (2018).
  27. Kosaka, T., Wu, J. Y. & Benoit, R. GABAergic neurons containing somatostatin-like immunoreactivity in the rat hippocampus and dentate gyrus. *Exp. Brain Res.* **71**, 388–398 (1988).
  28. Sloviter, R. S. & Nilaver, G. Immunocytochemical localization of GABA-, cholecystinin-, vasoactive intestinal polypeptide-, and somatostatin-like immunoreactivity in the area dentata and hippocampus of the rat. *J. Comp. Neurol.* **256**, 42–60 (1987).
  29. Maccaferri, G. & Dingledine, R. Control of feedforward dendritic inhibition by NMDA receptor-dependent spike timing in hippocampal interneurons. *J. Neurosci.* **22**, 5462–5472 (2002).
  30. Maccaferri, G. & McBain, C. J. The hyperpolarization-activated current (I<sub>h</sub>) and its contribution to pacemaker activity in rat CA1 hippocampal stratum oriens-alveus interneurons. *J. Physiol.* **497 ( Pt 1)**, 119–130 (1996).
  31. Czurkó, A., Huxter, J., Li, Y., Hangya, B. & Muller, R. U. Theta phase classification of interneurons in the hippocampal formation of freely moving rats. *J. Neurosci.* **31**, 2938–2947 (2011).
  32. Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* **13**, 407–420 (2012).
  33. Herreras, O. Local Field Potentials: Myths and Misunderstandings. *Front. Neural Circuits* **10**, 101 (2016).
  34. Ranck, J. B., Jr. Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. *Exp. Neurol.* **41**, 461–531 (1973).
  35. Woodson, W., Nitecka, L. & Ben-Ari, Y. Organization of the GABAergic system in the rat hippocampal formation: a quantitative immunocytochemical study. *J. Comp. Neurol.* **280**, 254–271 (1989).

36. Aika, Y., Ren, J. Q., Kosaka, K. & Kosaka, T. Quantitative analysis of GABA-like-immunoreactive and parvalbumin-containing neurons in the CA1 region of the rat hippocampus using a stereological method, the disector. *Exp. Brain Res.* **99**, 267–276 (1994).
37. Bezaire, M. J. & Soltesz, I. Quantitative assessment of CA1 local circuits: knowledge base for interneuron-pyramidal cell connectivity. *Hippocampus* **23**, 751–785 (2013).
38. Olbrich, H. G. & Braak, H. Ratio of pyramidal cells versus non-pyramidal cells in sector CA1 of the human Ammon's horn. *Anat. Embryol.* **173**, 105–110 (1985).
39. Sik, A., Penttonen, M., Ylinen, A. & Buzsáki, G. Hippocampal CA1 interneurons: an in vivo intracellular labeling study. *J. Neurosci.* **15**, 6651–6665 (1995).
40. Cobb, S. R., Buhl, E. H., Halasy, K., Paulsen, O. & Somogyi, P. Synchronization of neuronal activity in hippocampus by individual GABAergic interneurons. *Nature* **378**, 75–78 (1995).
41. Jonas, P., Bischofberger, J., Fricker, D. & Miles, R. Interneuron Diversity series: Fast in, fast out--temporal and spatial signal processing in hippocampal interneurons. *Trends Neurosci.* **27**, 30–40 (2004).
42. Verheugen, J. A., Fricker, D. & Miles, R. Noninvasive measurements of the membrane potential and GABAergic action in hippocampal interneurons. *J. Neurosci.* **19**, 2546–2555 (1999).
43. Eccles, J. C. *The Inhibitory Pathways of the Central Nervous System.* (1969).
44. Amaral, D. G. Emerging principles of intrinsic hippocampal organization. *Curr. Opin. Neurobiol.* **3**, 225–229 (1993).
45. Andersen, P., Bliss, T. V. & Skrede, K. K. Lamellar organization of hippocampal pathways. *Exp. Brain Res.* **13**, 222–238 (1971).
46. Klausberger, T. & Somogyi, P. Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science* **321**, 53–57 (2008).
47. Witter, M. P., Wouterlood, F. G., Naber, P. A. & Van Haften, T. Anatomical organization of the parahippocampal-hippocampal network. *Ann. N. Y. Acad. Sci.* **911**, 1–24 (2000).
48. Steward, O. Topographic organization of the projections from the entorhinal area to the hippocampal

- formation of the rat. *J. Comp. Neurol.* **167**, 285–314 (1976).
49. Macrides, F., Eichenbaum, H. B. & Forbes, W. B. Temporal relationship between sniffing and the limbic theta rhythm during odor discrimination reversal learning. *J. Neurosci.* **2**, 1705–1717 (1982).
  50. Lancet, D. Vertebrate olfactory reception. *Annu. Rev. Neurosci.* **9**, 329–355 (1986).
  51. Mombaerts, P. Odorant receptor gene choice in olfactory sensory neurons: the one receptor-one neuron hypothesis revisited. *Curr. Opin. Neurobiol.* **14**, 31–36 (2004).
  52. Brody, C. D. & Hopfield, J. J. Simple networks for spike-timing-based computation, with application to olfactory processing. *Neuron* **37**, 843–852 (2003).
  53. Schwerdtfeger, W. K., Buhl, E. H. & Germroth, P. Disynaptic olfactory input to the hippocampus mediated by stellate cells in the entorhinal cortex. *J. Comp. Neurol.* **292**, 163–177 (1990).
  54. Leitner, F. C. *et al.* Spatially segregated feedforward and feedback neurons support differential odor processing in the lateral entorhinal cortex. *Nat. Neurosci.* **19**, 935–944 (2016).
  55. Young, B. J., Otto, T., Fox, G. D. & Eichenbaum, H. Memory representation within the parahippocampal region. *J. Neurosci.* **17**, 5183–5195 (1997).
  56. Keene, C. S. *et al.* Complementary Functional Organization of Neuronal Activity Patterns in the Perirhinal, Lateral Entorhinal, and Medial Entorhinal Cortices. *J. Neurosci.* **36**, 3660–3675 (2016).
  57. Bilash, O. M., Chavlis, S., Johnson, C. D., Poirazi, P. & Basu, J. Lateral entorhinal cortex inputs modulate hippocampal dendritic excitability by recruiting a local disinhibitory microcircuit. *Cell Rep.* **42**, 111962 (2023).
  58. Olsen, G. M., Ohara, S., Iijima, T. & Witter, M. P. Parahippocampal and retrosplenial connections of rat posterior parietal cortex. *Hippocampus* **27**, 335–358 (2017).
  59. Deshmukh, S. S. & Knierim, J. J. Representation of non-spatial and spatial information in the lateral entorhinal cortex. *Front. Behav. Neurosci.* **5**, 69 (2011).
  60. Tsao, A., Moser, M.-B. & Moser, E. I. Traces of experience in the lateral entorhinal cortex. *Curr. Biol.* **23**, 399–405 (2013).
  61. Wilson, D. I. G. *et al.* Lateral entorhinal cortex is critical for novel object-context recognition.

- Hippocampus* **23**, 352–366 (2013).
62. Wang, C. *et al.* Egocentric coding of external items in the lateral entorhinal cortex. *Science* **362**, 945–949 (2018).
  63. Burwell, R. D. & Amaral, D. G. Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. *J. Comp. Neurol.* **398**, 179–205 (1998).
  64. Hjorth-Simonsen, A. & Jeune, B. Origin and termination of the hippocampal perforant path in the rat studied by silver impregnation. *J. Comp. Neurol.* **144**, 215–232 (1972).
  65. Hafting, T., Fyhn, M., Molden, S., Moser, M.-B. & Moser, E. I. Microstructure of a spatial map in the entorhinal cortex. *Nature* **436**, 801–806 (2005).
  66. Ohara, S. *et al.* Intrinsic Projections of Layer Vb Neurons to Layers Va, III, and II in the Lateral and Medial Entorhinal Cortex of the Rat. *Cell Rep.* **24**, 107–116 (2018).
  67. Tamamaki, N. & Nojyo, Y. Preservation of topography in the connections between the subiculum, field CA1, and the entorhinal cortex in rats. *J. Comp. Neurol.* **353**, 379–390 (1995).
  68. Boccara, C. N., Nardin, M., Stella, F., O’Neill, J. & Csicsvari, J. The entorhinal cognitive map is attracted to goals. *Science* **363**, 1443–1447 (2019).
  69. Barry, C., Ginzberg, L. L., O’Keefe, J. & Burgess, N. Grid cell firing patterns signal environmental novelty by expansion. *Proc. Natl. Acad. Sci. U. S. A.* **109**, 17687–17692 (2012).
  70. Sargolini, F. *et al.* Conjunctive representation of position, direction, and velocity in entorhinal cortex. *Science* **312**, 758–762 (2006).
  71. Rueckemann, J. W., Sosa, M., Giocomo, L. M. & Buffalo, E. A. The grid code for ordered experience. *Nat. Rev. Neurosci.* **22**, 637–649 (2021).
  72. Bush, D., Barry, C. & Burgess, N. What do grid cells contribute to place cell firing? *Trends Neurosci.* **37**, 136–145 (2014).
  73. Koenig, J., Linder, A. N., Leutgeb, J. K. & Leutgeb, S. The spatial periodicity of grid cells is not sustained during reduced theta oscillations. *Science* **332**, 592–595 (2011).
  74. Fenton, A. A. *et al.* Unmasking the CA1 ensemble place code by exposures to small and large

- environments: more place cells and multiple, irregularly arranged, and expanded place fields in the larger space. *J. Neurosci.* **28**, 11250–11262 (2008).
75. Singer, A. C., Karlsson, M. P., Nathe, A. R., Carr, M. F. & Frank, L. M. Experience-dependent development of coordinated hippocampal spatial activity representing the similarity of related locations. *J. Neurosci.* **30**, 11586–11604 (2010).
  76. Takács, V. T., Klausberger, T., Somogyi, P., Freund, T. F. & Gulyás, A. I. Extrinsic and local glutamatergic inputs of the rat hippocampal CA1 area differentially innervate pyramidal cells and interneurons. *Hippocampus* **22**, 1379–1391 (2012).
  77. McNaughton, B. L., Barnes, C. A. & O'Keefe, J. The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp. Brain Res.* **52**, 41–49 (1983).
  78. Kubie, J. L., Muller, R. U. & Bostock, E. Spatial firing properties of hippocampal theta cells. *J. Neurosci.* **10**, 1110–1123 (1990).
  79. Wilent, W. B. & Nitz, D. A. Discrete place fields of hippocampal formation interneurons. *J. Neurophysiol.* **97**, 4152–4161 (2007).
  80. Ego-Stengel, V. & Wilson, M. A. Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus* **17**, 161–174 (2007).
  81. Vanderwolf, C. H. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr. Clin. Neurophysiol.* **26**, 407–418 (1969).
  82. Buzsáki, G., Leung, L. W. & Vanderwolf, C. H. Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* **287**, 139–171 (1983).
  83. Igarashi, K. M., Lu, L., Colgin, L. L., Moser, M.-B. & Moser, E. I. Coordination of entorhinal-hippocampal ensemble activity during associative learning. *Nature* **510**, 143–147 (2014).
  84. Gattas, S., Elias, G. A., Janecek, J., Yassa, M. A. & Fortin, N. J. Proximal CA1 20-40 Hz power dynamics reflect trial-specific information processing supporting nonspatial sequence memory. *Elife* **11**, (2022).



85. Trimper, J. B., Galloway, C. R., Jones, A. C., Mandi, K. & Manns, J. R. Gamma Oscillations in Rat Hippocampal Subregions Dentate Gyrus, CA3, CA1, and Subiculum Underlie Associative Memory Encoding. *Cell Rep.* **21**, 2419–2432 (2017).
86. Feder, R. & Ranck, J. B., Jr. Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. II. Hippocampal slow waves and theta cell firing during bar pressing and other behaviors. *Exp. Neurol.* **41**, 532–555 (1973).
87. Goutagny, R., Jackson, J. & Williams, S. Self-generated theta oscillations in the hippocampus. *Nat. Neurosci.* **12**, 1491–1493 (2009).
88. Amilhon, B. *et al.* Parvalbumin Interneurons of Hippocampus Tune Population Activity at Theta Frequency. *Neuron* **86**, 1277–1289 (2015).
89. Whittington, M. A., Traub, R. D. & Jefferys, J. G. Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature* **373**, 612–615 (1995).
90. Whittington, M. A., Traub, R. D., Kopell, N., Ermentrout, B. & Buhl, E. H. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *Int. J. Psychophysiol.* **38**, 315–336 (2000).
91. Fries, P. Rhythms for Cognition: Communication through Coherence. *Neuron* **88**, 220–235 (2015).
92. Singer, W. Time as coding space? *Curr. Opin. Neurobiol.* **9**, 189–194 (1999).
93. Ryan, S. J. *et al.* Spike-timing precision and neuronal synchrony are enhanced by an interaction between synaptic inhibition and membrane oscillations in the amygdala. *PLoS One* **7**, e35320 (2012).

## CHAPTER 2: DIFFERENTIAL RECRUITMENT OF INHIBITORY INTERNEURONS: FIRING RATES

### 2.1 Abstract

Flexible, dynamic activity in the brain is essential to information processing. Neurons in the hippocampus are capable of conveying information about the continually evolving world through changes in their spiking activity. This information can be expressed through changes in firing rate and through the reorganization of spike timing in unique rhythmic profiles. Locally projecting interneurons of the hippocampus are in an ideal position to coordinate task-relevant changes in the spiking activity of the network, as their inhibitory influence allows them to constrain communication between neurons to rhythmic, optimal windows and facilitates selective responses to afferent input. During a context-guided odor-reward association task, interneurons and principal cells in the CA1 subregion of the rat hippocampus demonstrate distinct oscillatory profiles that correspond to correct and incorrect performance, despite similar firing rates during correct and incorrect trials (Rangel et al., eLife 5:e09849, 2016). Principal cells additionally contained information in their firing rates about task dimensions, reflective of highly selective responses to features such as single positions and odors. It remains to be determined whether interneurons also contain information about task dimensions in their firing rates. To address this question, we evaluated the information content for task dimensions in the firing rates of inhibitory neurons. Interneurons contained low, but significant information for task dimensions in their firing rates, with increases in information over the course of a trial that reflected the evolving availability of task dimensions. These results suggest that interneurons are capable of manifesting distinct rhythmic profiles and changes in firing rate that reflect task-relevant processing.

### 2.2 Introduction

Successful information processing in the brain is characterized by dynamic patterns of neuronal activity that reflect information about the changing world. In the hippocampus, a brain region important for learning and memory, these changes often manifest as selective responses to features of the

environment [1]. During associative memory processing, neurons in the CA1 region of the rat hippocampus have been shown to exhibit highly selective and reliable changes in firing rate in response to specific spatial locations within an environment, stimuli such as odors that are relevant for performance [2], and the conjunction of specific odors in particular spatial locations [3,4,5]. Metrics that exploit these selective increases in firing rate have been devised to quantify neuronal information content for particular task features [3, 6,7,8]. These analyses have revealed subsets of neurons in the hippocampus whose transient changes in firing rate reflect the availability of behaviorally relevant features in the environment, leading to the hypothesis that the hippocampus dynamically recruits appropriate neuronal ensembles in the service of memory [9,10,11,12,13].

Associative memory processing is additionally accompanied by dynamic shifts in the oscillatory profile of the local field potential (LFP) in the CA1 region of the rat hippocampus [3]. In a context-guided odor–reward association task, rats must learn that odors are differentially rewarded depending upon the context in which they are encountered. During intervals in which rats correctly associated odors with the contexts in which they were rewarded, we observed large amplitude changes in the theta (4–12 Hz), beta (15–35 Hz), low gamma (35–55 Hz), and high gamma (65–90 Hz) frequency ranges. A surprising indicator of performance was the reorganization of spike timing with respect to the ongoing oscillations, such that both interneurons and principal cells exhibited distinct profiles of engagement in each of the four rhythms depending upon the trial outcome (correct or incorrect). For example, while the largest proportion of interneurons exhibited consistent spike timing relationships to each of the four frequency ranges during correct performance, many exhibited spike timing relationships to only theta during incorrect performance despite similar overall firing rates between correct and incorrect trials. This finding revealed a previously underexplored task-relevant selectivity in the rhythmic domain, and highlighted the dynamic patterns of activity in the interneuron population as a hallmark of successful processing.

Having established that the selective reorganization of interneuron spike timing is related to successful performance, we wished to additionally evaluate the extent to which interneurons manifest selectivity for task dimensions in a manner similar to principal cells of the hippocampus. Specifically, we

tested whether interneurons modulate their firing rates for different task dimensions. We have previously reported that principal cells in the CA1 region of the hippocampus with distinct profiles of engagement in theta, beta, low gamma, and high gamma differentially represent task dimensions in their firing rates [3]. As the largest proportion of interneurons exhibited engagement in all four rhythms during correct performance, we examined the information content in the firing rates of this subpopulation for three key task dimensions: odors, positions, and odor–position combinations. Our preliminary results suggest that this interneuron population exhibited low, but significant information content for task dimensions in their firing rates, with increases in information for odors and odor–position conjunctions that reflected the availability of odor stimuli. This suggests that in addition to providing a task-relevant oscillatory framework for the hippocampal network, interneurons may convey information about task dimensions.

## **2.3 Materials and Methods**

All experimental procedures were carried out as previously described in Rangel, Rueckemann, Rivière et al. 2016 [3]. They are briefly described here.

### ***Behavioral Paradigm***

Rats performed a context-guided odor–reward association task in which odors were differentially rewarded depending on the spatial context in which they were encountered. In this task, rats were placed in a behavioral apparatus consisting of two arms, or *contexts*, oriented at 180° with respect to each other and separated by a central chamber (Fig. 2.1). Each context contained two odor ports. On each trial, rats were given access to one of the contexts, which were distinct from each other on the basis of their spatial location within the recording room. To further distinguish each context, vinyl or plastic contextual wraps of different colors and textures were used to cover each arm and face plate containing the odor ports. Rats learned to sample odors presented in each port by poking their snouts in an odor port. LED sensors within each odor port registered nose pokes (*poke onset*). Odors were delivered into the odor port 250 ms after the initiation of a nose poke (*odor onset*). Odors were presented in pairs that were consistent throughout

the experiment. One odor of a pair was assigned as correct in the first context, with the opposite odor assigned as correct in the second context. Odor positions were pseudorandomized and counterbalanced such that the rat had no prior knowledge of which odor port would contain the correct odor of a pair. *Correct trial:* In order to receive a water reward, rats were required to maintain a nose poke for 1500 ms in the odor port containing the correct odor of a pair. If an odor port contained the incorrect odor, the rat needed to remove his nose from the port before 1500 ms had elapsed from poke onset. Following an exit from the incorrect odor port, the rat was allowed to move to the adjacent port containing the correct odor, where he had to maintain a nose poke for 1500 ms to receive his reward. Upon completion of a nose poke in the correct odor port, a water reward was immediately delivered to an indentation in a tray positioned directly under the odor port. *Incorrect trial:* Holding a nose poke for 1500 ms in the odor port containing the unrewarded odor resulted in a white-noise buzz and no water reward on that trial. Rats encountered pairs of odors in blocks, with two blocks of odor pairs per half-session (four odors in a half-session, eight odors in a full session, 96 trials total). Contextual wraps with unique colors and textures covering each arm were switched between half-sessions. Data was analyzed during 1500-ms sampling intervals in correct odor ports (correct trials). Each half-session was analyzed separately.

During behavioral training, rats completed 80 trials each day until they could perform at 75% accuracy. Recordings began after rats achieved this performance criterion. Only those sessions in which the rat achieved at least 75% accuracy were included in the data analysis. Data during 1500-ms sampling intervals in incorrect odor ports (incorrect trials) were excluded from analysis due to an insufficient number of trials for each task dimension.

### ***Neural Recordings and Interneuron Identification.***

We performed high-density extracellular tetrode recordings from the CA1 region of the awake-behaving rat in order to obtain single cell and local field potential (LFP) activity. Signals were amplified 4000–8000 × and digitized at 40 kHz by an Omniplex Neural Acquisition system (Plexon). Local field potentials (LFPs) were digitally isolated with a band-pass filter from 1 to 400 Hz and spikes

were isolated with a band-pass filter from 400 to 8000 Hz. Putative interneurons were isolated according to firing rates and waveform characteristics. Waveform features, such as peak and valley voltage amplitudes and total peak-to-valley distance, were compared across tetrode wires in OfflineSorter (Plexon). Interneurons clustered according to mean firing rate ( $\geq 5$  Hz), mean width at half of the maximum amplitude of the waveform ( $< 150 \mu\text{s}$ ), and mean temporal offset from peak to trough ( $< 350 \mu\text{s}$ ) [14,15,16].

### ***Local Field Potential and Spike-Phase Coherence Analyses.***

A third-order Butterworth filter was used to band-pass filter the LFP in the theta (4–12 Hz), beta (15–35 Hz), low gamma (35–55 Hz), and high gamma (65–90 Hz) frequency ranges. The instantaneous phase was then calculated by taking the arctangent of the complex Hilbert transform of the filtered signal. For interneurons demonstrating spike-phase coherence to all four frequency ranges (theta<sub>4–12 Hz</sub>, beta<sub>15–35 Hz</sub>, low gamma<sub>35–55 Hz</sub>, and high gamma<sub>65–90 Hz</sub>) during correct trials (32 out of 67 total interneurons recorded, 6 rats, 53 half-sessions), the phase of the filtered LFP at the time of each spike was recorded for the 1.5 s odor sampling intervals leading up to reward delivery. Spike-phase relationships were assessed using a Rayleigh statistic, and interneurons were categorized as significantly phase coherent to a rhythm if exhibiting a  $p < 0.05$ . This criterion was previously used to identify interneurons with significant spike-phase coherence that, upon further characterization, also exhibited differences in the magnitude and phase of coherency across correct and incorrect trials [3].

### ***Quantifying Information.***

For interneurons demonstrating spike-phase coherence to all four frequency ranges (theta<sub>4–12 Hz</sub>, beta<sub>15–35 Hz</sub>, low gamma<sub>35–55 Hz</sub>, and high gamma<sub>65–90 Hz</sub>) during correct trials, the information contained in the firing rate of the interneuron for a task dimension (odors, positions, or odor–position conjunctions) was calculated according to the following equation:

$$I = \sum_{i=1}^n P_i \left( \frac{F_i}{F} \right) \log_2 \left( \frac{F_i}{F} \right)$$

Where  $i$  designates a variant of the task dimension (one of four possible odor port positions, four possible odors, or eight possible odor–position combinations as each odor is rewarded in more than one position),  $n$  is the total number of variants,  $P_i$  is the probability of the occurrence of variant  $i$ ,  $F_i$  is the mean firing rate during the occurrence of the variant  $i$ , and  $F$  is the overall mean firing rate of the cell. To determine whether calculated scores could be acquired by chance from the spiking behavior of a given interneuron, task conditions were randomly shuffled 1000 times and the observed information was considered significant if greater than the 95% confidence interval of the condition-shuffled scores.


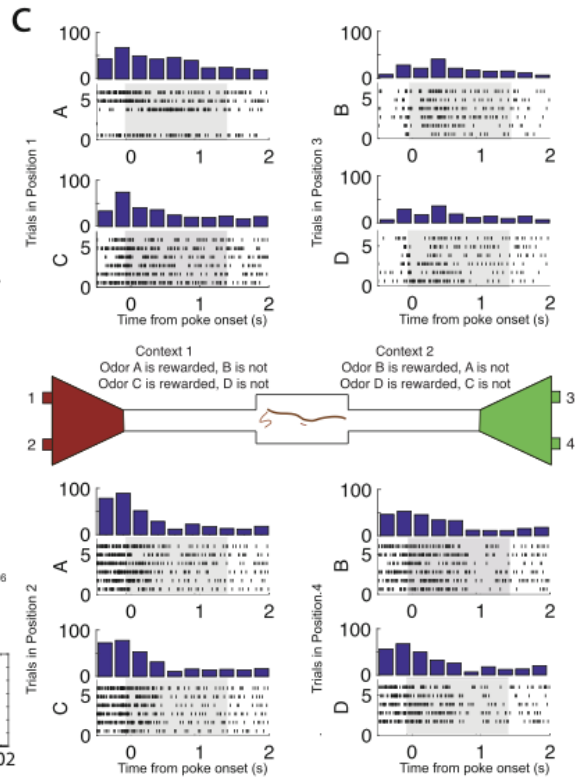
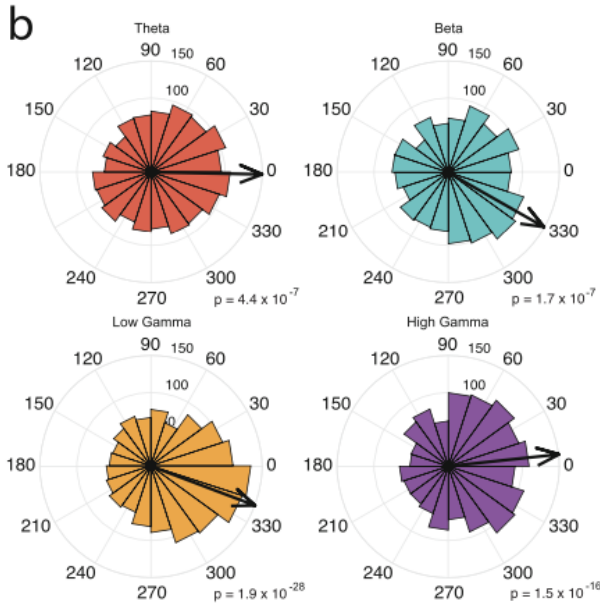

For the interneurons with significant information for a given task dimension, additional information scores were calculated for three non-overlapping 500 ms intervals that directly preceded and spanned the 1.5-s nose poke interval. These intervals included a 750–250-ms interval prior to nose poke onset (*before*), the 500 ms interval after odor onset (*odor*), and the last 500 ms of the nose poke (*end*). Differences in the median information across the three time intervals were assessed using a Friedman’s test, with post hoc pairwise comparisons performed using a Tukey’s Honest Significant Difference test.

## 2.4 Results.

All interneurons with significant spike-phase coherence to theta<sub>4–12 Hz</sub>, beta<sub>15–35 Hz</sub>, low gamma<sub>35–55 Hz</sub>, and high gamma<sub>65–90 Hz</sub> during correct trials (32 out of 67 total interneurons recorded, 6 rats, 53 half-sessions, see Sect. 2) also contained significant information for one or more task dimensions (odor, position, and odor–position) in their firing rates. An example of an interneuron with significant spike-phase coherence to all four rhythms and significant information for all three task dimensions is shown in Fig. 1. The median information for odors was 0.015 bits/spike (N = 21, interquartile range = 0.0194), the median information for position was 0.038 bits/spike (N = 47, interquartile range = 0.0720), and the median information for odor–position was 0.041 bits/spike (N = 51, interquartile range = 0.0609). The median of the average firing rates exhibited by this group of interneurons during correct trials was 24.57 Hz (interquartile range 12.45 Hz).

**Figure 2.1.** Spike-field coherence and firing profiles of a single interneuron in the CA1 region of the rat hippocampus. (a) *Upper:* Gabor spectrogram during the odor sampling interval of a single trial. *Middle:* Corresponding raw local field potential and the band-pass filtered local field potential in the theta (4–12 Hz), beta (15–35 Hz), low gamma (35–55 Hz), and high gamma (65–90 Hz) frequency ranges. *Lower:* Spike times of a single interneuron that demonstrated significant spike-field coherence to each of the four frequency ranges. (b) Circular histograms indicating the phases of theta<sub>4–12 Hz</sub> (orange), beta<sub>15–35 Hz</sub> (blue), low gamma<sub>35–55 Hz</sub> (yellow), and high gamma<sub>65–90 Hz</sub> (purple) at the time of each spike, with the direction of the mean resultant length vector (R) shown by the arrow in black. *P*-values were calculated using a Rayleigh statistic. (c) Spiking activity of the interneuron during the sampling intervals of rewarded odors on correct trials. Each row of tick marks represents spiking during a single trial. Bar graphs above tick marks indicate mean firing rates every 250 ms. The overall mean firing rate of this interneuron was 24.563 Hz. *Center:* Schematic of the context-guided odor–reward association task. Pairs of odors are differentially rewarded depending upon the context in which they are presented. (d) The spiking activity of this interneuron contains significant information (see Sect. 2) for odors, positions, and odor–position combinations. Information scores are indicated above distributions of information scores from shuffled data.

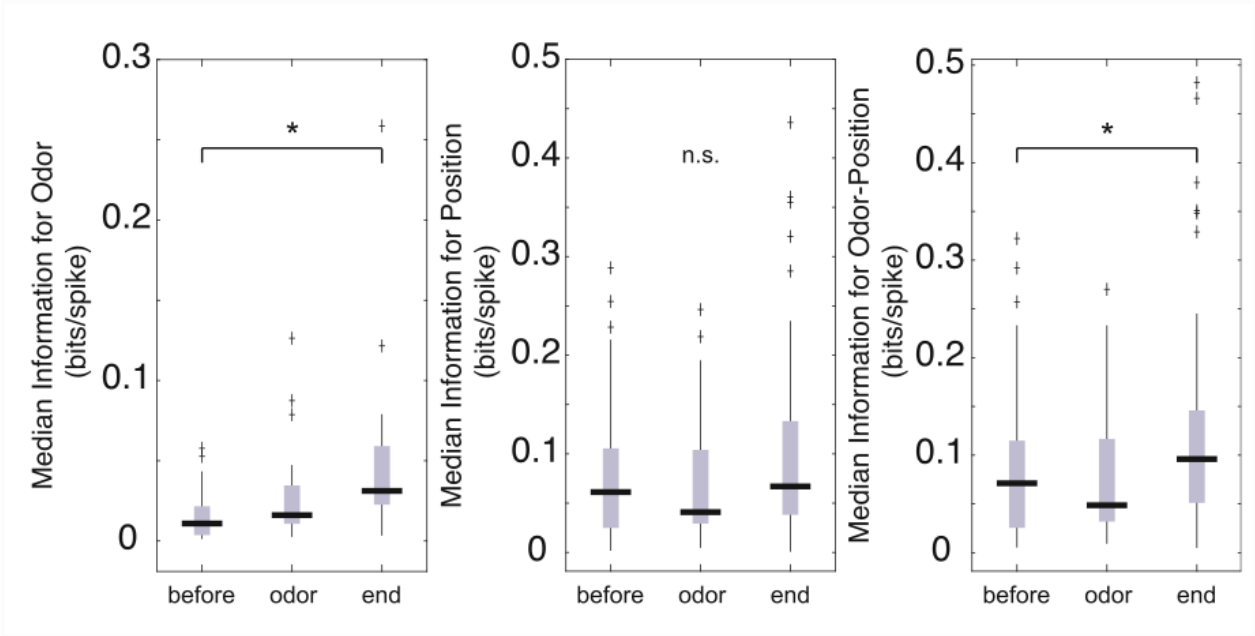




Although there were interneurons with significant spike-phase coherence to combinations of the four rhythms examined (e.g., theta only, theta and high gamma, etc.), few of these interneurons also exhibited significant information for task dimensions. For example, the population of interneurons with significant spike-phase coherence during correct trials to only theta<sub>4-12 Hz</sub> (14 interneurons, 6 rats, 18 half-sessions) produced just three instances of significant information for each task dimension. Given these low numbers, analyses were restricted to the population of interneurons coherent to all four of the rhythms examined during correct trials.

We then tested whether the information for a specific task dimension changed over the course of a trial (Fig. 2.2). We compared the information content during the *before*, *odor*, and *end* intervals. The interneurons coherent to all four rhythms exhibited an increase in odor information across the three intervals examined (Friedman's test: d.f. = 2,  $\chi^2 = 12.67$ ,  $p = 0.0018$ ). Post hoc comparisons revealed that this increase occurred at the end of the nose poke, but not during the interval immediately after odor delivery (Tukey's Honest Significant Difference test,  $p < 0.05$  for comparisons of the before interval to odor and end intervals). These interneurons also exhibited a similar increase in odor-position information across the three intervals examined (Friedman's test: d.f. = 2,  $\chi^2 = 11.76$ ,  $p = 0.0028$ ) that occurred only at the end of the nose poke (Tukey's Honest Significant Difference test,  $p < 0.05$ ). This interneuron group did not exhibit increases in position information across the three intervals examined (Friedman's test: d.f. = 2,  $\chi^2 = 3.87$ ,  $p = 0.1443$ ).

**Figure 2.2.** Information for odors and odor-positions evolves over the course of odor sampling. *Left:* Median information (bits/spike) for odors for interneurons exhibiting significant spike-phase coherence relationships to all four rhythms examined (theta<sub>4–12 Hz</sub>, beta<sub>15–35 Hz</sub>, low gamma<sub>35–55 Hz</sub>, and high gamma<sub>65–90 Hz</sub>) during a 500-ms interval prior to nose poke (*before*), a 500-ms interval directly after odor delivery (*odor*), and 500 ms prior to the end of the nose poke (*end*). Vertical gray bars indicate the interquartile range. The top vertical line indicates  $q3 + 1.5 \times (q3 - q1)$  and the bottom vertical line indicates  $q1 - 1.5 \times (q3 - q1)$ , where  $q1$  and  $q3$  are the 25th and 75th percentiles, respectively. Asterisks (\*) indicate a significant pairwise comparison using a Tukey’s Honest Significant Difference test,  $p < 0.05$ . *Middle:* Same as in A, for position information. *Right:* Same as in B, for odor–position information.



## 2.5 Discussion.

Our preliminary results suggest that interneurons may convey information about task dimensions in their firing rates. The low, but significant information for task dimensions in the interneuron population indicates that they have potentially subtle, but reliable changes in firing rate that account for changing task conditions. Notably, the median information for each dimension is roughly an order of magnitude lower than previously reported medians derived from the firing rates of the pyramidal cell population [3]. To illustrate the types of firing rate changes that might occur to produce information scores of this magnitude, we provide the following extreme case: an interneuron with a change in firing rate at only one of four positions, a mean firing rate of 24.5741 Hz (the median for this group of interneurons), and a position information value of 0.038 bits/spike (the median for this group of interneurons), would need to exhibit a firing rate increase of approximately 14 Hz. The extent to which this degree of selectivity impacts the hippocampal network remains to be determined. Although the interneurons generally have low information values, more analyses must be applied to better characterize the range and reliability of the firing rate changes driving the significant information scores in this study [8, 17].

Interneurons demonstrated an increase in firing rate selectivity over the course of the nose poke interval that reflected the availability of information about task dimensions. Specifically, the relatively stable degree of information for position before and during the odor sampling interval is consistent with the early availability of position information prior to the nose poke. In contrast, information about odors became available 250 ms after the initiation of a nose poke, and the interneurons demonstrate increases in information for these dimensions only at the end of the odor sampling interval. These results suggest that the selectivity of interneurons' firing rates for these dimensions is a product of task-relevant engagement.

A number of future analyses could enhance the interpretive power of these preliminary results. For instance, future characterizations of this data could employ information measures that are more sensitive to the direction of changes (increases or decreases) in firing rate demonstrated by interneurons over the course of a trial. In addition, it would be informative to quantify information during a number of additional intervals during the task, including incorrect trials, correct rejection trials, and reward

consumption intervals. Quantifying information during these additional intervals would help assess the degree to which the observed increases in information are related to a reward. Further exploration will facilitate the development of more nuanced and comprehensive hypotheses as to the mechanisms through which interneurons contribute to successful information processing in the hippocampus.

Leading hypotheses in the field have proposed that interneuron activity crucially shapes the oscillatory profile of the hippocampus, providing a temporal scaffolding within which pyramidal neurons can receive [18, 19], process, and successfully transmit behaviorally relevant information [20, 21]. Support for this view stems in part from the organization of the hippocampal network, where the locally projecting interneuron population densely innervates large numbers of principal cells [22, 23]. This widespread innervation, compounded with high firing rates [24], amplifies the impact of interneurons' inhibitory currents and places interneurons in a unique position to coordinate the simultaneous activity of large ensembles of pyramidal cells. In particular, periodically occurring inhibitory currents create alternating windows of suppression and relative excitability in the principal cell population, ensuring that only precisely timed inputs are able to elicit principal cell responses [21]. In this way, interneuron activity has the ability to simultaneously sculpt the oscillatory profile of the hippocampus while constraining the subset of inputs to which principal cells can respond [25]. The observed interneuron selectivity for task dimensions suggests that interneurons are capable of shaping the overall information content of the network through both the temporal coordination of principal cell excitability and reliable changes in firing rate.

The mechanisms whereby interneurons acquire the observed selectivity for task dimensions remain obscure. Here, we briefly consider two hypotheses for the emergence of this information content. On the one hand, interneurons might inherit information directly from the afferents that simultaneously recruit subsets of principal cells in a feedforward inhibitory manner. It is also possible that the observed firing rate selectivity for task dimensions in the interneuron population emerges directly from their interactions with various ensembles of pyramidal cells, each containing highly selective information for distinct task dimensions. Our finding contributes to a growing body of evidence that interneurons are

active facilitators of task-relevant processing through at least two forms of selective engagement: reorganization of spike timing into multiple rhythms and the conveyance of information about task dimensions through changes in firing rates [[14](#), [17](#), [26](#), [27](#)].

## 2.6 References.

1. J. O'Keefe, J. Dostrovsky, The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res.* 34(1), 171–175 (1971)
2. E.R. Wood, P.A. Dudchenko, H. Eichenbaum, The global record of memory in hippocampal neuronal activity. *Nature* 397(6720), 613–616 (1999)
3. L.M. Rangel, J.W. Rueckemann, P.D. Rivière, K.R. Keefe, B.S. Porter, I.S. Heimbuch, C.H. Budlong, H. Eichenbaum, Rhythmic coordination of hippocampal neurons during associative memory processing. *eLife* 5, e09849 (2016)
4. R.W. Komorowski, J.R. Manns, H. Eichenbaum, Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* 29(31), 9918–9929 (2009)
5. K.M. Igarashi, L. Lu, L.L. Colgin, M.-B. Moser, E.I. Moser, Coordination of entorhinal-hippocampal ensemble activity during associative learning. *Nature* 510(7503), 143–147 (2014)
6. W. Skaggs, B. McNaughton, K. Gothard, E.J. Markus, An information-theoretic approach to deciphering the hippocampal code. *Proc. IEEE* 1990, 1030–1037 (1993)

7. E.J. Markus, C.A. Barnes, B.L. McNaughton, V.L. Gladden, W.E. Skaggs, Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. *Hippocampus* 4(4), 410–421 (1994)
8. A.V. Olypher, P. Lánský, R.U. Muller, A.A. Fenton, Quantifying location-specific information in the discharge of rat hippocampal place cells. *J. Neurosci. Methods* 127(2), 123–135 (2003)
9. M.A. Wilson, B.L. McNaughton, Dynamics of the hippocampal ensemble code for space. *Science* 261(5124), 1055–1058 (1993)
10. K. Harris, J. Csicsvari, H. Hirase, G. Dragoi, G. Buzsáki, Organization of cell assemblies in the hippocampus. *Nature* 424(July), 552–556 (2003)
11. J.F. Guzowski, J.J. Knierim, E.I. Moser, Ensemble dynamics of hippocampal regions CA3 and CA1. *Neuron* 44(4), 581–584 (2004)
12. G. Buzsáki, Neural syntax : cell assemblies, synapsembles, and readers. *Neuron* 68, 362–385 (2010)
13. S. McKenzie, A.J. Frank, N.R. Kinsky, B. Porter, P.D. Rivière, H. Eichenbaum, Hippocampal representation of related and opposing memories develop within distinct, hierarchically organized neural schemas. *Neuron* 83(1), 202–215 (2014)
14. E. Stark, R. Eichler, L. Roux, S. Fujisawa, H.G. Rotstein, G. Buzsáki, Inhibition-induced theta resonance in cortical circuits. *Neuron* 80(5), 1263–1276 (2013)
15. J. Csicsvari, B. Jamieson, K.D. Wise, G. Buzsáki, Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron* 37(2), 311–322 (2003)
16. P. Barthó, H. Hirase, L. Monconduit, M. Zugaro, K.D. Harris, G. Buzsáki, Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* 92(1), 600–608 (2004)
17. W.B. Wilent, D.A. Nitz, Discrete place fields of hippocampal formation interneurons. *J. Neurophysiol.* 97(6), 4152–4161 (2007)



18. F. Pouille, M. Scanziani, Enforcement of temporal fidelity in pyramidal cells by somatic feed-forward inhibition. *Science* 293(5532), 1159–1163 (2001)
19. F. Pouille, M. Scanziani, Routing of spike series by dynamic circuits in the hippocampus. *Nature* 429(6993), 717–723 (2004)
20. G. Buzsáki, J.J. Chrobak, Temporal structure in spatially organized neuronal ensembles: a role for interneuronal networks. *Curr. Opin. Neurobiol.* 5(4), 504–510 (1995)
21. J. Cannon, M.M. McCarthy, S. Lee, J. Lee, C. Borgers, M.A. Whittington, N. Kopell, Neurosystems: brain rhythms and cognitive processing. *Eur. J. Neurol.* 39(5), 705–719 (2014)
22. T.F. Freund, G. Buzsáki, Interneurons of the hippocampus. *Hippocampus* 6(4), 347–470 (1996)
23. A. Sik, M. Penttonen, A. Ylinen, G. Buzsáki, Hippocampal CA1 interneurons: an in vivo intracellular labeling study. *J. Neurosci. Off. J. Soc. Neurosci.* 15, 6651–6665 (1995)
24. S. Fox, J. Ranck, Electrophysiological characteristics of hippocampal complex-spike cells and theta cells. *Exp. Brain Res.* 41(3), 399–410 (1981)
25. M.A. Whittington, R.D. Traub, N. Kopell, B. Ermentrout, E.H. Buhl, Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *Int. J. Psychophysiol.* 38(3), 315–336 (2000)
26. D. Nitz, B. McNaughton, Differential modulation of CA1 and dentate gyrus interneurons during exploration of novel environments. *J. Neurophysiol.* 91, 863–872 (2004)
27. S. Royer, B.V. Zemelman, A. Losonczy, J. Kim, F. Chance, J.C. Magee, G. Buzsáki, Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* 15(5), 769–775 (2012)

Chapter 2, in full, is a reprint of the material as it appears in *Advances in the Mathematical Sciences: AWM Research Symposium*, Los Angeles, CA, April 2017 (pp. 161-171). Springer International Publishing. Rivière, Pamela D; Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.

## CHAPTER 3: MODELING RELATIONSHIPS BETWEEN RHYTHMIC PROCESSES AND NEURONAL SPIKE TIMING

### **3.1 Abstract.**

Neurons are embedded in complex networks, where they participate in repetitive, coordinated interactions with other neurons. Neuronal spike timing is thus predictably constrained by a range of ionic currents that shape activity at both short (milliseconds) and longer (tens to hundreds of milliseconds) timescales, but we lack analytical tools to rigorously identify these relationships. Here, we innovate a modeling approach to test the relationship between oscillations in the local field potential (LFP) and neuronal spike timing. We use kernel density estimation to relate single neuron spike timing and the phase of LFP rhythms (in simulated and hippocampal CA1 neuronal spike trains). We then combine phase and short (3 ms) spike history information within a logistic regression framework (“phaseSH models”), and show that models that leverage refractory constraints and oscillatory phase information can effectively test whether—and the degree to which—rhythmic currents (as measured from the LFP) reliably explain variance in neuronal spike trains. This approach allows researchers to systematically test the relationship between oscillatory activity and neuronal spiking dynamics as they unfold over time and as they shift to adapt to distinct behavioral conditions.

### **3.2 Introduction.**

The neural code crucially depends on spike timing. In the brain, variations in spiking patterns over time convey information about the external world (1), instigate (and register) changes within the body (2), and engage distinct behavioral states (2-4). A major aim in neuroscience is the quantification of temporal regularity in these spike trains and the identification of its underlying generators. Establishing relationships between temporal patterns of spiking activity and their neural drivers is challenging, however, and requires methods that can 1) describe temporal fluctuations and 2) evaluate possible neural mechanisms responsible for these fluctuations. Individual neurons in the central nervous system are

embedded in complex networks, where they receive inputs from a variety of both distally located and neighboring neurons. Individual neurons additionally generate spontaneous intrinsic ionic currents that interact with these input sources. Currently, few quantitative methods are designed to systematically assess the contribution of specific rhythmic currents in shaping single neuron spike timing. Here, we provide a powerful statistical modeling approach that allows us to weigh the relative contributions of refractory and rhythmic current influences on neuronal spike timing, predict moment-to-moment variations, quantify the uncertainty of our predictions, and assess the goodness-of-fit of our models.

Neuronal activity obeys constraints at multiple timescales. Action potentials (spikes) are large-amplitude, short-lived (<1 ms) neuronal membrane voltage deflections. These are followed by transient periods of hyperpolarization, during which neurons are less likely or unable to emit further spikes (5-7). At longer timescales, regularly-occurring ionic currents create phasic temporal windows during which an individual neuron is more or less primed to produce spikes (8,9). These currents can arise spontaneously in the absence of synaptic inputs (10,11), or they may result from repetitive, coordinated network interactions. Notably, these longer-timescale ebbs and flows in subthreshold membrane potentials, whether intrinsically or synaptically generated, superimpose extracellularly to produce neural oscillations in the local field potential (LFP) (12,13). Oscillations in the LFP thus reflect a mixture of intrinsic and synaptic currents shaping single neuron activity, but the degree to which these combined currents influence spike timing remains uncertain.

In the hippocampus, a brain region extensively studied in the context of learning and memory, pyramidal cell and inhibitory interneuron spiking is often rhythmically organized in the theta (4-12 Hz) and low gamma (15-35 Hz) frequency ranges (14-17), and evidence suggests that interneurons powerfully contribute to periodicity in the hippocampal LFP (18,19). Rhythmic temporal structure in interneuron spike trains arises from both neuronal resonance (i.e. intrinsic cellular properties that give rise to preferential responses at specific frequency ranges and spontaneous membrane oscillations, (20,21)) as well as interactions with a range of intra-hippocampal (22,23), cortical (22-25), and subcortical (26-28) inputs. Thus, hippocampal interneuron spiking is temporally constrained at both short timescales (i.e.

milliseconds) due to refractoriness, and at longer theta and gamma rhythmic timescales (i.e. tens to hundreds of milliseconds) due to a combination of intrinsic and synaptic current fluctuations. Whereas the influence of the after-hyperpolarization on neuronal spiking is well-known, and the temporal structure of this influence has been characterized via statistical models (29,30), we currently lack the tools to evaluate the effect that longer-duration rhythmic constraints exert on spike timing. Moreover, there are few existing methods for assessing short and long timescale oscillatory influences together to provide a holistic view of the temporal constraints that can shape neuronal spiking. Here, we deploy a series of statistical models that leverage spike history information as well as LFP rhythms to model the rhythmic temporal structure of both simulated and hippocampal inhibitory interneuron spike trains.

Our approach is grounded in the use of generalized linear models (GLMs) and maximum likelihood estimation (31), which have been used extensively to account for the absolute and relative refractory periods following a spike occurrence (29,30,32,33). Typically, these models use logistic regression to predict the probability of spiking from the preceding spike history. Following previous work (29,30,32,33), we constructed a short history model (“SH model”), which predicted upcoming spiking according to spike history in the preceding 3 ms. Beyond capturing the absolute after-hyperpolarization constraint, we also modeled the relationship between spiking and a longer-duration history (250 ms, “LH model”) to detect patterns in spike emissions that cycled along hundreds of milliseconds at most. This LH model allowed us to flexibly capture structure in spike trains at exquisitely fine temporal resolution without prior assumptions about specific factors that contributed to the structure (e.g. refractoriness, spike-rate adaptation, rhythmicity at various frequency ranges). In this way, LH models behaved much like sliding-window autocorrelations or spectral analysis of spike trains (34), but applied in a predictive (rather than exclusively descriptive) setting. Despite their descriptive and predictive abilities, LH models do not by themselves disambiguate the various possible neural mechanisms underlying temporal structure in spike trains.

To explicitly test hypotheses regarding the presence of rhythmicity in spike trains, we developed a model that could represent rhythmic currents as measured from the LFP, and also quantify their

relationship to neuronal spike trains. In the work that follows, we present strategies for (1) representing spiking relationships to the phase of a particular LFP rhythm using Kernel Density Estimation (KDE), and (2) combining LFP phase information and neuronal refractory constraints on spiking. First, we used KDE to directly model the likelihood of LFP phases given observed spiking, and subsequently derived the posterior probability of spiking over time given a phase time series. Since phase models could not capture short timescale refractory constraints on spiking, and SH models were insensitive to long timescale structure, we innovatively combined phase and short history information in a logistic regression framework (“phaseSH models”). These comprehensive models were able to effectively capture and specify short (refractory) and long (rhythmic) timescale influences in simulated spike trains and hippocampal CA1 neurons, which addressed the major limitations of each model independently. Moreover, since our models specifically assess the probability of spiking relative to phase, they provide information about the reliability of these relationships over time that cannot be obtained from traditional measures of spike-phase relationship based on circular averages (e.g. the Rayleigh statistic). We additionally compared phaseSH model performance against that of phaseLH models (GLMs containing phase and long history as predictors), to test whether including long history provided additional information to better predict held-out spike trains. In our dataset, most spike trains were better predicted by the phaseSH models, suggesting that long history did not add information above and beyond phaseSH models in these cases. The model comparisons described in this work provide a new paradigm for rigorously assessing the contribution of rhythmic processes to single neuron spiking dynamics.

### **3.3 Results.**

#### ***Modeling temporal structure at short and long timescales***

To characterize the temporal structure of single neuron spiking activity, we constructed five distinct models that exhibited sensitivity to spiking fluctuations at either short or long timescales, or a combination of the two. Here, short timescales refer to millisecond-resolution variations in spiking driven by after-hyperpolarization currents, while long timescales refer to a variety of potential mechanisms that

could explain variations in spiking along tens to hundreds of milliseconds, including rhythmic synaptic currents. We applied these models to both simulated and electrophysiologically recorded datasets, with the recorded data drawn from previously published work that included CA1 pyramidal cells and inhibitory interneurons (35). This dataset was recorded during an associative memory paradigm, where rats learned that odors in a pair were differentially rewarded depending on the context in which they were encountered. To receive a reward, rats maintained their position at the odor port containing the correct odor of a pair for 1250 ms after delivery of the odor stimulus (see Materials and Methods). Our simulated data match the organizational structure of the CA1 recordings, with individual trials lasting 1250 ms, and each trial preceded by an additional 250 ms of spiking history for the early time samples in the trial (Fig. 1A). For the present work, all model parameters (for both simulated and recorded datasets) were learned with a dataset of training trials (50% of trials in the session), while the remaining trials were held out for performance evaluation.

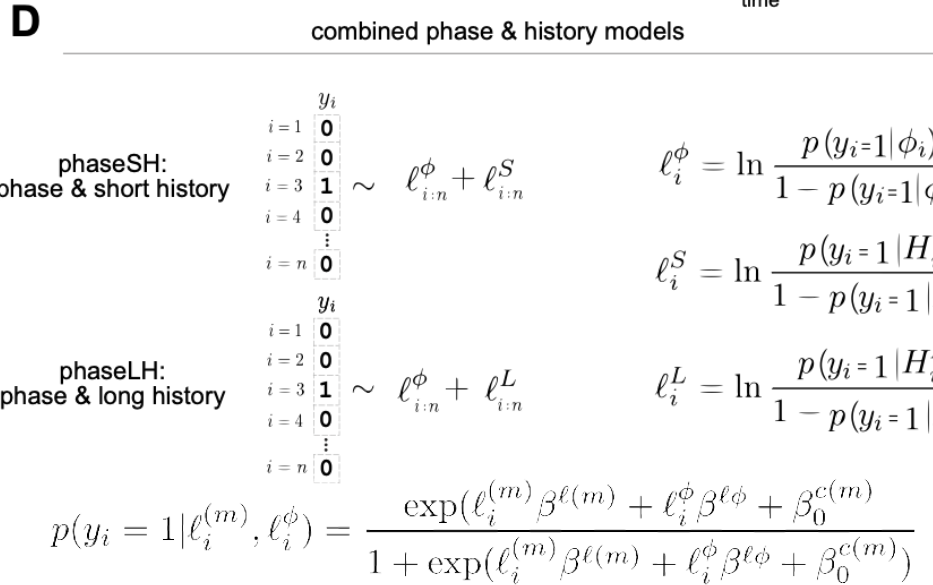
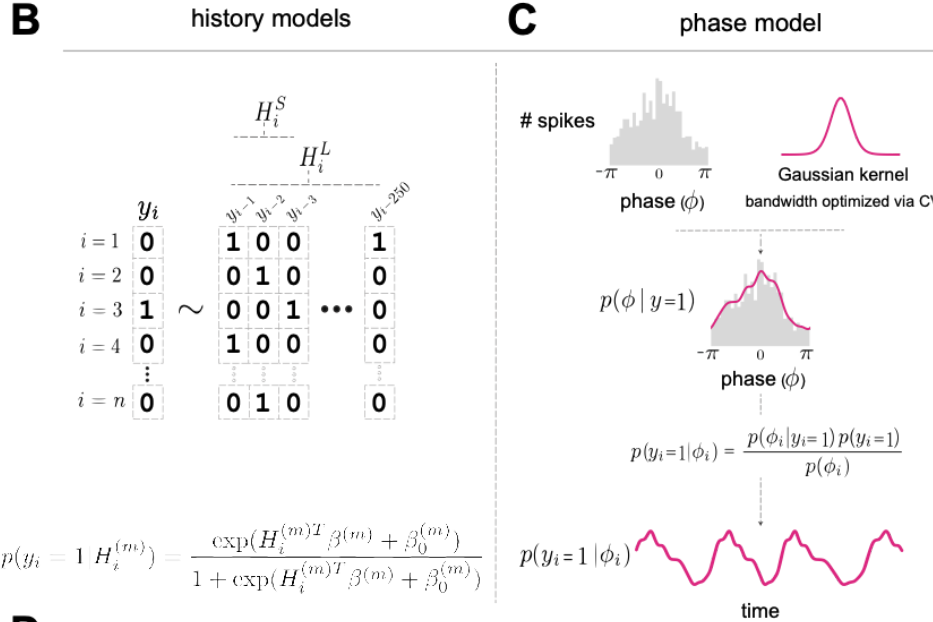

We first modeled short timescale constraints on neuronal spike timing imposed by absolute and relative refractory periods. To do this, we constructed a logistic regression model whose predictions relied solely on the short history (3 ms) of the spike train (“SH model”, Fig. 3.1B). We additionally constructed a long history model (“LH model”, 250 ms) that captured longer timescale constraints on neuronal spiking activity, including (but not limited to) modulation by oscillatory currents. This 250 ms period corresponded to the period of a 4 Hz rhythm, ensuring that the LH model could capture temporal modulation of spiking activity at frequencies greater than or equal to 4 Hz (Fig. 3.1B). With 250 parameters to learn, we anticipated that LH models would be vulnerable to fitting noise in the training set. We nevertheless sacrificed predictive generalizability in this model in order to (1) maximize the model’s ability to discover temporal structure at fine (and multiple) timescales, and (2) make predictions with fine temporal resolution. For instance, any neuron in the dataset that did not have a true relationship to the phase of the theta rhythm could still produce an LH model capable of capturing shorter timescale rhythmic fluctuations (e.g. low gamma, 35-55 Hz), neuronal adaptation dynamics, rebound spiking, as well as very fast membrane constraints (e.g. 1 ms refractory periods). Alternate models of long history

have opted to bin temporal epochs or smooth spike times by applying various temporal filters, effectively averaging spike numbers over longer periods of time to reduce overfitting at the expense of temporal resolution (42). Windowing methods impose a set duration for relative refractory periods, which could vary widely across neurons, and also limit the fastest frequency modulation they can identify. Allowing millisecond-to-millisecond variation in our LH model increases the size of the family of functions we can fit to the data. To strike a balance between this model’s flexibility and its ability to generalize to new data (i.e., robustness), we imposed an L2 penalty on the LH model’s regression weights (44), which reduced noise in coefficient estimates while allowing the model to discover, at fine temporal resolution, any systematic relationships present in the data regardless of the underlying generative mechanism.

To specifically test the hypothesis that CA1 neuronal spike timing was systematically related to the phase of an ongoing oscillation, we first estimated the instantaneous phase of this oscillation at all time samples, and selected phase estimates that co-occurred with spiking events. We subsequently estimated the distribution of phases given a spike occurrence using a kernel density estimator with a Gaussian kernel whose bandwidth was optimized via 5-fold cross validation, as described in the methods. We then computed the probability of observing a spike  $y_i = 1$  given the instantaneous phase  $\phi$  at every time sample  $i$ , or  $p(y_i = 1|\phi_i)$ , by applying Bayes’ transform to the kernel density estimator (Fig. 3.1C). Finally, to produce models that could capture temporal constraints on neural spiking activity at both short and long timescales, we constructed two logistic regression models that combined phase and history information: one containing phase and short history (“phaseSH models”), and another containing phase and long history (“phaseLH models”, Fig. 3.1D, see Materials and Methods).

**Figure 3.1. Data structure and model specification.** **(A)** Data are from a previously published study, in which rats successfully associated odors with a rewarded context. All analyzed intervals correspond to 1250 ms epochs (spikes in dark grey) with a preceding 250 ms serving solely as history regressors for the first time samples of the behavioral epochs (spikes in light grey). **(B)** Long (LH) and short history (SH) models, generated via logistic regression. The target consists of a neuron spike (or the lack thereof),  $y_i$ , for each time sample  $i$ . Predictors are either the 3ms (short history,  $H_i^S$ ) or the 250ms (long history,  $H_i^L$ ) preceding the current time sample. LH model coefficient estimates are uniquely subject to L2 regularization. **(C)** Phase model, generated via kernel density estimation (KDE). We perform cross validation to optimize the bandwidth of a Gaussian kernel. The resulting estimator represents the probability of a phase given that a spike has occurred, and this estimator can then be used to predict the probability of spiking over time, as a function of a phase time series. **(D)** Models combining a phase and history predictor. Each of these predictors is represented as the log-odds of the phase-based probabilities of spiking,  $l^\phi$ , and the log odds of history-based probabilities of spiking,  $l^{(m)}$ , respectively.





### ***Model performance on simulated data***

We first validated the ability of each model to capture temporal structure at short and long timescales in simulated spiking data. We generated four “types” of spike trains, where spiking obeyed ground truth probabilities that incorporated a combination of short timescale (refractoriness) and/or long timescale (rhythmic modulation) constraints on spiking: termed hereafter *atemporal*, *refractory non-rhythmic*, *non-refractory rhythmic*, and *refractory rhythmic* spike train types (Fig. 3.2A). A total of 50 simulated spike trains were created for each ground truth. The average target firing rate of these neuronal spike trains (6 Hz) approximated the average firing rate of a large proportion of CA1 inhibitory interneurons recorded during an associative memory paradigm ([Supplemental Fig. S3.1](#)).

Each model produced systematic predictions (and deviations from the ground truth) that varied according to simulated spike train type (Fig. 3.2A). In order to quantify deviations between model predictions and the ground truth of each type of spike train, we first leveraged an information-theoretic approach. This consisted of computing the average Kullback-Leibler (KL) divergence between the ground truth spiking probabilities and the empirically derived, model-based probabilities (see Materials and Methods: *Data Simulations*). Small ( $\sim 0$ ) KL divergences between the ground truth and model-based probabilities indicated little to no deviation from the ground truth, and thus demonstrated an ability to fit the temporal structure in the spike train. Larger KL divergences, on the other hand, indicated systematic deviations from the ground truth and demonstrated an inability to capture essential sources of variance in the spike train.

SH models incurred the smallest and least variable KL divergences for simulated spike trains exhibiting a refractory constraint but no phase modulation (*refractory non-rhythmic* spike trains), and produced systematically large KL divergences from the ground truth whenever spiking was rhythmically modulated (*non-refractory rhythmic* and *refractory rhythmic*; Fig. 3.2B, *middle and far right*). In cases where simulated spike trains did not have a refractory period, SH models predicted increased firing within milliseconds following a spike in the history (*atemporal* and *non-refractory rhythmic*; Fig. 3.2A, *second row, far left and middle right panels*). SH models thus excelled at capturing variance in spike trains

containing refractory constraints but no rhythmic phase modulation and, as expected, dramatically failed to capture variance in spike trains modulated by rhythmic phase.

Phase model predictions flexibly adapted to each simulated spike train's ground truth, capturing rhythmic phase modulation when it existed (*non-refractory rhythmic*, and *refractory rhythmic*), and basing predictions on the average spike rate when it did not exist (*atemporal*, and *refractory non-rhythmic*). These properties allowed the phase model to cleanly approximate the ground truth probabilities in the *atemporal* and *non-refractory rhythmic* spike trains, with very subtle deviations from the flat ground truth in the former case. The phase model incurred slightly larger KL divergences from the ground truth whenever the latter incorporated a refractory constraint (*refractory non-rhythmic*, and *refractory rhythmic*), but these KL divergences were relatively infrequent given the low firing rate of these simulated spike trains (Fig. 3.2B). At faster spike rates, however, phase models dramatically diverged from the ground truth for spike trains governed by refractory constraints (Supplemental Fig. S3.2A, *middle left* and *far right*). Phase models were thus adept at capturing variance in rhythmically modulated spike trains, but suffered from their inability to account for after-hyperpolarizations.

LH models were susceptible to spurious patterns in data whose ground truth was either completely uniform (*atemporal*) or only contained refractory constraints (*refractory non-rhythmic*), despite reducing noisier coefficient estimates through L2-regularization. These models consequently suffered KL divergences from the ground truth that were relatively larger than those incurred by phase and SH models for the non-rhythmic simulated spike trains (Fig. 3.2B, *far and middle left*). This model was, however, able to capture rhythmic temporal structure in spike trains whose ground truths incorporated phase modulation, albeit noisily (*non-refractory rhythmic*, *refractory rhythmic*; Fig. 3.2A). LH models, as formulated here, were not ideal for capturing rhythmic constraints on spiking.

PhaseSH models closely matched SH model predictions for simulated spike trains that had no phase modulation, performing slightly worse than SH model predictions (*atemporal*, and *refractory non-rhythmic*; Fig. 3.2B, *far and middle left*). This was likely a consequence of including the phase component which, while close to uniform, was not quite flat (Fig. 3.2A, *left boxed panel*, *fifth row of*

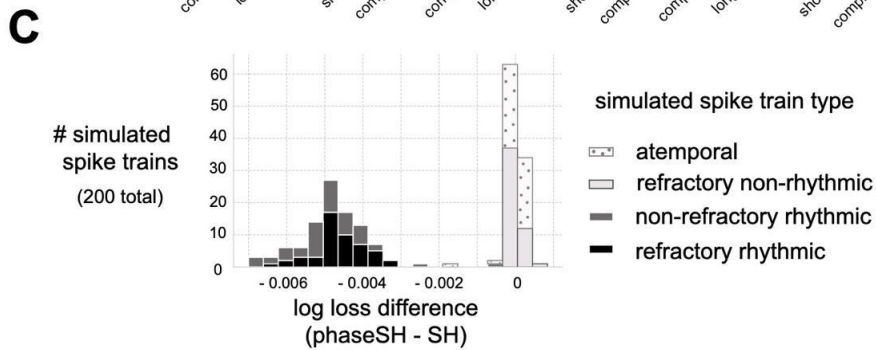
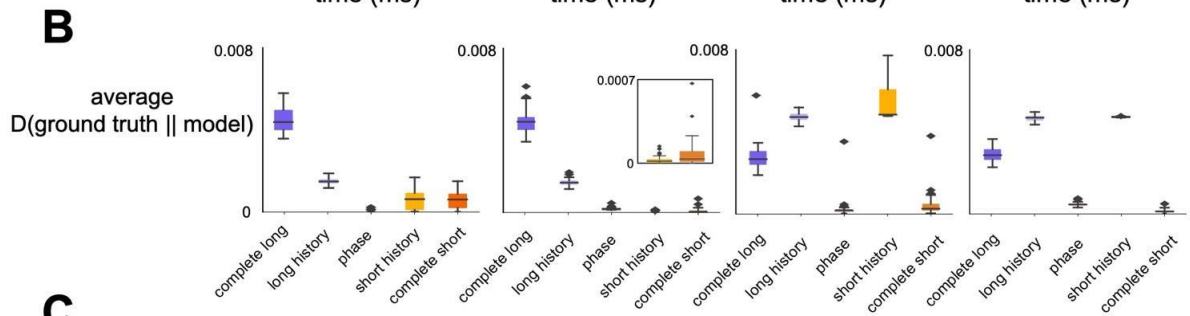
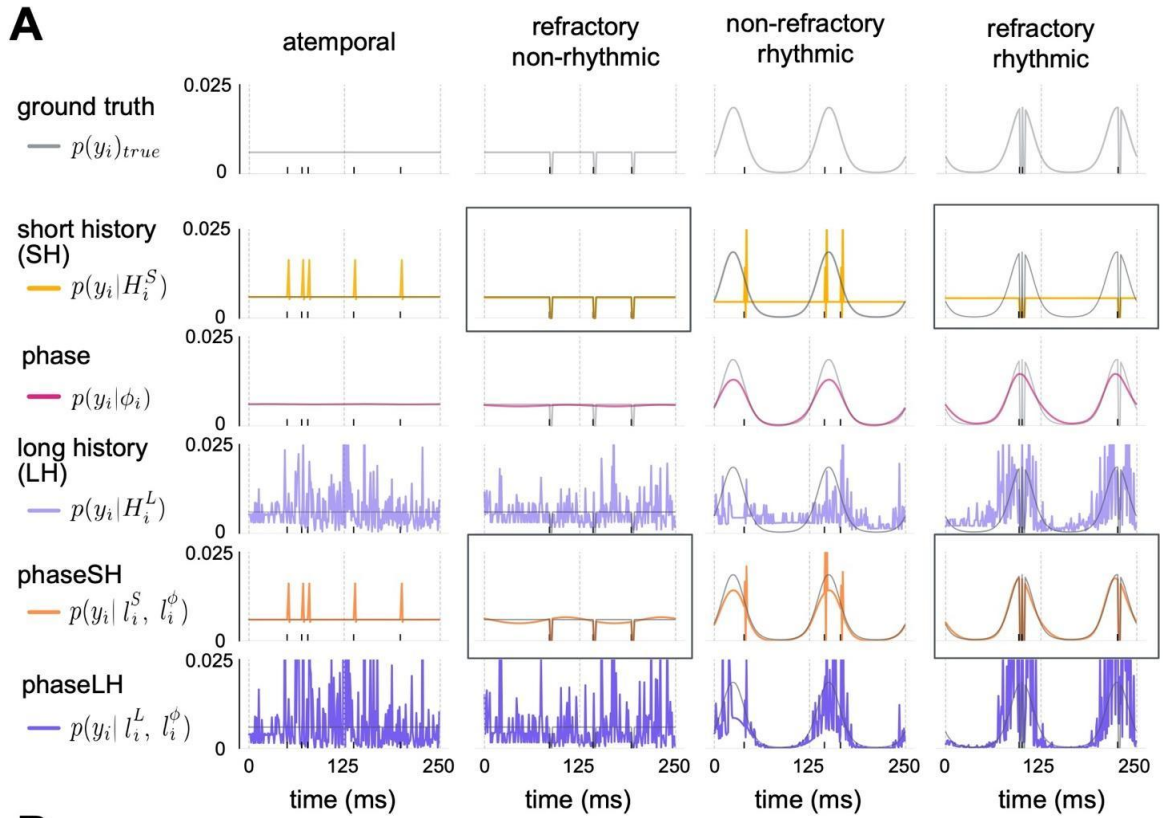
*models*). As with the SH models, in cases where there was no refractory constraint, the models tended to overestimate the probability of spiking immediately following a spike (*atemporal*, and *non-refractory rhythmic*; Fig. 3.2B, *far left, middle right*). PhaseSH models best approximated the ground truth probabilities for spike trains that exhibited both refractory and rhythmic constraints (*refractory rhythmic*; Fig. 3.2B, *far right*). PhaseSH models were thus well suited for capturing rhythmic constraints when they existed, readily captured variation in spike trains for simulated spike trains with both slow (~6 Hz) and fast (~40-60 Hz) average firing rates, (Supplemental Fig. S3.2A), and approximated SH model predictions when rhythmic constraints were not present in the data.

Lastly, phaseLH model predictions behaved similarly to LH models, capturing spurious patterns in *atemporal* and *refractory non-rhythmic* spike trains (Fig. 3.2B, *far and middle left*), but could approximate oscillatory spike distributions over time due to the combined contributions of LH and phase models (*non-refractory rhythmic*, and *refractory rhythmic*; Fig. 3.2A, *sixth row*). Noisy estimates, however, resulted in large and frequent KL divergences from the ground truth (Fig. 3.2B). In simulations, phaseLH models were thus less adept than phaseSH models at accurately describing rhythmic constraints in spike trains.

Using the phaseSH models, we hoped to identify neurons whose spike trains were specifically influenced by rhythmic currents and were thus systematically related to the phase of an LFP oscillation. Given the phase model's ability to produce reasonably smooth and accurate predictions even in cases where simulated spike trains did not exhibit phase modulation, we were not able to use phase model performance alone to identify the neurons we sought. Instead, we first noted that disparities in predictions between the phaseSH model and the SH model should lead to differences in likelihood-based measures of goodness-of-fit (42). We consequently took a model comparison approach and computed the log losses of phaseSH and SH models (evaluated on held out trial data, see Materials and Methods: *Evaluating Model Performance and Goodness-of-Fit*). We then subtracted these log losses to produce the log loss difference between models (see Materials and Methods: *Identifying Neurons Exhibiting Spike-Phase Relationships*). Log loss differences closer to zero resulted from similar performance between phaseSH and SH models,

indicating that oscillatory phase failed to provide information above and beyond what was already encoded in SH model predictions. This was precisely the pattern observed for *atemporal* and *refractory non-rhythmic* simulated spike trains, whose temporal structure was unrelated to oscillatory phase (Fig. 3.2C). In contrast, log loss differences much smaller than zero emerged when the phaseSH model's predictions held an advantage over SH model predictions, indicating that oscillatory phase information improved the likelihood-based goodness-of-fit. A similar pattern of results emerged when using KL divergences to compare phaseSH and SH models ([Supplemental Fig. S3.3](#)). Notably, this was exactly the result obtained for *non-refractory rhythmic* and *refractory rhythmic* simulated spike trains, which were systematically related to oscillatory phase (Fig. 3.2C, [Supplemental Fig. S3.2B](#)). Using likelihood-based model comparisons between the phaseSH and SH models, we were able to successfully recover the simulated spike trains whose spike probability was organized according to oscillatory phase.

**Fig. 3.2. Model results in a simulated dataset. (A)** Panels depict two cycles (250 ms) of held out spiking (black ticks) for individual simulated spike trains and their corresponding ground truth probabilities (grey, all rows), or probabilities generated from short history (yellow, 2<sup>nd</sup> row), phase (pink, 3<sup>rd</sup> row), long history (light purple, 4<sup>th</sup> row), phaseSH (short history and phase, orange, 5<sup>th</sup> row), and phaseLH (long history and phase, dark purple, 6<sup>th</sup> row) models for four simulated spike trains corresponding to different “types,” where a “type” consists of a unique combination of rhythmic and refractory constraints. LH and phaseLH model heights are truncated to standardize axes across panels. Ground truth probabilities are replicated on each panel (within column) to facilitate visual comparisons across true and empirical model probabilities. Some model probabilities lie exactly along the ground truth for the entire time series (such as the phase model for the atemporal neuron). Vertical dashed lines mark the edges of two cycles of data. Boxed probability distributions highlight the qualitative performance of SH models and phaseSH models on neurons with either only refractory constraints, or both refractory and rhythmic constraints. **(B)** Average Kullback-Leibler (KL) divergences of the ground truth from a model’s probabilities across 50 simulated spike trains per type. Inset for the refractory non-rhythmic spike train type shows values more narrowly distributed along zero for SH models than for phaseSH models. Black horizontal bars mark the median, and whiskers mark 25<sup>th</sup> and 75<sup>th</sup> percentiles of the distribution. **(C)** Number of simulated spike trains (N=200) by the difference in log losses between phaseSH and SH models.



### ***Model application to CA1 neurons***

We subsequently applied all models to CA1 inhibitory interneurons, where we aimed to identify interneurons in the dataset whose spike timing preferences were systematically related to the phase of the theta (4-12 Hz) rhythm. We additionally replicated this analysis using the phase of the low gamma (35-55 Hz) rhythm. For each rhythm separately, we trained each of the five models for every interneuron in the dataset, and subsequently designed a series of model comparisons to determine (1) whether phase-based models could more accurately predict spike timing in unseen trials relative to long history-based models, and (2) whether interneuron spiking could be reliably explained specifically by rhythmic currents, as measured from the LFP, across trials.

### ***Comparing phase-based and long history-based model performance***


To address this first question, we quantified the proportion of interneurons in this dataset whose phaseSH model reliably outperformed the phaseLH model. To test for reliability, we subtracted the phaseLH model log loss from the phaseSH model log loss for each of the 20 train-test splits of data, yielding a distribution of 20 log loss differences for each neuronal spike train. If, for a given spike train, the phaseSH model smoothly recapitulated all the relevant information embedded in the much more detailed but mechanism-agnostic phaseLH model, then we expected the distribution of log loss differences across the 20 folds to be significantly smaller than zero according to a one-tailed t-test, indicating that long history-based models overfit to the training set and were unable to generalize to held out data from the same neuronal spike train. With few notable exceptions, we found that the phaseSH models based on theta (Fig. 3.3A, *left*) and low gamma (Fig. 3.3A, *right*) phase information were able to account for variation in held out interneuron spike trains more readily than phaseLH models during the odor sampling epoch (theta: 91% interneuron spike trains,  $n = 122$  out of 134, one-tailed t-test,  $p < 0.001$ ; low gamma: ~89% interneuron spike trains,  $n = 120$  out of 134, one-tailed t-test,  $p < 0.001$ ). This trend was reversed when log losses were evaluated on the training trial data ([Supplemental Fig. S3.4](#)), where the phaseLH models held a clear advantage (unsurprisingly) over the phaseSH models.





**Figure 3.3. Model comparisons for odor sampling CA1 inhibitory interneuron data, for theta- and low gamma-phase based models separately.** (A) (top) Theta phase-based models: Differences between phaseSH and phaseLH model log losses (y-axis) evaluated on held out data, across 20 folds (train-test splits) of data for each CA1 interneuron spike train (x-axis). Black horizontal bar within boxplots marks the median, whiskers mark the 25<sup>th</sup> and 75<sup>th</sup> percentiles of the distribution of folds, diamonds mark outliers. A log loss difference of zero (marked by dotted red line) indicates equal performance across phaseSH and phaseLH models. Grey bars mark interneuron spike trains whose spike-phase distributions (drawn from the combined train and test trial data) significantly differed from uniform according to a Rayleigh test for non-uniformity ( $p\text{-val} < 0.001$ ). Pink bars mark the distributions of folds whose log loss differences significantly differed from zero ( $p < 0.001$ ), according to a one-tailed t-test. (bottom) Same as the top panel, but for differences between phaseSH and SH model log losses. Interneuron spike train identity in both top and bottom panels are sorted according to the median of the log loss difference between phaseSH and SH models across folds. (B) Low gamma-phase based models: all details same as in A.

**A**

assess whether  
phaseLH  
contains  
information  
beyond that  
encoded in  
phaseSH

**B**

test for  
rhythmicity



— Rayleigh p-val < 0.001

— log loss diff t-test p-val < 0.001

### ***Identifying neurons with reliable rhythmic engagement***

To address the second question, we aimed to assess how reliably rhythmic currents in the LFP accounted for variance in neuronal spike trains. To achieve this, it was necessary to compare model performance between phaseSH and SH models, a model comparison that we previously established as a suitable metric to assess reliable spike-phase relationships to rhythmic currents measured from the LFP (Fig. 3.2C, [Supplemental Fig. S3.2B](#)). We compared the log losses for these models for the same 20 train-test splits used previously (Fig. 3.3A), and compared the distribution of log loss differences for each interneuron spike train against the null hypothesis that there was either no difference across the model log losses (which would indicate no effect of rhythmic currents on spike train temporal structure) or that this difference was larger than zero (which would indicate that the SH model outperformed the phaseSH model). This metric resulted in ~80% ( $n = 107$  out of 134, one-tailed t-test,  $p < 0.001$ ) of interneuron spike trains exhibiting log loss differences that were reliably smaller than zero, suggesting that theta phase robustly explained variance in these neurons' spike trains (Fig. 3.3B, *left*). We also performed this analysis for phaseSH models leveraging low gamma phase information and found that 44% of interneuron spike trains ( $n = 59$  out of 134, one-tailed t-test,  $p < 0.001$ ) exhibited log loss differences that were reliably smaller than zero (Fig. 3.3B, *right*).

We additionally compared the results of the log loss differences between phaseSH and SH models against the results of a Rayleigh test for non-uniformity in circular distributions (Fig. 3.3B, *bottom grey bars*). Few discrepancies between the metrics emerged when identifying neuronal relationships to the theta rhythm, amounting to approximately 8% of interneuron spike trains ( $n = 11$  out of 134; where the Rayleigh test identified 9 out of 11 neurons missed by the model comparison approach, and the latter identified 2 out of 11 neurons missed by the Rayleigh test; Fig. 3.3B, *left*). Discrepancies between the metrics also arose when assessing low gamma relationships. Here, the disparities amounted to approximately 10% of interneuron spike trains ( $n = 14$  out of 134; Rayleigh test: identified 10 out of 14; model comparisons: identified 4 out of 14; Fig. 3.3B, *right*). The exceptions highlight a crucial difference

between the Rayleigh test and the phaseSH model. The Rayleigh statistic exclusively accounts for *number* of spikes pooled across all trials, irrespective of the total number of time samples available for the recording. The phaseSH model, on the other hand, accounts for firing *rates* (the number of spikes *and* the number of total time samples), and our cross-validation approach implicitly assesses trial to trial reliability of spiking. With our modeling approach, a cell is identified as significantly entrained to a particular rhythm if its baseline firing rate is robust and its spike timing is reliably reproducible across trials. The Rayleigh statistic and phaseSH models can thus offer complementary information: the Rayleigh statistic reveals *whether* a cell emits spikes biased towards particular phase angles, and phaseSH models assess the *predictability* and *reliability* of spikes occurring at particular phase angles.

This effect was particularly evident when we tested model performance on CA1 pyramidal cell data. Although CA1 pyramidal cells in this dataset expressed high firing rates during odor sampling intervals for particular positions, odors, and odor position combinations (35), there were only six trials available for every odor position combination. Given the high degree of selectivity for certain stimulus combinations, we analyzed pyramidal cell spike trains for each odor position combination separately, yielding only six trials per pyramidal cell spike train: a much smaller number of available trials for training and testing than even some of the sparsely firing inhibitory interneurons (Supplemental Fig. S3.6). This resulted in larger discrepancies between the Rayleigh statistic and theta phase-based model comparisons, particularly when attempting to identify neurons whose spiking was related to the theta rhythm. While the Rayleigh statistic identified 35% of pyramidal cell spike trains as related to the theta rhythm ( $n = 15$  out of 43 pyramidal cell spike trains,  $p < 0.001$ ), the model comparison approach identified 12% of pyramidal cell spike trains as reliably related to the theta rhythm ( $n = 5$  out of 43, one-tailed t-test,  $p < 0.001$ ). These discrepancies were dramatically reduced when identifying neurons whose spiking was related to the low gamma rhythm. Here, the Rayleigh statistic and the low gamma phase-based model comparisons yielded closer agreement, where the Rayleigh statistic identified  $\sim 7\%$  of pyramidal cell spike trains as related to low gamma ( $n = 3$  out of 43,  $p < 0.001$ ), while the model comparison approach identified  $\sim 5\%$  of pyramidal cell spike trains as reliably related to low gamma ( $n =$

2 out of 43, one-tailed t-test,  $p < 0.001$ ). In each case, the phase-based model comparison approach tended to yield more conservative estimates of spike-LFP relationships than those produced by the Rayleigh statistic, given the phaseSH model's ability to account for both the *probability* of spiking over time as well as the *rhythmic temporal structure* of spiking.

To further verify that population-level metrics presented in Fig. 3.3 were accurately selecting for *rhythmic* temporal structure in spike trains, we closely evaluated models and their performance on a subset of sample CA1 interneuron spike trains strategically drawn from the population to represent the median log loss difference at the 0<sup>th</sup> (most negative median log loss difference, marking the most predictive theta phaseSH model), 25<sup>th</sup>, 49<sup>th</sup>, 51<sup>st</sup>, 75<sup>th</sup>, and 100<sup>th</sup> percentiles of the distribution of all CA1 interneuron median log loss differences ([Supplemental Fig. S3.7](#)). This selection procedure ensured close inspection of a wide range of neuronal spike trains, some of which should be reliably leveraging phase information specifically to achieve accurate phaseSH model predictions above and beyond the average firing probabilities represented by their SH models. With this subset of neurons, we specifically asked: how much did phase information contribute to the performance of each of the models that included it? To address this question, we trained a hundred versions of each model whose predictions were either partially or completely based on phase data (phase, phaseSH, and phaseLH models), with phase values shuffled randomly within a single cycle while keeping spike history intact. We then computed the log losses evaluated on the held out data for the phase-permuted models, as well as those from the unperturbed models.

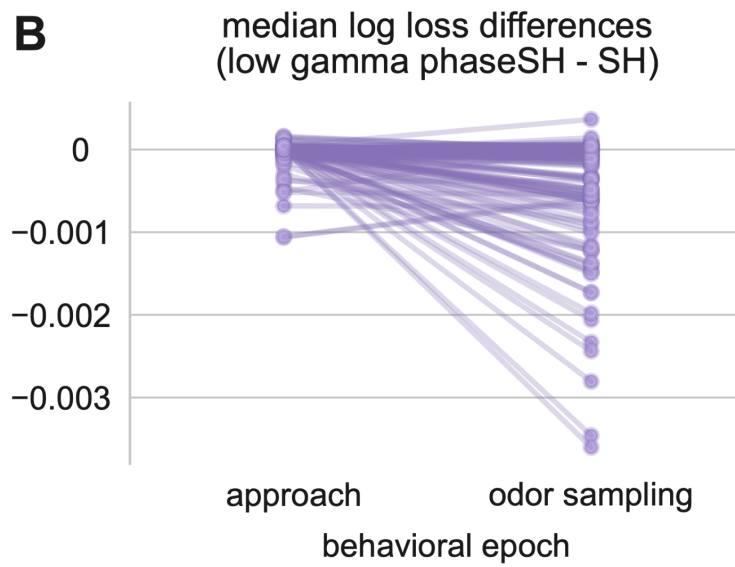
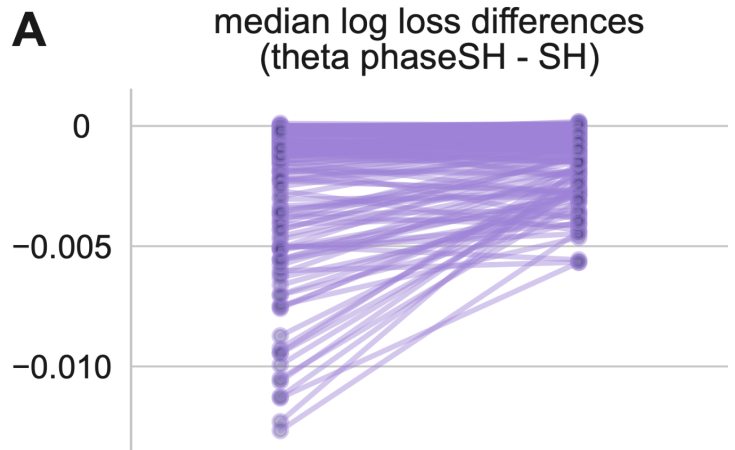
For interneuron spike trains drawn from the 0<sup>th</sup> to the 51<sup>st</sup> percentile of the distribution of median log loss differences, it was clear that models using intact phase information produced log loss values that were far smaller than the most extreme values of the distribution of log losses from the models trained on phase-permuted data ([Supplemental Fig. S3.7](#)). The log loss values for LH and SH models are included for reference. This predictive accuracy metric supported the results from the population metrics in Fig. 3.3B, suggesting that interneuron spike trains whose log loss differences between phaseSH and SH

models are significantly smaller than zero have spike trains whose temporal structure is reliably related to the phase of the theta rhythm.

### ***Characterizing rhythmic engagement profiles across behavioral epochs***

The temporal structure of neuronal spiking varies across distinct behaviors (2-4). Here, we additionally apply model comparisons between phaseSH and SH models to evaluate hypothesized rhythmic current influences across different behavioral epochs during the associative memory task. We specifically aimed to visualize shifts in rhythmic current influences across the running approach to the odor port and the subsequent odor sampling epoch when the CA1 region was likely engaging associative memory processes. The pattern of model comparisons reveals a qualitative shift between the approach and odor sampling epochs (explored quantitatively in [Supplemental Fig. S3.5](#)). While interneuron spike trains appear to be strongly related to rhythmic currents in the theta frequency range, this degree of relationship tends to shrink during the odor sampling epoch (Fig. 3.4A). On the other hand, relationships to low gamma-range rhythmic currents tend to increase during the odor sampling epoch relative to their influence during the approach epoch (Fig. 3.4B). Model comparisons between phaseSH and SH models can reveal differences in the degree of influence that rhythmic currents exert across distinct behavioral epochs.

**Figure 3.4. PhaseSH models trained on data from distinct behavioral epochs capture shifts in CA1 interneuron spike-LFP relationships.** (A) Each line represents the median log loss difference between the theta phaseSH model and the SH model trained on a single interneuron's spike train in either the running approach or odor sampling epoch. A difference in height between the approach and odor sampling behavioral epochs thus indicates a change in a spike train's degree and reliability of entrainment to the theta rhythm across trials. (B) Same as in A, for the low gamma phaseSH models. Note the scale differences across A and B. More negative values indicate that the phaseSH model outperformed the SH model.





### ***Characterizing rhythmic engagement profiles that change over time***

We next sought to leverage the temporal resolution of the phaseSH models to characterize the time course of single neuron relationships to rhythmic processes, such as the theta (4-12 Hz) and low gamma (35-55 Hz) rhythms, as reflected in the local field potential (LFP). In the hippocampus, running epochs coincide with the dominant presence of theta oscillations (45, 46), and running speed is positively correlated with increasing interneuron firing rates (14, 17). Accordingly, we hypothesized that behavioral epochs during which rats were running towards odor ports (approach epochs) would likely coincide with high rates of CA1 interneuron spike phase relationships to the theta rhythm relatively *consistently* throughout the approach interval. In contrast, the odor sampling epoch of the associative memory task requires parsing multiple input streams to integrate the relevant contextual and odor information that should guide the rat's successful behavior. We thus hypothesized that some interneurons would be more likely to exhibit *shifts* in spike phase relationships to distinct CA1 rhythms during this odor sampling, associative memory window.

To characterize the time course of spike phase relationships during the approach and odor sampling epochs, we again compute differences in model log losses on held out data to ask whether the relative predictive accuracies between the theta and low gamma phaseSH models varied over the course of associative memory processing within the odor sampling epoch, for each interneuron independently. The millisecond resolution of phaseSH models allowed us to track whether any given held-out spike was better predicted by the probabilities generated from the theta or low gamma phaseSH models. We could then aggregate these log losses within arbitrarily small or wide window sizes (in this case 50 ms duration with no overlap) and take the median of these across results from the 20 train test splits of the data within the odor sampling window to obtain an estimate of the relative strength of the relationship to either of the two rhythms.

We present the median log loss differences (computed over time) between the following models: theta phaseSH minus SH, theta phaseSH minus low gamma phaseSH, and low gamma phaseSH minus SH. Across the population, this analysis captured the relatively weaker low gamma rhythmic current

influences on spiking during the approach epoch (Supplemental Fig. S3.8A-C left, Fig. 4, Supplemental Fig. S3.5). During the odor sampling epoch, however, transient periods emerged during which low gamma phaseSH models outperformed theta phaseSH models (Supplemental Fig. S3.7B, right). This pattern of low gamma rhythmic current influence bore out when comparing the theta phaseSH model against the SH model: along the diagonal of the matrix, the accuracy of the theta phaseSH model tended to suffer (Supplementary Fig. S3.8A, right), while model comparisons between the low gamma phaseSH and SH models revealed slight increases in the influence of low gamma rhythmic currents on spike timing (Supplementary Fig. S3.8C, right).



**Figure 3.5. Model comparisons over time reveal rapid shifts in spike-phase relationships to LFP rhythms during an odor sampling epoch.** Models were fit to spiking and rhythmic phase data from behavioral epochs when rats were either running towards the odor port (left column, approach) or odor sampling (right column). **(A)** Rows correspond to three unique interneuron identities, matched across approach and odor sampling behavioral epochs. Columns within each heatmap correspond to non-overlapping 50 ms-long time windows spanning a continuous interval of 1250 ms. Color scale corresponds to log loss differences between the theta phaseSH and SH model computed on held out data for each of the 20 train test splits. Warmer red colors indicated that theta phaseSH models outperformed SH models. The average firing rates across all 20 held out sets during the approach are as follows, for each neuron in order of rows from top to bottom: 35.5 Hz, 36.6 Hz, 42.4 Hz. For the odor sampling epoch, the average firing rates are as follows: 20.3 Hz, 20.9 Hz, 26.2 Hz. **(B)** Same as in **A**, for the same three neurons, but for log loss differences between the theta phaseSH model and the low gamma phaseSH model. Warmer red colors indicated better performance by theta phaseSH models, while cooler blue colors indicated better performance by low gamma phaseSH models. **(C)** Same as in **A**, but for log loss differences between low gamma phaseSH models and SH models. Warmer red colors indicated better performance for low gamma phaseSH models. All color maps are displayed on the same scale.

**approach**

**odor sampling**



**A**

log loss difference (theta phaseSH - SH)





**B**

log loss difference (theta phaseSH - low gamma phaseSH)



**C**

log loss difference (low gamma phaseSH - SH)



time windows  
(50 ms duration, no overlap)

time windows  
(50 ms duration, no overlap)

**Figure 3.6. Model comparisons over time on held-out data**, with significance assessed per time window using a two-tailed one-sample t test. As in Fig. 5, models were fit to spiking and rhythmic phase data from behavioral epochs when rats were either running toward the odor port (left column, approach) or odor sampling (right column). Model performance was tested on the held-out trials of each train-test split. A: rows correspond to three unique interneuron identities, matched across approach and odor sampling behavioral epochs. Columns within each heatmap correspond to nonoverlapping 50-ms long time windows spanning a continuous interval of 1,250 ms. Colors correspond to log loss differences between models combining theta phase and short history (phaseSH) and SH model computed on held-out data for each of the 20 train test splits that were significantly different from a distribution centered around zero according to a two-tailed one-sample t test ( $n = 20$  log loss differences per window,  $P < 0.00004$ ). Pink bars in heatmaps correspond to time windows whose distribution of log loss differences indicate better theta phaseSH model predictions relative to SH model predictions; blue indicates SH model predictions significantly outperformed theta phaseSH model predictions. B: same as in A, for the same three neurons, but for significant log loss differences between the theta phaseSH model and the low gamma phaseSH model. Pink in these heatmaps indicates that the theta phaseSH model significantly outperformed the low gamma phaseSH model; blue indicates that low gamma phaseSH models significantly outperformed the theta phaseSH models. C: same as in A, but for significant log loss differences between low gamma phaseSH models and SH models. Blue in these heatmaps indicates low gamma phaseSH models significantly outperformed SH models; pink indicates that SH models outperformed low gamma phaseSH models. White bars indicate no significant difference in model performance.

**approach**

**odor sampling**

**A**

log loss difference (theta phaseSH - SH)

two-tailed t-test  $p < 4e-05$





■ **theta** phaseSH  
better

**B**

log loss difference (theta phaseSH - low gamma phaseSH)

two-tailed t-test  $p < 4e-05$

interneuron  
identity




■ **low gamma**  
phaseSH better  
■ **theta** phaseSH  
better

**C**

log loss difference (low gamma phaseSH - SH)

two-tailed t-test  $p < 4e-05$



■ **low gamma**  
phaseSH better

time windows

time windows

(50 ms duration, no overlap)

(50 ms duration, no overlap)

To further explore the range of trends in rhythmic influence, we selected a subset of neuronal spike trains with dynamic outcomes in the model comparisons over time. Notably, one of the three sample interneuron spike trains exhibited relatively few time windows with significant theta or low gamma rhythmic current influences during the approach or the odor sampling epoch (Fig. 3.5A-C & Fig. 3.6A-C, *left, middle rows, rows*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window), but it did exhibit what appeared to be a relatively sustained relationship to low gamma rhythmic currents during the odor sampling epoch (Fig. 3.5B & Fig. 3.6B, *right, middle rows*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window). The other two sample interneuron spike trains exhibited relatively strong and consistent theta relationships during the approach (Fig. 3.5A & B, Fig. 3.6A & B, *left, top & bottom rows*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window), and weaker low gamma relationships during this interval (Fig. 3.5C & Fig. 3.6C, *left*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window). During the odor sampling epoch, however, the theta phaseSH models for these neurons decreased in predictive utility over the course of the odor sampling epoch (Fig. 3.5A & B, Fig. 3.6A & B, *right, top & bottom rows*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window). Surprisingly, when the theta rhythmic influence decayed in one of these spike trains, the low gamma rhythmic influence acquired prominence towards the end of the epoch (Fig. 3.5B & Fig. 3.6B, *right, top row*, two-tailed one-sample t-test,  $p < 0.00004$ ,  $n = 20$  log loss differences per time window).

Model comparisons evaluated over time can thus test the hypothesis that specific rhythmic currents shape spike timing at fine time resolution, and the phaseSH model's failures *and* successes in predictive utility have the potential to reveal rapid shifts in spike timing.

### **3.4 Discussion.**

Neuronal spike timing and its rapid dynamics equip organisms to successfully adjust to a shifting landscape of sensory information and behavioral demands. With the present work, we offer a modeling approach that can (1) capture short-lived intrinsic constraints on spiking as well as rhythmic current

influences that unfold over relatively longer timescales, and (2) evaluate neuronal recruitment by a constellation of distinct rhythmic currents, as measured from the local field potential (LFP), over time. Specifically, models including the immediate spiking history and information about the phase of an LFP oscillation—termed “phaseSH” models in this work—can recover short and rhythmic long timescale constraints on cellular spiking. Importantly, we prescribe approaches to identify neurons in large datasets that exhibit systematic relationships to LFP rhythms. By innovatively applying long-established statistical modeling methods and their associated performance metrics, these models can make informed (and verifiable) predictions over spiking at fine time resolution, rather than collapsing estimates across long windows of time and assuming consistent temporal structure during these periods. With our statistical models, it is possible to explicitly test neuronal participation in rhythmic processes at exquisite temporal resolution.

In this work, we propose a novel application of kernel density estimation (KDE) to estimate the conditional probability of spiking given the phase of an LFP oscillation and we apply these predictions over time. To specifically test the hypothesis that spike timing is related to the phase of an oscillation (and is thus likely influenced by rhythmic currents), it is imperative to construct models that avoid imposing distributional biases (e.g. normality) on their relationship. Biased models may fail to recover various features (e.g. variance, complexity) of the true underlying relationship and result in hypothesis tests using unrepresentative relationships. With KDE, we can instead empirically estimate the relationship from training data. We first estimate the likelihood of theta (or low gamma) phases given spike occurrences using KDE (Fig. 1C). We subsequently compute the posterior probability of observing a spike given observed phase values, and we apply this probability mapping over the duration of the phase timeseries to produce probabilities of spiking that unfold over time (Fig. 1C). Crucially, we leverage cross validation to select the optimal kernel bandwidth used to convolve the distribution of phase observations that co-occurred with spikes. Our use of KDE to estimate the conditional probability of spiking given phase, and our ability to produce an optimized estimator, presents an innovative contribution to the problem of characterizing spike-phase relationships.



Our goal was to apply KDE-based “phase” model predictions over time. Phase model predictions suffered, however, from the inability to account for short timescale refractory constraints. For this reason, we additionally constructed logistic regression models using the 3 ms short (“SH”) history of spiking, and combined the predictions from this model with those from the phase model (Fig. 3.1D). This combination yielded the “phaseSH” model, which was capable of adjusting predictions according to refractory constraints while simultaneously testing whether particular rhythmic currents explained variance in spike trains (Fig. 3.2). PhaseSH models captured ground truth rhythmic dynamics in simulated spike trains and they did so more accurately than models whose predictions were based on the 250 ms long (“LH”) history of spiking (Fig. 3.2B).

Simulations also revealed a key property of phaseSH models: they produced flat predictions over time whenever spikes were not systematically organized according to the phase time series being tested. This property allowed us to leverage model comparisons between phaseSH and SH models in order to identify simulated (Fig. 3.2C, [Supplemental Fig. S3.B](#)) and CA1 neuronal spike trains (Fig. 3.3B, [Supplemental Fig. S3.5](#), [S3.6](#)) that exhibited reliable rhythmic relationships.

Alternative statistical approaches to quantifying rhythmic temporal structure in spike trains exist. Of these, the Rayleigh statistic and mean resultant length vectors are prominently used in the field (35,47). While a useful metric to distinguish the density of spiking concentrated along some phases (example in Fig. 3.1C) from a uniform distribution, the Rayleigh statistic requires pooling spikes over large windows of time. In contrast, our cross validated model comparison approach allowed us to quantify the predictive utility of phaseSH models on held out data, enabling assessment of its predictions both across trials (Fig. 3.3) and within trials over arbitrarily small time samples down to millisecond resolution (Fig. 3.4). These features allow us to examine the trial-to-trial *reliability* of spike-phase relationships, as well as identify rapid shifts in spike-phase relationships over the course of behavior. A strength of the Rayleigh statistic, however, is that it will not take baseline firing rate into account, unlike our KDE and logistic regression models. This means that its estimates of spike-phase relationships will not distinguish

between neurons exhibiting vastly different firing rate behavior, as long as each neuron provides the same total number of spikes. A sparsely firing neuron recorded over a long session may produce the same number of spikes, and obtain the same Rayleigh statistic, as a frequently firing neuron recorded over a much shorter session. In contrast, KDE probabilities implicitly account for the baseline probability of firing (Fig. 3.1C). A sparsely firing neuron recorded over a long session will thus produce a smaller amplitude KDE curve than a frequently firing neuron recorded over a shorter session, for the same number of spikes and same degree of phase relationship. This distinction crucially disambiguates *rhythmic* neurons—where rhythmicity requires firing at consistent phase bins for a large fraction of consecutive cycles—from neurons that exhibit systematic phase relationships but which spike so rarely as to coincide with a smaller, and more sporadic (e.g. non-consecutive), fraction of the rhythm’s many cycles. In this regard, the Rayleigh statistic and our modeling approach provide complementary information that together fully captures systematic relationships to LFP rhythms, as well as the likelihood of a neuron’s *rhythmicity*.

When we applied cross validated model comparisons between phaseSH and SH models to multiple behavioral contexts and neuron types, we revealed a range of dynamic relationships to theta and low gamma rhythms in the hippocampus. Specifically, we identified changing degrees of CA1 interneuron entrainment to theta and low gamma rhythms across high velocity and odor sampling epochs (Fig. 3.4), and even changing entrainment over the course of the odor sampling interval at fine (50 ms) time scales (Fig. 3.5, Fig. 3.6, [Supplemental Fig. S3.8](#)). Given the dynamism we describe in the hippocampus, and the often highly task-selective responses of hippocampal neurons, we recommend thoughtful consideration of the isolated intervals of interest. Specifically, eliciting many trials with stereotyped behavior is necessary to best leverage the ability of our method to characterize reliable rhythmic entrainment when cells are active and the evolution of rhythmic entrainment over time for particular processing states.

Oscillatory currents prominently manifest within a variety of brain regions. Sensory stimuli can evoke a cascade of oscillatory interactions, which often reflect tight excitatory and inhibitory control of spiking in piriform cortex during the presentation of odors (48), and produce exquisite gamma oscillatory

dynamics in visual cortices in response to attended stimuli (49, 50). In motor systems, tight coordination can also be observed throughout the basal ganglia in response to utilized cues (51) but also in travelling waves through motor cortices during movement preparation and execution (52-54). Aberrant oscillatory patterns have been observed in a number of neural disorders (e.g. disrupted gamma in Alzheimer's disease (55,56), disrupted beta in Parkinson's disease (57)) and are commonly used as markers of altered interactions and disease progression. However, in all cases, the relationship between these oscillatory currents and cell spiking activity has been poorly understood due to a limited ability to assess spiking relationships to ongoing oscillations at the time scales over which they evolve (as opposed to averages over seconds-long epochs). Our modeling approach can be applied broadly in these contexts to precisely characterize how and *when* large scale, coordinated currents engage cells to support dynamic information processing.

### **3.5 Materials and Methods.**

#### ***Experimental Design: Behavioral Paradigm***

All procedures involving animals were carried out according to guidelines set forth by the National Institute of Health, with approval from the Institutional Animal Care and Use Committee (IACUC) in Boston University (approval number: 13-057). Experimental procedures have been previously described at length in (35) and will be briefly reviewed here.

We aimed to elicit multiple stereotyped and behaviorally controlled instances of “associative memory processing” during a window of time in which rats had to integrate information about odor cues and the context in which they were presented to produce an appropriate, rewarded response.

#### ***Odor sampling epoch.***

The data analyzed in this study were recorded from the CA1 subregion of the hippocampus in rats performing a context-guided associative memory task (35,36). Rats learned that odors in a pair were differentially rewarded as a function of the context in which they appeared. Contexts were differentiated according to their spatial position in the room as well as the texture and coloring of material wrapped

around the context's surface. During a given trial, rats were allowed access to one context at a time, and odor delivery within an odor port (start time of each trial) was triggered 250 ms after the rat poked his nose into that odor port. Odor reward contingencies remained constant throughout the experiment, and odor delivery locations were pseudorandomized across odor ports within a context. Two pairs of odors were used for each session, with the first block (48 consecutive trials) featuring odors A and B, and the second block consisting of odors C and D, for a total of 96 trials per session. Interneuron data recorded from each block of odors was analyzed separately, given the potential differences in spiking temporal structure driven by distinct contexts. Pyramidal cell data was analyzed separately for each odor-position combination since these neurons have previously been shown to spike selectively for specific positions and odor-position combinations in this task (35). **Correct trials.** In order to receive a water reward, rats maintained a nose poke in the correct odor port (the one containing the rewarded odor within that context) for 1250 ms following the odor onset, after which a water droplet was delivered to a well directly below the odor port. If the rats initially poked their nose into the incorrect port, they had up to 1250 ms following odor onset to remove their snout, after which they were able to poke into the correct odor port and maintain the nose poke. **Incorrect trials.** If the rat failed to end the nose poke 1250 ms following odor onset in the incorrect odor port, a buzzer would sound and no reward would be delivered. Recordings began only after rats had successfully achieved a 75% performance criterion. Incorrect trials are consequently infrequent, and we excluded them from this study. For our phase modeling efforts, all intervals were restricted to the 1250 ms intervals following odor onset in which rats maintained their nose in the odor port containing the *correct* odor prior to receiving a reward. We refer to this interval as the “odor sampling epoch.” The first time point assessed by our models begins at the time of odor presentation, with the preceding 250 ms serving solely as spike history for the first 250 time samples of the odor-sampling interval. **Approach epoch.** We also separately fit models to data drawn from a behaviorally distinct epoch: the running approach to the odor port. To create these approach epochs, we selected the 1500 ms interval with the highest median velocity preceding a *correct* odor sampling trial. As

we did for the odor sampling epoch, phase models are fit to data from a continuous 1250 ms period, with the preceding 250 ms serving as spike history for the first 250 time samples of the approach intervals.

### ***Data Acquisition***

We carried out high-density extracellular tetrode recordings from the CA1 subregion of the hippocampus and extracted single-unit and local field potential (LFP) data. The Omniplex Neural Data Acquisition system (Plexon) amplified (4000-8000 x) and digitized (40 kHz) signals. ***Single units.*** Putative interneurons were isolated according to waveform features: mean firing rate ( $\geq 5$  Hz), mean width at half-maximum of waveform ( $<150$   $\mu$ s), and mean temporal offset from peak to trough ( $<350$   $\mu$ s) (18, 37). ***Local Field Potentials.*** LFPs were bandpass-filtered in the theta (4-12 Hz) range using a third-order Butterworth filter. We then obtained instantaneous phase estimates by computing the arctangent of the filtered signal's complex Hilbert transform. Frequency range selection was based on the observable frequency bands in average spectrograms during the odor sampling epoch (35).

### ***Data Preparation***

In the present study, we examined interneuron spiking and CA1 local field potentials (LFPs) for single correct trials, each consisting of 1500 1-ms time samples (250 ms before odor stimulus delivery, 1250 ms after odor stimulus delivery). Time samples occurring *within* trials are consecutively sampled (sampling frequency = 1000 Hz), while time samples *between* trials are separated by longer (and variable) stretches of time. Each interneuron spike train was analyzed with respect to the simultaneously recorded CA1 LFP, acquired from the same tetrode that collected single-unit spikes. To test the generalizability of trained models, we partitioned trial labels into train (50%) and test sets (50%). We generated a total of twenty train-test splits for each neuronal spike train. Consecutive time samples within a trial retained their temporal structure. ***History Variables.*** During the odor sampling epoch, the odor that the rat used to make a decision (to stay or move to the next odor port) was made available 250 ms after the nose poke. In the first 250 ms of the trial, the rat's behavior is the same (stationary) as its behavior for the rest of a correct trial, the only difference being the availability of the odor. Given this, we use the first 250 ms of the trial

as the history for the first time samples post-odor release. The 250 ms at the very beginning of each trial were only included in the models as spiking history regressors; they never served as target variables (Fig. 1A). Fitting history models to data from the approach epochs proceeded identically.

### ***Data Simulations***

We produced synthetic data to simulate four distinct “types” of spike trains, each with unique ground-truth dependencies that governed spike probability. These simulated spike trains exhibited: (1) *no temporal modulation / “atemporal”* with no refractory period and no relationship with respect to a simulated phase time series oscillating at 8 Hz, (2) *short timescale temporal modulation / “refractory non-rhythmic”* with a refractory period (lasting 3 ms post-spike) and no relationship with respect to a simulated phase time series oscillating at 8 Hz, (3) *long timescale modulation / “non-refractory rhythmic”* with no refractory period but a relationship with respect to a simulated phase time series oscillating at 8 Hz, and (4) *both short and long timescale modulation / “refractory rhythmic”* with a refractory period (lasting 3 ms post-spike) and a relationship with respect to a simulated phase time series oscillating at 8 Hz.

We simulated forty-eight time series, each spanning 1500 ms (identical in length and approximate number to trials in our behavioral task, where the first 250 ms of each trial serve exclusively as history predictors for the early time samples of each odor sampling interval, Fig. 1A). For each time series, we created a time-matched vector of probability densities indicating which phases were most likely to coincide with spikes. For simulated spike trains that had no phase relationship, each cycle of the probability vector consisted of a flat, uniform distribution whose value corresponded to the average probability of spiking. For simulated spike trains governed by a phase relationship, each cycle of the probability vector consisted of a von Mises distribution peaking at 0 radians, with an inverse variance of 2 (equivalent to a Gaussian kernel

bandwidth of 0.5). Phase cycles were each 125 ms long, simulating an 8 Hz oscillation. This probability vector thus represented the probability of spiking given the phase of an 8 Hz oscillation. To equip simulated spike trains with a refractory period, we forced the model probabilities generated above to adjust to 0.00001 (essentially zero) in the 3 time samples immediately following a sampled spike (see sampling procedure below).

To sample spiking events from these ground truth probabilities, we used a single-trial binomial distribution (software: Python Numpy, `random.binom`), which produced a spike (1) or not a spike (0) according to the probability value at that particular time sample. Once the simulated data had been generated, each simulated spike train underwent history and phase model-fitting as described in the sections below. To test the degree of divergence between model-based probabilities and ground truth probabilities, we calculated the Kullback-Leibler (KL) divergence,

$$D(p_i||q_i) = p_i \log_2 \frac{p_i}{q_i} + (1 - p_i) \log_2 \frac{(1-p_i)}{(1-q_i)} \quad (1)$$

where  $p_i$  and  $q_i$  are the ground truth and modeled probability of spiking at time sample  $i$ , respectively. This value can be interpreted as the cost, in bits, incurred when modeling the distribution of spikes using the model-based probabilities instead of the ground truth.

We computed the KL divergence at each time sample to compare the ground truth to a given model's predictions. We then computed the average of all the KL divergences across time, to provide an overall sense for the degree of misalignment between the ground truth and model-based predictions throughout the entirety of the time series. A smaller average KL divergence indicates increased alignment between the ground truth and the model.

### ***Constructing History Models***

We constructed models that used the history of spiking over long or short timescales. For short history models (“SH models”), we reasoned that using 3 ms of history would capture both the absolute and initial segments of the relative refractory periods that neurons undergo immediately after emitting an action potential. In contrast, the long history model (“LH model”) can reflect various temporal constraints on spiking over a more extended past. Consequently, for neurons in the hippocampus, the history should also encode the periodic structure of spiking due to oscillatory current modulation, if it exists. In order to test whether long history could capture the periodic distribution of spikes over time, we set the long history interval to begin 250 ms prior to current time sample  $i$ , equivalent to a full period of the lowest frequency (4 Hz) within the theta 4-12 Hz frequency range, and additionally able to capture any rhythmic current influences at faster frequencies.

In order to generate probabilities of spiking given the long and short histories separately, we fit logistic regression models, according to the expression below,

$$p\left(y_i = 1 \mid H_i^{(m)}\right) = \frac{\exp\left(H_i^{(m)T} \beta^{(m)} + \beta_0^{(m)}\right)}{1 + \exp\left(H_i^{(m)T} \beta^{(m)} + \beta_0^{(m)}\right)} \quad (2)$$

where  $y_i$  corresponds to spiking at time sample  $i$  and  $H_i^{(m)}$  is a vector that represents history, long (with dimensions: 1 x 250) or short (with dimensions: 1 x 3), for that time sample. In models that commit to including either long history or short history parameters,  $m$  is replaced by  $L$  for long history, or  $S$  for short

history (Fig. 3.1A). For ease of notation, we define the vectors  $\tilde{H}_i^{(m)} = \left[1 \ H_i^{(m)}\right]^T$  and  $\tilde{\beta}^{(m)} = \left[\beta_0^{(m)} \ \beta^{(m)}\right]^T$

, such that  $\tilde{H}_i^{(m)T} \tilde{\beta}^{(m)} = H_i^{(m)T} \beta^{(m)} + \beta_0^{(m)}$ . To estimate the coefficient vector  $\tilde{\beta}^{(m)}$ , we find the vector

that minimizes the normalized negative log likelihood, NNLL, given by



$$\text{NNLL}(\tilde{\beta}^{(m)}) = \frac{1}{n} \sum_{i=1}^n -y_i \left( \frac{\tilde{H}_i^{(m)T} \tilde{\beta}^{(m)}}{\beta} \right) + \log \left( 1 + e^{\frac{\tilde{H}_i^{(m)T} \tilde{\beta}^{(m)}}{\beta}} \right) \quad (3)$$

For short history models, we fit a coefficient vector  $\tilde{\beta}^{(S)}$  containing four parameters ( $\tilde{\beta}^{(S)} = \beta_0^{(S)}, \beta_1^{(S)}, \beta_2^{(S)}, \beta_3^{(S)}$ ), with one parameter for the baseline spike rate and one for each millisecond of short history. In contrast, long history models require fitting a coefficient vector  $\tilde{\beta}^{(L)}$  containing 251 parameters ( $\tilde{\beta}^{(L)} = \beta_0^{(L)}, \beta_1^{(L)}, \dots, \beta_{250}^{(L)}$ ). This large number of parameters renders the LH model a high-variance model, susceptible to large variations in coefficient estimates as a function of small changes to the training dataset, and weakly generalizable to held-out data (38). To mitigate this statistical disadvantage while retaining the flexibility and mechanistic agnosticism we prize in this model, we minimize the NNLL in this model with an added L2 penalty,

$$\tilde{\beta}^{(L)} = \arg \min_{\tilde{\beta}} \left\{ \text{NNLL}(\tilde{\beta}) + \lambda \|\tilde{\beta}\|_2^2 \right\}, \quad \|\tilde{\beta}\|_2^2 = \sum_{p=0}^{250} \tilde{\beta}_p^2 \quad (4)$$

where the penalty term  $\lambda \|\tilde{\beta}\|_2^2$  works to *increase* the NNLL. To minimize the NNLL in light of this penalty, the coefficient vector is optimized such that the most useful predictors are assigned large coefficients, while remaining predictors are assigned coefficients approaching zero, thus attenuating the effect of noise in parameter estimation. The regularization strength parameter,  $\lambda$  is set to equal 1 by default; we did no optimization over this parameter. LH model fits were performed using Python's `sklearn.linear_model.Logistic Regression` method, and SH model fits ensured that the penalty argument was not active.

### ***Learning Phase Representations***

The local field potential (LFP) is a complex signal, manifesting a number of oscillatory and non-oscillatory components at different times. Others have successfully leveraged the Fast Fourier Transform of spike-triggered averages of the LFP as a predictor of synchronous spiking activity across multiple, simultaneously recorded units (41). Spike-triggered averages of the LFP may comprise multiple nested oscillatory and non-oscillatory components, and will thus represent a complex signal. In our work, we aim to tease apart the variety of oscillatory components present in the LFP, and use these to assess the *timing* of single neuron spiking relative to the oscillatory process. Systematic spike timing relative to rhythmic phase may indicate a neuron's engagement in the processes that give rise to the corresponding oscillatory LFP feature. Moreover, we wish to identify the potential reorganization of spike timing relative to the phase of different rhythms present in the LFP over time. For this reason, we have focused on modeling spike timing according to rhythmic phase.

Raw circular variables like oscillatory phase are inappropriate regressors in a generalized linear model (GLM). To generate a useful phase representation, we first noted that the spike-phase relationship can often be approximated by eye from a histogram of spike counts over phase. Leveraging this visual representation of the relationship, we applied kernel density estimation (KDE) by first repeating the distribution of spikes over phase to create three identical instances of the distribution and then convolving the three concatenated distributions with a Gaussian kernel. This procedure yields a smooth function representing the conditional probability of observing a phase  $\phi$  given spiking (after dividing by three to account for the triple counting of spikes),  $p(y = 1)$  (Fig. 3.1C), which approximates the effect of convolving the data with a von Mises kernel.

The Gaussian kernel's bandwidth (the width at half-maximum amplitude of the Gaussian kernel) crucially determines the ultimate shape of the kernel density estimator. In order to learn the optimal kernel bandwidth within a given interneuron dataset, we applied 5-fold cross validation on the training subset of trials (software package: Python sklearn, GridSearchCV, Kernel Density method, grid size = 1000 (39)). For the CA1 pyramidal cell datasets, we applied leave-one-out cross validation (instead of 5-fold cross

validation) on the training subset of trials. This change ensured maximal use of the small amount of pyramidal cell spiking data available for kernel bandwidth selection.

On each fold of the cross-validation procedure, we sampled observations from the training set to generate a distribution of spiking over phase, and the optimal bandwidth (out of several possible bandwidths) was selected if its resulting estimator maximized the log likelihood over validation sets. It is important to explore the cross validated performance of a set of candidate kernel bandwidths that are sufficiently narrow or wide to capture fine or broad fluctuations of spiking as a function of phase. Increasing the total number of candidate kernel bandwidths will increase the computational time required to select the final bandwidth, but it will potentially result in a more accurate estimator. To estimate spike-phase relationships in this work, we explored the cross validated log likelihoods of 20 candidate bandwidths ranging from the equivalent of 6% to 40% of a phase cycle at kernel half-maximum. This range was selected to span a wide assortment of candidate bandwidths and smoothly capture slow or fast fluctuations of spiking over phase. This range is similar to the set of inverse variances deployed in (40), who previously developed a successful method for estimating spiking relationships to oscillatory phase that differs substantively from the method currently applied here.

After estimating the probability of phases conditioned on spiking, we applied Bayes' Theorem to obtain the posterior probability of spiking given phase for the training set:

$$p(y = 1 | \phi)_{train} = \frac{p(\phi|y=1)_{train} p(y=1)_{train}}{p(\phi)} \quad (6)$$

where the prior over spiking activity in the training set,  $p(y = 1)_{train}$ , was computed as the average number of spikes divided by total time samples in the training set, and the prior over phase,  $p(\phi)$ , was uniform,  $\frac{1}{2\pi}$ . The probability of spiking given held out phase data,  $p(y = 1 | \phi)_{test}$ , was computed in the same way, but with likelihood  $p(\phi | y = 1)_{test}$ , derived from the held-out distribution of spikes over

phase. Importantly, we used the prior over spiking from the training set, to prevent the phase model from using any information in the held out set of spikes. These steps ensured that the held-out set is a truly independent, and unbiased test of model performance.

### ***Models Combining Phase & History Information***

We constructed models that account for both short and longer timescale temporal constraints by combining phase and history information. We fit logistic regression models to regressor combinations that included either short history and phase (“phaseSH models”), or long history and phase (“phaseLH models”):

$$p(y_i = 1 | l_i^{(m)}, l_i^\phi) = \frac{\exp(l_i^{(m)}\beta^{(lm)} + l_i^\phi\beta^{l\phi} + \beta_0^{c(m)})}{1 + \exp(l_i^{(m)}\beta^{(lm)} + l_i^\phi\beta^{l\phi} + \beta_0^{c(m)})} \quad (7)$$

In these models, we avoid re-fitting history coefficients by taking the log-odds of the probabilities over spiking generated by the long history,  $l_i^{(L)}$ , or short history,  $l_i^{(S)}$ . To include a phase representation that was suitable for logistic regression, we took the log-odds,  $l_i^\phi$ , of the  $p(y_i = 1|\phi_i)$ , for each time sample  $i$  derived from the training data. For models combining phase and history information, only three coefficients are fit: the intercept,  $\beta_0^{c(m)}$  (with superscript  $c(m)$  to signal this coefficient’s applicability only to these combined models, as distinct from the history-only models specified in Eq. 3), and coefficients  $\beta^{(lm)}$  and  $\beta^{l\phi}$ . Notably, these models will only fit these three parameters irrespective of the dimensionality of the originating history term, given that only the log-odds of the probability of spiking given history—and not the original history vector itself—is included in the models. The log-odds are expressed below,

$$l_i^{(m)} = \log \frac{p(y_i=1|H_i^{(m)})}{1-p(y_i=1|H_i^{(m)})} \quad (8)$$

$$l_i^\phi = \log \frac{p(y_i=1|\phi_i)}{1-p(y_i=1|\phi_i)} \quad (9)$$

We then generated predictions over spiking using the held out dataset with each of the models: short history only  $p(H_i^{(S)})$ , long history only  $p(H_i^{(L)})$ , phase  $p(y_i = 1 | \phi_i)$ , phase and short history  $p(l_i^{(S)}, l_i^\phi)$ , and phase and long history  $p(l_i^{(L)}, l_i^\phi)$ . We refer to these models as SH, LH, phase, phaseSH, and phaseLH models, respectively, throughout the manuscript.

### ***Evaluating Model Performance and Goodness-of-Fit***

To assess the likelihood-based goodness-of-fit (42) of models on the held out set of spikes, we computed the average log loss  $LL$  of each model,

$$LL = -\frac{1}{n} \sum_{i=1}^n [y_i \log_2 \hat{y}_i + (1 - y_i) \log_2 (1 - \hat{y}_i)] \quad (10)$$

where  $i$  corresponds to each 1 ms time sample,  $y_i$  represents a spiking event in the held out data,  $1-y_i$  represents the absence of a spiking event in the held out data,  $\hat{y}_i$  corresponds to the predicted probability of a spike occurring according to the model, and  $1 - \hat{y}_i$  corresponds to the model's predicted probability that a spike will not occur at time sample  $i$ . The log loss represents the degree of disparity between a model's predicted spiking probability and the actual occurrence of a spike or its absence. As a consequence, smaller values indicate improved predictions.

### ***Identifying Neurons Exhibiting Spike-Phase Relationships***

The phase model (as described in section *Learning Phase Representations*) can fit a relatively flat line to spiking data that is uniform over all phases. This makes it difficult to assess whether the phase model performed well because it approximated the average probability of spiking or whether it performed well by approximating rhythmic structure. To identify cells with strong phase relationships, we computed the log loss (eq. 10) of the phaseSH and SH models, and subtracted the log loss of the SH model from that of the phaseSH model. A *negatively* valued log loss difference indicates that including phase provides information beyond that encoded in the SH model (phaseSH log loss < SH log loss). In contrast, a log loss difference approaching zero indicates that phase is not adding much more information than the SH model contains by itself (phaseSH log loss  $\approx$  SH log loss). We computed this log loss difference for each of the twenty train-test splits of data for each neuron, and then performed a one-tailed t-test ( $p < 0.001$ ) to assess whether the distribution of log loss differences was significantly smaller than zero, indicating that phaseSH models significantly outperformed SH models. This method effectively disambiguates neurons with systematic relationships to the phase of an LFP oscillation from those whose spiking is unrelated to the specified oscillatory phase.

### ***Characterizing Variations in Spike-Phase Relationships Over Time***

One substantial advantage of phaseSH models is their ability to make predictions over spiking at fine temporal resolution. To leverage this ability in individual neurons, we first fit phaseSH and SH models for each train-test split of the data in a behavioral epoch of interest, we bin the behavioral epoch into arbitrarily small window sizes (here, 50ms, with no overlap), and then compute the mean log loss differences between the models for each train-test split of data ( $n = 20$  unique log loss differences, one for each train-test split of data). This yields, for each 50ms window, a distribution of log loss differences between the models we wish to compare (e.g. theta phaseSH to SH, or theta phaseSH to low gamma phaseSH, or low gamma phaseSH to SH). For each time window, we then used a two-tailed one-sample t-test to assess whether the distribution of log loss differences (in that window) was significantly different from zero at a Bonferroni-corrected alpha of 0.00004 (that is, 0.001 divided by the total number of time

windows, which corresponds to the number of t-tests performed per neuron). The alpha was corrected to account for the multiple comparisons performed which, uncorrected, would dramatically inflate the likelihood of a Type I error.

### ***Rayleigh Test for Non-Uniformity of Circular Distributions***

To compare the results of the widely used Rayleigh test for non-uniformity against our model comparison approach to identify spike-phase relationships (described above), we computed the Rayleigh statistic for each neuronal spike train. The mean resultant length vector,  $r$ , is defined by,

$$r = \sqrt{\left(\frac{1}{\eta} \sum_{j=1}^{\eta} \cos(\phi_j)\right)^2 + \left(\frac{1}{\eta} \sum_{j=1}^{\eta} \sin(\phi_j)\right)^2} \quad (11)$$

where  $\eta$  corresponds to the total number of spikes in a spike train, and  $\phi_j$  represents a phase angle that co-occurred with spike  $j$ . The mean resultant length vector can then be used to compute the Rayleigh statistic,  $Z$ , where  $Z = r^2 / \eta$ .

Under this framework, the probability that a neuronal spike train generates a  $Z$  statistic of a particular magnitude given that the spike train was in fact drawn from a uniform distribution of spikes over phase is given by,

$$p_{Ray} = \exp(\sqrt{1 + 4\eta + 4(\eta^2 - R^2)} - (1 + 2\eta)) \quad (12)$$

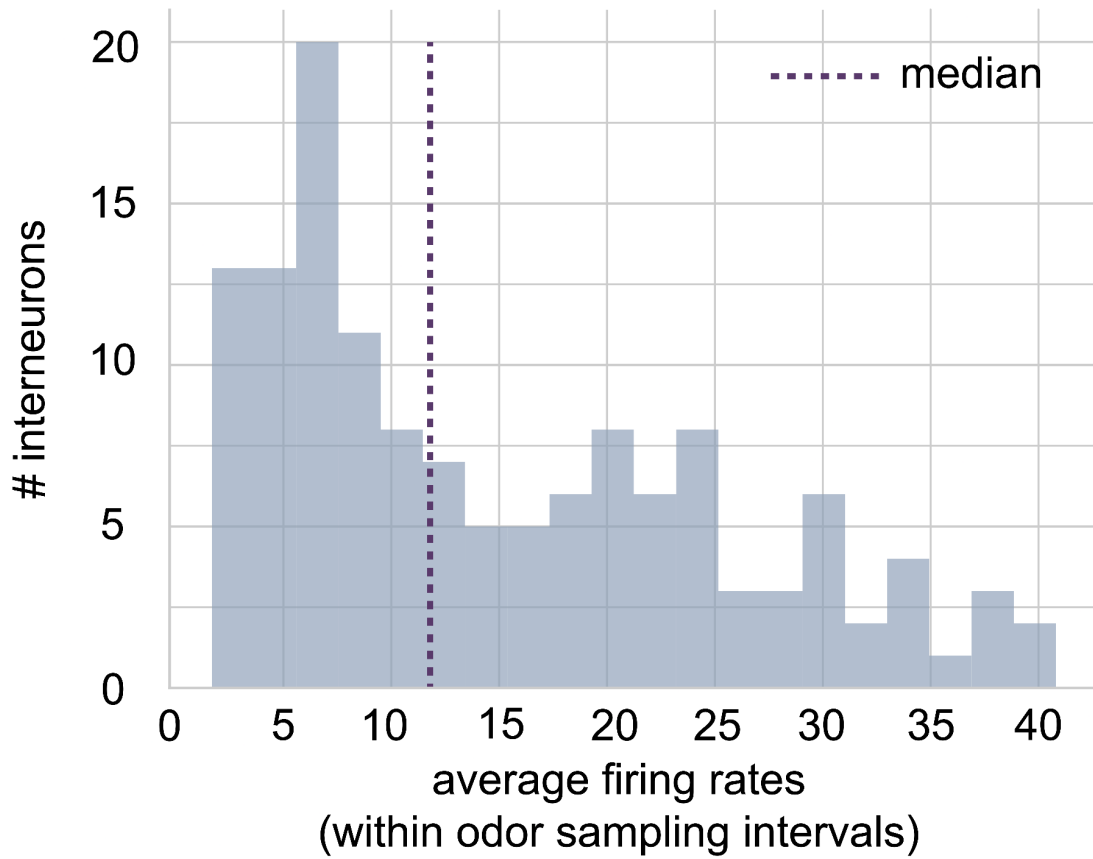
where, as in eq. (11),  $\eta$  corresponds to the total number of spikes in a given spike train, and  $R$  refers to the resultant length vector, expressed as  $R = \eta r$ .

### ***Code Availability***

All code underlying the analysis methods described above is freely available at the following:  
<https://doi.org/10.5281/zenodo.6837168>.



Supplementary Figures occupy the remaining pages of this chapter.



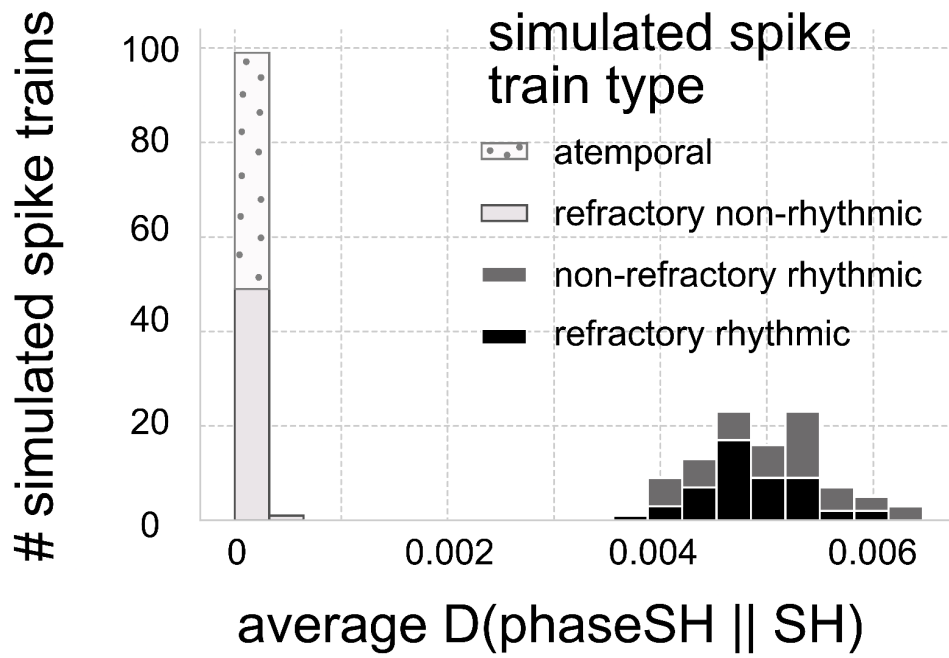


**Supplementary Figure S3.1. Average firing rates across the CA1 interneuron population during odor sampling intervals.** Firing rates were computed over all 1250 ms odor sampling intervals within a session for a given CA1 interneuron. The median average firing rate is 11.8 Hz. A large proportion of the population clusters near 6 Hz, below the median.

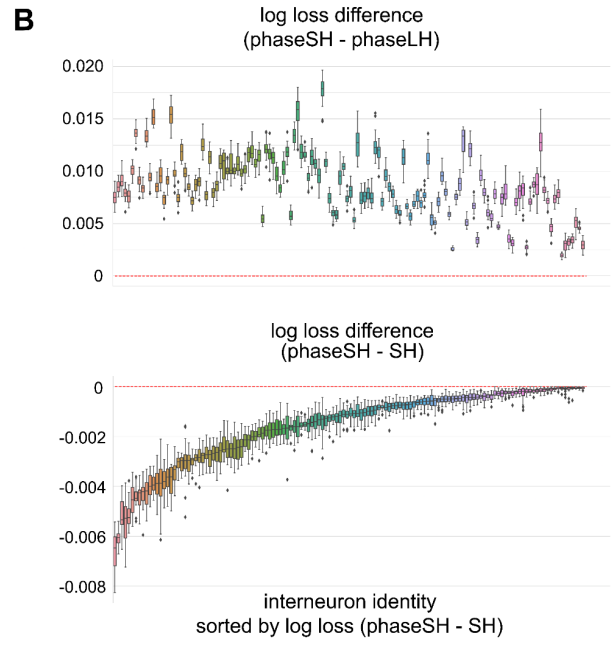
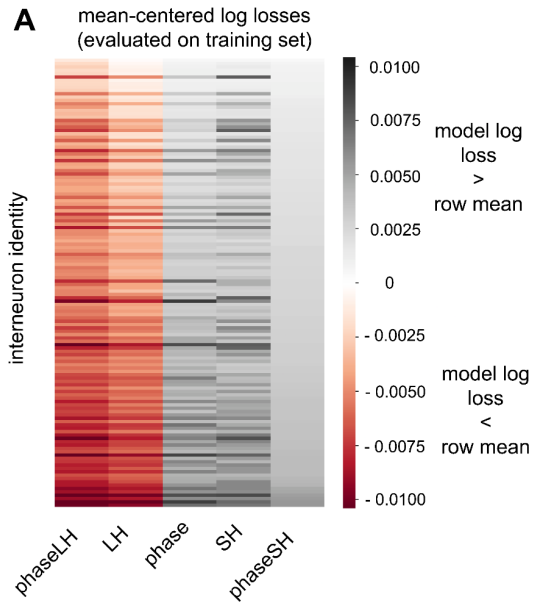
**Supplementary Figure S3.2. Model results in a simulated dataset of spike trains with high average firing rates (40-60 Hz).** (A) Average Kullback-Leibler (KL) divergences of the ground truth from a model's probabilities across 50 simulated spike trains per type. Note the large average divergences incurred by the phase model for spike trains governed by a refractory period (middle left and far right subpanels). (B) Number of simulated spike trains (N=200) by the difference in log losses between phaseSH and SH models. Log loss differences between phaseSH and SH models separate rhythmic simulated spike trains from non-rhythmic ones.

**A****B**

**Supplemental Figure. S3.3. Average Kullback-Leibler (KL) divergence between phaseSH and SH model probabilities separate rhythmic from non-rhythmic simulated spike trains.** Histogram depicts the number of simulated spike trains (with a ~6 Hz average firing rate) with a particular average KL divergence between phaseSH and SH model probabilities. Bars are colored according to the ground truth (“spike train type”) of the contributing spike train (50 spike trains per type).



**Supplemental Figure S3.4. Population log losses evaluated on training trials.** (a) Rows along the heatmap correspond to unique CA1 interneurons; each column corresponds to one of the five models. Color scale indicates the log loss of a model trained and tested on a single fold of a given interneuron's data. Warmer values signal log losses that are smaller than the corresponding row mean. Interneurons were sorted according to their phaseSH model log loss. (b) (top) Differences between phaseSH and phaseLH model log losses (y-axis) evaluated on the same data that was used to fit model coefficients (training trials), across 20 train-test splits of data for each CA1 interneuron (x-axis). Black horizontal bar in boxplots marks the median, whiskers mark the 25th and 75th percentiles of the distribution of train-test splits. The red horizontal line lies along zero, the value that would indicate equal performance across models. (bottom) Same as top panel of B, but for differences between phaseSH and SH model log losses. Interneuron identity in both top and bottom panels are sorted according to the median log loss difference between phaseSH and SH models across the 20 folds.




**Supplementary Figure S3.5. Model comparisons for approach epochs. (A)** Theta phase-based models: Differences between phaseSH and SH model log losses (y-axis) evaluated on held-out data, across 20 folds (train-test splits) of data for each CA1 interneuron (x-axis). Black horizontal bar within boxplots marks the median, whiskers mark the 25th and 75th percentiles of the distribution of folds, diamonds mark outliers. A log loss difference of zero (marked by dotted red line) indicates equal performance across models. Grey bars mark interneurons whose spike-phase distributions (drawn from the combined train and test trial data) significantly differed from uniform according to a Rayleigh test for non-uniformity ( $p\text{-val} < 0.001$ ). Pink bars mark the distributions of folds whose log loss differences significantly differed from zero ( $p < 0.001$ ), according to a one-tailed t-test. (bottom) Interneuron identities are sorted according to the median of the log loss difference between models across folds. **(B)** Low gamma-phase based models: all details same as in A.




**A**

theta (4 - 12 Hz)  
log loss difference  
(phaseSH - SH)

**B**

low gamma (35 - 55 Hz)  
log loss difference  
(phaseSH - SH)




■ Rayleigh p-val < 0.001    ■ log loss diff t-test p-val < 0.001

interneuron identity  
sorted by median log loss difference  
(phaseSH - SH)


**Supplementary Figure S3.6. Model comparisons for CA1 pyramidal cell activity during odor sampling.** (A) (top) Theta phase-based models: Differences between phaseSH and phaseLH model log losses (y-axis) evaluated on held out data, across 20 folds (train-test splits) of data for each CA1 pyramidal cell spike train (x-axis). Black horizontal bar within boxplots marks the median, whiskers mark the 25th and 75th percentiles of the distribution of folds, diamonds mark outliers. A log loss difference of zero (marked by dotted red line) indicates equal performance across complete short and complete long models. Grey bars mark pyramidal cell spike trains whose spike-phase distributions (drawn from the combined train and test trial data) significantly differed from uniform according to a Rayleigh test for non-uniformity ( $p\text{-val} < 0.001$ ). Pink bars mark the distributions of folds whose log loss differences significantly differed from zero ( $p < 0.001$ ), according to a one-tailed t-test. (bottom) Same as the top panel, but for differences between phaseSH and SH model log losses. Pyramidal cell spike train identity in both top and bottom panels are sorted according to the median of the log loss difference between phaseSH and SH models across folds. (B) Low gamma-phase based models: all details same as in A.

**A**

**theta (4 - 12 Hz)**  
log loss difference  
(phaseSH - phaseLH)



log loss difference  
(phaseSH - SH)




■ Rayleigh p-val < 0.001

**B**

**low gamma (35 - 55 Hz)**  
log loss difference  
(phaseSH - phaseLH)



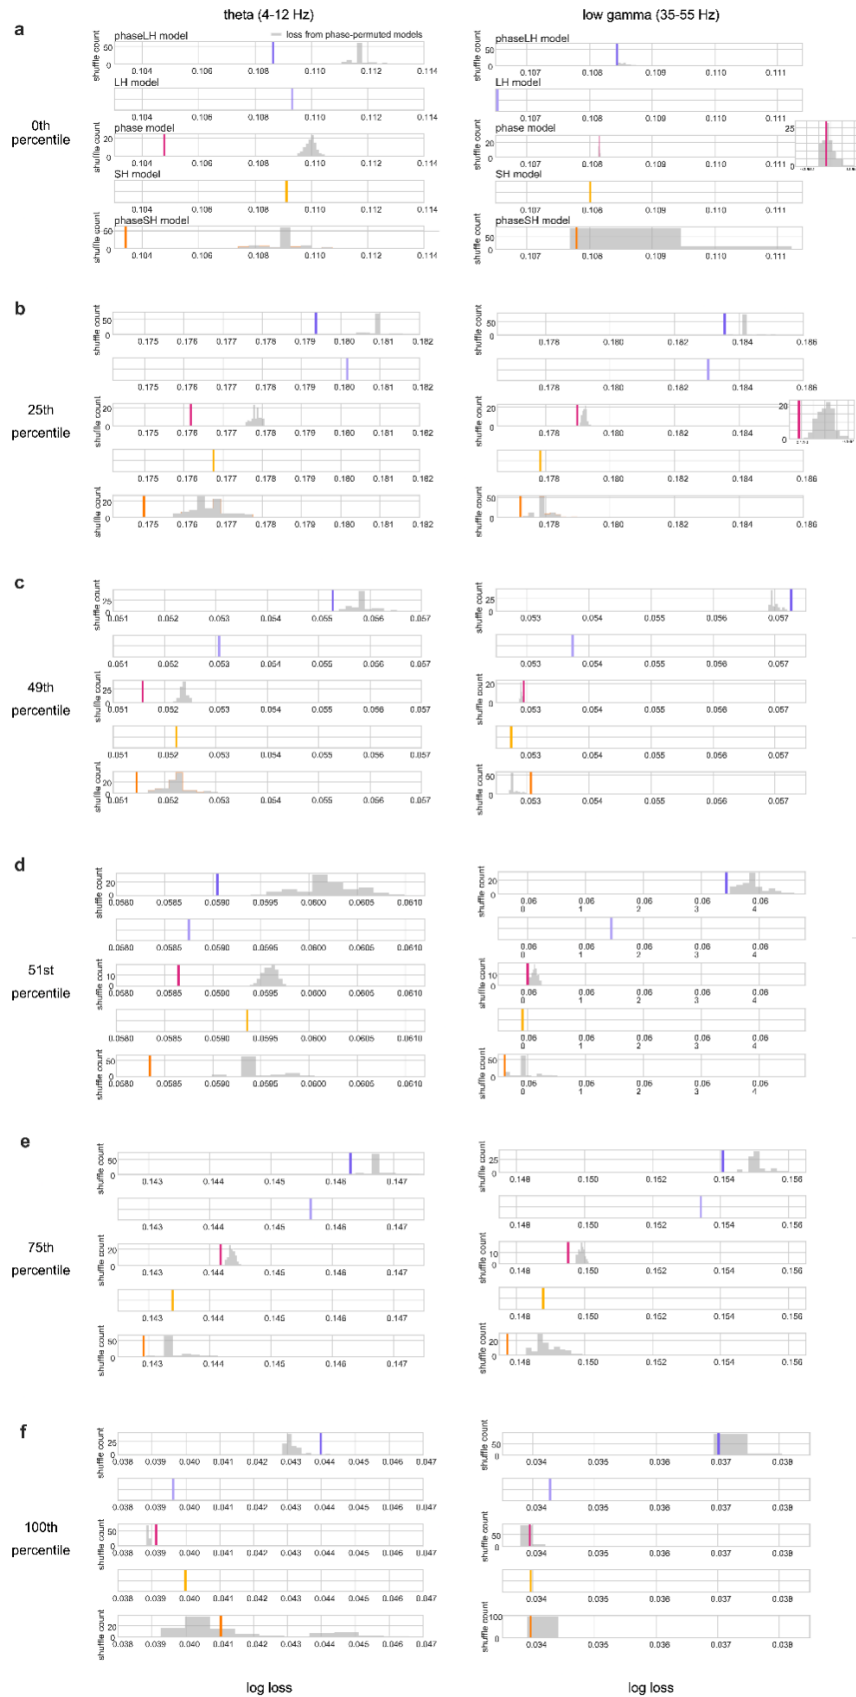
log loss difference  
(phaseSH - SH)



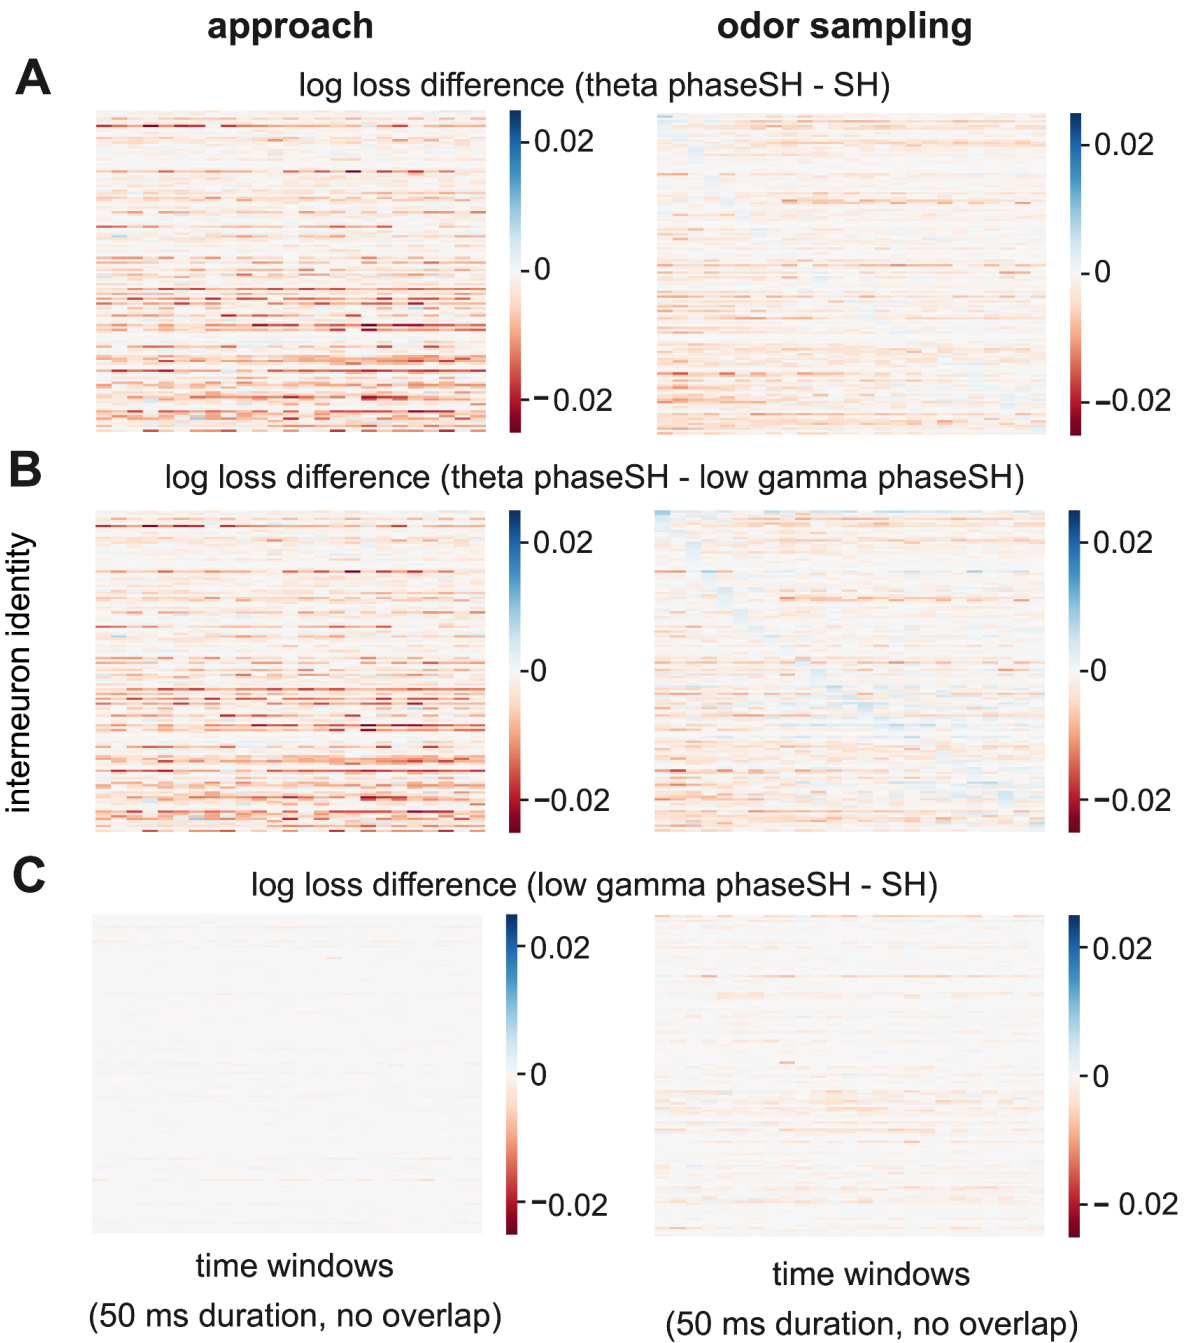
■ log loss diff t-test p-val < 0.001

pyramidal cell identity  
sorted by median log loss difference  
(phaseSH - SH)

**Supplemental Figure S3.7. Model log loss permutation tests.** (A) Model log loss performance (vertical tick marks) for phaseLH history (dark purple), LH (light purple), phase (pink), SH (yellow), and phaseSH (orange) models relative to distributions of 100 log losses drawn from corresponding models trained on permuted phase predictors (grey histogram). Phase observations were randomly shuffled within cycles, while spike times (and consequently spike history) remained intact. The data for phase-permuted models were sourced from the same training trials as intact-phase models, and were tested on the same held out set of trials. Left column corresponds to models including theta phase information, while right column corresponds to models including low gamma phase information for the same interneuron. (B - F) Same as in A, but for distinct CA1 interneuron spike trains. Interneurons were selected according to whether the median log loss difference between their theta phaseSH and SH models corresponded to the 0th, 25th, 49th, 51st, 75th, and 100th percentiles, respectively, of the distribution of all CA1 interneuron median log loss difference, in the bottom subpanel of Fig. 3.3A.



**Supplemental Figure S3.8. Model comparisons over time across the population across behavioral states.** Models were fit to spiking and rhythmic phase data from behavioral epochs when rats were either running towards the odor port (left column, approach) or odor sampling (right column). (A) Rows correspond to three unique interneuron spike train identities, matched across approach and odor sampling behavioral epochs, and are sorted according to the time window with the most positive value for the log loss difference (theta phaseSH - low gamma phaseSH) model comparison. Columns within each heatmap correspond to non-overlapping 50 ms-long time windows spanning a continuous interval of 1250 ms. Color scale corresponds to log loss differences between the theta phaseSH and SH model computed on held out data for each of the 20 train test splits. Warmer red colors indicated that theta phaseSH models outperformed SH models. (B) Same as in A, for log loss differences between the theta phaseSH model and the low gamma phaseSH model. Warmer red colors indicated better performance by theta phaseSH models, while cooler blue colors indicated better performance by low gamma phaseSH models. (C) Same as in A, but for log loss differences between low gamma phaseSH and SH models. Warmer red colors in this heatmap indicated better performance for low gamma phase SH models.



### 3.6 References

1. Egea-Weiss A, Renner A, Kleineidam CJ, Szyszka P. High precision of spike timing across olfactory receptor neurons allows rapid odor coding in *Drosophila*. *iScience*. 29(4), 76-83, 2018. <http://doi.org/10.1016/j.isci.2018.05.009>
2. Marder E, Bucher D. Central pattern generators and the control of rhythmic movements. *Current Biology*. 11(23), R986-R996, 2001. [http://doi.org/10.1016/S0960-9822\(01\)00581-4](http://doi.org/10.1016/S0960-9822(01)00581-4)
3. Getting PA. Mechanisms of pattern generation underlying swimming in *Tritonia*. II. Network reconstruction. *Journal of Neurophysiology*. 49(4), 1017-1035, 1983.
4. von Reyn CR, Breads P, Peek MY, Zheng GZ, Williamson WR, Yee AL, Leonardo A, Card GM. A spike timing mechanism for action selection. *Nature Neuroscience*. 17, 962-970 (2014). <https://doi.org/10.1038/nn.3741>
5. Hodgkin AL, Huxley AF. Currents carried by sodium and potassium ions through the membrane of the giant axon of *Loligo*. *J. Physiol*. 116, 449-472, 1952a. [doi: 10.1113/jphysiol.1952.sp004717](https://doi.org/10.1113/jphysiol.1952.sp004717)
6. Hodgkin AL, Huxley AF. The components of membrane conductance in the giant axon of *Loligo*. *J. Physiol*. 116, 473-496, 1952b. [doi: 10.1113/jphysiol.1952.sp004718](https://doi.org/10.1113/jphysiol.1952.sp004718)
7. Hodgkin AL, Huxley AF. The dual effect of membrane potential on sodium conductance in the giant axon of *Loligo*. *J. Physiol*. 117, 497-506, 1952c. [doi: 10.1113/jphysiol.1952.sp004719](https://doi.org/10.1113/jphysiol.1952.sp004719)
8. Burchell TR, Faulkner HJ, Whittington MA. Gamma frequency oscillations gate temporally coded afferent inputs in the rat hippocampal slice. *Neuroscience Letters*. 255(3), 151-154, 1998. [https://doi.org/10.1016/S0304-3940\(98\)00676-4](https://doi.org/10.1016/S0304-3940(98)00676-4)
9. Volgushev M, Chistiakova M, Singer W. Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience*. 83(1), 15-25, 1998. [https://doi.org/10.1016/S0306-4522\(97\)00380-1](https://doi.org/10.1016/S0306-4522(97)00380-1)
10. Leung LWS, Yim CYC. Intrinsic membrane potential oscillations in hippocampal neurons *in vitro*. *Brain Research*. 553(2), 261-274, 1991. [https://doi.org/10.1016/0006-8993\(91\)90834-I](https://doi.org/10.1016/0006-8993(91)90834-I)



11. Wilson CJ. The mechanism of intrinsic amplification of hyperpolarizations and spontaneous bursting in striatal cholinergic interneurons. *Neuron*. 45(4), 575-585, 2005. <https://doi.org/10.1016/j.neuron.2004.12.053>
12. Logothetis NK. The underpinnings of the BOLD functional magnetic resonance imaging signal. *J. Neurosci*. 23(10), 3963-3971, 2003. <https://doi.org/10.1523/JNEUROSCI.23-10-03963.2003>
13. Buzsáki G, Anastassiou CA, Koch C. The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nature Reviews Neuroscience*. 13(6), 407-420, 2012. <https://doi.org/10.1038/nrn3241>
14. Ranck Jr JB. Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats: Part I Behavioral correlates and firing repertoires. *Experimental Neurology*. 41(2), 462-531, 1973. [https://doi.org/10.1016/0014-4886\(73\)90290-2](https://doi.org/10.1016/0014-4886(73)90290-2)
15. Buzsáki G, Eidelberg E. Phase relations of hippocampal projection cells and interneurons to theta activity in the anesthetized rat. *Brain Research*. 266(2), 334-339, 1983. [https://doi.org/10.1016/0006-8993\(83\)90665-0](https://doi.org/10.1016/0006-8993(83)90665-0)
16. Varga C, Golshani P, Soltesz I. Frequency-invariant temporal ordering of interneuronal discharges during hippocampal oscillations in awake mice. *Proc. Natl. Acad. Sci. U.S.A.* 109(40), E2726-E2734, 2012. <https://doi.org/10.1073/pnas.1210929109>
17. Czurkó A, Huxter J, Li Y, Hangya B, Muller RU. Theta phase classification of interneurons in the hippocampal formation of freely moving rats. *J. Neurosci*. 31(8), 2938-2947, 2011. <https://doi.org/10.1523/JNEUROSCI.5037-10.2011>
18. Stark E, Eichler R, Roux L, Fujisawa S, Rotstein HG, Buzsáki G. Inhibition-induced theta resonance in cortical circuits. *Neuron*. 80(5) 1263-1276, 2013. <https://doi.org/10.1016/j.neuron.2013.09.033>
19. Amilhon B, Huh CYL, Manseau F, Ducharme G, Nichol H, Adamantidis A, Williams S. Parvalbumin interneurons of hippocampus tune population activity at theta frequency. *Neuron*. 86(5), 1277-1289, 2015. <https://doi.org/10.1016/j.neuron.2015.05.027>

20. Maccaferri G, McBain CJ. The hyperpolarization-activated current ( $I_h$ ) and its contribution to pacemaker activity in rat CA1 hippocampal stratum oriens-alveus interneurons. *J. Physiol.* 497(1), 119-130, 1996. <https://doi.org/10.1113/jphysiol.1996.sp021754>
21. Pike FG, Goddard RS, Suckling JM, Ganter P, Kasthuri N, Paulsen O. Distinct frequency preferences of different types of rat hippocampal neurones in response to oscillatory input currents. *J. Physiol.* 529(1), 205-213, 2000. <https://doi.org/10.1111/j.1469-7793.2000.00205.x>
22. Wierenga CJ, Wadman WJ. Excitatory inputs to CA1 interneurons show selective synaptic dynamics. *J. Neurophysiol.* 90, 811-821, 2003. <https://doi.org/10.1152/jn.00865.2002>
23. Kajiwara R, Wouterlood FG, Sah A, Boekel AJ, Baks-te Bulte LTG, Witter MP. 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. <https://doi.org/10.1002/hipo.20385>
24. Basu J, Srinivas KV, Cheung SK, Taniguchi H, Huang ZJ, Siegelbaum SA. A cortico-hippocampal learning rule shapes inhibitory microcircuit activity to enhance hippocampal information flow. *Neuron.* 79(6), 1208-1221, 2013. <https://doi.org/10.1016/j.neuron.2013.07.001>
25. Basu J, Zaremba JD, Cheung SK, Hitti FL, Zemelman BV, Losonczy A, Siegelbaum SA. Gating of hippocampal activity, plasticity, and memory by entorhinal cortex long-range inhibition. *Science.* 351(6269), aaa5694, 2016. DOI: [10.1126/science.aaa5694](https://doi.org/10.1126/science.aaa5694)
26. Bland SK, Bland BH. Medial septal modulation of hippocampal theta cell discharges. *Brain Research.* 375(1), 102-116, 1986. [https://doi.org/10.1016/0006-8993\(86\)90963-7](https://doi.org/10.1016/0006-8993(86)90963-7)
27. Smythe JW, Colom LV, Bland BH. The extrinsic modulation of hippocampal theta depends on the coactivation of cholinergic and GABA-ergic medial septal inputs. *Neuroscience & Biobehavioral Reviews.* 16(3), 289-308, 1992. [https://doi.org/10.1016/S0149-7634\(05\)80203-9](https://doi.org/10.1016/S0149-7634(05)80203-9)
28. Lee MG, Chrobak JJ, Sik A, Wiley RG, Buzsáki G. Hippocampal theta activity following selective lesion of the septal cholinergic system. *Neuroscience.* 62(4), 1033-1047, 1994. [https://doi.org/10.1016/0306-4522\(94\)90341-7](https://doi.org/10.1016/0306-4522(94)90341-7)

29. Truccolo W, Eden UT, Fellows MR, Donoghue JP, Brown EN. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *J. Neurophysiol.* 93(2), 1074-1089, 2005. <https://doi.org/10.1152/jn.00697.2004>
30. Paninski L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Comput. Neural Syst.* 15(4), 243-262, 2004. [doi:10.1088/0954-898X/15/4/002](https://doi.org/10.1088/0954-898X/15/4/002)
31. Brillinger DR. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol. Cybern.* 59, 189-200, 1988. <https://doi.org/10.1007/BF00318010>
32. Pillow J. Likelihood-based approaches to modeling the neural code. In *Bayesian Brain: Probabilistic Approaches to Neural Coding*, edited by Poya K, Ishii S, Pouget A, Rao RP, Eds. MIT Press, 2007, pp. 53-70.
33. Banerjee A, Dean HL, Pesaran B. Parametric models to relate spike train and LFP dynamics with neural information processing. *Front. Comput. Neurosci.* 6, 51, 2012. <https://doi.org/10.3389/fncom.2012.00051>
34. Rivlin-Etzion M, Ritov Y, Heimer G, Bergman H, Bar-Gad I. Local shuffling of spike trains boosts the accuracy of spike train spectral analysis. *Journal of Neurophysiology.* 95(5), 3245-3256, 2006. <https://doi.org/10.1152/jn.00055.2005>
35. Rangel LM, J. W. Rueckemann JW, Rivière PD, Keefe KR, Porter BS, Heimbuch IS, Budlong CH, Eichenbaum H. Rhythmic coordination of hippocampal neurons during associative memory processing. *eLife.* 5, e09849, 2016. DOI: [10.7554/eLife.09849](https://doi.org/10.7554/eLife.09849)
36. Rivière PD, Rangel LM. Spike-field coherence and firing rate profiles of CA1 interneurons during an associative memory task. In *Advances in the Mathematical Sciences, AWMRS 2017, Association for Women in Mathematics Series*, edited by Deines A, Ferrero D, Graham E, Im M, Manore C, Price C, Springer, 2018, pp. 161-171. [https://doi.org/10.1007/978-3-319-98684-5\\_10](https://doi.org/10.1007/978-3-319-98684-5_10)
37. Barthó P, Hirase H, Monconduit L, Zugaro M, Harris KD, Buzsáki G. Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* 92(1), 600-608, 2004. <https://doi.org/10.1152/jn.01170.2003>

38. James G, Witten D, Hastie T, Tibshirani R. An introduction to statistical learning. New York: Springer, 2013, vol. 112.
39. Hall P, Wand MP. On the accuracy of binned kernel density estimators. *Journal of Multivariate Analysis*. 56, 165-184, 1996. <https://doi.org/10.1006/jmva.1996.0009>
40. Johnson TD, Coleman TP, Rangel LM. A flexible likelihood approach for predicting neural spiking activity from oscillatory phase. *Journal of Neuroscience Methods*. 311, 307-317, 2019. <https://doi.org/10.1016/j.jneumeth.2018.10.028>
41. Goldberg JA, Rokni U, Boraud T, Vaadia E, Bergman H. Spike synchronization in the cortex-basal ganglia networks of parkinsonian primates reflects global dynamics of the local field potentials. *Journal of Neuroscience*. 24(26), 6003-6010.
42. Pillow JW. Time-rescaling methods for the estimation and assessment of non-Poisson neural encoding models. In *Advances in Neural Information Processing Systems*, edited by Bengio Y, Schuurmans D, Lafferty J, Williams C, Culotta A, Cambridge, MA: MIT Press, 2009, pp. 1473-1481.
43. Pillow J, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*. 454, 995-999, 2008. <https://doi.org/10.1038/nature07140>
44. Hoerl AE, Kennard RW. Ridge Regression: Biased estimation for nonorthogonal problems. *Technometrics*. 12(1), 55-67, 1970. <https://doi.org/10.1080/00401706.1970.10488634>
45. Vanderwolf CH. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and clinical neurophysiology*. 26(4), 407-418, 1969. [https://doi.org/10.1016/0013-4694\(69\)90092-3](https://doi.org/10.1016/0013-4694(69)90092-3)
46. Sławińska U, Kasicki S. The frequency of rat's hippocampal theta rhythm is related to the speed of locomotion. *Brain Research*. 796(1-2), 327-331, 1998. [https://doi.org/10.1016/S0006-8993\(98\)00390-4](https://doi.org/10.1016/S0006-8993(98)00390-4)

47. Siapas AG, Lubenov EV, Wilson MA. Prefrontal phase locking to hippocampal theta oscillations. *Neuron*. 46(1), 141-151, 2005. <https://doi.org/10.1016/j.neuron.2005.02.028>
48. Poo C, Isaacson JS. Odor representations in olfactory cortex: "sparse" coding, global inhibition, and oscillations. *Neuron*. 62(6), 850-861 (2009). <https://doi.org/10.1016/j.neuron.2009.05.022>
49. Fries P, Reynolds JH, Rorie AE, Desimone R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*. 291(5508), 1560-1563, 2001. DOI: [10.1126/science.1055465](https://doi.org/10.1126/science.1055465)
50. Landau AN, Schreyer HM, van Pelt S, Fries P. Distributed attention is implemented through theta-rhythmic gamma modulation. *Current Biology*. 25(17), 2332-2337, 2015. <https://doi.org/10.1016/j.cub.2015.07.048>
51. Leventhal DK, Gage GJ, Schmidt R, Pettibone JR, Case AC, Berke JD. Basal ganglia beta oscillations accompany cue utilization. *Neuron*. 73(3), 523-536, 2012. <https://doi.org/10.1016/j.neuron.2011.11.032>
52. Rubino D, Robbins KA, Hatsopoulos NG. Propagating waves mediate information transfer in the motor cortex. *Nature Neuroscience*. 9(12), 1549-1557, 2006. <https://doi.org/10.1038/nn1802>
53. Kim S, Quinn CJ, Kiyavash N, Coleman TP. Dynamic and succinct statistical analysis of neuroscience data, *invited paper, Proceedings of the IEEE*. 102(5), 683-698, 2014. DOI: [10.1109/JPROC.2014.2307888](https://doi.org/10.1109/JPROC.2014.2307888)
54. Takahashi K, Kim S, Coleman TP, Brown KA, Suminski AJ, Best MD, Hatsopoulos NG. Large-scale spatiotemporal spike patterning consistent with wave propagation in motor cortex. *Nature Communications*. 6, 7169, 2015. doi: [10.1038/ncomms8169](https://doi.org/10.1038/ncomms8169)
55. Klein AS, Donoso JR, Kempter R, Schmitz D, Beed P. Early cortical changes in gamma oscillations in Alzheimer's disease. *Front. Syst. Neurosci.* 10, 83, 2016. <https://doi.org/10.3389/fnsys.2016.00083>
56. Mably AJ, Colgin LL. Gamma oscillations in cognitive disorders. *Current Opinion in Neurobiology*. 52, 182-187, 2018. <https://doi.org/10.1016/j.conb.2018.07.009>

57. McCarthy MM, Moore-Kochlacs C, Gu X, Boyden ES, Han X, Kopell N. Striatal origin of the pathologic beta oscillations in Parkinson's disease. *Proc. Natl. Acad. Sc.* 108(28) 1160-11625, 2011. <https://doi.org/10.1073/pnas.1107748108>

### **Acknowledgments**

The authors would like to thank Sean T. Trott and James B. Priestley, for extensive and insightful discussions regarding statistical models and their implementation, as well as all members of the Neural Crossroads Laboratory for their support, encouragement, and discussions of this work.

**Grants:** This research was partially funded by:

Frontiers of Innovation Scholars Program, FISP (LMR, TPC)

Hellman Fellowship (LMR)

Kavli Institute for Brain and Mind (LMR, PDR, TPC, GS)

National Institutes of Health grant R01MH1105142 (AC – Andrea Chiba, TPC, LMR)

National Institutes of Health grant R03MH120406 (LMR, TPC)

National Institutes of Mental Health MH052090 (HE – Howard Eichenbaum)

National Science Foundation Graduate Research Fellowship Program (PDR)

National Science Foundation DMS-1042134 (LMR)

National Science Foundation Science of Learning Center grant SBE 0542013, Temporal

Dynamics of Learning Center (PDR, LMR)

**Disclosures:** Authors declare that no conflict of interests exist.

Chapter 3, in full, is a reprint of the material as it appears in the *Journal of Neurophysiology*. Rivière, Pamela D; Schamberg, Gabriel; Coleman, Todd P; Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.

## CHAPTER 4: BEHAVIORALLY DEPENDENT HIPPOCAMPAL INTERNEURON ENTRAINMENT TO MULTIPLE RHYTHMS

### 4.1 Abstract

Inhibition orchestrates neuronal responses to inputs, and sculpts the oscillatory dynamics of brain regions. Characterizing inhibitory interneuron spike timing with respect to rhythms consequently provides insight into the mechanisms of network organization in response to behavioral demands. Relatively few studies have assessed the dynamism of interneuron oscillatory spiking across multiple rhythms, changing behaviors, and varying sensory inputs. To address this gap, we statistically modeled the relationship between CA1 interneuron spike timing and the rhythmic phase of local field potential (LFP) oscillations as rats performed a context-guided associative memory paradigm. CA1 interneurons organized into groups with preferential engagement in distinct phases of the theta rhythm, and theta phase preference predicted the phase of their entrainment to higher frequency low gamma (35-55 Hz) and high gamma (65-90 Hz) rhythms. Notably, patterns of engagement in low gamma—but not theta or high gamma—varied across associative memory processing and goal-oriented locomotive behavioral states. Within associative memory processing epochs, distinct olfactory and visuospatial inputs provoked changes to theta, low gamma, and high gamma spike-phase relationships, with predominant shifts involving entrainment to a particular rhythmic phase, or the complete lack thereof. Interneuron spike-phase relationships can consequently be flexible—rather than fixed—when multiple behaviors and inputs are considered. Our results suggest that inhibition may dynamically shape hippocampal oscillatory dynamics to meet changing behavioral demands.

### 4.2 Introduction

Continuously evolving environments impose extreme demands on organisms. To mitigate the metabolic costs involved in foraging for food, and the dangers inherent to wandering into predators' territories, organisms must learn to associate the features of an environment (e.g. what it looks like, what it smells like) with the positive or adverse outcomes they predict. This enables organisms to accurately infer the future presence of rewards or dangers while minimizing risks. Hippocampal subregions have been implicated in the formation of precisely these types of associations<sup>1-3</sup>, with the CA1 subregion serving as a primary output node from the hippocampal trisynaptic network<sup>4-6</sup>. To effectively support

associations between sensory, visuospatial, and affective features of an environment, this subregion must parse a wide array of input streams from cortical<sup>7-10</sup>, subcortical<sup>8,11-13</sup>, and intrahippocampal<sup>8,9</sup> structures. In the present work, we leveraged an associative memory paradigm to assess how various combinations of rewarded olfactory and visuospatial inputs recruited putative inhibitory interneurons recorded from the hippocampal CA1 subregion.

Inhibition plays a key organizing role in neural circuit dynamics<sup>14</sup>. In the CA1, inhibitory interneurons bias the array of effective inputs that ultimately shape excitatory principal cells' outputs<sup>8,15-17</sup>, and in doing so, they help sculpt the region's oscillatory profile<sup>18,19</sup>. The two major CA1 input streams—CA3 Schaffer collaterals and entorhinal cortex perforant path fibers—each recruit feedforward inhibition as they attempt to entrain principal cells in the network<sup>7,20-22</sup>. A diverse array of inhibitory interneurons carefully orchestrates pyramidal cells' responses to these inputs, by targeting inhibition onto distinct pyramidal cell dendritic and somatic compartments<sup>8,9,16,17,23,24</sup>. Previous studies have shown that the activity of these distinct interneuron types can work to enhance the influence of each afferent pathway. For instance, CA1 interneurons targeting pyramidal cells' distal dendrites are selectively recruited by subcortical neuromodulatory inputs, a process that transiently blocks perforant path inputs while simultaneously privileging Schaffer collaterals synapsing onto pyramidal cells' proximal dendrites<sup>8</sup>. Transient blocks of certain inputs need not constitute a competitive mechanism: in fact, enforcing a temporal sequence where Schaffer collateral inputs precede perforant path inputs results in *enhanced* pyramidal cell voltage responses to the latter<sup>9</sup>. Inhibition can consequently work cooperatively with various afferent streams to scaffold the temporal relationships that shape hippocampal output.

The coordinated interplay of extrinsic inputs, local inhibitory recruitment, and pyramidal cells' responses results in systematic patterns of ionic currents, which superimpose to produce rhythmic fluctuations of the CA1 local field potential (LFP)<sup>25,26</sup>. LFP oscillations are remarkably dynamic, with the dominant frequencies within a region evolving over time and as a function of the animal's behavioral state<sup>19,27-32</sup>. Notable frequency ranges in the CA1 include the theta (4-12 Hz) oscillation, which covaries with rodent locomotion<sup>30-32</sup> and sniffing<sup>28</sup>, a behavioral affordance that provides rodents with access to



high-fidelity olfactory information<sup>33</sup>. Optogenetic silencing of parvalbumin positive (PV+) interneurons in the CA1 disrupts theta LFP oscillations, suggesting an important inhibitory contribution to this hippocampal rhythm<sup>18</sup>. Very fast rhythmic fluctuations also emerge in the CA1. These faster frequencies can be subdivided into distinct ranges, which differ slightly depending on the study<sup>34–36</sup>, but generally fall into lower (35-55 Hz) and higher (65-90 Hz) gamma bands. Notably, many studies find that slower and faster gamma frequency bands manifest along distinct phases of the theta cycle<sup>34,35,37,38</sup>, a result that aligns with studies showing that distinct inputs influence the CA1 sequentially within the timescale of a theta cycle<sup>9,39</sup>. Inhibitory activity is additionally implicated in the generation of gamma oscillations<sup>25,40–42</sup>, and the diversity of interneuron types<sup>23,43</sup> may result in differential recruitment of interneuron classes into distinct gamma frequencies<sup>36</sup>.

A combination of intrinsic resonance properties and preferential recruitment by distinct neuromodulatory, perforant path, and Schaffer collateral afferents may render distinct CA1 interneuron types more or less likely to participate in slow versus high gamma oscillations. A number of studies have attempted to map the morphological and molecular properties of interneurons to their electrical activity patterns. This work has successfully correlated interneuron expression of molecular markers (such as parvalbumin, somatostatin, among others) to these cells' tendency to spike according to particular phases of the hippocampal theta oscillation<sup>43–48</sup>. Valuable extensions of this work additionally characterized interneurons according to their spike-phase relationships relative to faster oscillations—often treating gamma as a single, broad oscillatory band spanning frequencies within the 25-90 Hz<sup>49,50</sup> range, or above<sup>44,50</sup>. These studies have identified some degree of systematicity between molecularly and morphologically characterized interneuron types and their spike-phase relationships to faster gamma rhythms. With a finer grained separation of gamma frequency bands as they manifest during complex behaviors, it is possible that we may observe interneurons preferentially participating in distinct gamma oscillatory processes.

In this study, we were concerned with assessing CA1 inhibitory interneurons' spike-phase relationships across theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) rhythms in

response to distinct combinations of olfactory and visuospatial inputs during associative memory processing. To this end, we evaluated the spike timing patterns of putative CA1 interneurons recorded during a context-guided associative memory task, where odorants' reward values depended on the spatial context in which they were experienced<sup>1,27</sup>. Similarly to previous work, in the present study we were able to cluster inhibitory interneurons according to their spike-phase relationships to the theta (4-12 Hz) oscillation<sup>43,45</sup>. Unlike the bulk of prior studies, however—which explored these relationships either under anesthesia<sup>43</sup> or during locomotive behaviors<sup>45</sup>—we examined interneuron theta spike-phase relationships during an odor sampling, associative memory processing window. To date, it remained unclear from prior work whether interneuron theta spike-phase relationships constitute a relatively stable intrinsic neuronal property, or whether distinct input combinations can flexibly recruit interneurons into distinct theta spike-phase relationships. Here, we leveraged various combinations of odorants and visuospatial information to show that spike timing relationships to the theta rhythm are not necessarily a fixed property of inhibitory interneuron types, and can be flexibly recruited into either theta-entrained or theta-insensitive states in response to different combinations of inputs.

We additionally explored the relationship between theta spike-phase relationships and interneuron entrainment to higher frequency oscillations. We found systematic relationships between inhibitory spike timing organization according to particular theta phases and the likelihood of entrainment to either low gamma (35-55 Hz) or high gamma (65-90 Hz) rhythms. In particular, we observed that the likelihood of entrainment to low gamma phase differed as a function of behavioral demand (associative memory processing versus goal-oriented locomotion), while the likelihood of entrainment to high gamma frequencies given theta phase preference remained relatively stable across behavioral states. Surprisingly, relationships between theta phase preference and high frequency relationships were not necessarily simultaneous or nested in nature: entrainment across rhythms could be systematic but occur at orthogonal times over the course of associative memory processing.

This work provides a unique perspective on the degree to which distinct olfactory and visuospatial inputs flexibly or stably recruit interneurons into theta spike-phase relationships, a property

that has often been discussed as a fixed intrinsic feature of interneurons that facilitates classification of interneurons into types. We learn here that inhibitory interneuron dynamics during an important evolutionarily-honed skill—that of leveraging arbitrary stimuli combinations to predict the likely future presence of rewards—can be flexibly recruited into a constellation of rhythmic spike timing patterns.

## 4.2 Results

To characterize inhibitory interneuron spike-phase relationships during associative memory processing, we leveraged a context-guided associative memory paradigm<sup>1,27</sup>. Odors in a pair were differentially rewarded depending on the spatial context in which they were presented, and rats had to sustain a nose poke in the odor port containing the rewarded odor for 1500 ms—an epoch referred to as “odor sampling”—in order to receive a water reward (**Figure 4.1A**). During this task, we performed high-density extracellular tetrode recordings of the CA1 subregion of the hippocampus, obtaining a sample of 31 putative inhibitory interneurons (see STAR Methods). To specifically examine CA1 interneuron spike timing as a function of olfactory and visuospatial inputs, we split interneuron spike trains and local field potentials (LFPs) according to the odor-context combination present during the odor sampling epoch. Within an experimental session, rats experienced four consecutive blocks of odor pairs (24 trials per block, with a maximum of 12 rewarded trials per odor in a pair), resulting in a maximum of 8 spike trains per inhibitory interneuron. We excluded odor blocks where the performance criterion ( $\geq 75\%$  accuracy) was not met. To assess the effects of behavioral state on inhibitory dynamics, we additionally evaluated interneuron spike trains during a time-matched (1500 ms) locomotive approach epoch preceding corresponding odor sampling epochs.

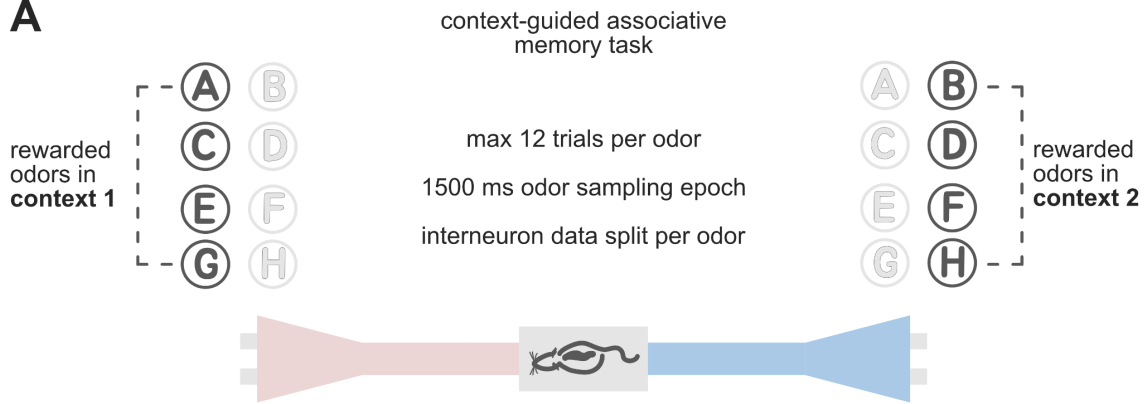

Next, we applied kernel density estimation to model<sup>51</sup> the relationships between interneuron spike timing and the phase of the theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) LFP oscillations (**Figure 4.1B**). Each interneuron spike train was summarized by these three rhythmic “phase models”, which describe the probability of spiking given the phase of the respective CA1 oscillation. We fit rhythmic phase models to a total of 206 (odor sampling) and 207 (approach) interneuron spike trains,

rather than the expected maximum of 248 spike trains. These discrepancies stem from our spike train exclusion criteria: we excluded spike trains corresponding to low accuracy odor blocks, as well as spike trains recorded during periods in which interneurons did not emit the minimum number of spikes necessary to fit the rhythmic phase models (see Methods).

### ***CA1 Interneurons' Theta Spike-Phase Relationships Predict Entrainment to Higher Frequencies.***

After having fit the rhythmic phase models to each spike train, we first examined whether CA1 interneuron spike timing was organized according to stereotyped phases of the theta oscillation during associative memory processing, as previously reported under anesthesia<sup>43</sup> and awake, locomotive<sup>45</sup> epochs. We used the circular mean of phase angles that co-occurred with spiking, as well as the spike-phase distributions' ability to meet the Rayleigh statistic for non-uniformity ( $p < 0.05$ ), to cluster interneuron spike trains (**Figure 4.2A**). This classification scheme yielded four distinct clusters of phase models, with peaks in probability of spiking along the trough (yellow, 76 out of 206 spike trains, 37.1%), ascending (blue, 63 out of 206 spike trains, 30.7%), and the early descending (orange, 12 out of 206 spike trains, 5.85%) phase of the theta cycle. The distances between the centroids of each theta cluster were as follows: trough (yellow) to ascending phase (blue): 1.91 rad, 109.48 deg; trough (yellow) to early descending phase (orange): 1.79 rad, 102.95 deg; and ascending (blue) to early descending (orange): 2.57 rad, 147.56 deg (**Figure 4.2B, top**). We additionally identified a cluster of spike trains with a uniform distribution (hot pink, 55 out of 206, 26.7%) of spikes along the theta cycle, termed “theta-insensitive”<sup>45</sup>, “tonic”<sup>52</sup>, or “theta-off”<sup>53</sup> interneurons in earlier work.

**Figure 4.1. Behavioral variables and model fitting procedure. (A)** The context-guided associative memory task takes place on a linear maze whose distinct compartments, context 1 (pale pink) and context 2 (blue), are separated by a rectangular start box (depicted in grey, with schematic rat). The rat waits in the start box until granted access to either one of the two contexts on a given trial. Two odor ports lie at the end of each context (grey squares). Rats learn that odors in a pair (e.g. A, B) are differentially rewarded depending on the context in which they are presented (e.g. A rewarded in context 1, B rewarded in context 2). Each odor port emits a single odor on a given trial, and the odor delivered per odor port is pseudo randomized across trials. To obtain a reward, rats must maintain a nose poke in the odor port containing the rewarded odor for a duration of 1500 ms. The available odor pair changes every 24 trials. Here, we analyze inhibitory interneuron spiking and CA1 local field potential (LFP) data only for correct trials, resulting in a maximum of 12 trials per odor-context combination, with a maximum of 8 unique odor-context combinations available. A single interneuron contributes at most 8 spike trains. **(B)** Model-fitting procedure<sup>26</sup>, consisting of identifying rhythmic phase values that coincided with inhibitory interneuron spiking, and convolving this distribution with a Gaussian kernel that has been adapted to a circular basis and whose bandwidth was selected via 7-fold cross validation<sup>26</sup>. Lastly, we compute the conditional probability of spiking given phase using Bayes' theorem to obtain the phase model for a given interneuron's spike train.

**A****B**

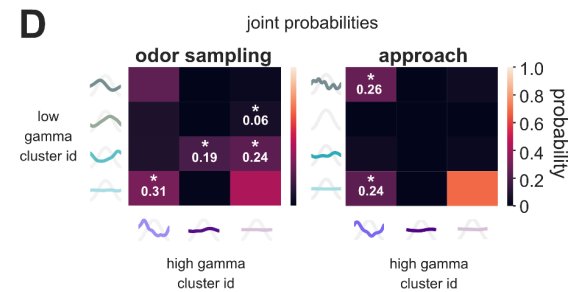
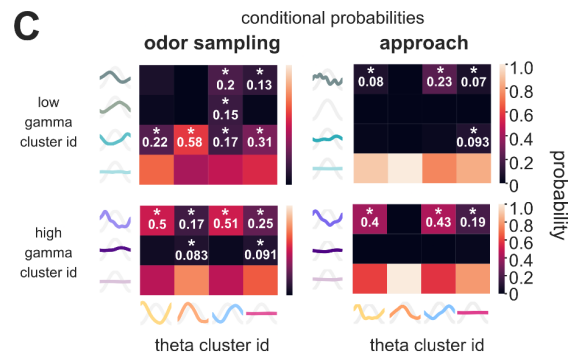
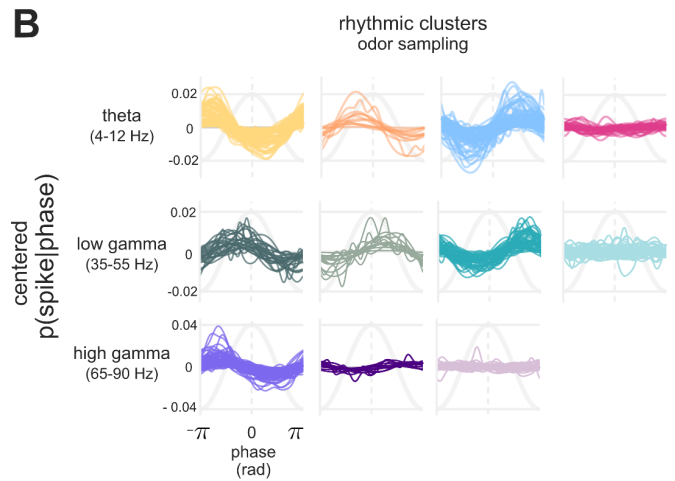

Interestingly, when we performed this clustering procedure (using distinct phase boundaries, see Methods) for the low gamma (35-55 Hz) and high gamma (65-90 Hz) phase models, we also identified distinct spike timing preferences relative to each of these rhythms (**Figure 4.2B, middle & bottom; Supplementary Figure S4.1B**). According to these schemes, four low gamma clusters emerged, with peaks along the ascending (26 out of 206 spike trains, 12.6%), early descending (14 out of 206 spike trains, 6.80%), trough (51 out of 206, 24.7%) phases of the low gamma cycle, as well as a low gamma-insensitive group (115 out of 206 spike trains, 55.8%; **Figure 4.2B, middle; Supplementary Figure S4.1B, top**). In the high gamma rhythm, three clusters appeared, with peak firing probabilities along the ascending (86 out of 206 spike trains, 41.7%) and descending (11 out of 206 spike trains, 5.34%) phases of the high gamma cycle, and a high gamma-insensitive cluster (109 out of 206 spike trains, 52.9%; **Figure 4.2B, bottom; Supplementary Figure S4.1B, bottom**).

Interneuron spike timing patterns may depend on the behavioral state of the animal. Applying the same clustering scheme as in the odor sampling epoch, we assessed the distributions of rhythmic phase models fit to data from the locomotive approach epoch (**Supplementary Figure S4.1C & D**). We found relatively similar proportions of spike trains sorting into the theta phase clusters (trough: 90 out of 207, 43.5%; ascending phase: 62 out of 207, 29.9%; early descending phase: 12 out of 207, 5.80%; and theta-insensitive: 43 out of 207, 20.8%), as well as the high gamma phase clusters (early descending phase: 70 out of 207, 33.8%; ascending phase: 1 out of 207, 0.483%; and high gamma-insensitive: 136 out of 207, 65.7%). Approach epoch models tended to differ, however, in how they distributed along the phase of the low gamma cycle, with spike timing more frequently insensitive to the phase of the low gamma cycle than observed during odor sampling, with the exception of spike timing preferences relative to the descending phase (descending phase: 25 out of 207, 12.1%; early ascending phase: 4 out of 207, 1.93%; late ascending phase: 4 out of 207, 1.93%; and low gamma-insensitive: 174 out of 207, 84.1%). This difference in the interneuron spiking timing relationships to low gamma were apparent despite observing no statistically significant differences in the amplitude of the low gamma oscillation across the odor sampling and approach epochs (two-tailed paired t-test comparing **mean** low gamma amplitudes

across LFP wires in the odor sampling v. approach epochs:  $t(29) = -1.1202$ ,  $p = 0.2718$ ). Interneuron rhythmic engagement along the phases of the low gamma oscillation may represent a unique consequence of either olfactory input processing specifically, or associative memory processing more generally.



**Figure 2. Interneuron spike trains cluster according to spike-phase relationships with the theta (4-12 Hz) CA1 oscillation. (A)** Each observation corresponds to a CA1 interneuron spike train. The x-axis describes the circular mean of the spike train's spike times relative to theta (4-12 Hz) phase angles. The y-axis corresponds to the sum of the Kullback-Leibler divergences (sum kld) of each spike train's theta phase model against its own average probability of spiking. The sum kld is not used for clustering, but provides a visual intuition for the distribution of rhythmic entrainments to the theta oscillation across different interneuron spike trains. An idealized theta cycle is illustrated in pale grey. Three clusters comprise spike-phase distributions that met the Rayleigh criterion for non-uniformity ( $p < 0.05$ ) and whose circular means lay within the following edges:  $(2.85, -2]$ ,  $(-2, 0]$ , and  $(0, -2.85]$ , corresponding to the yellow (theta trough), orange (theta ascending phase), blue (theta descending phase) clusters, respectively. Cluster centroids lie at:  $-2.65$  rad,  $-0.952$  rad, and  $1.74$  rad, respectively. The fourth cluster (hot pink) comprises spike-phase distributions that did not meet the Rayleigh criterion for non-uniformity ( $p \geq 0.05$ ). **(B)** Mean-centered phase models fit to spike-phase relationships relative to the theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) rhythms filtered from the CA1 local field potential (LFP). Low gamma cluster centroids lie at:  $-0.753$  rad,  $0.770$  rad, and  $2.31$  rad, respectively. High gamma cluster centroids lie at:  $-1.73$  rad, and  $2.13$  rad, respectively. Each interneuron spike train is described by these three rhythmic phase models. Phase models for each rhythm have been clustered according to their ability to meet the Rayleigh criterion for non-uniformity, and their peaks along their respective rhythm's cycle. Pale grey sinusoids in each subpanel represent the corresponding rhythm's idealized cycle trajectory. X axes for all subpanels span the range between  $-\pi$  to  $\pi$ . **(C)** Conditional probabilities (color axis) of spike train assignment to a given high-frequency cluster (low gamma, top; high gamma, bottom) given membership in a theta cluster for the odor sampling (left subpanels) and time-matched running approach (right subpanels) epochs. Asterisks and values mark the theta-high frequency combinations whose conditional probabilities were statistically greater than expected by chance as evaluated through a permutation test. Entries lacking annotations in the matrix correspond to conditional probabilities that were likely to occur by chance ( $p < 0.05$ ). **(D)** Joint probabilities (color axis) of spike trains' assignment to low gamma (rows) and high gamma (columns) cluster identities. As in **C**, asterisks and values mark the combinations of cluster assignments that were unlikely to occur by chance, as evaluated through a permutation test ( $p < 0.05$ ).



Perforant path and Schaffer collateral pathways providing input to the region exert their influence along distinct phases of the theta cycle<sup>9</sup>, and high frequency oscillations have also been reported to distribute along different phases of the theta oscillation<sup>34,35,38</sup>. We consequently asked whether a given interneuron spike train's theta phase relationship systematically predicted its engagement and phase preference with respect to low gamma and high gamma oscillations. To achieve this, we computed the conditional probability of belonging to a low gamma cluster identity as a function of membership in a theta cluster (**Figure 4.2C, top left**), and separately, the conditional probability of membership in a high gamma cluster given membership in a theta cluster, during the odor sampling epoch (**Figure 4.2C, bottom left**). While the largest likelihoods were assigned to membership in low gamma-insensitive or high gamma-insensitive groups regardless of theta phase preference, other systematic relationships emerged, with theta-trough spike trains organizing along low gamma-descent phases. Theta-descending interneuron spike trains were capable of organizing along all low gamma phase clusters. Theta-trough and theta-descending spike trains were particularly likely to organize along high gamma-ascending phases. On the other hand, theta-ascending spike trains tended overwhelmingly to organize their spike trains significantly along the low gamma-descent (**Figure 4.2C, left**), with smaller likelihoods of entrainment to the high gamma-ascending phase. Notably, theta-insensitive spike trains had a significant likelihood of organizing according to low gamma-ascent and -descent, as well as the high gamma rhythm's ascent, with significance assessed via permutation tests (**Supplementary Figure S4.2A**).

We performed the same analysis with the approach epoch rhythmic models and found that theta phase similarly predicted high gamma relationships, but results differed in the predictions for low gamma relationships (**Figure 4.2C, right; Supplementary Figure S4.2B**). Here, theta-ascent spike trains ceased to engage entirely in low gamma spike-phase relationships, while theta-trough and -descent spike trains shifted spike timing to the low gamma rhythm's ascending phase, again reinforcing the notion that odor sampling epochs recruit interneuron spiking into low gamma relationships in a way that differs from locomotive approach epochs.

The ability of interneurons to organize their spike timing according to either low gamma or high gamma frequencies may result from distinct patterns of inputs, or represent an intrinsic constraint borne from unique resonance properties<sup>48-50</sup>, or a combination of these factors. If this were the case, we would expect interneuron spike trains to organize along one of the high frequency bands, but remain insensitive to the other one. To determine whether this prediction holds true, we computed interneuron spike trains' joint probabilities of belonging to a given low gamma and high gamma cluster identity, during the odor sampling (**Figure 4.2D, left**) and approach (**Figure 4.2D, right**) epochs. For each of the following proportions, note that the denominator reflects the intersection, in terms of number of spike trains, of evaluated cluster pairs (see Methods). We first found that a majority of spike trains were insensitive to both rhythms (**Figure 4.2D, left & right**, see bottom right entry in matrix). The next-largest proportions exhibiting joint membership, approximately 30%, participated only in the high gamma frequency band while remaining insensitive to low gamma frequencies during the odor sampling epoch, and this proportion dropped modestly to 24% during the approach epoch. Although low gamma-early descent relationships emerged during odor sampling epochs, these relationships disappeared during the approach. Surprisingly, however, 26% of spike trains participated preferentially in *both* high gamma and the low gamma-ascent phase during the approach epoch. These findings reveal diverse interneuron engagement that depends on the animal's behavioral state: insensitivity to fast frequencies altogether, engagement in one fast frequency to the exclusion of the other, and—although less frequent—a combination of low gamma and high gamma relationships.

***Temporal coincidence and anti-correlation of interneuron rhythmic spike-phase relationships.***

We subsequently worked to understand whether interneurons engaged in theta, low gamma, and high gamma frequencies simultaneously, or whether these relationships were systematically predictive of each other, but occurred at distinct times over the course of odor sampling. During odor sampling, we qualitatively observe intervals within a single trial where high frequencies manifest in the absence of large amplitude theta oscillations (**Figure 4.3A**, see interval between 200-600 ms), as well as nested within

theta oscillations (**Figure 4.3A**, see interval between 1000 ms to the end of the LFP trace). Whenever fast frequencies *did* coincide with theta oscillations, low gamma frequencies appeared to concentrate along the rising, peak, and falling phases of the theta cycle. On the other hand, high gamma frequencies appeared to manifest discretely near the peaks and troughs of theta cycles (**Figure 4.3B**). These results align with previous hippocampal silicon probe recordings acquired under anesthesia<sup>35</sup> as well as tetrode recordings performed during rat locomotive epochs<sup>34</sup>.

To assess the relative timing between interneuron entrainment to theta, low gamma, and high gamma oscillations, we fit new phase models to data drawn from non-overlapping 250 ms windows spanning the duration of the odor sampling epoch (1500 ms). We were then able to assess each spike train's degree of entrainment (see Methods) to these rhythms over the course of associative memory processing. In **Figure 4.3C**, we show an example interneuron's entrainment time course for its spike trains in response to odors G (*top*) and H (*bottom*), respectively. This interneuron exhibited simultaneous entrainment to theta and low gamma frequencies when odor G was present, but organized its spikes only to the low gamma oscillation when odor H was present (**Figure 4.3C**). To quantify the simultaneity or temporal anti-correlation of interneuron spike trains' rhythmic relationships, we computed a Pearson correlation between measures of entrainment to each rhythm obtained for each non-overlapping time window, where degree of entrainment was measured using the sum of the Kullback-Leibler divergences (sum kld, see Methods). This was done for each pairwise combination of rhythms: theta and low gamma, theta and high gamma, and low gamma and high gamma (**Figure 4.3D**). Each distribution of Pearson correlation coefficients ( $r$ ) yielded two modes, one located near perfect correlations, and another concentrated along more negative correlations. To examine whether a given interneuron could exhibit simultaneity of entrainment *and* anti-correlated entrainment to multiple rhythms, we counted the number of interneurons whose odor-context spike trains yielded a mixture of strongly anti-correlated ( $r < -0.05$ ) and strongly correlated ( $r > 0.05$ ) entrainment to pairs of rhythms, and we restricted this analysis to only those interneurons with significant entrainment (Rayleigh test,  $p < 0.05$ ) to both rhythms at some point during the odor sampling epoch (theta-low gamma: 20 out of 30 interneurons engaged in both rhythms;

theta-high gamma: 29 out of 30 interneurons; low gamma-high gamma: 5 out of 19). We found relatively few interneurons with odor-context spike trains whose multi-rhythmic engagement could be both simultaneous and anti-correlated in time (theta-low gamma: 3 out of 20 interneurons; theta-high gamma: 10 out of 29; low gamma-high gamma: 5 out of 19). In short, few interneurons—recorded across different odor-context conditions—could participate in multiple rhythms orthogonally, unsystematically, or simultaneously over the course of associative memory processing, and this appeared to be particularly frequent in interneurons with theta and high gamma rhythmic entrainment.

***Interneurons can flexibly shift in and out of entrainment to rhythmic phase in response to distinct olfactory and visuospatial inputs.***

There is a substantial body of work that attempts to identify systematicity between interneurons' morphological, transcriptomic, and spike timing profiles<sup>43,50</sup>, which has facilitated theta-phase interneuron classification schemes in studies that lack access to morphological and molecular information<sup>45</sup>. Seemingly implicit in these classification schemes is that interneurons' theta phase preferences may represent a stable neuronal property, either by virtue of intrinsic resonance<sup>54</sup>, or connectivity patterns, or a combination of these. Given the heterogeneous patterns of entrainment to distinct frequency bands observed when we fit models across behavioral states (**Figure 4.2**), and over time within odor sampling epochs (**Figure 4.3**), we sought to characterize explicitly whether an interneuron's activity could be organized differently within a rhythmic frequency band as a function of the active set of environmental stimuli during associative memory processing.

To address this question, we counted the number of interneurons whose various odor-context responses had been sorted into multiple theta clusters (**Figure 4.4A**). We found that 16 out of 30 interneurons reorganized their spike timing within the theta frequency band, as a function of different olfactory and visuospatial inputs. With few exceptions, cluster reassignments tended to reflect entrainment to a specific theta phase and theta-insensitive states (14 out of 16 reorganizing interneurons), rather than shifting theta phase preferences. To obtain a quantitative measure of the interchangeability of

phase models across odor-context conditions, we assessed the predictive accuracy of an interneuron's phase models (**Figure 4.4B, left**) against spike-phase distributions drawn from other odor-context conditions ("other", **Figure 4.4B, right**). To set a baseline expectation for model predictive accuracy, we also tested each phase model against the spike-phase distribution used during model fitting ("own"), which should on average display more accurate predictions (here, log likelihoods closer to zero).

An interneuron that exhibited a theta cluster reassignment (**Figure 4.4B, top left**) incurred larger prediction errors (more negative log likelihoods) when its various phase models were tested against data sourced from the other conditions than when the models were tested against the data that generated them (**Figure 4.4B, top right**). In contrast, an interneuron whose theta phase models were consistently sorted into the same cluster identity exhibited overlapping distributions of log likelihoods when tested against their own data, and against data from other odor-context conditions, indicating stable theta-phase spiking relationships across odor-context combinations (**Figures 4.4A & B, bottom**).

We took the same approach to assess whether interneurons exhibited flexibility and stability in low gamma and high gamma entrainment. We found that a majority of interneurons shifted their low gamma (19 out of 30 interneurons) and high gamma (21 out of 30 interneurons) spike-phase relationships as a function of the active set of stimuli. Similar to our observations for the theta rhythm, the majority of shifts involved transitions between entrainment to a given phase and loss of entrainment altogether (low gamma: 15 out of 19 interneurons; high gamma: 21 out of 21 interneurons, **Figures 4.4C & E**). An interneuron could be insensitive to a fast frequency for most input combinations yet participate rhythmically for a single odor-context combination (**Figure 4.4D**). The reverse could also be true: an interneuron could participate rhythmically in a fast frequency for the majority of input combinations, with varying degrees of entrainment, yet remain insensitive in response to a few sets of stimuli (**Figure 4.4F**).

In an effort to quantify this variability in interneurons' rhythmic responses to odor-context conditions across our recorded population ( $n = 30$ ), we used a linear mixed effects model to predict interneurons' average cross-validated log likelihood per condition using as fixed effects both rhythm (with levels theta, low gamma, and high gamma) and cross validation test type (with levels 'own' or


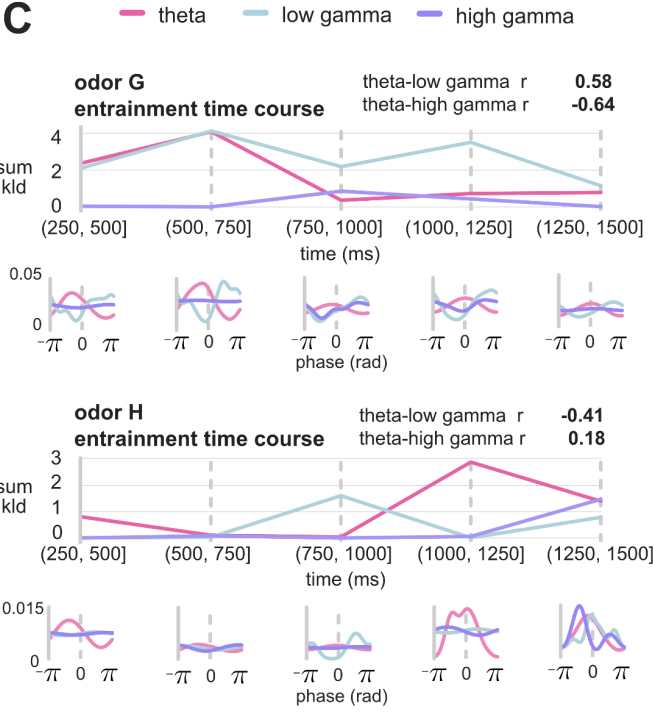
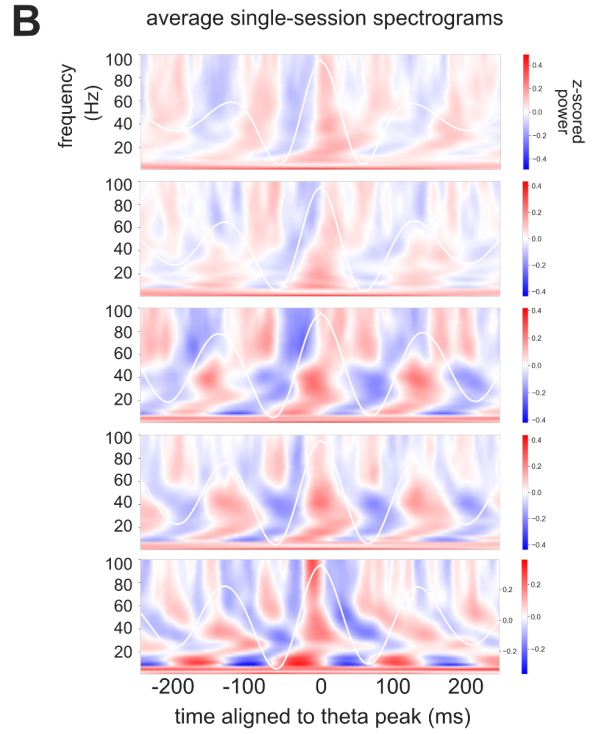

‘other’, referring to whether the kernel density estimator was tested against the data that generated it or the data sourced from the remaining odor-context conditions, respectively). To account for interneuron-level variability (e.g. differences in baseline rates of spiking), the linear mixed effects model was also specified with random intercepts for each interneuron. This model failed to explain substantially more variance in average cross-validated log likelihood than a model omitting test type ( $\chi^2(1) = 0.979$ ,  $p=0.32$ ). Our results indicate that interneurons tended to—on average—engage in relatively stable spike-phase relationships for different combinations of olfactory and visuospatial inputs, as reflected in the pattern of kernel density estimator predictive accuracies across test type (**Figure 4.4G**).

### 4.3 Discussion

CA1 inhibitory interneurons comprise a heterogeneous class of cells whose properties render them effective sources of rhythmic temporal organization<sup>25</sup>, promoting network stability during excitatory input processing<sup>14</sup>. In the hippocampus, inhibitory interneurons are recruited through a combination of feedforward excitatory inputs<sup>7-9</sup>, as well as dense feedback projections from their local pyramidal cell targets<sup>21</sup>. This intricate pattern of connectivity, along with changes to the predominant afferent sources of excitatory drive into the network, may give rise to variable interneuron spike timing engagement in rhythmic circuit interactions. In this study, we leveraged a context-guided associative memory paradigm, where rats learned that odors in a pair were differentially rewarded according to spatial context<sup>1,27</sup>. We were consequently able to examine the spike timing responses of CA1 inhibitory interneurons in response to distinct combinations of olfactory and visuospatial inputs, as well as compare activity patterns across behavioral states during the same task. By combining a complex behavioral paradigm and statistical modeling techniques, we show that putative CA1 interneurons can flexibly participate in rhythmic network processes as a function of behavioral state, and in response to distinct combinations of behaviorally relevant olfactory and visuospatial inputs during associative memory processing.



**Figure 4.3. Theta, low gamma, and high gamma frequencies can entrain interneurons simultaneously and orthogonally.** **(A)** Single trial local field potential (LFP) from an odor sampling epoch. Illustrates periods of high frequency content with weak to no theta prevalence (200 to 500 ms), and periods during which high frequencies appear nested within visible theta cycles (1000 to 1500 ms). **(B)** Example average z-scored power spectrograms from four rats' LFPs, with time aligned to the peak of theta cycles, obtained during the odor sampling epoch. Corresponding average theta-filtered signals are superimposed in white in each spectrogram. The first and second spectrograms were computed from two different simultaneously recorded LFPs from the same rat. The remaining three spectrograms were contributed by three different rats. **(C)** Sample entrainment time course for a single interneuron's spike trains in response to odor G (top) and odor H (bottom middle). Entrainment for each rhythm is measured as the divergence between the rhythm's phase model and a uniform model predicting spikes according to the neuron's average probability of spiking. Each rhythmic phase model is fit to odor sampling data spanning 250 ms consecutive, non-overlapping windows, starting with the window immediately following odor delivery, 250 ms following the nose poke; each window's edges are printed along the time axis. Below each spike train's entrainment time course, we show the corresponding, color-coded phase models from which the entrainment metric was derived. Note the second window for each spike train, where odor A recruits the interneuron into both theta (hot pink) and low gamma (light teal), while odor F recruits the interneuron into low gamma only. **(D)** Distributions of Pearson  $r$  correlations reflecting the relationship between the sum of the Kullback-Leibler divergences (sum kld) computed from the same spike-train's theta, low gamma, and high gamma phase models, for all spike trains (grey) and spike trains excluding those whose models were insensitive to at least one of the examined rhythms (red).




To assess CA1 interneuron selectivity to particular odor-context combinations, we first segmented interneuron spike trains and local field potential (LFP) phase estimates according to the odor-context combination available during odor sampling. We then modeled the relationship between interneuron spike timing and the phase of the theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) LFP oscillations during odor sampling and goal-oriented locomotive epochs within the context-guided associative memory task (**Figure 4.1**)<sup>27</sup>.

In addition to identifying well-established classes of theta phase relationships<sup>43,45</sup>, we found that low gamma and high gamma spike timing relationships also reliably clustered according to their respective phases (**Figures 4.2A & B; Supplementary Figures S4.1B-D**). Previous juxtacellular recording studies identified interneuron spike-phase relationships to the oscillatory phase of broadly defined gamma frequency bands (25-90 Hz)<sup>44,49,50</sup>. Our findings build on this work by characterizing gamma-band interneuron spike timing across behavioral states, when gamma oscillations manifest along dissociable sub-ranges that likely reflect distinct underlying generators<sup>36,57</sup>. Inhibitory interneurons have been particularly implicated in gamma-rhythmic generation<sup>25</sup>. Stimulation of CA3 Schaffer collateral afferents produces an evoked response in CA1 that is accompanied—in the absence of further CA3 input—by gamma oscillations centered approximately at 40 Hz<sup>32</sup>. GABA<sub>A</sub> receptor blockage abolishes this response, leading to the hypothesis that a pyramidal-interneuron CA1 network oscillator was sustaining the accompanying gamma oscillations. The CA1 additionally comprises a variety of interneuron classes exhibiting distinct intrinsic properties that may differentially contribute to the generation of slower or faster gamma rhythms. Two distinct populations, differentiable by the duration of their membrane excitatory postsynaptic potentials (EPSPs), have been reported in the region<sup>24</sup>. Interneurons with slower EPSPs tended toward less spike timing precision than those sporting shorter duration EPSPs, whose spike timing was very fast and precise, a difference that was attributable to each group's NMDA receptor activation kinetics<sup>24</sup>. Distinct *phase* preferences along fast oscillatory cycles could reflect a variety of influences, but perhaps most tantalizing is the contribution of gap junctions between interneurons<sup>58</sup>. Gap junction connectivity could ensure spike time-locking to the same gamma


cycle<sup>59</sup>, but slight conduction delays (introduced either by the influence of chemical synapses or the spatial distribution of gap junctions<sup>60</sup>) across interneurons might produce reliable, short phase lags that manifest as distinct low gamma and high gamma phase clusters in our data. A combination of distinct laminar distributions of afferent connectivity, electrochemical interneuron-interneuron interactions, and intrinsic ion channel and receptor variation may thus explain different interneurons' predilection for spike-phase relationships relative to distinct slower and faster gamma frequency ranges, as well as the emergence of the distinguishable phase-based clusters observed here.

**Figure 4.4. Flexibility and stability in CA1 interneuron oscillatory responses to olfactory and visuospatial inputs.** (A) Counts of interneurons whose spike trains were sorted into more than one theta cluster identity. (B) Schematic of the maze, with color-coded odor identities (top). Red-tinged colors (left) are only rewarded in the left-side, pink context; blue-tinged colors (right) are only rewarded in the right-side, blue context. Left column contains phase models corresponding to two different interneurons' theta phase models, where each phase model is color-coded according to the combination of olfactory and visuospatial inputs present. The interneuron in the top left exhibited theta cluster reassignments, while the interneuron in the bottom left was stably assigned to a single theta cluster identity across all odor-context combinations. Right column panels correspond to log likelihood distributions computed by testing a phase model trained on data from a given condition (x-axis) on data from its own training set ('test condition = own', dark grey) or against data from all remaining conditions ('test condition = other', pale grey). Less negative log likelihood values reflect more accurate predictions. (C) Same as in A, but for the low gamma cluster identities. (D) Sample interneuron with low gamma entrainment exclusively for one odor-context combination; testing this model against data from all other conditions yields the largest prediction errors (log likelihood distributions that do not overlap with and are more negative than the distribution of log likelihoods obtained from testing the model on the data that generated it). Results are jittered laterally for the conditions with flat phase models. (E) Same as in A and C, but for high gamma cluster reassignments. (F) Sample interneuron, whose phase models were sorted into multiple high gamma cluster identities. (G) Distributions of average log likelihood across the 30 interneurons, for each rhythm, for models tested against the data that generated them ('own'), and data from other odor-context conditions ('other') for a given interneuron. Mean log likelihoods correspond to the linear mixed effects model intercepts for each rhythm by test type.

**A** types of **theta** cluster shifts  
odor sampling




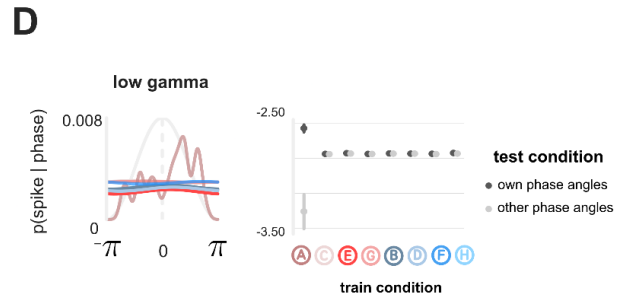
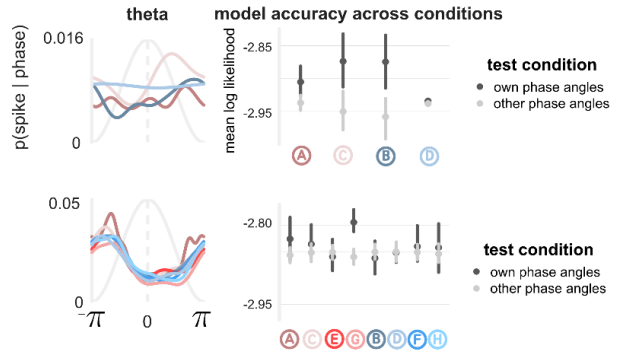


**C** types of **low gamma** cluster shifts



**E** types of **high gamma** cluster shifts



**G** distributions of average cross-validated log likelihood across the population



In the present study, interneuron entrainment across distinct phases of the low gamma oscillation emerged more frequently during odor sampling periods relative to locomotive approach epochs, despite a lack of significant difference in low gamma amplitudes across behavioral epochs (**Figure 4.2C & D**). The emergence of interneuron spike timing recruitment across diverse low gamma phases may reflect lateral entorhinal cortex (LEC) coordination with hippocampal CA1 cell assemblies. During odor sampling in a variant of the olfactory associative memory paradigm, Igarashi et al. (2014) observed concurrent increases in LEC and CA1 low gamma-band power, in the 20-40 Hz range<sup>29</sup>. Recent work from our group has provided evidence supporting a causal relationship between LEC and CA1 low gamma coherence: rhythmically stimulating the LEC resulted in optimal recruitment of CA1 neuronal spike timing particularly within subsets of the low gamma frequency range identified here<sup>61</sup>. LEC influence on interneuron spike timing may also emerge indirectly through the trisynaptic pathway. The LEC receives direct projections from the rodent olfactory bulb<sup>62</sup>, and specifically disrupting odorant information—while maintaining intact respiration—abolishes a rhythmic current sink in the outer third of the dentate gyrus (DG) molecular layer<sup>63</sup>, the location where LEC fibers synapse onto DG apical dendrites<sup>64,65</sup>. It is possible that (transformed) odorant information percolates through the trisynaptic pathway and converges onto both CA1 radiatum interneurons and local CA1 pyramidal cells by way of CA3 Schaffer collaterals. The absence of a difference in low gamma amplitudes across odor sampling and locomotion may be attributable to volume conduction during locomotive epochs, as previously documented<sup>35</sup>. On the other hand, odor sampling low gamma may arise locally in the CA1 pyramidal layer through the multiplexed contribution of LEC and CA3 inputs interacting with local interneuron and pyramidal cell microcircuits. Local generation of low gamma oscillations, as opposed to volume conduction, may specifically emerge during associative memory processing, given that power in this frequency band increases in the CA1 over the course of olfactory associative learning<sup>29,66</sup>. Although hypothetical at this point, future work could explicitly examine the possibility of behavioral state-specific local generation of low gamma frequency power.

A spike train's entrainment—or its insensitivity—to particular theta phases systematically predicted its pattern of low or high gamma phase activity (**Figure 4.2C**). Moreover, the same interneuron spike train—although rare—could engage in both low and high gamma oscillatory phase, while the majority engaged selectively in either low gamma or high gamma oscillations (**Figure 4.2D**). We investigated whether the systematicity in these observed cross-rhythmic spike-phase profiles was a consequence of nested theta-gamma rhythms—in which case entrainment to each rhythm should appear simultaneously—or whether high-frequency relationships could emerge orthogonally from theta entrainment over the course of associative memory processing. When we fit models to discrete windows over the course of odor sampling, we found that temporal relationships were broadly distributed across the population of spike trains, some of which exhibited orthogonal theta-low gamma, and theta-high gamma relationships, while others exhibited similar entrainment trajectories across theta and higher frequencies (**Figure 4.3**). In particular, high gamma oscillatory entrainment tended to emerge both simultaneously and orthogonally from theta relationships across interneurons' various odor-context responses. While theta-gamma nested relationships have received substantial attention in the field<sup>34,35,37,67</sup>, and the theta-gamma LFP relationships we observe align with previous reports<sup>35,38</sup>, less is known about gamma rhythmic dynamics that emerge independently of theta oscillations, or which at least coincide with clear reductions in the power of the theta oscillation<sup>27,67,68</sup>. These periods may reflect processes distinct from those generating strong theta-gamma coupling in the CA1 network. Moreover, their prevalence may vary topographically: a recent study leveraging an olfactory association task showed that the power of faster frequencies appears to be inversely proportional to theta power across the CA1 proximo-distal axis<sup>66</sup>. As with the questions relating to local generation versus volume conduction of low gamma in the CA1 pyramidal layer, addressing the functional role of theta-independent gamma oscillations—as well as their underlying mechanisms—will require further work.

Lastly, we find occasional flexibility, as well as overwhelming stability, in interneurons' rhythmic spike timing responses to varying odor-context combinations. Interneurons in our sample could shift between entrainment and insensitivity to each of the rhythms examined. They nevertheless maintained



their phase preference across odor-context conditions when they *were* entrained, with few exceptions (**Figure 4.4A**). Differential recruitment of interneuron activity is most frequently explored in changes to firing rates. Indeed, interneurons' firing rates can be modulated according to the active subset of pyramidal cells<sup>69</sup>, they can vary across the animal's spatial location in a manner reminiscent of principal cell place field dynamics<sup>70-72</sup>, and they can change as a function of olfactory inputs in associative memory tasks<sup>28,68</sup>. Our work reveals an additional mechanism—oscillatory spike timing variations—whereby interneurons reflect information about relevant constellations of environmental features.

Inhibitory dynamics have long been considered essential stabilizing components in networks that receive overwhelming excitatory drive<sup>14</sup>. Although it is clear that our population of CA1 interneurons exhibited overwhelming stability in spike-phase relationships, it might appear surprising that some in this sample displayed flexibility for some odor-context conditions. The functional consequence of this state-dependent flexible engagement in multiple rhythms remains to be causally probed, and offers exciting new directions towards understanding the intrinsic and microcircuit parameters that determine interneuron rhythmic spike timing. Notably, the present study illuminates the importance of assessing inhibitory dynamics—indeed, any neuronal dynamic—within diverse behavioral conditions. Our work presents an early and essential first step in ascertaining the role of inhibition in coordinating inputs and local excitation to adaptively meet a wide range of behavioral demands.

#### **4.4 Methods**

##### **Behavioral Paradigm.**

All protocols involving animals were performed in accordance with National Institutes of Health (NIH) guidelines, and with the prior approval of the Institutional Animal Care and Use Committee (IACUC) in Boston University (approval number 13-057). The context-guided associative memory task and apparatus have been previously described in detail (see Rangel et al. 2016), and will be briefly described here.

##### ***Apparatus.***

The behavioral apparatus comprises a central start box flanked by two diametrically opposed 45 cm stems that make up each spatial context. Contexts were differentiable according to the spatial position in the room, and by the color and texture of wraps secured over the surface of the arms. Each stem widened into a trapezoidal area, also covered with the context wrap, at the end of which sat a face plate containing two circular apertures, termed odor ports (**Figure 1A**). Rats were behaviorally shaped to insert (“poke”) their noses into the odor ports, a process that triggered the release of a given odor 250 ms after poke onset. The odor port containing the rewarded odor was counterbalanced across trials.

### ***Odor sampling epoch***

***Correct trials.*** Rats learned that odors in a pair were differentially rewarded depending on the context in which they were presented. To trigger the release of a water reward (from a well located underneath the odor port), rats learned to sustain a 1500 ms nose poke in the port containing the rewarded odor of the pair. If the rat first poked its snout into the port containing the unrewarded odor, the rat had up to 1500 ms to remove his snout, after which he could initiate a 1500 ms nose poke in the correct odor port. A single correct trial consisted of a 1500 ms interval (“termed odor sampling”) that led to a reward.

***Incorrect trials.*** Nose pokes lasting 1500 ms in the port containing the unrewarded odor triggered a white noise buzz instead of a water reward. We did not include analysis of these trials because rats were overtrained when we began electrophysiological recordings, resulting in very few incorrect trials per odor-context combination.

***Approach epoch.*** To compare results from odor sampling windows against those from locomotive behavioral states, we selected time-matched (1500 ms) epochs prior to each correct odor sampling trial, when the rat displayed the highest velocity during stem traversals towards the odor ports.

***Task Structure.*** Rats performed the task for a total of 96 trials, which consisted of four consecutive 24-trial blocks in which rats deployed associations for a single odor pair. Because each odor in a pair was only rewarded in a single context, each odor contributes a maximum of 12 correct trials to the data set. We excluded data from odors that had fewer than four correct trials. Since there were a total

of four odor pairs over the course of an experimental session, and given that we only analyze correct trials, we obtained interneuron responses for a maximum of eight odor-context combinations.

### **Electrophysiological Recordings.**

We performed high-density extracellular tetrode recordings, with independently moveable tetrodes implanted into the CA1 subregion of the hippocampus. This allowed us to isolate single units as well as local field potentials (LFPs). Signals were amplified (4,000-8,000 times) and digitized (40 kHz) using Plexon's Omniplex Neural Data Acquisition system.

**Single Units.** We isolated putative CA1 inhibitory interneurons according to waveform shape features (mean width at half maximum: 150 us; mean temporal offset from peak to trough: 350 us), as well as an average firing rate criterion ( $\geq 5$  Hz)<sup>73</sup>. Inhibitory interneurons possess extremely large dendritic and axonal arborizations, spanning various CA1 layers<sup>23,35</sup>. It may consequently be possible to record the same interneuron over multiple days, even as tetrode wires are gradually advanced. To ensure extremely conservative estimates of interneuron counts, if the same tetrode within a given rat contained an interneuron across multiple sessions, we selected only one of the sessions for which the tetrode recorded putative interneuron activity.

**Local Field Potentials (Estimating Instantaneous Phase).** We used a third-order Butterworth filter to bandpass-filter LFPs in the theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) frequency ranges. We subsequently took the arctangent of the filtered signal's complex Hilbert transform to compute instantaneous phase estimates for each corresponding frequency band. Frequency ranges were selected according to the observed frequency ranges in a previous study leveraging the same dataset<sup>27</sup>.

### **Dataset Preparation.**

Our goal was to characterize CA1 interneuron spike timing relationships in response to different combinations of odorants and visuospatial (context) information. We consequently split each interneuron's data according to the available odor-context combinations. Each interneuron could thus

contribute a maximum of eight spike trains. In some cases, the full eight odor-context combinations were unavailable as a function of the following exclusion criteria: performance for a given 24-trial block fell below 75% accuracy; the accuracy criterion for the 24-trial block was met, but one of the odors in the pair had fewer than 4 correct trials; or the interneuron exhibited fewer than 7 spikes for a given odor-context combination. After applying these exclusion criteria, the dataset consisted of 206 odor sampling interneuron spike trains (from  $n = 30$  out of 31 interneurons, one interneuron did not spike during odor sampling), and 207 approach spike trains (from 31 out of 31 interneurons). One interneuron didn't emit spikes at all during the odor sampling epoch, but did spike during a subset of the approach epochs preceding odor sampling. Those spike trains are consequently analyzed during approach epochs, but having no spikes to fit models to, are not included in the odor sampling epoch spike trains.

### **Kernel Density Estimation and Phase Models.**

Phase models, representing the conditional probability of spiking given oscillatory phase, were fit according to the procedure described in Rivière et al. (2022). For a given spike train and the CA1 LFP recorded on the same tetrode, we selected all phase angles  $x_a$  that coincided with a spike, where  $x_a \in X$ , and  $X = \{x_1, \dots, x_A\}$ . This distribution of phase angle observations was convolved with a Gaussian kernel  $K(x; h)$  of the form,

$$K(x; h) = \frac{1}{\sqrt{2\pi}} \exp\left(\frac{-x^2}{2h^2}\right), \quad (\text{equation 1})$$

where the smoothing parameter, or bandwidth,  $h$  controls the width or spread of the kernel, and whose bandwidth (also termed “smoothing”) parameter  $h$  was selected via 7-fold cross-validation<sup>74,75</sup> (software package: Python sklearn, GridSearchCV, Kernel Density method, grid size = 1,000)<sup>51,76</sup>. Candidate bandwidths spanned the equivalent of 6% to 40% of a phase cycle at the kernel's half-maximal point,

which ensured that phase models could capture the potentially varied distributional features of spike densities within a phase cycle (e.g. multiple modes, dispersion)<sup>77</sup>.

The goal is to estimate an unobservable probability density function  $p(x)$  that generates the empirically observed distribution of phase angles that coincide with spikes. To estimate this function, we convolve the distribution of phase angles  $X$  with  $K(x; h)$ , according to the following expression,

$$\hat{p}(x; h) = \frac{1}{Ah} \sum_{a=1}^A K\left(\frac{x-x_a}{h}\right) = p(\text{phase} | \text{spike} = 1) , \quad (\text{equation 2})$$

a probability density function (or kernel density estimator<sup>76</sup>) generated from the mixture of  $A$  Gaussian probability densities, where each individual Gaussian's mean corresponds to the location of each observation,  $x_a$ . These Gaussians have been adapted to treat  $-\pi$  and  $\pi$  as identical points<sup>51</sup>. The resulting kernel density estimator<sup>76</sup> yielded the probability of observing a phase angle given spiking. Lastly, we applied Bayes' Theorem to this estimator to obtain the probability of spiking given phase, termed “phase model” (**Figure 1B**),

$$p(\text{spike} = 1 | \text{phase}) = \frac{p(\text{phase} | \text{spike} = 1) p(\text{spike})}{p(\text{phase})} , \quad (\text{equation 3})$$

where  $p(\text{spike})$  corresponds to the average probability of spiking, and  $p(\text{phase})$  represents the uniform probability of observing any one phase angle; that is,  $\frac{1}{2\pi}$ .

### **Computing the sum of Kullback-Leibler divergences (sum kld).**

To examine the degree of entrainment to rhythmic phase, we compare a spike train's phase model against a model representing the spike train's uniform average probability of spiking (**Supplementary**

**Figure 1A).** Probability distributions are routinely compared using the Kullback-Leibler divergence (kld)<sup>78</sup>, expressed below as,

$$kld(p_i || q_i) = p_i \log_2\left(\frac{p_i}{q_i}\right) + (1 - p_i) \log_2\left(\frac{1-p_i}{1-q_i}\right), \quad (\text{equation 4})$$

where  $i$  corresponds to a given phase bin (out of a 1000 phase bins used to fit the kernel density estimator),  $p$  refers to the empirically obtained phase model, and  $q$  represents the average probability of spiking. The kld is a non-symmetrical measure, meaning that  $kld(p_i || q_i) \neq kld(q_i || p_i)$  unless  $p_i = q_i$ , in which case  $kld = 0$ . Our formulation commits to treating the phase model, which is empirically estimated, as our proxy for the (unobservable) ground truth distribution underlying the spike train. We believe this a reasonable decision given that we have previously shown, in simulations, that phase models accurately approximate known ground truths across spike trains driven by both rhythmic (e.g. those related to phase) and uniform underlying distributions<sup>51</sup>.

To summarize the Kullback-Leibler divergences across all 1000 phase bins over which we fit the phase models, we sum across all computed  $kld_i$  (**Supplementary Figure S4.1A**),

$$sum\ kld = \sum_i^{1000} kld_i, \quad (\text{equation 5})$$

which provides a numerical estimate of the degree of a spike train's degree of entrainment with respect to the phase of an oscillation. We wish to emphasize that this provides information beyond that offered by the Rayleigh test for non-uniformity of circular distributions<sup>79</sup>: while the latter considers only the number of spikes that organize around particular phase angles, the sum kld implicitly takes *spike rate* (and consequently, time) into account. This enables the sum kld to distinguish between spike trains that are truly *rhythmic*—where spikes systematically are emitted across all cycles, or nearly every cycle of an oscillation—from spike trains that are systematically related to particular phases of an oscillation, but are

so sparse as to skip several cycles of the oscillation<sup>51</sup>. For this reason, we take the sum kld to approximate *degree of rhythmic entrainment*.

### **Clustering Procedure.**

To organize spike trains according to their theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) phase preferences, we leveraged the Rayleigh test for non-uniformity of circular distributions<sup>79</sup>, as well as the circular mean of the spike-phase distributions (**Figure 4.2A, Supplementary Figure S4.1B & C**). First, we separated spike-phase distributions according to their ability to significantly deviate (or fail to deviate) from a uniform distribution according to the Rayleigh statistic ( $p < 0.05$ ). The spike-phase distributions that did not meet this criterion constituted the theta-insensitive, low gamma-insensitive, and high gamma-insensitive clusters, for each rhythm respectively.

The remaining spike-phase distributions were clustered according to fixed phase boundaries, set for each rhythm separately. Boundaries for theta clusters were set as follows: trough (2.85 to -2] rad; ascent (-2 to 0] rad; and descent (0 to 2.85] rad, where parentheses (“(”, “)”) and brackets indicate non-inclusive and inclusive edges, respectively. Boundaries for low gamma clusters were set as follows: ascent (-2 to 0.25] rad; descent (0.25 to 1.25] rad; trough (1.25 to -2] rad. Lastly, boundaries for high gamma clusters were set as follows: ascent  $[-\pi$  to 0) rad; and descent  $[0$  to  $\pi$ ) rad. For each rhythm, cluster boundaries were found for odor sampling epochs, and subsequently applied identically to cluster spike-phase distributions from approach epochs (**Supplementary Figures S4.1C & D**).

### **Conditional & Joint Probabilities of Cluster Membership.**

One of our goals was to examine whether a spike train’s theta phase preference systematically predicted the likelihood of that spike train’s entrainment to higher frequencies. To address this question, we computed the conditional probability of observing spike train membership in low gamma or high

gamma cluster identities given spike trains' theta phase preferences (**Figure 4.2C**), for odor sampling and approach epochs separately. We computed the conditional probabilities according to the following,

$$p(\text{low gamma}_l | \text{theta}_k) = \frac{n(\text{low gamma}_l \& \text{theta}_k)}{n(\text{theta}_k)}, \quad (\text{equation 6})$$

where  $n(\text{low gamma}_l \& \text{theta}_k)$  represents the number of spike trains that were sorted into both low gamma cluster identity  $l$  and theta cluster identity  $k$ . The same procedure was applied to compute conditional probabilities of high gamma cluster membership given theta cluster membership,

$$p(\text{high gamma}_m | \text{theta}_k) = \frac{n(\text{high gamma}_m \& \text{theta}_k)}{n(\text{theta}_k)}, \quad (\text{equation 7})$$

where  $n(\text{high gamma}_m \& \text{theta}_k)$  represents the number of spike trains that were sorted into both high gamma cluster identity  $m$  and theta cluster identity  $k$ .

We additionally asked how likely spike trains were to participate in both low gamma and high gamma relationships. To this end, we computed the joint probability of membership in a given low gamma and high gamma cluster identity,

$$p(\text{low gamma}_l \& \text{high gamma}_m) = \frac{n(\text{low gamma}_l \& \text{high gamma}_m)}{[n(\text{low gamma}_l) + n(\text{high gamma}_m)] - n(\text{low gamma}_l \& \text{high gamma}_m)}, \quad (\text{equation 8})$$

where  $n(\text{low gamma}_l \& \text{high gamma}_m)$  represents the number of spike trains that were assigned to both low gamma cluster identity  $l$  and high gamma cluster identity  $m$ .



### **Permutation Tests to Assess Chance Conditional and Joint Probabilities of Cluster Membership.**

To test for the likelihood of obtaining the conditional and joint probabilities described above by chance, we performed a permutation test. For each spike train and its corresponding LFP, we identified all LFP theta phase cycles, binned each cycle according to the cluster edges (see “Clustering Procedure”), and we randomly jittered spike times *within* theta cluster bins. We repeated the permutation process twenty times, to generate a distribution of twenty conditional (**Supplementary Figure S4.2A & B**) and joint (**Supplementary Figure S4.2C & D**) probabilities computed from shuffled data. We then computed all conditional and joint probabilities using these twenty new shuffled datasets. This permutation procedure avoids disrupting the theta spike-phase relationships, while selectively asking whether high frequency relationships could have been obtained by chance. We then computed the fraction of conditional or joint probabilities (depending on the analysis) that exceeded the value of the conditional or joint probability obtained from intact data. If this fraction was smaller than 0.05, we considered the result statistically significant.

### **Phase Amplitude Spectrograms.**

To assess how high frequency oscillations distributed along the theta cycle, we computed spectrograms using the complex Morlet wavelet transform<sup>80</sup>. Time points corresponding to theta peaks were estimated from the filtered LFP trace. Spectrograms in **Figure 4.3B** were generated using a wavelet transform (Python, `scipy.signal, 'cwt'` method, wavelet order = 4, peak Morlet frequency = 250 Hz)<sup>38</sup> of the 1000 ms surrounding (symmetrically) each identified theta peak within a session. We subsequently compute z-scored power spectrograms, and report the average of all individual trial, z-scored spectrograms.

### **Simultaneity of spike train rhythmic entrainment across frequency bands.**

To examine the extent to which interneuron spiking was simultaneously entrained in theta and higher frequency bands, we first segmented odor sampling trials into 250 ms windows, starting at the time

of odor delivery up to the end. This yielded five non-overlapping windows. We then fit each rhythmic phase model to the distributions of phase angles that coincided with spiking within each of these windows, producing three rhythmic phase models per window, per interneuron spike train. For each of these phase models, we computed the degree of entrainment (sum kld). Next, for each spike train, we assessed the degree to which pairs of rhythmic phase models (e.g. theta and low gamma, theta and high gamma, low gamma and high gamma) tended to exhibit similar degrees of entrainment across the five non-overlapping windows. To achieve this, we computed Pearson correlation coefficients between the sum klds for a pair of phase models, for five observations, one per window. Correlation coefficients close to 1 indicated mirrored time courses of entrainment across a pair of rhythms. Of note, correlation coefficients close to 1 could not distinguish between entrainment trajectories that were similar because they were both quite flat (e.g. not different from a uniform distribution of spiking across phase) or trajectories that shared similarly large degrees of entrainment. On the other hand, correlation coefficients close to -1 indicated opposing entrainment time courses (e.g. when one rhythm's sum kld was low, the other rhythm's sum kld was likely larger). Lastly, correlation coefficients closer to 0 reflected unsystematic trajectories of entrainment across pairs of rhythms.

### **Selectivity of rhythmic entrainment as a function of olfactory & visuospatial inputs.**

One of our primary goals was to assess the degree to which interneuron rhythmic entrainment was stably or flexibly recruited across distinct combinations of olfactory and visuospatial inputs (odor-context combinations). We formulated this question as one where we assessed each interneuron's kernel density estimators' interchangeability across odor-context combinations. For instance, a kernel density estimator fit to phase angles drawn from odor-context A would prove "interchangeable" with an estimator from odor-context C if its prediction accuracies on spike-phase data from condition C were indistinguishable from predictions on its own data (e.g. data from condition A).

Prediction accuracy was measured via cross-validated log likelihoods (Python, sklearn, KernelDensity estimator, `score\_samples` method) where a given condition's kernel density estimator was

tested on the data that generated it (“own”) and, in separate tests, against data drawn from all other conditions (“other”). The mean log likelihood computation is formulated as below,

$$\text{mean log likelihood} = \frac{1}{A} \sum_{a=1}^A - \ln(\hat{p}(x_a; h)) \quad (\text{equation 9})$$

where we evaluate the kernel density estimator  $\hat{p}(x_a; h)$ —representing the probability of observing a phase angle given spiking,  $p(\text{phase} | \text{spike} = 1)$ —at each phase angle  $x_a$ . More negative mean log likelihoods indicate worse performance than less negative log likelihoods.

To assess the degree to which odor-context combinations selectively entrained spike-phase relationships across the population of interneurons recorded during odor sampling, we fit a linear mixed effects model predicting average log likelihood per condition, per interneuron, with fixed effects of rhythm (levels: theta, low gamma, and high gamma) and test type (levels: “own” and “other”). Test type refers to whether the log likelihood resulted from a kernel density estimator predictions (using equation 9) on the spike-phase data from the condition the estimator was trained on, while “other” referred to log likelihoods resulting from predictions on spike-phase data from the remaining odor-context conditions.

We additionally included random intercepts for interneuron identity, to account for different baseline firing rates across neurons in the population. The model was specified programmatically (R, lme4 package, `lmer` method) as below,

$$\text{mean log likelihood} \sim \text{rhythm} + \text{test type} + (1 | \text{interneuron identity}) \quad (\text{equation 10})$$

where the notation  $(1 | \text{interneuron identity})$  specifies a unique intercept (“1”) given a unique interneuron identity<sup>81</sup>. We additionally performed a log likelihood ratio test<sup>82</sup> to determine the degree to

which test type explained significant variance in average log likelihoods. The reduced model was specified similarly to equation 10, but omitted the fixed effect of test type,

$$\text{mean log likelihood} \sim \text{rhythm} + (1|\text{interneuron identity}) . \quad (\text{equation 11})$$

We determined the significant differences across the full and reduced models by comparing the ratio of the log likelihoods against the null hypothesis that the ratio should equal 1 under a chi squared distribution:

$$\text{log likelihood ratio} = -2 (\log \text{likelihood}_{\text{reduced}} - \log \text{likelihood}_{\text{full}}) \quad (\text{equation 12})$$

#### **Data and Code Availability.**

All code and data are made publicly accessible at [Github repository](#).

#### **Acknowledgements.**

We thank Chris Heyman, all members of the Neural Crossroads Laboratory, Michael Preston, and Alessandro D'Amico for their valuable feedback and comments on this manuscript. We additionally would like to thank Dr. Douglas Nitz and Dr. Jon Rueckemann, for their detailed analysis suggestions; and a special thanks to Dr. Sean Trott for his extensive assistance in linear mixed effects model specification and implementation.

#### **4.5 References**

1. Komorowski, R. W., Manns, J. R. & Eichenbaum, H. Robust conjunctive item-place coding by hippocampal neurons parallels learning what happens where. *J. Neurosci.* **29**, 9918–9929 (2009).
2. Rajji, T., Chapman, D., Eichenbaum, H. & Greene, R. The role of CA3 hippocampal NMDA receptors in paired associate learning. *J. Neurosci.* **26**, 908–915 (2006).

3. Gilbert, P. E. & Kesner, R. P. Role of the rodent hippocampus in paired-associate learning involving associations between a stimulus and a spatial location. *Behav. Neurosci.* **116**, 63–71 (2002).
4. 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).
5. Amaral, D. G. Emerging principles of intrinsic hippocampal organization. *Curr. Opin. Neurobiol.* **3**, 225–229 (1993).
6. Amaral, D. G. & Witter, M. P. The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience* **31**, 571–591 (1989).
7. Bilash, O. M., Chavlis, S., Johnson, C. D., Poirazi, P. & Basu, J. Lateral entorhinal cortex inputs modulate hippocampal dendritic excitability by recruiting a local disinhibitory microcircuit. *Cell Rep.* **42**, 111962 (2023).
8. Leão, R. N. *et al.* OLM interneurons differentially modulate CA3 and entorhinal inputs to hippocampal CA1 neurons. *Nat. Neurosci.* **15**, 1524–1530 (2012).
9. Ang, C. W., Carlson, G. C. & Coulter, D. A. Hippocampal CA1 circuitry dynamically gates direct cortical inputs preferentially at theta frequencies. *J. Neurosci.* **25**, 9567–9580 (2005).
10. Empson, R. M. & Heinemann, U. The perforant path projection to hippocampal area CA1 in the rat hippocampal-entorhinal cortex combined slice. *J. Physiol.* **484 ( Pt 3)**, 707–720 (1995).
11. Nitz, D. A. & McNaughton, B. L. Hippocampal EEG and unit activity responses to modulation of serotonergic median raphe neurons in the freely behaving rat. *Learn. Mem.* **6**, 153–167 (1999).
12. Varga, V. *et al.* Fast synaptic subcortical control of hippocampal circuits. *Science* **326**, 449–453 (2009).
13. Colom, L. V., Ford, R. D. & Bland, B. H. Hippocampal formation neurons code the level of activation of the cholinergic septohippocampal pathway. *Brain Res.* **410**, 12–20 (1987).
14. Sadeh, S. & Clopath, C. Inhibitory stabilization and cortical computation. *Nat. Rev. Neurosci.* **22**, 21–37 (2021).
15. Valero, M., Navas-Olive, A., de la Prida, L. M. & Buzsáki, G. Inhibitory conductance controls place

- field dynamics in the hippocampus. *Cell Rep.* **40**, 111232 (2022).
16. Royer, S. *et al.* Control of timing, rate and bursts of hippocampal place cells by dendritic and somatic inhibition. *Nat. Neurosci.* **15**, 769–775 (2012).
  17. Pouille, F. & Scanziani, M. Routing of spike series by dynamic circuits in the hippocampus. *Nature* **429**, 717–723 (2004).
  18. Amilhon, B. *et al.* Parvalbumin Interneurons of Hippocampus Tune Population Activity at Theta Frequency. *Neuron* **86**, 1277–1289 (2015).
  19. Csicsvari, J., Jamieson, B., Wise, K. D. & Buzsáki, G. Mechanisms of gamma oscillations in the hippocampus of the behaving rat. *Neuron* **37**, 311–322 (2003).
  20. Hu, H., Gan, J. & Jonas, P. Interneurons. Fast-spiking, parvalbumin<sup>+</sup> GABAergic interneurons: from cellular design to microcircuit function. *Science* **345**, 1255263 (2014).
  21. Takács, V. T., Klausberger, T., Somogyi, P., Freund, T. F. & Gulyás, A. I. Extrinsic and local glutamatergic inputs of the rat hippocampal CA1 area differentially innervate pyramidal cells and interneurons. *Hippocampus* **22**, 1379–1391 (2012).
  22. Gulyás, A. I., Megias, M., Emri, Z. & Freund, T. F. Total number and ratio of excitatory and inhibitory synapses converging onto single interneurons of different types in the CA1 area of the rat hippocampus. *J. Neurosci.* **19**, 10082–10097 (1999).
  23. Freund, T. F. & Buzsáki, G. Interneurons of the hippocampus. *Hippocampus* **6**, 347–470 (1996).
  24. Maccaferri, G. & Dingledine, R. Control of feedforward dendritic inhibition by NMDA receptor-dependent spike timing in hippocampal interneurons. *J. Neurosci.* **22**, 5462–5472 (2002).
  25. Whittington, M. A., Traub, R. D., Kopell, N., Ermentrout, B. & Buhl, E. H. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *Int. J. Psychophysiol.* **38**, 315–336 (2000).
  26. Buzsáki, G., Anastassiou, C. A. & Koch, C. The origin of extracellular fields and currents--EEG, ECoG, LFP and spikes. *Nat. Rev. Neurosci.* **13**, 407–420 (2012).
  27. Rangel, L. M. *et al.* Rhythmic coordination of hippocampal neurons during associative memory

- processing. *Elife* **5**, e09849 (2016).
28. Symanski, C. A., Bladon, J. H., Kullberg, E. T., Miller, P. & Jadhav, S. P. Rhythmic coordination and ensemble dynamics in the hippocampal-prefrontal network during odor-place associative memory and decision making. *Elife* **11**, (2022).
  29. Igarashi, K. M., Lu, L., Colgin, L. L., Moser, M.-B. & Moser, E. I. Coordination of entorhinal-hippocampal ensemble activity during associative learning. *Nature* **510**, 143–147 (2014).
  30. Buzsáki, G. Theta oscillations in the hippocampus. *Neuron* **33**, 325–340 (2002).
  31. Buzsáki, G., Leung, L. W. & Vanderwolf, C. H. Cellular bases of hippocampal EEG in the behaving rat. *Brain Res.* **287**, 139–171 (1983).
  32. Leung, L. S. Generation of theta and gamma rhythms in the hippocampus. *Neurosci. Biobehav. Rev.* **22**, 275–290 (1998).
  33. Kepecs, A., Uchida, N. & Mainen, Z. F. The sniff as a unit of olfactory processing. *Chem. Senses* **31**, 167–179 (2006).
  34. Zhang, L., Lee, J., Rozell, C. & Singer, A. C. Sub-second dynamics of theta-gamma coupling in hippocampal CA1. *Elife* **8**, (2019).
  35. Lasztóczy, B. & Klausberger, T. Layer-specific GABAergic control of distinct gamma oscillations in the CA1 hippocampus. *Neuron* **81**, 1126–1139 (2014).
  36. Colgin, L. L. Do slow and fast gamma rhythms correspond to distinct functional states in the hippocampal network? *Brain Res.* **1621**, 309–315 (2015).
  37. Bragin, A. *et al.* Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J. Neurosci.* **15**, 47–60 (1995).
  38. Alexander, A. S., Rangel, L. M., Tingley, D. & Nitz, D. A. Neurophysiological signatures of temporal coordination between retrosplenial cortex and the hippocampal formation. *Behav. Neurosci.* **132**, 453–468 (2018).
  39. Schomburg, E. W. *et al.* Theta phase segregation of input-specific gamma patterns in entorhinal-hippocampal networks. *Neuron* **84**, 470–485 (2014).

40. Cardin, J. A. *et al.* Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* **459**, 663–667 (2009).
41. Buzsáki, G. & Wang, X.-J. Mechanisms of gamma oscillations. *Annu. Rev. Neurosci.* **35**, 203–225 (2012).
42. Bartos, M. *et al.* Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *Proc. Natl. Acad. Sci. U. S. A.* **99**, 13222–13227 (2002).
43. Klausberger, T. *et al.* Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* **421**, 844–848 (2003).
44. Varga, C. *et al.* Functional fission of parvalbumin interneuron classes during fast network events. *Elife* **3**, (2014).
45. Czurkó, A., Huxter, J., Li, Y., Hangya, B. & Muller, R. U. Theta phase classification of interneurons in the hippocampal formation of freely moving rats. *J. Neurosci.* **31**, 2938–2947 (2011).
46. Forro, T., Valenti, O., Lasztocki, B. & Klausberger, T. Temporal organization of GABAergic interneurons in the intermediate CA1 hippocampus during network oscillations. *Cereb. Cortex* **25**, 1228–1240 (2015).
47. Somogyi, P. & Klausberger, T. Defined types of cortical interneurone structure space and spike timing in the hippocampus. *J. Physiol.* **562**, 9–26 (2005).
48. Klausberger, T. *et al.* Spike timing of dendrite-targeting bistratified cells during hippocampal network oscillations in vivo. *Nat. Neurosci.* **7**, 41–47 (2004).
49. Tukker, J. J., Fuentealba, P., Hartwich, K., Somogyi, P. & Klausberger, T. Cell type-specific tuning of hippocampal interneuron firing during gamma oscillations in vivo. *J. Neurosci.* **27**, 8184–8189 (2007).
50. Varga, C., Golshani, P. & Soltesz, I. Frequency-invariant temporal ordering of interneuronal discharges during hippocampal oscillations in awake mice. *Proc. Natl. Acad. Sci. U. S. A.* **109**, E2726–34 (2012).
51. Rivière, P. D., Schamberg, G., Coleman, T. P. & Rangel, L. M. Modeling relationships between



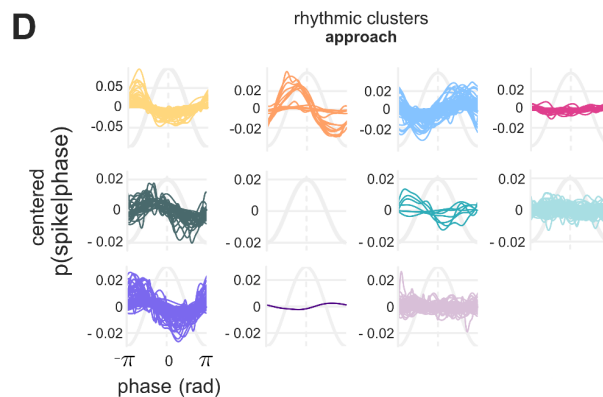
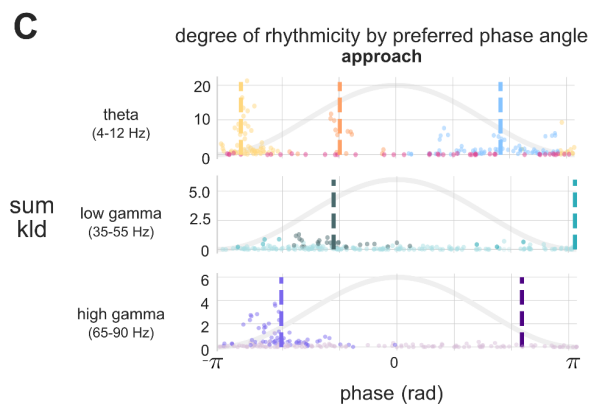
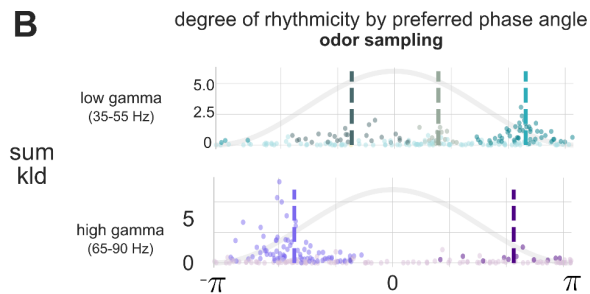

- rhythmic processes and neuronal spike timing. *J. Neurophysiol.* **128**, 593–610 (2022).
52. Colom, L. V. & Bland, B. H. State-dependent spike train dynamics of hippocampal formation neurons: evidence for theta-on and theta-off cells. *Brain Res.* **422**, 277–286 (1987).
  53. Nitz, D. & McNaughton, B. Differential modulation of CA1 and dentate gyrus interneurons during exploration of novel environments. *J. Neurophysiol.* **91**, 863–872 (2004).
  54. Hutcheon, B. & Yarom, Y. Resonance, oscillation and the intrinsic frequency preferences of neurons. *Trends Neurosci.* **23**, 216–222 (2000).
  55. Zemankovics, R., Káli, S., Paulsen, O., Freund, T. F. & Hájos, N. Differences in subthreshold resonance of hippocampal pyramidal cells and interneurons: the role of h-current and passive membrane characteristics. *J. Physiol.* **588**, 2109–2132 (2010).
  56. Whittington, M. A. & Traub, R. D. Interneuron diversity series: inhibitory interneurons and network oscillations in vitro. *Trends Neurosci.* **26**, 676–682 (2003).
  57. Ahmed, O. J. & Mehta, M. R. Running speed alters the frequency of hippocampal gamma oscillations. *J. Neurosci.* **32**, 7373–7383 (2012).
  58. Fukuda, T. & Kosaka, T. Gap junctions linking the dendritic network of GABAergic interneurons in the hippocampus. *J. Neurosci.* **20**, 1519–1528 (2000).
  59. Traub, R. D. *et al.* Gap junctions between interneuron dendrites can enhance synchrony of gamma oscillations in distributed networks. *J. Neurosci.* **21**, 9478–9486 (2001).
  60. Fricker, D. & Miles, R. Interneurons, spike timing, and perception. *Neuron* **32**, 771–774 (2001).
  61. Johnson, T. D., Keefe, K. R. & Rangel, L. M. Stimulation-induced entrainment of hippocampal network activity: Identifying optimal input frequencies. *Hippocampus* **33**, 85–95 (2023).
  62. Schwerdtfeger, W. K., Buhl, E. H. & Germroth, P. Disynaptic olfactory input to the hippocampus mediated by stellate cells in the entorhinal cortex. *J. Comp. Neurol.* **292**, 163–177 (1990).
  63. Karalis, N. & Sirota, A. Breathing coordinates cortico-hippocampal dynamics in mice during offline states. *Nat. Commun.* **13**, 467 (2022).
  64. Wilson, R. C. & Steward, O. Polysynaptic activation of the dentate gyrus of the hippocampal

- formation: an olfactory input via the lateral entorhinal cortex. *Exp. Brain Res.* **33**, 523–534 (1978).
65. Witter, M. P. The perforant path: projections from the entorhinal cortex to the dentate gyrus. *Prog. Brain Res.* **163**, 43–61 (2007).
66. Gattas, S., Elias, G. A., Janecek, J., Yassa, M. A. & Fortin, N. J. Proximal CA1 20-40 Hz power dynamics reflect trial-specific information processing supporting nonspatial sequence memory. *Elife* **11**, (2022).
67. Tort, A. B. L., Komorowski, R. W., Manns, J. R., Kopell, N. J. & Eichenbaum, H. Theta-gamma coupling increases during the learning of item-context associations. *Proc. Natl. Acad. Sci. U. S. A.* **106**, 20942–20947 (2009).
68. Rivière, P. D. & Rangel, L. M. Spike-Field Coherence and Firing Rate Profiles of CA1 Interneurons During an Associative Memory Task. in *Advances in the Mathematical Sciences* 161–171 (Springer International Publishing, 2018).
69. Dupret, D., O’Neill, J. & Csicsvari, J. Dynamic reconfiguration of hippocampal interneuron circuits during spatial learning. *Neuron* **78**, 166–180 (2013).
70. Wilent, W. B. & Nitz, D. A. Discrete place fields of hippocampal formation interneurons. *J. Neurophysiol.* **97**, 4152–4161 (2007).
71. Kubie, J. L., Muller, R. U. & Bostock, E. Spatial firing properties of hippocampal theta cells. *J. Neurosci.* **10**, 1110–1123 (1990).
72. Ego-Stengel, V. & Wilson, M. A. Spatial selectivity and theta phase precession in CA1 interneurons. *Hippocampus* **17**, 161–174 (2007).
73. Barthó, P. *et al.* Characterization of neocortical principal cells and interneurons by network interactions and extracellular features. *J. Neurophysiol.* **92**, 600–608 (2004).
74. Marron, J. S. *An Asymptotically Efficient Solution to the Bandwidth Problem of Kernel Density Estimation.* (University of North Carolina Institute of Statistics, 1983).
75. Jones, M. C. & Henderson, D. A. Maximum likelihood kernel density estimation. 1–29 (2001).
76. Hall, P. & Wand, M. P. *On the Accuracy of Binned Kernel Density Estimators.* (1993).

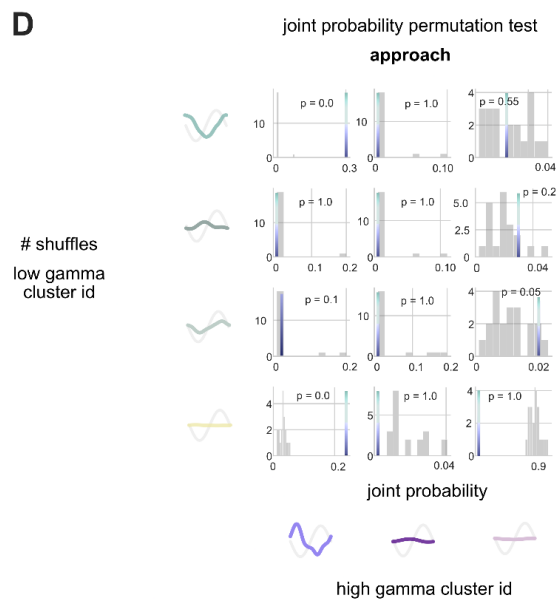
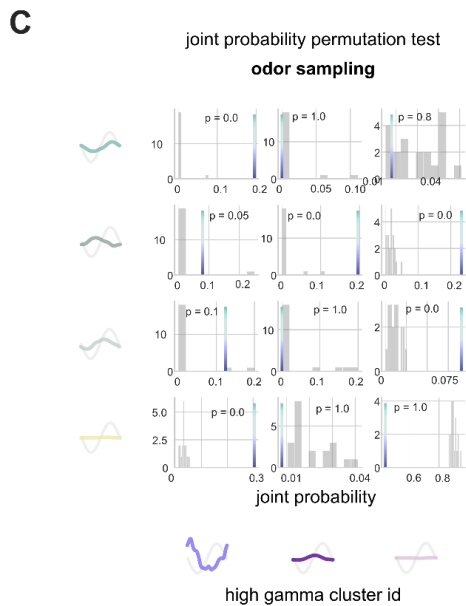
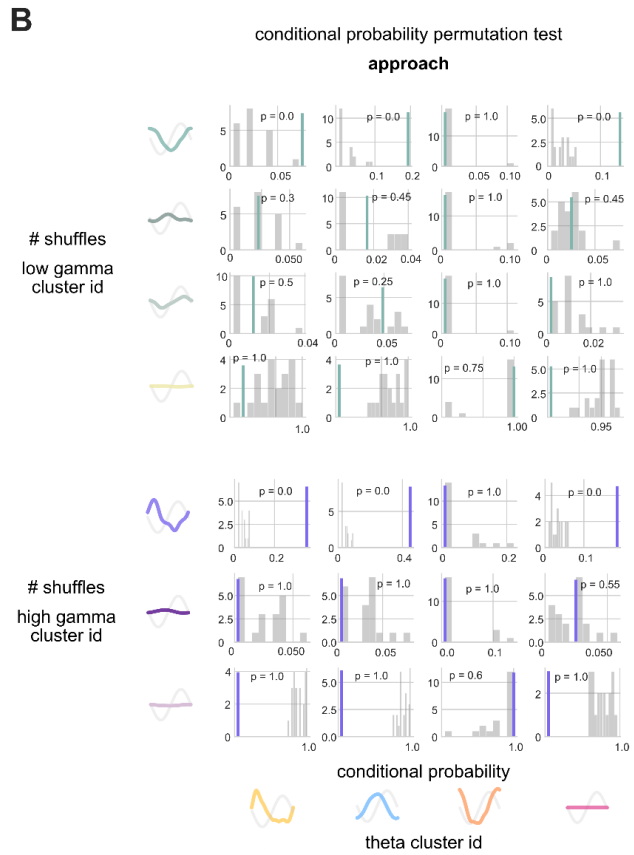
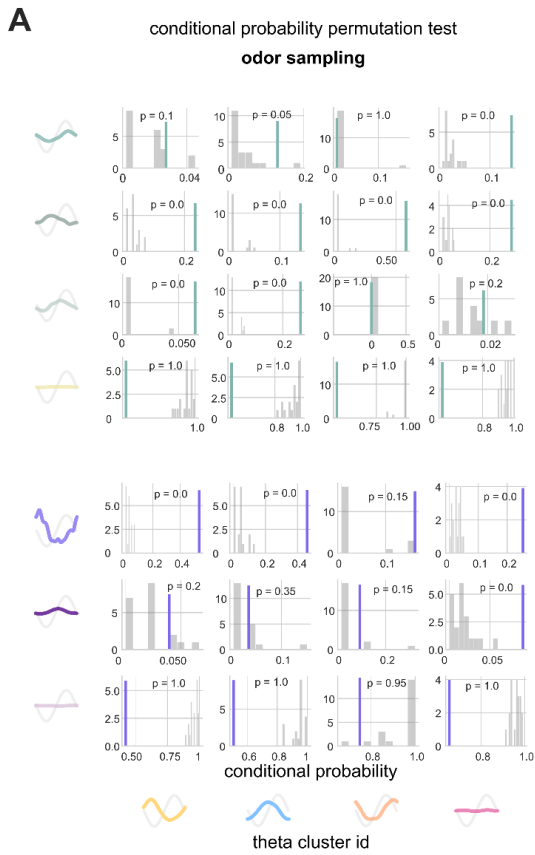
77. Johnson, T. D., Coleman, T. P. & Rangel, L. M. A flexible likelihood approach for predicting neural spiking activity from oscillatory phase. *J. Neurosci. Methods* **311**, 307–317 (2019).
78. Cover, T. M. & Thomas, J. A. *Elements of Information Theory*. (John Wiley & Sons, 2012).
79. Mardia, K. V. & Jupp, P. E. *Directional Statistics*. (John Wiley & Sons, 2009).
80. Cohen, M. X. *Analyzing Neural Time Series Data: Theory and Practice*. (MIT Press, 2014).
81. Winter, B. *Statistics for Linguists: An Introduction Using R*. (Routledge, 2019).
82. Hosmer, D. W. & Lemeshow, S. *Goodness of Fit Test for the Multiple Logistic Regression Model*. (1979).

Supplementary figures make up the remainder of this chapter.

**Supplementary Figure S4.1.** Schematic illustration and equation describing the computation for the sum of the Kullback-Leibler divergences (kld) between a given spike-train's rhythmic phase model and the spike-train's own uniform average probability of spiking. Phase model is in pale teal, uniform model is a flat black horizontal line, and individual kld values are represented by black vertical dotted lines. **(B)** Clustering schemes for low gamma, and high gamma for models fit on interneuron spike trains recorded during the odor sampling epoch. **(C)** Rhythmic phase models fit to spike trains recorded during the locomotive approach epoch. Theta model centroids (top subpanel) during the approach emerged as follows: trough (yellow dotted line): -2.73 rad; ascending phase (orange dotted line): -0.987 rad; descending phase (blue dotted line): 1.83 rad. Low gamma model centroids (middle) subpanel: ascending (leftmost dotted line): -1.10 rad; trough (rightmost dotted line): 3.12 rad. High gamma (bottom subpanel) centroids: ascending phase (left dotted line): -2.02 rad, descending phase: 2.18 rad. **(D)** Mean-centered rhythmic phase models fit to interneuron spike trains and LFPs drawn from the locomotive approach epoch. Rows correspond to models fit to the phase of theta, low gamma, and high gamma, respectively. Note that no interneuron spike trains were sorted into the low gamma early descent phases.



**Supplementary Figure S4.2. Permutation test results for conditional and joint probabilities of rhythmic entrainment across frequency bands.** (A) Results of the permutation tests examining the likelihood of randomly observing conditional probabilities of low gamma (top, rows) or high gamma (bottom, rows) cluster membership given theta phase relationship (columns) during odor sampling. Each subpanel corresponds to a combination of low gamma given theta cluster identity (top), or high gamma given theta cluster identity (bottom). The colored vertical line in each subpanel represents the conditional probability computed from the intact data, while the grey histogram corresponds to the conditional probabilities computed from the shuffled data. The probability of obtaining a conditional probability greater than those drawn from the distribution of shuffles is printed on each subpanel, and constituted the basis for significance ( $p < 0.05$ ). (B) Same as in A, but for the approach epoch. (C) Results of the permutation tests examining the likelihood of randomly observing joint probabilities of low gamma (rows) and high gamma (columns) cluster membership, during the odor sampling epoch. The features of each subpanel match all features described in A. (D) Same as in C, but for the approach epoch.



Chapter 4, in full, is being prepared for submission to *Current Biology*. Rivière, Pamela D;  
Rangel, Lara M. The dissertation author was the primary investigator and author of this paper.



## CHAPTER 5: DISCUSSION

Throughout this dissertation, I explore hippocampal inhibitory interneurons' spike rates and spike timing in response to evolving constellations of inputs, as they manifest across distinct behavioral states and stimuli. In Chapter 1, I demonstrate that a population of hippocampal CA1 inhibitory interneurons reflect behaviorally relevant stimulus information in their firing rates during associative memory processing. This result suggests that inputs differentially recruit inhibition as animals perform the task. To quantify and predict inhibitory interneuron spike timing dynamics, in Chapter 2 I innovatively apply statistical tools to model the relationship between interneuron spiking and the phase of two important CA1 oscillations, the theta (4-12 Hz) and low gamma (35-55 Hz) frequency bands. I leverage each model's cross-validated prediction accuracy to reveal moment-to-moment variations in interneuron spike-phase relationships across locomotive and associative memory processing epochs. In some cases, interneurons demonstrated dynamism in their oscillatory spike timing *within* associative memory processing windows, flexibly shifting their spike timing preferences from theta to low gamma frequencies. With the work in Chapter 3, I go on to characterize the specific rhythmic phase at which CA1 inhibitory interneurons relate to theta (4-12 Hz), low gamma (35-55 Hz), and high gamma (65-90 Hz) oscillations. I find that an interneuron's theta phase preference (or lack thereof), predicts its engagement in faster oscillations, and these multi-rhythmic relationships vary as a function of the animal's behavioral state. Moreover, if an interneuron exhibited a rhythmic phase preference, it could still shift to firing uniformly relative to a given rhythm when other odor-context combinations were present, indicating variable likelihoods of entrainment during associative memory processing. In sum, this work offers evidence in support of flexible hippocampal interneuron recruitment, and motivates future research efforts reevaluating neural dynamics across the wide range of behavioral states that organisms express as they make their way in the world.