

December 2022

Deciphering the Firing Patterns of Hippocampal Neurons During Sharp-Wave Ripples

Kourosh Maboudi Ashmankamachali
University of Wisconsin-Milwaukee

Follow this and additional works at: <https://dc.uwm.edu/etd>



Part of the [Biomedical Engineering and Bioengineering Commons](#), and the [Neuroscience and Neurobiology Commons](#)

Recommended Citation

Maboudi Ashmankamachali, Kourosh, "Deciphering the Firing Patterns of Hippocampal Neurons During Sharp-Wave Ripples" (2022). *Theses and Dissertations*. 3039.
<https://dc.uwm.edu/etd/3039>

This Dissertation is brought to you for free and open access by UWM Digital Commons. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of UWM Digital Commons. For more information, please contact scholarlycommunicationteam-group@uwm.edu.

**DECIPHERING THE FIRING PATTERNS OF HIPPOCAMPAL
NEURONS DURING SHARP-WAVE RIPPLES**

by

Kourosh Maboudi

A Dissertation Submitted in

Partial Fulfillment of the

Requirements for the Degree of

Doctor of Philosophy

in Psychology

at

The University of Wisconsin-Milwaukee

December 2022

ABSTRACT

DECIPHERING THE FIRING PATTERNS OF HIPPOCAMPAL NEURONS DURING SHARP-WAVE RIPPLES

by

Kourosh Maboudi

The University of Wisconsin-Milwaukee, 2022
Under the Supervision of Dr. Kamran Diba

The hippocampus is essential for learning and memory. Neurons in the rat hippocampus selectively fire when the animal is at specific locations - place fields - within an environment. Place fields corresponding to such place cells tile the entire environment, forming a stable spatial map supporting navigation and planning. Remarkably, the same place cells reactivate together outside of their place fields and in coincidence with sharp-wave ripples (SWRs) - dominant electrical field oscillations (150-250 Hz) in the hippocampus. These offline SWR events frequently occur during quiet wake periods in the middle of exploration and the follow-up slow-wave sleep and are associated with spatial memory performance and stabilization of spatial maps. Therefore, deciphering the firing patterns during these events is essential to understanding offline memory processing.

I provide two novel methods to analyze the SWRs firing patterns in this dissertation project. The first method uses hidden Markov models (HMM), in which I model the dynamics of neural activity during SWRs in terms of transitions between distinct states of neuronal ensemble activity. This method detects consistent temporal structures over many instances of SWRs and, in contrast to standard approaches, relaxes the dependence on positional data during the behavior

to interpret temporal patterns during SWRs. To validate this method, I applied the method to quiet wake SWRs. In a simple spatial memory task in which the animal ran on a linear track or in an open arena, the individual states corresponded to the activation of distinct group of neurons with inter-state transitions that resembled the animal's trajectories during the exploration. In other words, this method enabled us to identify the topology and spatial map of the explored environment by dissecting the firings occurring during the quiescence periods' SWRs. This result indicated that downstream brain regions may rely only on SWRs to uncover hippocampal code as a substrate for memory processing.

I developed a second analysis method based on the principles of Bayesian learning. This method enabled us to track the spatial tunings over the sleep following exploration of an environment by taking neurons' place fields in the environment as the prior belief and updating it using dynamic ensemble firing patterns unfolding over time. This method introduces a neuronal-ensemble-based approach that calculates tunings to the position encoded by ensemble firings during sleep rather than the animal's actual position during exploration. When I applied this method to several datasets, I found that during the early slow-wave sleep after an experience, but not during late hours of sleep or sleep before the exploration, the spatial tunings highly resembled the place fields on the track. Furthermore, the fidelity of the spatial tunings to the place fields predicted the place fields' stability when the animal was re-exposed to the same environment after ~ 9 h. Moreover, even for neurons with shifted place fields during re-exposure, the spatial tunings during early sleep were predictive of the place fields during the re-exposure. These results indicated that early sleep actively maintains or retunes the place fields of neurons, explaining the representational drift of place fields across multiple exposures.

© Copyright by Kourosh Maboudi, 2022
All Rights Reserved

To
Woman, Life, Freedom
The brave women and men who are fighting for freedom around the globe,
especially in Iran

TABLE OF CONTENTS

Abstract	ii
TABLE OF CONTENTS	vi
LIST OF FIGURES	ix
ACKNOWLEDGEMENTS	xi
1 Introduction.....	1
1.1 Sleep.....	1
1.2 Sleep benefits memory	2
1.3 The hippocampus and its role in memory	3
1.4 Patterns of neural activity in the hippocampus across wake and sleep	4
1.5 Theta sequences	6
1.6 Sharp-wave ripples.....	7
1.7 Interferential studies on the mnemonic role of SWRs	8
1.8 The spike content of sharp-wave ripples.....	9
1.8.1 Awake replay supports memory storage and planning	10
1.8.2 Sleep reactivations and replays	11
1.9 Project goal: temporal patterns of firing during the sharp-wave ripples	13
1.9.1 Hidden Markov models to detect SWRs spatiotemporal structures.....	14
1.9.2 Bayesian learning of spatial tunings during sleep	15
2 Uncovering temporal structure in hippocampal output patterns.....	17
2.1 Introduction.....	17
2.2 Results.....	19
2.2.1 Awake population burst events	19
2.2.2 Learning hidden Markov models from PBE data	20
2.2.3 What do the learned model parameters tell us about PBEs?	23
2.2.4 HMM-congruent PBEs capture sequence replay	26
2.2.5 Modeling internally generated activity during open field behavior.....	32
2.2.6 Extra-spatial information	35
2.3 Discussion	39
2.3.1 Graded, non-binary replay detection.....	41
2.3.2 Learning, model congruence and replay quality	41
2.3.3 Application to extra-spatial behaviors	42
2.3.4 Future possibilities	43
2.3.5 Conclusions.....	44

2.4	Materials and methods	45
2.4.1	Experiment paradigm/neural data recording	45
2.4.2	Population burst events	47
2.4.3	Hidden Markov model of PBE activity.....	48
2.4.4	Ordering states for visualization	49
2.4.5	Surrogate datasets and shuffle methods	50
2.4.6	Calculating sparsity and connectivity of the model parameters.....	50
2.4.7	Model connectivity and sequences.....	51
2.4.8	Latent state place fields.....	52
2.4.9	Decoding position from latent state sequences	52
2.4.10	Bayesian replay detection	53
2.4.11	Replay detection via PBE model congruence	54
2.4.12	Human scoring and detection comparison	55
2.4.13	HMM model quality across sessions.....	56
3	Retuning of hippocampal place representations during sleep	57
3.1	Introduction.....	59
3.2	Results.....	62
3.2.1	POST but not PRE ripples align with place fields on the maze.....	62
3.2.2	Learned tunings reveal the stability of spatial representations in sleep.	64
3.2.3	Sleep dynamics predict representational consistency and drift.....	68
3.2.4	Representations during high score preplays do not align with maze place fields	71
3.2.5	Sleep representations are driven by awake ripples and theta oscillations.....	74
3.3	Discussion	77
3.4	Methods.....	80
3.4.1	Behavioral task and data acquisition.....	80
3.4.2	Place field calculation	82
3.4.3	LFP analysis and brain state detection	82
3.4.4	Detection of ripple events	83
3.4.5	Bayesian learned tunings	83
3.4.6	Fidelity of the learned tunings across epochs	86
3.4.7	Learned tuning's dynamics	86
3.4.8	Ripple event replay scores	88
3.4.9	Place fields' overlap with decoded posterior	89
3.4.10	Multiple regression analyses.....	89
4	Conclusion and future direction	91

References..... 98

LIST OF FIGURES

Figure 1.1: Difference in the number of correctly recalled syllables at various sleep or wake retention intervals.....	2
Figure 1.2: Slow oscillatory stimulation enhances declarative memory performance.....	3
Figure 1.3: Behavioral sequence of place cell activation during a traversal on a linear track.	5
Figure 1.4. Millisecond sequences of place cells firings during theta cycles (theta sequences).....	6
Figure 1.5: Increased BOLD activity in different neocortical and limbic brain regions, coincident with the sharp-wave ripples in the hippocampus.	8
Figure 1.6: Disrupting sharp-wave ripples through electrical stimulation of the hippocampus leads to impairment in spatial memory.	10
Figure 1.7: Example of population burst events co-occurring with hippocampal sharp-wave ripples.	11
Figure 1.8: Enhanced coactivation of neurons during the sleep following an experience.....	12
Figure 1.9: A Markov chain model.....	15
Figure 1.10. A hidden Markov model to model temporal patterns of firing on a maze.....	16
Figure 2.1: A hidden Markov model of ensemble activity during PBEs.	22
Figure 2.2: A hidden Markov model of ensemble activity during population burst events.....	24
Figure 2.3: Latent states capture positional code.	25
Figure 2.4: Replay events can be detected via HMM congruence.....	28
Figure 2.5: Comparing HMM congruence and Bayesian decoding in replay detection..	30
Figure 2.6. Modeling PBEs in open field.....	33
Figure 2.7: Extra-spatial structure.....	37
Figure 2.8: Temporal structure during a sleep period following object-location memory task.	39
Figure 3.1: Schematic for Bayesian learning of hippocampal spatial tunings during offline states..	61
Figure 3.2: Bayesian learned tunings in theta and offline states.....	63
Figure 3.3: Learned tunings during ripples are stable for hours in POST.	65
Figure 3.4: Spatial representations decay over the course of sleep.....	67
Figure 3.5: POST ripple tunings predict maze re-exposure place fields.....	70
Figure 3.6: Place field fidelities do not strictly correlate with replay score.....	73
Figure 3.7: A diversity of pre-existing representations impacts the tunings in POST sleep.....	76
Figure 3.8: A L-ratio measure to quantify the degree of overlap in the spike feature space between pairs of units.....	87
Figure 4.1: Sequential content during PBEs within each behavioral epoch.	93

Figure 4.2: The states during POST but not PRE represented maze locations.	94
Figure 4.3. Examples of methods to calculate replay scores and their statistical significance.	96
Figure 4.4: Agreement of each algorithmic replay score method with human scores..	97

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisor, Kamran Diba, who gave me the opportunity to join his lab as a Ph.D. student. You have been always available to discuss the current state of my work, even during the weekends. No matter how complicated the topic of discussion was, either computational and experimental, you were always ready for brainstorming and coming up with great ideas. Thank you so much for your guidance and advice along the path. Thank you to my dissertation committee members for your interest, enthusiasm, and patience. Thanks particularly to Caleb Kemere who I had the honor of collaborating for my dissertation projects. Your energy and enthusiasm helped me push the limits. I want to acknowledge my committee members in Milwaukee, Karyn Frick, Fred Helmstetter, and Rodney Swain for your valuable insights and feedback regarding my projects. Thank you to the directors of Graduate Studies of Psychology Department in past and present, Susan Lima, who although is no longer with us, I never forget how helpful and sympathetic she was towards me, and Han Joo Lee who always promptly and kindly responded to my many questions.

Thank you to my lab mates both past and present: Dr. Nat Kinsky, Dr. Utku Kaya, Dr. Bapun Giri, Rachel Wahlberg, Pho Hale, and Dr. Laurel Watkins de Jong. You guys have been great friends and colleagues and your advice and feedback have been always helpful. Thank you to Mary Norat, Nikki Vladimerou, and Amy Harms and other staff and faculties in the University of Michigan Anesthesiology Department that assisted me with my research after that our lab moved from Milwaukee to Ann Arbor.

This also would have not been possible without the support of my family and friends. Thank you to my friends in Ann Arbor: Mehrdad, Bahador, Negin, Amir, Ghoncheh, and Ahmad for all the fun that we had together in Ann Arbor and your support. Obviously, I would have not made it this far without the love and support from my parents, Siavash Maboudi and Kobra Ghavidel, and my beloved sisters, Maria and Gandom whose love and support encouraged me to follow my interests, for which I could not be more grateful.

1 Introduction

1.1 Sleep

A significant portion of every animal's life is devoted to sleeping, a state of body and brain observed throughout the development and preserved across various species, from fruit flies to humans. Sleep is not simply a state of inactivity but impacts body and brain physiology. Sleep balances heart rate, blood pressure, and respiration rate (Snyder et al., 1964), and cell repair and growth are most significant during sleep (Bellesi et al., 2013; Guzmán-marín et al., 2003). While the wake behavior of the animal and the constellation of stimuli experienced by the animal every day requires strengthened synapses between neurons in the brain, sleep decreases the net synaptic strength and restores cellular homeostasis (Tononi & Cirelli, 2014). This global downscaling of the synapses benefits brain memory function through an increased signal-to-noise ratio for salient memories and increases the capacity of the brain to generate synapses for new learnings.

Sleep cycles through distinct stages with transitions regulated by the brain neuromodulatory systems in the hypothalamus and brain stem. While slow wave sleep (SWS) is recognized with low levels of acetylcholine (ACh), norepinephrine (NE), serotonin (5-HT), and histamine (HA), rapid eye movement (REM) is known with a high acetylcholine tone (Diekelmann & Born, 2010; Krishnan et al., 2016). This differential neuromodulatory influence leads to distinct patterns of neural activity in the thalamocortical system; low amplitude theta waves, similar to active exploration and high attentive states, occur during REM sleep. However, SWS is abundant with slow oscillations (<1 Hz) in the neocortex, thalamocortical spindles, and hippocampal SWRs.

1.2 Sleep benefits memory

Early studies on humans indicated the effect of sleep on memory retention. A classic study by Jenkins and Dallenbach showed that subjects forget slower while asleep (Jenkins & Dallenbach, 1924; Rasch & Born, 2022); the subjects correctly recalled more meaningless syllables while sleeping during the retention interval than when they were awake (**Figure 1.1**).




Figure 1.1: Difference in the number of correctly recalled syllables at various sleep or wake retention intervals. Retention of nonsense syllables was tested after intervals of different length when the participants were asleep (Sleep) or awake (Wake) during the interval. The rate of correctly remembered syllables were higher during Sleep than Wake condition (figure reproduced from Rasch & Born, 2022).

This study suggested that sleep maintains learned information over time by either passively protecting it from interference due to subsequent learnings during the wake or actively strengthening the original memory traces (Rasch & Born, 2022). The latter view was further supported by studies that showed that the consolidation depends on the composition of sleep, with differential outcomes for sleep rich of REM sleep or SWS. While SWS-rich sleep benefits declarative memories (word-pair, spatial memory), REM-rich sleep improved procedural, implicit (mirror tracing), and perceptual memories (Karni et al., 1994; Rasch & Born, 2022). Enhancing SWS by inducing slow oscillations during early sleep after learning using transcranial

electrical stimulations at 0.75 Hz enhanced the retention of memory for word pairs (**Figure 1.2**; Marshall et al., 2006). The number of recalled words was higher for the Stimulation than the Sham group with 5 Hz stimulation.




Figure 1.2: Slow oscillatory stimulation enhances declarative memory performance. Performance on the declarative paired-associate memory task after sleep in two groups of subjects with stimulations at either < 1 Hz (Stimulation) or 5 Hz (Sham) during the early sleep are shown. Performance is expressed as the difference between the number of correct words reported at retention test and learning (figure reproduced from Marshall et al., 2006)

Sleep is associated with the reactivation of memory traces formed during an experience (Peigneux et al., 2004). In an object-location task paired with an experimental odor, presenting the odor during the following sleep to induce reactivations led to enhanced performance in a follow-up retention test (Diekelmann et al., 2011). The reactivation of memory traces during sleep was extensively studied for spatial memories in rodents (Nádasdy et al., 1999; Wilson & McNaughton, 1994). These studies mainly used the correlation between the spike trains of different neurons as a proxy for reactivation.

1.3 The hippocampus and its role in memory

The hippocampus is involved in forming long-term memories of events in humans (Scoville & Milner, 1957) and spatial memory in rodents (Morris et al., 1986). A case study on a patient

named H.M. with a lesion in the hippocampus and surrounding temporal lobe structures as a last resort to relieve frequent epileptic seizures provided the initial evidence in support of the necessity of the hippocampus for processing episodic memories in humans. The hippocampus receives highly processed information from many areas of the brain and is in an ideal position to link events across space and time. According to the two-stage model of memory (Buzsáki, 1989), upon a new experience, memory traces are initially formed through development of new synapses between hippocampal neurons. However, these initial memory traces are vulnerable to disruption and forgetting. To convert the short-term to long-term memory, the hippocampus needs to gradually transfer the learned information to neocortex. To do this, the hippocampus reactivates and spreads the information by effectively discharging neocortical neurons.

1.4 Patterns of neural activity in the hippocampus across wake and sleep

When an animal begins to explore a new environment, each hippocampal place cell develops one or more firing fields at specific locations within the environment (J. O'Keefe & Dostrovsky, 1971; J. O'Keefe & Speakman, 1987). A set of place cells collectively represent an environment in a unique way. These representations are stable across long-term periods which means that these representations are retrieved and not created de novo every time the animal explores a familiar environment (Thompson & Best, 1989).

Patterns of neuronal activity unfolding at different timescales are involved in forming ensembles underlying spatial memories. Hippocampal ensemble activity tends to be sequential (Buzsáki & Tingley, 2018). Sequential activity is thought to underlie numerous fundamental operations of the hippocampus. When the animal takes a trajectory within the environment, place cells activate in the same order as the animal traverses their place fields along the trajectory (**Figure 1.3**).

Other continuous task variables, besides position, map onto sequential activations in the hippocampus, as well; a study showed that during a delay prior to a choice, each moment was represented by the activation of a distinct set of neurons (Pastalkova et al., 2008). This sequential activation was unique and predictive of the upcoming choice of the animal. Another study showed that in a behavioral task in which the animal was trained to continuously increase sound frequency within a specific range towards a target (through a joystick), hippocampal neurons developed selective firing fields that spanned the entire frequency range (Aronov et al., 2017). These studies indicated that the hippocampus is well-purposed to represent (and process) the continuous dimensions of the experiences through developing sequential firing fields.




Figure 1.3: Behavioral sequence of place cell activation during a traversal on a linear track.

A sequence of place cells activation when a rat takes a trajectory from left to right on a linear track. Spike trains of different place cells are displayed in different colors. The speed of the animal at each moment is displayed in the bottom. On the top LFP trace signifies theta oscillations (6-10 Hz) during the traversal and SWRs before and after the traversal. The behavioral sequence is replayed during the SWRs, marked with rectangles. The two other rectangles on the sides show the magnified versions (figure reproduced from Diba & Buzsáki, 2007).

In addition to the behavioral time scale sequences, the activation of place cells is further entrained by dominant oscillations that are observed in the hippocampal local field potential (LFP), leading to the emergence of millisecond time-scale sequences: theta sequences (Dragoi & Buzsáki, 2006) and SWR-associated replays (Diba & Buzsáki, 2007; Foster & Wilson, 2006). The speed-up sequences that are observed during the theta sequences and SWR-associated

replays are suggested to be critical to spatial memory processing, as they might strengthen memory-related representations (Drieu et al., 2018; Sadowski et al., 2016) and facilitate transfer of learned information to cortical structures (Maingret et al., 2016; Ji & Wilson, 2007).

1.5 Theta sequences

During the active exploration of an environment, theta oscillations (6-10 Hz) increase in power. These oscillations are internally generated in the hippocampus and triggered by medial septum projections to the hippocampus (Buzsaki, 2002). Place cells with overlapping place fields sequentially activate during each theta cycle, with the same order as in the behavioral sequence (**Figure 1.4**). The millisecond time scale activation of place cells during the theta sequences depends on the theta oscillations (Drieu et al., 2018; Wang et al., 2015). A recent study using T-maze indicated that theta sequences observed on the stem of the T-maze constantly cycle between the representation of either arm, suggesting a cognitive role for the theta sequences for representation and evaluation of the animal's future choices (Kay et al., 2020).




Figure 1.4. Millisecond sequences of place cells firings during theta cycles (theta sequences). Top, partially overlapping place fields of three example place cells plotted as colored ellipses. Middle, the behavioral sequence of place cells activation on the track. Bottom, theta LFP with cycles separated using dotted lines. Intermingled with the behavioral sequence, theta sequences are observed within each theta cycle. Each theta sequence is a partial replication of the behavioral sequence in a temporally-compressed manner (figure reproduced from Wikenheiser & Redish, 2015).

Inactivation of theta oscillations by injecting muscimol into the medial septum reduces the number of observed theta sequences (Wang et al., 2015). Theta sequence compression allows cell assemblies to fire with brief delays that allow strengthening of the synapse according to the Hebbian spike-timing dependent plasticity (STDP) (Magee & Johnston, 1997).

1.6 Sharp-wave ripples


A sharp wave ripple (SWR) is a complex of oscillations in the hippocampal LFP that are predominantly observed during SWS and quiet wake periods when the animal is immobile, drinking, eating, or grooming. The SWR complexes in cornu Ammonis 1 (CA1) hippocampal subregion consist of a negative deflection in the LFP in the stratum radiatum layer - which reflects the inputs from the CA3 to CA1 - and ripple oscillations (150-250 Hz) in the CA1 pyramidal layer (Buzsaki, 2015). The SWRs are intrinsically generated in the hippocampus and are initiated by powerful population bursts in the recurrent network of CA3. The powerful SWR bursts invade the CA1, subiculum and deep layers of the entorhinal cortex (Buzsaki, 1998).

A series of studies supported the role of SWRs in memory processing. First, hippocampal neurons that encode learned information tend to coactivate in a time-compressed manner during the SWRs. The proximal discharge of the neurons during the SWRs is suitable for induction of long-term potentiation, as supported by in vitro studies (Sadowski et al., 2016). Second, the rate of sharp wave ripples increases in frequency and amplitude during SWS after a learning compared with a baseline prior to the learning (Eschenko et al., 2008). Third, the hippocampal SWRs tend to synchronize with thalamocortical spindles (5-15 Hz) (Maingret et al., 2016) and cortical ripples (Khodagholy et al., 2017) around the slow oscillation, promoting hippocampal-cortical communication underlying memory consolidation. Accordingly, a monkey study showed

brain-wide activation during SWRs (Logothetis et al., 2012), indicating SWRs' role as a conduit for communication between the hippocampus and other brain regions, including the cortex (

Figure 1.5).

Figure 1.5: Increased BOLD activity in different neocortical and limbic brain regions, coincident with the sharp-wave ripples in the hippocampus. Average ripple-triggered activation map show an increased activation in neocortex and limbic regions, but a suppression in the subcortical areas (diencephalon, mesencephalon and metencephalon) (d) (figure reproduced from Logothetis et al., 2012).



1.7 Interferential studies on the mnemonic role of SWRs

A series of studies on rats showed that disrupting SWRs during the quiet wake or SWS leads to impairment in spatial memory and spatial working memory performance (Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Jadhav et al., 2012). A study using a W-maze when the animals were required to alternate between different arms of the maze showed that disrupting quiet wake SWRs caused impairment in learning (Jadhav et al., 2012). While this study found intact place field activity following SWR disruption, a more recent study found that optogenetic

silencing of place cells' activity during quiet wake SWRs destabilized place fields when the animal was re-exposed to the same environment after an hour of rest (Roux et al., 2017). These findings on the mnemonic role of quiet wake SWRs were complemented by a gain-of-function study indicating that elongating the SWRs through optogenetic stimulation of the CA1 pyramidal layer enhances W-maze spatial alternation performance (Fernández-Ruiz et al., 2019).

Similarly, disrupting SWRs during SWS impeded acquisition of memory for learnings before and after sleep. Girardeau et al. showed that in a task in which rats learned the constantly-baited arms of an eight-arm radial maze over multiple days, the learning rate was slower when the SWRs were disrupted during the first hour of sleep following the training every day (**Figure 1.6**, Girardeau et al., 2009). Optogenetic silencing of SWRs during SWS disrupted learning for an object-location association task following sleep (Norimoto et al., 2018). However, the effect of disturbing SWRs during SWS on the stability of spatial representation is ambiguous. A study showed that the SWR disruption during SWS does not interfere with the formation of stable spatial representations (Kovacs et al., 2016). This finding was challenged by another study from the same lab where selective suppression of SWR-associated reactivations of a neuronal assembly specific to an environment impaired spatial memory performance and destabilized place fields (Gridchyn et al., 2020).

1.8 The spike content of sharp-wave ripples

Neurons within the hippocampus synchronously fire during SWRs (**Figure 1.7**). The neuronal population bursts (population burst events: PBEs) happen in the absence of sensorimotor cues or perceptual references to create place cells' activation patterns as observed during the behavior. The time scale of firings during the sharp-wave ripple can induce long-term potentiation, a

phenomenon associated with memory-related plasticity at the cellular level (Sadowski et al., 2016). We use SWRs and coinciding PBEs interchangeably throughout the document.




Figure 1.6: Disrupting sharp-wave ripples through electrical stimulation of the hippocampus leads to impairment in spatial memory. Left, a memory task in which rats were trained to find food rewards on an eight-arm radial maze in which the same three arms were baited every day. Right, the performance in three groups of animals; SWR suppression (red), and two control groups with either random delayed stimulation (blue) or lack of stimulation (black). Although all groups showed an increase in performance over days of learning, rats with SWR suppression took more days to perform above chance level (gray shading) (figure reproduced from Girardeau et al., 2009).

1.8.1 Awake replay supports memory storage and planning

During the quiet wake SWRs, behavioral sequences of place field activation re-occur in a time-compressed manner. These SWR-associated replay events represent the behavioral sequences in a forward (the order in which they were experienced) (Diba & Buzsáki, 2007; Pfeiffer & Foster, 2013) or reverse (Foster & Wilson, 2006) manner (**Figure 1.3**). While reverse replays are associated with maintaining the representation of previously taken paths in the environment (Ambrose et al., 2016; Foster & Wilson, 2006), the forward replays predict the immediate subsequent trajectory taken by the animal, supporting planning (Pfeiffer & Foster, 2013; Xu et al., 2019). A shift from the reverse replay-based retrospective evaluation to forward-replay-based planning happens as animals learn to correctly alternate in a W-maze task (Shin et al., 2019).

However, the forward replays' role in planning has been challenged recently by (Gillespie et al., 2021), which showed that the replay content is decoupled from the immediate subsequent choice of the animal, supporting that the awake replays only serve as a mechanism for consolidation and storage of memories and not planning.




Figure 1.7: Example of population burst events co-occurring with hippocampal sharp-wave ripples. On the top, hippocampal LFP filtered at ripple frequency band (100-250 Hz) and in the bottom, spike trains of simultaneously recorded units are shown. Example SWR-associated population burst events (PBEs) are highlighted in blue shades.

The replays do not simply reflect the current experience of the animal due to elevated excitability of the neurons. Instead, they tend to represent remote less-experienced sequences than current familiar sequences (Gupta et al., 2010; Karlsson & Frank, 2009). Moreover, they might represent the never-experienced trajectories, which could be imagined by the animal but restricted by the experimenter (Gupta et al., 2010).

1.8.2 Sleep reactivations and replays

Early rodent studies showed that place cells with overlapping place fields displayed an increased tendency to coactivate during sleep after exploration of an environment than the sleep period before the experience (**Figure 1.8**, Wilson & McNaughton, 1994). In a follow-up study,

Kudrimoti et al. showed that compared with pre-maze sleep, spike content in post-maze SWS, and not REM sleep, better explained the distribution of correlation between the maze firings of pairs of neurons (Kudrimoti et al., 1999). These coactivations that mainly occurred during sleep SWRs last for several hours after a novel experience (Giri et al., 2019). SWR-associated reactivation of assemblies that were formed during the exploration of an environment impacts the reinstatement of those assemblies when the animal is re-exposed to the same environment (van de Ven et al., 2016), suggesting that reactivation of assemblies during sleep is essential for their stabilization. Moreover, studies showed coordinated reactivation in the hippocampus and neocortex (Ji & Wilson, 2007), supporting that the SWS hippocampal reactivations distribute the memory traces across brain systems, facilitating their long-term storage.




Figure 1.8: Enhanced coactivation of neurons during the sleep following an experience. Individual neurons are represented as dots on the perimeter of each circle. Correlations between pairs of neurons' spike trains are shown with lines. Warmer colors correspond to higher correlations and the thick lines indicate cell pairs that were correlated during RUN and also correlated during either PRE or POST (figure reproduced from Wilson & McNaughton, 1994).

Sleep reactivations in the form of extended sequences, similar to awake replays, have been reported as well (Grosmark & Buzsáki, 2016; Lee & Wilson, 2002; Nádasdy et al., 1999).

However, a study showed that the sleep reactivation mainly reflect random sequences of place field activation (or random trajectories in the environments) than behavioral sequences (Stella et al., 2019).

Replay versus preplay. It is still unclear whether sequential replays are essential to memory processing. The level of sequential replay during post-maze sleep is significantly lower than during the behavior (Grosmark & Buzsáki, 2016). A series of studies reported that the sequences observed during a novel experience are pre-played during the preceding sleep (Dragoi & Tonegawa, 2011; Farooq et al., 2019; Grosmark & Buzsáki, 2016). The novel experience instead strengthens millisecond time scale coordination between neurons and adjusts neuronal participation within a consistent sequential framework (Farooq et al., 2019). Therefore, one hypothesis is that sequential replays are formed de novo upon a novel experience, and their formation depends on integrity of memory-elicited plasticity via the activation of NMDA receptors, as supported by Silva et al., 2015. An alternative hypothesis, supported by preplay studies, is that the replays are just epiphenomena reflecting solely the preexisting hardwired connectivity within the hippocampal network, casting doubt on their role in memory processing.

1.9 Project goal: temporal patterns of firing during the sharp-wave ripples

Although interferential studies have indicated the necessity of SWRs in spatial memories, it is not well known what patterns of neural activity occur during these events during quiet wake and sleep. The structure of neural activity during the SWRs is critical to understand neural correlates of memory processing.

Methods have been developed to analyze the activity during the SWRs. These methods mainly rely on calculating the correlation between the place cell activations during SWRs and those observed during the behavior. The animal's position during the exploration can be accurately estimated from the spike trains of a population of place cells (Wilson & McNaughton, 1993). Accordingly, in a standard method based on Bayesian decoding, the animal's position is decoded in successive snapshots of population firings within each SWR event, given the place fields

measured during the behavior. Then, the replay content of each SWR is assessed by calculating the extent to which the virtual trajectory through the decoded positions is similar to the animal's actual trajectory on the track.

Although Bayesian decoding has been proven successful in interpreting the activity patterns during the quiet wake SWRs, the same might not apply to the neural activity during the sleep SWRs. The main reason is that the decoding approach assumes that the spike patterns during the SWRs take trajectories through a state space with states limited to the explored maze locations. This assumption brings limitations, as remote and recent experiences are represented together during sleep, and each transforms over the course of sleep. Therefore, the ensembles during sleep might be quantitatively and qualitatively different from those on the maze, and examining the SWR activity through the lens of maze neural activity could be faulty. Hence, the standard Bayesian decoding methods are unsuitable for detecting temporally rich and dynamic representations during sleep.

In this project, our goal was to provide new insights into the temporal coordination between the firings of place cells and their contribution to spatial memory processes. To this end, we developed two analysis methods that will be explained in chapters 2 and 3 of this dissertation.

1.9.1 Hidden Markov models to detect SWRs spatiotemporal structures

We first developed an unsupervised method based on hidden Markov models (HMMs) that without reference to positions recorded during the exploration, identifies the sequential activities within SWRs in terms of distinct ensembles of neurons and their activation sequences. We validated this method by analyzing the quiet wake SWRs and showed that an HMM-based method can detect sequences pertaining to the current and remote experiences of the animal. We show that we can use such a method to analyze rich and dynamic neural activities during sleep.

Hidden Markov models. Dynamic systems can be described in terms of distinct states and the probability of transitions between each pair of states. A Markovian assumption asserts that the future state at each point depends only on the system's current state and is independent of the prior history of transitions. An example Markov chain model that consisted of two states, A and B, is shown in **Figure 1.9**.



Figure 1.9: A Markov chain model. Top left, a model with two states A and B and inter-state transitions that are demonstrated by arrows. The model at each successive time point can transition to another state or stay in the same state (self-transition). Top right, the corresponding transition probability matrix consisting of between-state transition probabilities on the off-diagonal (lighter) elements and self-transitions on the diagonal (darker). Bottom, an example sequence of states generated by the model.

Patterns of firings within hippocampal neuronal networks can be modeled using Markov models (**Figure 1.10**). The only difference with the system in **Figure 1.9** is that the states are assumed to be unknown, hence latent or hidden. Forward-backward algorithms have been designed to estimate the individual states in terms of each neuron's participation probability and calculate the transition probabilities between the different states (Rabiner, 1989).

1.9.2 Bayesian learning of spatial tunings during sleep

The second approach is based on the principles of Bayesian learning to update the neurons' place fields based on the sleep ensemble firings. We assume that the maze place fields do not remain consistent during sleep, but they transform, and this transformation is evident from the ensemble with which they cofire during sleep. This methodology enabled us to track spatial tunings of

single neurons moment by moment during sleep following a novel experience. In contrast to sleep prior to exploration, we found that a group of neurons maintain consistent spatial tunings during early but not late SWS. Furthermore, the place fidelity of sleep spatial tunings across neurons was predictive of the neurons' place field stability upon delayed exposure to the same environment.

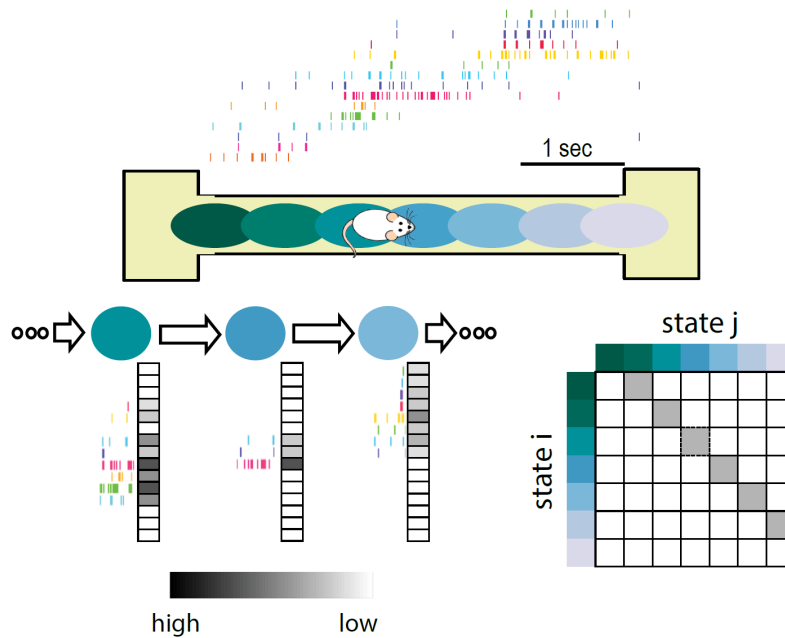


Figure 1.10. A hidden Markov model to model temporal patterns of firing on a maze. Top, place cells sequential activation when a rat runs between two ends of a linear track. Overlaid ellipses on the track show distinct states corresponding to different track locations. Bottom left, the temporal evolution of the system between three example states. Each state is characterized using a unique distribution of participation probabilities across neurons (observation probabilities). Bottom right, transition probability matrix. The off-diagonal elements show the successive transitions between nearby states (locations). While the states here are mapped onto different track locations, in general, there are no assumptions about the task variable that is represented by the states.

2 Uncovering temporal structure in hippocampal output patterns

(Presented from Maboudi et al., 2018)

Summary

Place cell activity of hippocampal pyramidal cells has been described as the cognitive substrate of spatial memory. Replay is observed during hippocampal sharp-wave-ripple-associated population burst events (PBEs) and is critical for consolidation and recall-guided behaviors. PBE activity has historically been analyzed as a phenomenon subordinate to the place code. Here, we use hidden Markov models to study PBEs observed in rats during exploration of both linear mazes and open fields. We demonstrate that estimated models are consistent with a spatial map of the environment, and can even decode animals' positions during behavior. Moreover, we demonstrate the model can be used to identify hippocampal replay without recourse to the place code, using only PBE model congruence. These results suggest that downstream regions may rely on PBEs to provide a substrate for memory. Additionally, by forming models independent of animal behavior, we lay the groundwork for studies of non-spatial memory.

2.1 Introduction

Large populations of neurons fire in tandem during hippocampal sharp-waves and their accompanying CA1 layer ripple oscillations (Buzsaki, 1986). By now, multiple studies have shown that during many sharp-wave ripple-associated PBEs, hippocampal 'place cells' (John

O'Keefe, 1976) in compressed sequences that reflect the firing order determined by the sequential locations of their individual place fields (Diba & Buzsáki, 2007; Foster & Wilson, 2006; Lee & Wilson, 2002; Nádasdy et al., 1999). While the firing patterns during active exploration are considered to represent the brain's global positioning system and provide a substrate for spatial and episodic memory, instead it is the synchronized activity during PBEs that is most likely to affect cortical activity beyond the hippocampus (G. Buzsaki, 1989; Carr et al., 2011; Diekelmann & Born, 2010). Likewise, widespread activity modulation is seen throughout the brain following these sharp-wave ripple population bursts (Logothetis et al., 2012).

The literature on PBEs has largely focused on developing templates of firing patterns during active behavior and evaluating the extent to which these templates' patterns are reprised during subsequent PBE. But what if the fundamental mode of the hippocampus is not the re-expression of place fields, but rather the PBE sequences during SWR? PBE sequences are enhanced during exploration of novel environments (Sen Cheng & Frank, 2008; Foster & Wilson, 2006), they presage learning-related changes in place fields (Dupret et al., 2010), and appear to be critical to task learning (Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Jadhav et al., 2012). Here, we examine the information provided by CA1 and CA3 pyramidal neurons, the output nodes of the hippocampus, through the looking glass of PBE firing patterns.

We developed a technique to build models of PBE sequences strictly outside of active exploration and independent of place fields and demonstrate that this nevertheless allows us to uncover spatial maps. Furthermore, these models can be used to detect congruent events that are consistent with replay but without any explicit place cell template. Our technique therefore

provides new possibilities for evaluating hippocampal output patterns in single-trial and other fast learning paradigms, where a reliable sequential template pattern is not readily available. Overall, our work suggests that a sequence-first approach can provide an alternative view of hippocampal activity that may shed new light on how memories are formed, stored, and recalled.

2.2 Results

2.2.1 Awake population burst events

We began by analyzing the activity of large numbers of individual neurons in areas CA1 and CA3 of the dorsal hippocampus as rats navigated linear mazes for water reward (linear track: $n=3$ rats, $m=18$ sessions; previously used by (Diba & Buzsáki, 2007)). Using pooled multiunit activity, we detected PBEs during which many neurons were simultaneously active. The majority of these events occurred when animals paused running (speed <5 cm/s, corresponding to $54.0\% \pm 20.1\%$ sd of events) to obtain reward, groom, or survey their surroundings, and were accompanied by SWR complexes, distinguished by a burst of oscillatory activity in the 150–250 Hz band of the CA1 LFP. Because we are interested in understanding internally generated activity during PBEs, we included only these periods without active behavior, ensuring that theta sequences would not bias our results. While we identified active behavior using a speed criterion, we found similar results when we instead used a theta-state detection approach (not shown). We did not add any other restrictions on behavior, LFP, or the participation of place cells. We found that inactive PBEs occupied an average of 1.8% of the periods during which animals were on the linear track (16.9 ± 15.1 s of 832.6 ± 390.5 s). In comparison, classical Bayesian approaches to understand PBE activity require the 34.8% of time

animals are running (speed >10 cm/s) on the track (254.4 ± 106.6 s of 832.6 ± 390.5 s) to build models of place fields.

2.2.2 Learning hidden Markov models from PBE data

Activity during PBEs is widely understood to be internally-generated in the hippocampal-entorhinal formation, and likely to affect neuronal firing in downstream regions (Buzsaki, 1989; Logothetis et al., 2012; Yamamoto & Tonegawa, 2017). Given the prevalence of PBEs during an animal’s early experience, we hypothesized that the neural activity during these events would be sufficient to train a machine learning model of sequential patterns—a hidden Markov model—and that this model would capture the relevant spatial information encoded in the hippocampus independent of exploration itself.

Hidden Markov models have been very fruitfully used to understand sequentially structured data in a variety of contexts. A hidden Markov model captures information about data in two ways. First, it clusters observations into groups (‘states’) with shared patterns. In our case, this corresponds to finding time bins in which the same sets of neurons are co-active. This is equivalent to reducing the dimension of the ensemble observations into a discretized latent space or manifold. Second, it models the dynamics of state transitions. This model is Markovian because it is assumed that the probability to transition to the next state only depends on the current state. Critically, these operations of clustering and sequence modeling are jointly optimized, allowing the structure of ensemble firing corresponding to each of the final states to combine information over many observations. Given the role of the hippocampus in memory, in our HMMs, the unobserved latent variable presumably corresponds to the temporal evolution of a memory trace that is represented by co-active ensembles of CA1 and CA3 neurons. The full

model will correspond to the structure which connects all the memory traces activated during PBEs.

The parameters of our model that are fit to data include the observation model (the cluster descriptions, or predicted activity of each excitatory neuron within the CA1/CA3 ensemble for a given state), the state transition model (the probability that the CA1/CA3 ensemble will transition from a start state to a destination state in the next time bin), and the initial state distribution (the probability for sequences to start in each given state). In prior work using HMMs to model neural activity, a variety of statistical distributions have been used to characterize ensemble firing during a specific state (the observation model). We opted for the Poisson distribution to minimize the number of parameters per state and per neuron (see Materials and methods). We used the standard iterative EM algorithm (Rabiner, 1989) to learn the parameters of an HMM from binned PBE data (20 ms bins). **Figure 2.1** depicts the resultant state transition matrix and observation model for an example linear-track session.

Using separate training- and test-datasets (cross-validation) mitigates over-fitting to training data, but it is still possible for the cross-validated goodness-of-fit to increase with training without any underlying dynamics, e.g., if groups of neurons tend to activate in a correlated fashion. Does the model we have learned reflect underlying sequential structure of memory traces beyond pairwise co-firing? To answer this question, we cross-validated the model against both real ‘test’ data and against surrogate ‘test’ data derived from shuffling each PBE in two ways: one in which the binned spiking activity was circularly permuted across time for each neuron independently of the other neurons (‘temporal shuffle’, which removes co-activation), and one in which the order of the binned data was scrambled coherently across all neurons

(‘time-swap’, which maintains co-activation). Note that the second shuffle preserves pairwise correlations while removing the order of any sequential patterns that might be present. Using five-fold cross-validation, we compared learned models against both actual and surrogate test data and found that the model likelihood was significantly greater for real data (vs. temporal shuffle, $p < 0.001$, vs. time-swap, $p < 0.001$, $n = 18$ sessions, Wilcoxon signed-rank test).

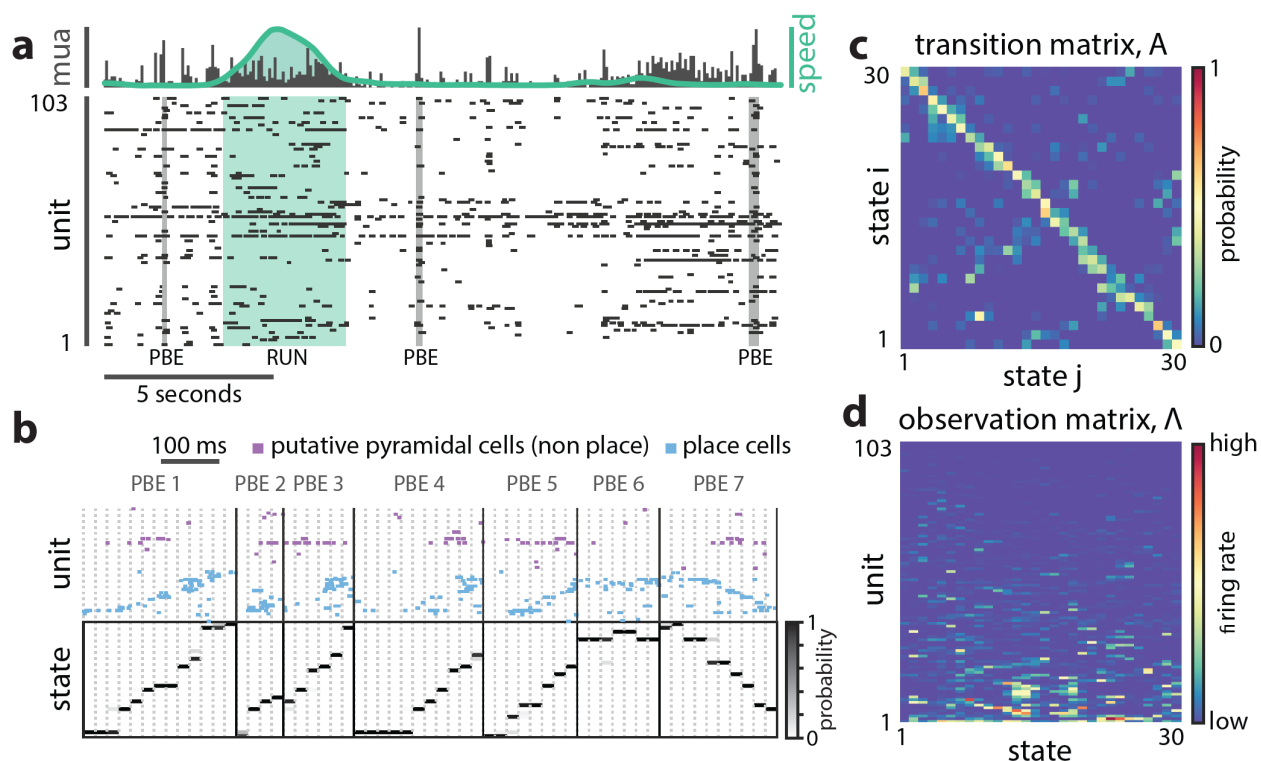


Figure 2.1: A hidden Markov model of ensemble activity during PBEs. A hidden Markov model of ensemble activity during PBEs. (a) Examples of three PBEs and a run epoch. (b) Spikes during seven example PBEs (top) and their associated (30 state HMM-decoded) latent space distributions (bottom). The place cells are ordered by their place fields on the track, whereas the non-place cells are unordered. The latent states are ordered according to the peak densities of the lsPFs (see Materials and methods). (c) The transition matrix models the dynamics of the unobserved internally-generated state. The sparsity and banded-diagonal shape are suggestive of sequential dynamics. (d) The observation model of our HMM is a set of Poisson probability distributions (one for each neuron) for each hidden state. Looking across columns (states), the mean firing rate is typically elevated for only a few of the neurons and individual neurons have elevated firing rates for only a few states.

2.2.3 What do the learned model parameters tell us about PBEs?

To begin to understand what structure we learn from PBE activity, we compared our HMMs (trained on real data) against models trained on multiple different surrogate datasets (**Figure 2.2a,b**). These surrogate datasets were obtained from actual data following: (1) temporal shuffles and (2) time-swaps, as above, and (3) by producing a surrogate PBE from independent Poisson simulations according to each unit's mean firing rate within the original PBEs. First, we investigated the sparsity of the transition matrices using the Gini coefficient (see Materials and methods). A higher Gini coefficient corresponds to higher sparsity. Strikingly, the actual data yielded models in which the state transition matrix was sparser than in each of the surrogate counterparts ($p < 0.001$, **Figure 2.2c**), reflecting that each state transitions only to a few other states. Thus, intricate yet reliable details are captured by the HMMs. Next, we quantified the sparsity of the observation model. We found that actual data yielded mean firing rates which were highly sparse (**Figure 2.2d**), indicating that individual neurons were likely to be active during only a small fraction of the states. Using a graph search algorithm (see Materials and methods), we simulated paths through state space generated by these transition matrices, and found that this increased sparsity accompanied longer trajectories through the state space of the model. Thus, the state transition matrices we learn are suggestive of dynamics in which each sparse state is preceded and followed by only a few other, in turn, sparse states, providing long sequential paths through state space-consistent with spatial relationships in the environment in which the animal was behaving, but generated from PBEs. The increased sparsity of the observation model and transition matrix in the example session was representative of a significant increase over all remaining sessions ($p < 0.05$, $n = 18$ sessions, Wilcoxon signed-rank tests, **Figure 2.2e,f**).

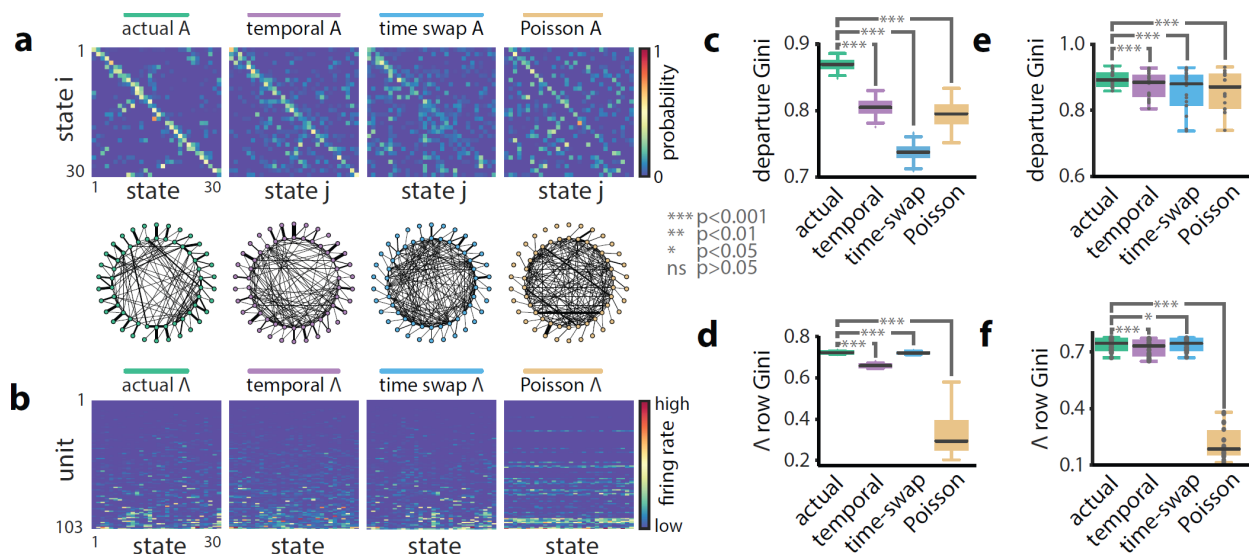


Figure 2.2: A hidden Markov model of ensemble activity during population burst events. Models of PBE activity are sparse. We trained HMMs on neural activity during PBEs (in 20 ms bins), as well as on surrogate transformations of those PBEs. **(a)** (top) The transition matrices for the actual and surrogate PBE models with states ordered to maximize the transition probability from state i to state $i+1$. (bottom) Undirected connectivity graphs corresponding to the transition matrices. The nodes correspond to states (progressing clockwise, starting at the top). The weights of the edges are proportional to the transition probabilities between the nodes (states). The transition probabilities from state i to every other state except $i+1$ are shown in the interior of the graph, whereas for clarity, transition probabilities from state i to itself, as well as to neighboring state $i+1$ are shown between the inner and outer rings of nodes (the nodes on the inner and outer rings represent the same states). **(b)** The observation matrices for actual and surrogate PBE models show the mean firing rate for neurons in each state. For visualization, neurons are ordered by their firing rates. **(c)** We quantified the sparsity of transitions from one state to all other states using the Gini coefficient of rows of the transition matrix for the example session in (a). Actual data yielded sparser transition matrices than shuffles. **(d)** The observation models—each neuron’s expected activity for each state—learned from actual data for the example session are significantly sparser than those learned after shuffling. This implies that as the hippocampus evolves through the learned latent space, each neuron is active during only a few states. **(e)** Summary of transition matrix sparsity and **(f)** observation model sparsity with corresponding shuffle data pooled over all sessions/animals. (***: $p < 0.001$, *: $p < 0.05$; single session comparisons: $n=250$ realizations, Welch’s t-test; aggregated comparisons - $n=18$ sessions, Wilcoxon signed-rank test).

These observations indicate that PBEs inform an HMM about extant spatial relationships within the environment. So, next we asked how the firing patterns of neurons during actual behavior project into the learned latent spaces. To observe the evolution of the latent states during behavior, we used our model to determine the most likely sequence of latent states corresponding to decode the neural activity observed in 100 ms bins during epochs that displayed strong theta oscillations (exclusive of PBEs) when rats were running (speed >10 cm/s; see Materials and methods). If the learned model was distinct from ensemble patterns during behavior, we might

expect the resulting state space probability distributions at each point in time to be randomly spread among multiple states. Instead, we found distributions that resembled sequential trajectories through the latent space (**Figure 2.3a**) in parallel with the physical trajectories made by the animal along the track, further demonstrating that the latent state dynamics learned from PBEs corresponds to an internalized model of physical space.

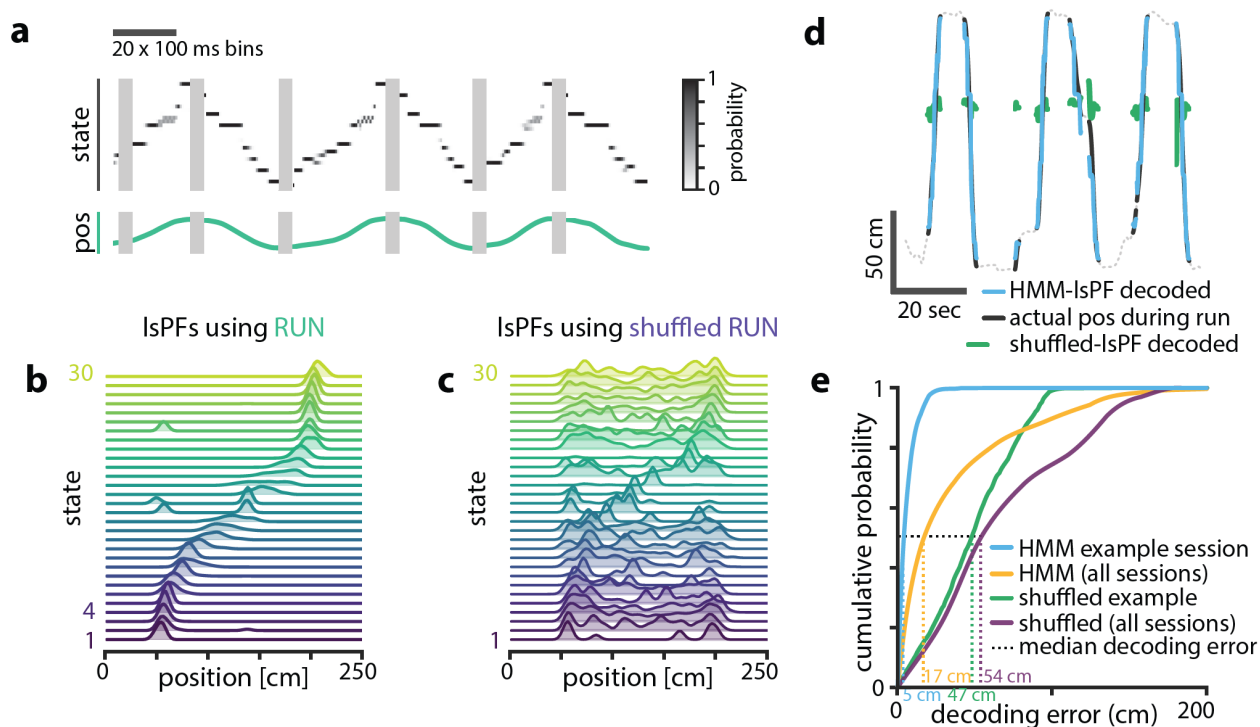


Figure 2.3: Latent states capture positional code. (a) Using the model parameters estimated from PBEs, we decoded latent state probabilities from neural activity during periods when the animal was running. An example shows the trajectory of the decoded latent state probabilities during six runs across the track. (b) Mapping latent state probabilities to associated animal positions yields latent-state place fields (IsPFs) which describe the probability of each state for positions along the track. (c) Shuffling the position associations yields uninformative state mappings. (d) For an example session, position decoding during run periods through the latent space gives significantly better accuracy than decoding using the shuffled tuning curves. The dotted line shows the animal's position during intervening non run periods. (e) The distribution of position decoding accuracy over all sessions ($n=18$) was significantly greater than chance. ($p<0.001$).

To better understand the relationship between the latent space and physical space, we used the latent state trajectories decoded during running to form an estimate of the likelihood of each state as a function of location on the track (see Materials and methods). These latent state place fields

(lsPFs, **Figure 2.3b**) in many ways resembled neuronal place fields and similarly tiled the extent of the track. This spatial localization went away when we re-estimated the lsPFs with shuffled positions (**Figure 2.3c**). To quantify how informative the latent states were about position, we used the lsPFs to map decoded state sequences to position during running periods (**Figure 2.3d**). In our example session, decoding through the latent space resulted in a median accuracy of 5 cm, significantly greater than the 47 cm obtained from shuffled lsPFs ($p < 0.001$, Wilcoxon signed-rank test, **Figure 2.3d**). When we evaluated decoding error over our full set of sessions, we observed a similar result ($p < 0.001$, Wilcoxon signed-rank test, **Figure 2.3e**). As our method required discretizing the state space, a potential caveat is that the number of latent states is a relevant parameter, which we arbitrarily chose to be 30. However, latent-state place fields were informative of position over a wide range of values of this parameter. Note that decoding into the latent space and then mapping to position resulted in slightly higher error than simply performing Bayesian decoding on the neural activity during behavior. This suggests that the latent space we learn from PBEs may not capture all the information about space that is present in hippocampal activity during behavior, though this may also reflect the limited number of PBEs from which we can learn.

2.2.4 HMM-congruent PBEs capture sequence replay

We and others have previously described how the pattern of place cell firing during many PBEs recapitulates the order in which they are active when animals run on the track (**Figure 2.4a**). We employed the versatile and widely-used Bayesian decoding method to ascribe a replay score to sequential patterns during PBEs. Briefly, for each PBE, we used place-field maps to estimate a spatial trajectory (an *a posteriori* distribution of positions) in 20 ms bins. We generated surrogate

data via a column-cycle shuffle (i.e., a circular shift across positions for each time bin (Davidson et al., 2009)) of the *a posteriori* distributions during PBEs. The real and surrogate trajectories were scored (see Materials and methods), and we defined replay events as those for which the score of the actual trajectory was larger than a threshold fraction of the null distribution generated by the surrogate scores. Using this approach, we found that 57% of PBEs (1064 of 1883) were identified as replay beyond a threshold of 99% (median across datasets 54.2%, interquartile range=32.8–61.0%). Thus, as has been reported many times (Davidson et al., 2009; Diba & Buzsáki, 2007; Foster & Wilson, 2006; Karlsson & Frank, 2009), only a fraction of PBEs (but many more than expected by chance) represent statistically significant replay. Given that we use all PBEs for model learning and our models capture the structure of the environment and the patterns expressed by place cells during exploration, we were interested in understanding whether we could also use our latent-space models to find these replay events. Indeed, for many events when we decode trajectories through state space, they resemble the sequential patterns observed when we decode position using Bayesian techniques and the place cell map (**Figure 2.4a**, left). However, given previous evidence for replay of environments not recently experienced (Gupta et al., 2010; Karlsson & Frank, 2009), we hypothesized that some PBEs might contain ensemble neural activity which is unstructured and thus unrelated to the learned model, and that these would correspond to the ‘non-replay’ events found using traditional methods.

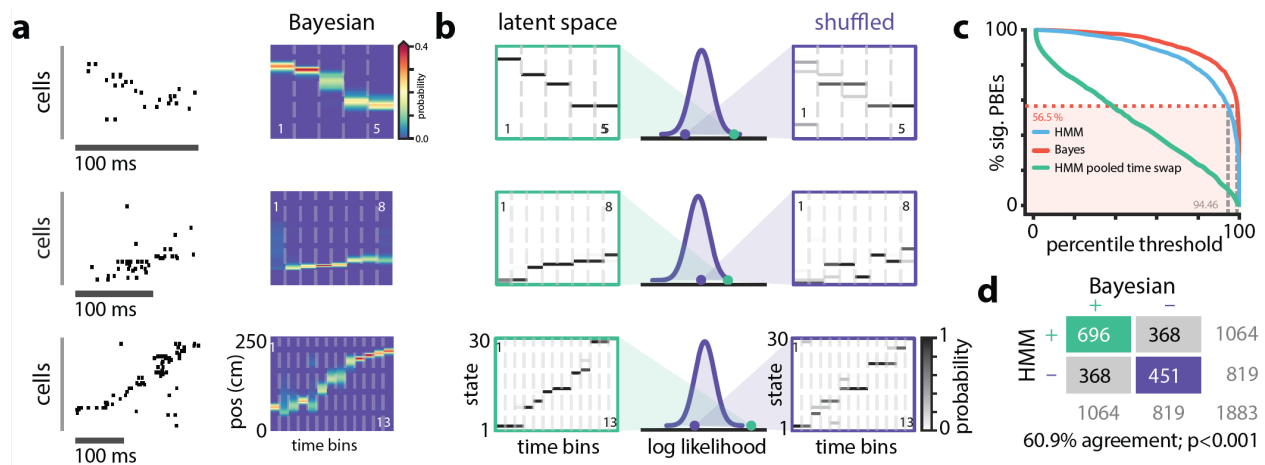


Figure 2.4: Replay events can be detected via HMM congruence. (a) Example PBEs decoded to position using Bayesian decoding. (b) (left) Same examples decoded to the latent space using the learned HMM. (right) Examples decoded after shuffling the transition matrix, and (middle) the sequence likelihood using actual and shuffled models. (c) Effect of significance threshold on the fraction of events identified as replay using Bayesian decoding and model congruent events using the HMM approach. (d) Comparing Bayesian and model-congruence approaches for all PBEs recorded, we find statistically significant agreement in event identification (60.9% agreement, $n=1883$ events from 18 sessions, $p < 0.001$, Fisher's exact test two sided).

To assess how well the pattern of ensemble activity during individual PBEs related to the overall state-space model learned from PBE activity ('congruence'), we developed a statistical approach for identifying the subset of strongly structured PBEs. Specifically, rather than comparing real and surrogate PBEs, we compared the goodness-of-fit for each event to a null distribution generated via a computationally-efficient manipulation of the transition matrix of the model (Figure 2.4b); we row-wise shuffled the non-diagonal elements of the transition matrix to assess whether an individual PBE is a more ordered sequence through state space than would be expected by chance. Maintaining the diagonal avoids identifying as different from chance sequences which consist of few repeated states, marked by transitions between state i and itself. As described above, the fraction of events identified as replay using Bayesian decoding is strongly tied to how the null-distribution is generated (i.e., what shuffle is used), some secondary criteria (e.g., number of active cells, unit cluster quality, peak firing rate, trajectory 'jumps', etc.), and the value of the significance threshold arbitrarily chosen to be 90%, 95%, or 99% of

shuffles in different reports. When we combined across datasets, we found that our transition matrix shuffle yielded a null distribution for which a 99% confidence interval identified slightly fewer PBEs as significant than the column-cycle shuffle did for Bayesian decoding (**Figure 2.4c**). To make a principled comparison of Bayesian- and HMM-based replay detection schemes, we fixed the Bayesian-based significance threshold at 99% but selected the significance threshold for the HMM-congruence null distribution so that the fraction of replay events detected would be the same between the two schemes. Following this approach, we found that model-congruent/incongruent PBEs largely overlapped with the replay/non-replay events detected using Bayesian decoding of the place cell map (**Figure 2.4d**). Thus, using only the neural activity during PBEs, without access to any place cell activity, we are remarkably able to detect the sequential patterns typically described as ‘replay’ based only on their consistency with the structure of other PBE activity.

There were, however, also differences between the Bayesian and HMM-congruent approaches, including events that reached significance in one but not the other formalism. We wanted to understand where and why these approaches differed in identifying significant sequences. When we examined individual PBEs, we found sequences for which both Bayesian and model-congruence replay detection approaches appeared to malfunction (**Figure 2.5a**). This was not a failure of the choice of significance threshold, as for both techniques we found what appeared to be false-negatives (patterns which looked like replay but were not significant) as well as false-positives (patterns which looked noisy but were identified as replay). Thus, in order to quantitatively compare the two approaches, we asked eight humans to visually examine all the PBEs in our database. They were instructed to label as replay PBEs in which the animal’s

Bayesian decoded position translated sequentially without big jumps ((Silva et al., 2015); see Materials and methods).

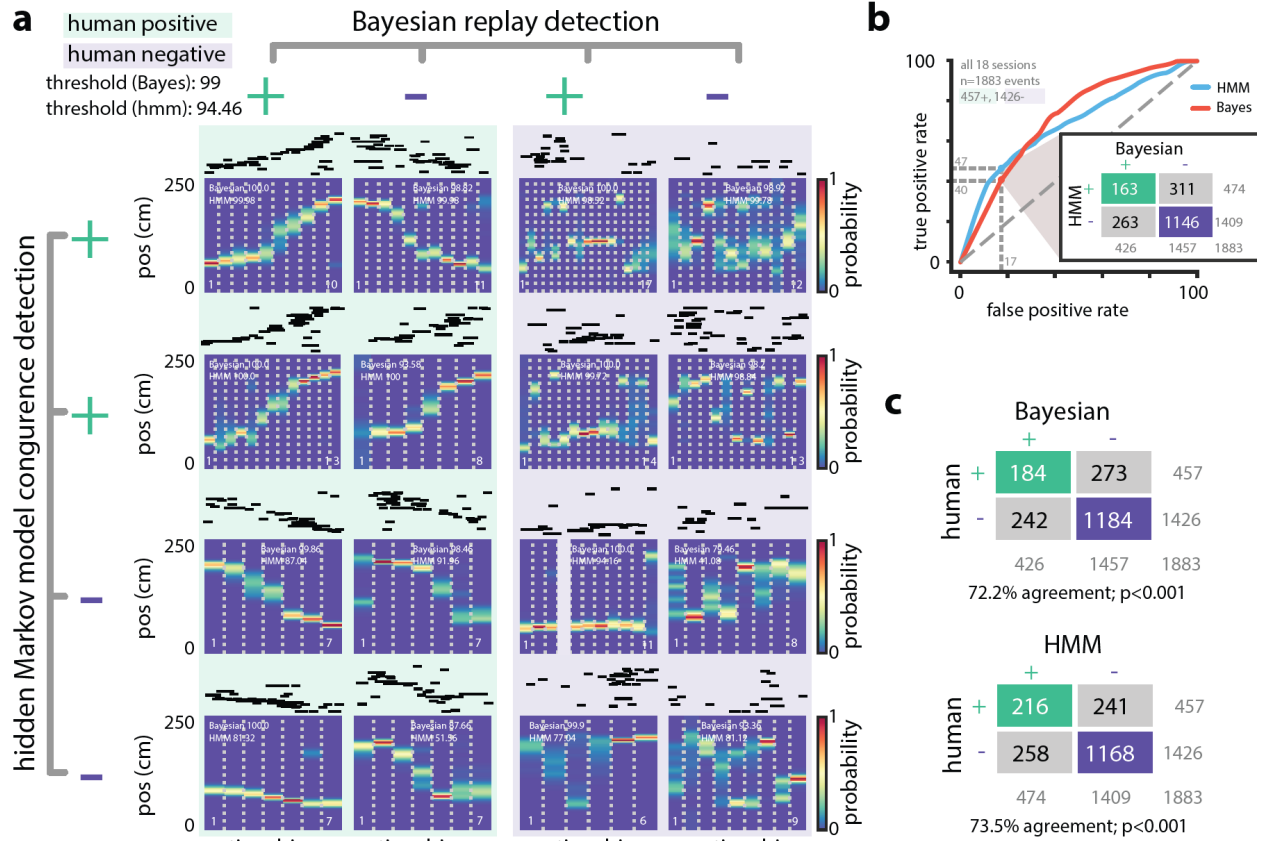


Figure 2.5: Comparing HMM congruence and Bayesian decoding in replay detection. (a) Eight examples from one session show that Bayesian decoding and HMM model-congruence can differ in labeling of significant replay events. For each event, spike rasters (ordered by the location of each neuron's place field) and the Bayesian decoded trajectory are shown. '+' ('-') label corresponds to significant (insignificant) events. (left) Both methods can fail to label events that appear to be sequential as replay and (right) label events replay that appear non-sequential. (b) We recruited human scorers to visually inspect Bayesian decoded spike trains and identify putative sequential replay events. Using their identifications as labels, we can define an ROC curve for both Bayesian and HMM model-congruence which shows how detection performance changes as the significance threshold is varied. (inset) Human scorers identify 24% of PBEs as replay. Setting thresholds to match this value results in agreement of 70% between Bayesian and HMM model-congruence. (c) Using the same thresholds, we find $\approx 70\%$ agreement between algorithmic and human replay identification. (All comparison matrices, $p < 0.001$, Fisher's exact test two-tailed.).

We marked each event as a 'true' community replay if it was identified by a majority of scorers (six individuals scored $n=1883$ events, two individuals scored a subset of $n=1423$ events). We calculated an ROC curve which compared the rate of true positive and false positive detections

as the significance thresholds for Bayesian and model-congruence approaches were varied (**Figure 2.5b**). A perfect detector would have an AUC of unity. We did not find a significant difference between the AUC of Bayesian decoding and model-congruence ($p=0.14$, bootstrap, see Materials and methods). If we select thresholds such that our algorithms yield a similar fraction of significant vs. total events as the 24% denoted by our human scorers, we find that both Bayesian and model-congruence yield agreement of $\approx 70\%$ labeled events with each other and with human scorers (**Figure 2.5c**).

Thus, congruence with an HMM trained only on PBEs appears to work as reliably as Bayesian decoding in detecting sequential reactivation of linear track behaviors. However, when we examined individual sessions, we noticed that performance was quite variable. Given that our models are learned only from PBEs, we reasoned that the statistics or structure of the PBEs within each session might yield models which vary in quality depending on the number of recorded units, the number of PBEs detected, and their self-consistency across events. We created a model quality metric by comparing cross-validated learning statistics to models which were learned from shuffled events (see Materials and methods). We found that the performance of model-congruence detection was tied to model quality ($R^2=0.17$, $F=2.9$, $n=18$ sessions). Model quality, in turn, was highly correlated with the number of PBEs during the session ($R^2=0.96$, $F=392.6$, $n=18$ sessions). Not surprisingly, the performance of Bayesian decoding relative to human scorers was independent of the quality of the HMM, or the number of PBEs, as the place field model is learned from ensemble neural activity during running. Thus, we find an intriguing contrast—when there is an abundance of PBEs (indicating novelty, learning, hippocampus-dependent planning, etc.), even in the absence of repeated experience, replay detection based on PBE activity is highly effective. Conversely, when there are few PBEs (i.e.,

scenarios in which PBEs are uncorrelated with cognitive function), but an abundance of repeated behavioral trials, Bayesian decoding of these limited events proves more effective.

2.2.5 Modeling internally generated activity during open field behavior

The linear track environment represents a highly-constrained behavior. We therefore asked whether the HMM approach could generalize to more complex environments and behavioral tasks. (Pfeiffer & Foster, 2013, 2015) had previously recorded activity of CA1 neurons in rats as they explored in a 2 m × 2 m open field arena for liquid reward. Briefly, animals were trained to discover which one of 36 liquid reward wells would be the ‘home’ well on a given day. Then, they were required to alternate between searching for a randomly rewarded well and returning to the home well. Using the place cell map in this task and Bayesian decoding, many PBEs were decoded to trajectories through two-dimensional space that were predictive of behavior and shaped by reward. Using this same dataset, we trained HMMs on neural activity during PBEs in the open field. Here, we used the same PBE detected previously (Pfeiffer & Foster, 2013, 2015) which occupied an average of $2.53 \pm 0.42\%$ of the periods during which animals were behaving (77.91 ± 21.16 s out of 3064.86 ± 540.26 s). Given the large number of units available in this dataset and the increased behavioral variability in the open field environment compared to the linear track, we chose to estimate HMMs with 50 latent states. The transition matrix and observation model from a sample session are shown in **Figure 2.6a,b**. Despite the complex and varied trajectories displayed by animals, the HMM captured sequential dynamics in PBE activity, as in the 1D case, when we compared learned models against both actual and surrogate test data, we found that the model likelihood was significantly greater for real data ($p < 0.001$, Wilcoxon signed-rank test).

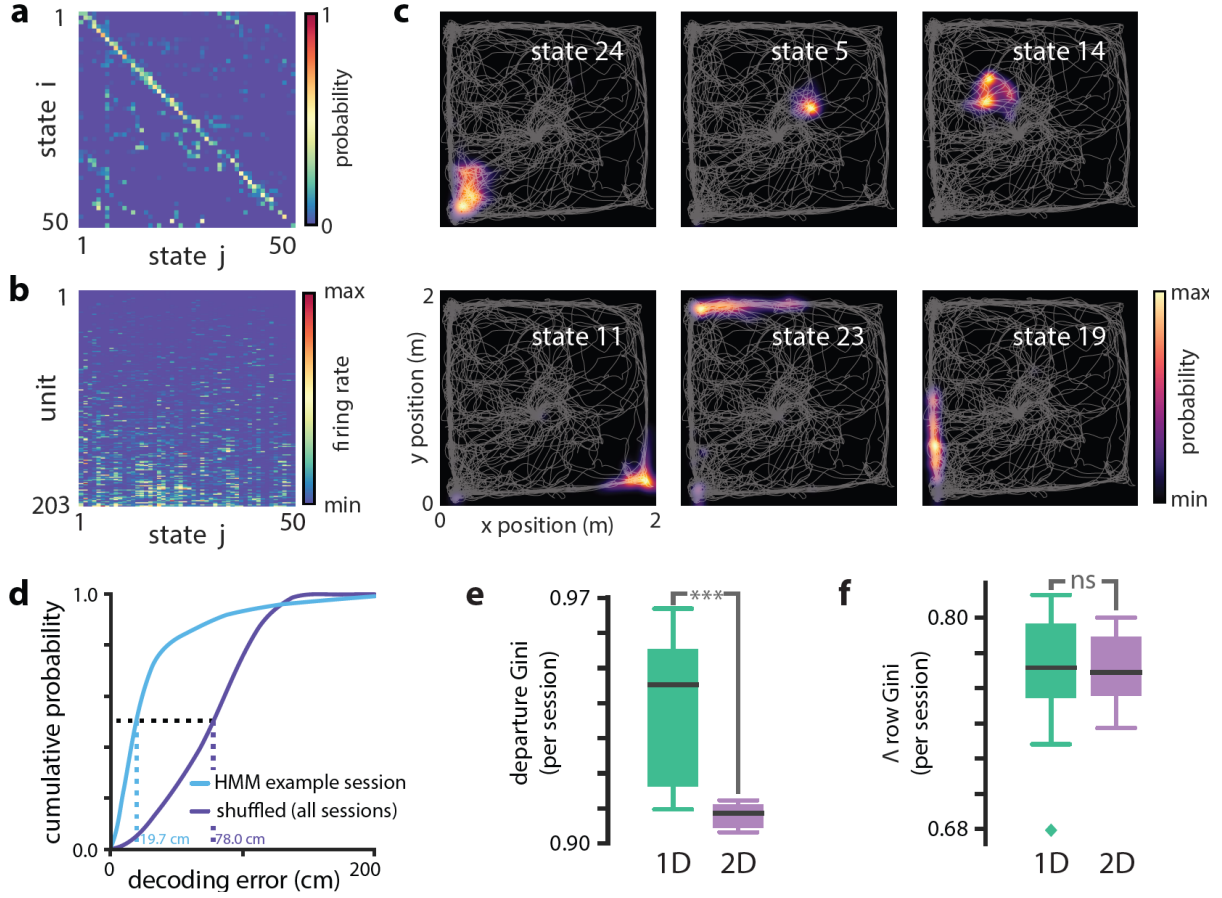


Figure 2.6. Modeling PBEs in open field. (a) The transition matrix estimated from activity detected during PBEs in an example session in the open field. (b) The corresponding observation model (203 neurons) shows sparsity similar to the linear track. (c) Example latent state place fields show spatially-limited elevated activity in two dimensions. (d) For an example session, position decoding through the latent space gives significantly better accuracy than decoding using the shuffled latent state place fields. (e) Comparing the sparsity of the transition matrices (mean Gini coefficient of the departure probabilities) between the linear track and open field reveals that, as expected, over the sessions we observed, the open field is significantly less sparse ($p < 0.001$), since the environment is less constrained. (f) In contrast, there is not a significant difference between the sparsity of the observation model (mean Gini coefficient of the rows) between the linear track and the open field. Note that the linear track models are sparser than in Figure 2.2 due to using 50 states rather than 30 to match the open field.

In the case of the linear track, we linked sparsity of the transition matrix to the sequential nature of behaviors in that environment. An unconstrained, two-dimensional environment permits a much richer repertoire of behavioral trajectories. However, behavior is still constrained by the structure of space—arbitrary teleportation from one location to another is impossible. We found that learning from PBEs in the open field yielded transition matrices (Figure 2.6a) that were significantly sparser than models learned from shuffled data ($p < 0.05$, Wilcoxon signed-rank test,

n=8 sessions). However, consistent with increased freedom of potential behaviors, when we compared the sparsity of models learned from open field PBEs with 50-state models learned from PBEs in linear tracks, the open field transition matrices were less sparse ($p < 0.001$, Mann–Whitney U test comparing 8 and 18 sessions, **Figure 2.6e**). Likewise, when we examined the observation model for the open field, we found that the activity across states for individual neurons was significantly more sparse than in models learned from shuffled data ($p < 0.05$, Wilcoxon signed-rank test, n=8 sessions). However, the sparsity of linear track and open field observation models were not significantly different ($p = 0.44$, Mann–Whitney U test).

Do the latent states learned from PBEs capture spatial information in a 2D environment? We used the PBE-trained model to decode run data, as in the linear track case. We found that the latent states corresponded with specific locations in the open field, as we expected (**Figure 2.6c**). Moreover, we were able to decode animals' movements with significantly greater than chance accuracy by converting decoded latent states to positions using the lsPF ($p < 0.001$, **Figure 2.6d**). Finally, we examined model-congruency for PBEs detected in the open field. Previously, it was reported that 27.3% (815 of 2980, n=8 sessions) were identified as 'trajectory events' (Pfeiffer & Foster, 2015). We chose a significance threshold to match this fraction and found that there was significant overlap between the events detected through Bayesian and model-congruence techniques ($p < 0.01$, Fisher's exact test). Thus, an HMM of the activity during population bursts captures the structure of neural activity in two dimensional environments during complex tasks and can be used to decode events consistent with trajectories through that environment.

2.2.6 Extra-spatial information

As described earlier, while we observed a similar fraction of events to be similar by HMM-congruence and Bayesian decoding, there was not an exact event-to-event correspondence. An intriguing potential explanation is that the latent space represented in PBE sequential firing and captured by the HMM is richer than simply the spatial structure of the present environment. In most hippocampal ensemble recording experiments, maze or open field tasks are structured to intentionally map memory elements to spatial behavior, and thus this potential richness is difficult to test. We used two sample datasets to explore the potential of the HMM to capture extra-spatial richness in the PBE sequences.

First, we considered the possibility that in the awake behaving animal, PBE activity might be sequential reactivation of environments other than the one being explored ('remote replay'). We reasoned that we could enhance the model's representation of remote environments by filtering out local replay from the training data. We evaluated how the model-quality of our HMM changed as progressively more sequences labeled as replay by Bayesian decoding were removed from the training data. In the linear track sessions we considered, we found that refining the training data resulted in models that lowered in quality at different rates as the threshold for Bayesian replay was decreased (**Figure 2.7**). Most, but not all, models dropped precipitously in quality: >50% when we removed events detected as Bayesian replay at a 95% threshold, as would be expected if the HMM represented only the local environment. In many outlier sessions in which model quality decreased more slowly, the initial (baseline) model quality was low. Intriguingly, however, in at least one outlier session where model quality decreased slowly with refinement (blue line, **Figure 2.7a**), the initial model quality was still high, and we further noted

that position decoding using lsPF yielded relatively high error (blue dot, **Figure 2.7b**). Thus, we wondered whether this and similar sessions might have contained non-local or extra-spatial PBEs that were captured by the HMM.

In order to validate the concept of model-training refinement, we considered a dataset in which multiple environments were explored on the same day and remote replay was previously observed (Karlsson & Frank, 2009). These data consisted of a series of short exploratory sessions in which an animal first explored a novel maze (E1) and then was placed in a familiar one (E2). We identified awake PBEs during the familiar E2 session and used them to train an HMM. When we refined this model by removing Bayesian-significant local replay events from the training data, we found that the model quality decreased comparatively slowly (**Figure 2.7a**, green line), indicating that the HMM was capturing more than the local spatial structure. In contrast, when we used place fields from E1 to identify Bayesian-significant remote replay events and removed these from the training data, we found that the model quality decreased rapidly as with the general linear track cases (**Figure 2.7a**, red line). When we examined individual events in detail in this data, we found many examples in which HMM-significant, Bayesian non-significant PBEs decoded to extended state sequences which turned out to correspond to reactivation of the remote track (two are shown in **Figure 2.7c-1**). If we imagine that in this experiment data were only recorded during exploration of the familiar environment, classical Bayesian decoding would treat these events as noise, as shown in the bottom half of the two examples. In contrast, our HMM-based analysis finds these events to be significant, as shown in the top half of the two examples. Thus, by combining classical Bayesian decoding and HMM-congruence, we are able to identify a signature of when an HMM trained on PBEs captures sequential structure distinct

from that dictated by the local environment. Additionally, in these cases, we show that specific non-local reactivation events can be identified.

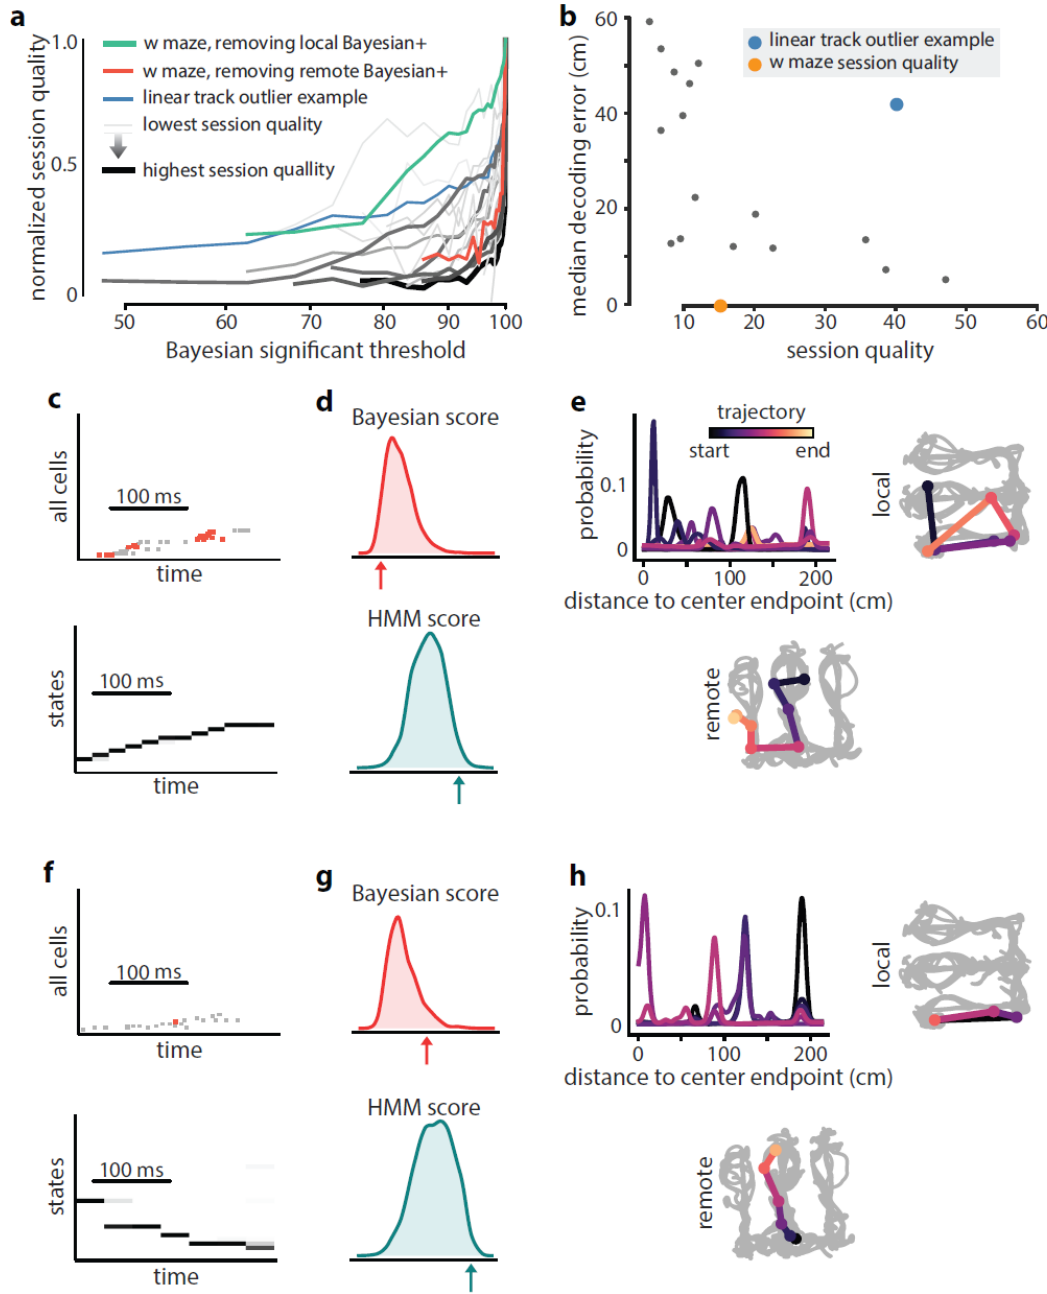


Figure 2.7: Extra-spatial structure. (a) The normalized session quality drops as local-replay events above the Bayesian significance threshold are removed from the data. Each trace corresponds to one of the 18 linear track sessions, with the stroke width and the stroke intensity proportional to the baseline (all-events) session quality. The blue line identifies a session in which model quality drops more slowly, indicating the potential presence of extra-spatial information. The reduction in session quality for a W maze experiment with known extra-spatial information

is even slower (green). When, instead, Bayesian-significant remote events are removed, rapid reduction in session quality is again revealed (red). **(b)** The lsPF-based median decoding errors are shown as a function of baseline session quality for all 18 linear track sessions. The blue dot indicates the outlier session from panel (a) with potential extra-spatial information: this session shows high decoding error combined with high session quality. Session quality of the W maze session is also indicated on the x-axis (decoding error is not directly comparable). **(c–h)** Two example HMM-congruent but not Bayesian-significant events from the W maze session are depicted to highlight the fact that congruence can correspond to remote replay. **(c)** Spikes during ripple with local place cells highlighted (top panel) and the corresponding latent state probabilities (bottom panel) decoded using the HMM show sequential structure (grayscale intensity corresponds to probability). **(d)** In this event, the Bayesian score relative to the shuffle distribution (top panel) indicates that the event is not-significant, whereas the HMM score relative to shuffles indicates (bottom panel) the ripple event is HMM-congruent. **(e)** Estimates of position using local place fields show jumpy, multi-modal a posteriori distribution over space in 1D (top left panel) and 2D (top right panel; distribution modes and time is denoted in color). Bayesian decoding using the remote environment place fields (bottom panel) indicates that the sample event is a remote replay. Note that in a typical experiment, only the local place fields would be available. **(f–h)** Same as (c–e), but for a different ripple event.

Finally, we considered the potential of our methodology for uncovering temporal patterns in PBE activity under scenarios where complex behavior does not permit identification of well-defined place-fields or in the absence of behavior, such as during sleep. As we have emphasized, a remarkable aspect of learning HMMs from PBE activity is that the model can be built entirely without behavioral data, so can our model capture significant sequential information outside of immobility periods during quiet waking? To demonstrate this potential, we examined HMMs trained on PBEs in sleep following the learning phase of an object-location memory task when animals explored three objects in an open field (see Material and methods). Previous studies have demonstrated that subsequent recall of this memory is hippocampus-dependent, and requires consolidation in post-task sleep (Inostroza & Born, 2013; Prince et al., 2014). However, while this task involves spatial exploration of objects in an arena, whether the subsequent post-task sleep contains sequential structure and whether object memory is contained in this code has remained elusive (Larkin et al., 2014). In order to assess the presence of sequential structure in the PBEs, we first used cross validation to generate a distribution of sequence HMM-congruence scores. For each set of test PBEs, we also generated surrogates by shuffling time bins across events (pooled time-swap). Using our HMM-congruence score which explicitly tests for sequences through state space, the large difference between actual and shuffled score

distributions indicates evidence for significant sequential structure in the PBEs ($p < 0.001$, Mann–Whitney U test, **Figure 2.8**). While more work is needed to evaluate the mnemonic relevance of these HMM-congruent sequences, these data support the notion that the HMM can uncover sequential activity in sleep away from the task environment. This approach further demonstrates the utility of the HMM approach as an initial analysis of a novel dataset, or as a way of comparing the sequential content encoded in PBEs during different periods.




Figure 2.8: Temporal structure during a sleep period following object-location memory task. Using cross validation, we calculate the HMM-congruence score (which ranges from 0 to 1) for test PBEs. For each event, we also calculate the score of a surrogate chosen using a pooled time-swap shuffle across all test events. The distribution of scores of actual events is significantly higher than that of the surrogate data ($p < 0.001$, Mann–Whitney U test).

2.3 Discussion

Increasing lines of evidence point to the importance of hippocampal ensemble activity during PBEs in guiding on-going behavior and active learning. Despite being the strongest output patterns of the hippocampus, however, this activity has been assumed to be interpretable only in the context of other theta-associated place cell activity expressed during behavior. Our findings demonstrate that over the course of a behavioral session, ensemble activity during PBEs alone is sufficient to form a model which captures the spatial relationships within an environment. This suggests that areas downstream of the hippocampus might be able to make use solely of PBE

activity to form models of external space. In an extreme view, place cell activity might merely subserve the internal mechanisms in the hippocampus which generate PBE sequences. To the extent that animals might wish to use the spatial code obtained from PBEs to identify their current location, we show that this can be done after translating ensemble activity into the latent states of the model. Do the PBEs contain ‘full information’ about the environment? Bayesian decoding of location from place cell activity results in lower error than location estimates generated using the latent states and lsPFs. This suggests that the manifold defined by the HMM may not capture all the dimensions of information represented during exploration. However, it is possible that with more PBE data, we would learn a more refined state space. Thus, the difference between the latent space represented during behavior and within PBEs may be an interesting focus of future study.

When we examined the transition matrices we learned from PBEs, we found that they were marked by significant sparsity. This sparsity results from the sequential patterns generated during PBEs. Latent variable models have previously been used to analyze the structure of hippocampal place cell activity (Chen et al., 2012, 2014; Dabaghian et al., 2014). In these studies, the learned transition matrices were mapped to undirected graphs which could be analyzed using topological measures. It is intriguing that similar structure is apparent in PBE activity. For example, we observed that transition matrices learned from PBEs associated with linear track behavior were significantly sparser than those learned from the open field, which we hypothesize is a consequence of the greater freedom of behavior in the latter (a topological difference). Whether hippocampal PBE activity must always be sequential, i.e., evolve through a sparsely-connected latent space, is an open and interesting question, as are differences between the latent state space dynamics learned during PBEs and those learned from place cell activity.

2.3.1 Graded, non-binary replay detection

Remarkably, evaluating the congruence or likelihood of test data against our HMM provided a highly novel method to detect events that are consistent with replay, without a need to access the ‘play’ itself. In the process of evaluating the potential of HMMs for detecting replay, we developed an approach to compare different replay-detection strategies. Our results highlight how the data do not readily admit to a strict separation between ‘replay’ and ‘non-replay’ events. While it is possible that with additional shuffles or other restrictions (Feng et al., 2015), automated performance might be rendered closer to human-labeling, even human scorers had variation in their opinions. This calls into doubt judgments of memory-related functions which build on a binary distinction between replay and non-replay sequences. Model congruence, either as a raw statistical likelihood or weighted against a shuffle distribution, seems to be a very reasonable metric to associate with individual PBEs. Moreover, evaluating congruence with an HMM does not require access to repeated behavioral sequences, which may be infeasible under widely-used single- or few-trial learning paradigms or when the events involve replay of a remote internalized environment. Given these benefits, along with computational efficiency, we would suggest that future analyses of the downstream impact of hippocampal reactivation regress effects against this measure rather than assuming a binary distinction.

2.3.2 Learning, model congruence and replay quality

Not surprisingly, the rate of PBEs had a large effect on our ability to measure model congruence. Interestingly, it has been noted that the density of PBE is higher during early exposure to a novel environment (S Cheng & Frank, 2011; Frank et al., 2004; Kemere et al., 2013; Kudrimoti et al., 1999). This might suggest that for the animal, PBE activity could be an important source for

generating models of the world when the animal is actively learning about the environment. If as hypothesized, replay is a form of rehearsal signal generated by the hippocampus to train neocortical modules (G. Buzsaki, 1989), then indeed the brain’s internal machinery may also be evaluating whether a given sequential PBE pattern is congruent and consistent with previously observed PBEs. In later sessions, as animals have been repeatedly exposed to the same environments, downstream regions will have already witnessed many PBEs from which to estimate the structure of the world. Overall, our approach provides a novel viewpoint from the perspective of hippocampal PBEs. An interesting future line of inquiry would be to assess the extent to which a model built on PBEs during first experience of a novel environment is slower or faster to converge to the final spatial map than models built on theta-associated place activity.

2.3.3 Application to extra-spatial behaviors

We have analyzed data gathered in experiments in which rats carried out simple spatial navigation tasks. Thus, to some extent it is not surprising that when we decoded ensemble activity during behavior, we found that spatial positions the animal is exploring are strongly associated with the latent states.

We anticipate that our approach for calculating lsPFs would be equally useful in tasks in which the hippocampal map is organized around time (Eichenbaum, 2014) or other continuous variables (e.g. sound frequency (Aronov et al., 2017)). Our two proof-of-concept analyses, however, suggest that it should be possible to use HMMs to infer the presence of extra-spatial sequential reactivation in PBEs. For example, we showed that there is significant sequential structure during sleep after an animal explores novel objects in an environment. We anticipate that careful experimental design and further algorithmic development would allow for the

conjunctive coding of object identity and spatial locations to be detected in the latent states we learn from PBEs, with model-congruence providing a tool to study sequential hippocampal reactivation in these types of tasks.

Conjunctive, non-spatial information might be one source of the apparent variability that results in many PBEs not being detected as replay using traditional Bayesian decoding. Another proposed source of this variability is reactivation of other environments. Our second proof-of-concept analysis suggests that HMMs learned from PBEs can, in fact, capture the spatial structure of environments beyond the one the animal is currently exploring. It appears that it should be possible to use only the PBEs and information about the place-cell map of the local environment to refine the training set for remote replay activity and learn the structure of a remote environment. While we used Bayesian decoding to detect putative local replays, we anticipate related approaches might use an HMM or other approaches to model local place cell activity.

2.3.4 Future possibilities

It has been previously observed that the rate of hippocampal reactivations in PBEs during awake behavior is much higher than during sleep (Grosmark & Buzsáki, 2016), but the reasons for this are not well understood. One hypothesis is that many sleep PBEs contain the reactivation of contexts other than those measured during a behavioral experiment. Another hypothesis is that sleep activity involves remodeling of dynamic network architectures (Buhry et al., 2011; Tononi & Cirelli, 2014). Our approach has the potential to illuminate some sources of variability during sleep. While we have given preliminary evidence that information about a remote context can be present in PBEs along with the local context, further work is required to understand how our

model's ability to capture this structure scales with the number of different contexts. With sufficient data, our HMM approach should be able to learn disjoint sets of latent states (or 'sub-models') which would capture these separate contexts and allow us to test this possibility. Alternatively, sleep PBEs could yield models which represent a known behavioral context but are markedly different (e.g., less sparse) than those learned from awake PBEs. This might support the network remodeling function of sleep. In the latter case, we might imagine that only a small subset of sleep PBEs—corresponding to learning-related replay—would be congruent with a model learned from awake PBE data.

2.3.5 Conclusions

We have demonstrated a new analytical framework for studying hippocampal ensemble activity which enables primacy of PBEs in model formation. We use an unsupervised learning technique commonly used in the machine learning field to study sequential patterns, the hidden Markov model. This contrasts with existing approaches in which the model—estimated place fields for the ensemble—is formed using the theta-associated place cell activity. We find that our PBE-first approach results in a model which still captures the spatial structure of the behavioral tasks we studied. Additionally, we demonstrate that we can use model-congruence as a tool for assessing whether individual PBEs contain hippocampal replay. Finally, we present proofs-of-concept that this analytical approach can detect the presence of sequential reactivation in experimental scenarios in which existing approaches are insufficient. Thus, the use of unsupervised learning of latent variable models—specifically HMMs and statistical congruence as a marker of individual events—bears much promise for expanding our ability to understand how PBEs enable the cognitive functions of the hippocampus.

2.4 Materials and methods

2.4.1 Experiment paradigm/neural data recording

We analyzed neural activity recorded from the hippocampus of rats during periods in which they performed behavioral tasks in different paradigms. First, we considered data from animals running back and forth in a linear track 150 or 200 cm long. As previously reported using these same data (Diba & Buzsáki, 2007), we recorded neural activity using chronically-implanted silicon probes to acquire the activity of hippocampal CA1/CA3 neurons. From these experiments, we chose sessions during which we observed at least 20 place cells during active place-field exploration, and at least 30 PBEs (see below). Place cells were identified as pyramidal cells which had (i) a minimum peak firing rate of 2 Hz, (ii) a maximum mean firing rate of 5 Hz, and (iii) a peak-to-mean firing rate ratio of at least 3, all estimated exclusively during periods of run (as defined before, that is, when the animal was running >10 cm/s). This selection yielded $n=18$ sessions with 41–203 neurons (36–186 pyramidal cells). All procedures were approved by the Institutional Animal Care and Use Committee of Rutgers University and followed US National Institutes of Health animal use guidelines (protocol 90–042).

A second dataset used tetrodes to record a large number (101–242) of putative pyramidal neurons in area CA1 during two sessions each in four rats. Briefly, as was previously reported using these data (Pfeiffer & Foster, 2013, 2015), rats explored an arena in which there were 36 reward sites. In each session, one site was designated as ‘home’. During a session, rats would repeatedly alternate between retrieving a random reward site in one of the remaining 35 locations and retrieving a reward at the home location. All procedures were approved by the Johns

Hopkins University Animal Care and Use Committee and followed US National Institutes of Health animal use guidelines (protocols RA08M138, RA11M16, and RA14M48).

In order to investigate remote replay, we used data from an experiment in which this phenomenon has been previously reported (Karlsson & Frank, 2009). Briefly, rats were implanted with multi-electrode microdrives with tetrodes targeting CA1 and CA3. They were trained to carry out a continuous-alternation task in an initially novel ‘w’-shaped maze (E2) for liquid reward for multiple daily run sessions interspersed by rest-periods in an enclosed box. After they learned the task, they were introduced to a novel w-maze (E1) in a different orientation in which they had two run sessions followed by a run in the now-familiar E2. For our proof-of-concept analysis (Figure 2.7), we used data from the second day of the novel maze (i.e., third and fourth exposures) in animal ‘Bon’.

Finally, we recorded neural activity during an object-location memory task using a 32-channel silicon probes (Buzsaki32, Neuronexus, MI) equipped with light fibers lowered to area CA1 of the dorsal hippocampus. The animal was previously infused with AAV-CamKIIa-ArchT-GFP for the purpose of another experiment. Putative pyramidal cells and interneurons were distinguished based on their spike waveforms and spike auto-correlograms. On the day before the recordings, the animal was repeatedly exposed to an empty test chamber on four successive six minute blocks, interleaved by three minute rest periods in the home cage. On the recording day, the first of these six-minute blocks was again the empty test chamber, but on the remaining blocks, the animal was exposed to a fixed configuration of three different novel objects placed in the northeast, center and southeast corners of the box. These blocks were again interleaved with three minute rest periods in the home cage. The test chamber was a $60 \times 60 \text{ cm}^2$ box with a local

cue (8.5 in. \times 11 in. sheet printout) placed on one test wall. Following the last acquisition exposure, the animal was returned to its home cage for a four hour extended sleep period. The subsequent day, one of the objects in the box was displaced and the animal was reintroduced into the box to test for interactions with the displaced versus non-displaced objects. All procedures were approved by the Institutional Animal Care and Use Committee of the University of Wisconsin-Milwaukee and followed US National Institutes of Health animal use guidelines (protocol 13–14 #28)

2.4.2 Population burst events

To identify PBEs in the linear track data, a spike density function (SDF) was calculated by counting the total number of spikes across all recorded single and multi-units in non-overlapping 1 ms time bins. The SDF was then smoothed using a Gaussian kernel (20 ms standard deviation, 60 ms half-width). Candidate events were identified as time windows with a peak SDF of at least three standard deviations above the mean calculated over all the session. The boundaries of each event were set to time points of crossing the mean, preceding and following the peak. Events during which animals were moving (average movement speed of >5 cm/s) were excluded from all further analyses to prevent possible theta sequences from biasing our results. For analysis, we then binned each PBE into 20 ms (non-sliding) time bins. Spikes from putative interneurons (mean firing rate when moving >10 Hz) were excluded, as were events with duration less than four time bins or with fewer than four active pyramidal cells. For the open field data, we used the previously reported criteria (Pfeiffer & Foster, 2013) for identifying PBEs prior to binning (10 ms standard deviation kernel, minimum of 10% of units active, duration between 50 ms and 2000 ms).

2.4.3 Hidden Markov model of PBE activity

We trained HMMs complete sequence on the PBEs. In an HMM, an unobserved discrete latent state evolves through time according to a first order Markov process. The temporal evolution of the latent state is described by the $M \times M$ matrix A , whose elements $\{a_{ij}\}$ signify the probability after each time bin of transitioning from state i to state j , $a_{ij} = \Pr(q_{t+1} = j | q_t = i)$. The number of states, M , is a specified hyperparameter. We found that our results were insensitive to the value of M through a wide range of values from 20 to 100. During each time bin of an event, the identity of the latent state influences what is observed via a state-dependent probability distribution. We modeled the N -dimensional vector of binned spiking from our ensemble of N neurons at time t , O_t , as a Poisson process. Specifically, for each state, i , we model neuron n as independently firing according to a Poisson process with rate λ_{ni} .

$$\Pr(O_t | q_t = i) = \prod_{n=1}^N \Pr(o_{nt} | q_t = i) \propto \prod_{n=1}^N \lambda_{ni}^{o_{nt}} \exp(-\lambda_{ni})$$

where o_{nt} is the number of spikes observed from neuron n at time t . The final parameter which specifies our model is the probability distribution of the initial state for a given event, $\pi_i = \Pr(q_i = i)$. Thus, our model is specified by parameters $\theta = \{A, \Lambda, \pi\}$, where $\Lambda = \{\lambda_{ni}\}$ is an $N \times M$ matrix and $\pi = \{\pi_i\}$ is an M -dimensional vector.

To learn model parameters, we follow the well-known iterative EM procedure (Rabiner, 1989), treating each training PBE as an observation sequence. In order to regularize the model, we impose a minimum firing rate for each neuron of 0.001 (0.05 Hz) during the M-step of EM. For a given PBE (i.e., observation sequence) with K bins, we use the ‘forward-backward algorithm’

(Rabiner, 1989) to calculate the probability distribution of the latent state for each time bin, $Pr(qt|O_1, \dots, O_t, \dots, O_K)$. For a particular time bin, t , in a given sequence, the forward-backward algorithm allows information from all observation bins, previous and subsequent, to affect this state probability distribution (as well the observation bin at time t). The forward-backward algorithm also efficiently calculates the ‘score’, or likelihood of the complete sequence, $Pr(O_1, \dots, O_K)$. All HMMs learned in this work used five-fold cross validation, that is, the PBEs were divided into five randomly selected fifths (‘folds’), and then each fold was evaluated as a test set, with the model trained using the remaining four folds. We define the model likelihood of an HMM as the product of the scores of each event using this five-fold cross validation. To initially evaluate model learning, we compared model likelihoods calculated using real and shuffled test data. Models which have learned to properly represent the data should show significant increases. To quantify the presence of PBE sequences in a model we used a model quality metric as described below.

2.4.4 Ordering states for visualization

For visualization, we wanted to order the states to maximize the super diagonal of the transition matrix. We used a greedy approach which typically yields this solution. We started by assigning the first index to the state with the highest initial probability and added states based on the most probable state transitions. The undirected connectivity graphs were then generated from this transition matrix, averaging the strength of reciprocal connections, a_{ij} and a_{ji} .

2.4.5 Surrogate datasets and shuffle methods

In order to analyze the HMMs that we learned, we compared them against different types of surrogate datasets obtained by shuffling the neural activity during PBEs in distinct ways. (1) Temporal shuffle: within each event, the binned spiking activity was circularly permuted across time for each unit, independently of other units. This goal of this shuffle is to disrupt unit co-activation, while maintaining the temporal dynamics for each unit. (2) Time-swap shuffle: within each event, the order of the binned columns of neural activity was randomly permuted across time, coherently across units. The goal of this shuffle is to change the temporal dynamics of ensemble activity, while maintaining unit co-activation. (3) Poisson surrogate ‘shuffle’: we estimated each unit’s mean firing rate across all PBEs, and then produced surrogate PBEs from independent Poisson simulations according to each unit’s mean firing rate. (4) Pooled time-swap shuffle: the order of the binned columns of neural activity was randomly permuted across all pooled events, coherently across units. This shuffle has been previously used in Bayesian replay detection (Davidson et al., 2009).

2.4.6 Calculating sparsity and connectivity of the model parameters

Sparsity of the transitions from individual states (departure sparsity) was measured by calculating the Gini coefficient of corresponding rows of the transition matrix (Hurley et al., 2009). The Gini coefficient is a measure of how variable the values of this probability distribution are, with equality across states corresponding to a coefficient of zero (minimal sparsity), and a singular distribution with a probability-one transition to a single other state corresponding to a coefficient of one (maximal sparsity). The sparsity of the full transition matrix was calculated by averaging the Gini coefficient across rows. For analyses of PBE models

from linear tracks, we computed the mean sparsity across states for each of the 250 surrogate datasets, and these means were used to generate the box plots of Figure 2.2c. Note that for the actual data, we generate a distribution by randomly initializing the model 250 times and calculating the mean sparsity over all initializations. For analyses of models learned from PBEs in open fields (and the linear track comparison with 50 states), we created 50 surrogates/random initializations. To compare across sessions, we calculated the mean sparsity by averaging over all 250 surrogate datasets to obtain a single mean sparsity per session, so that $n=18$ per-session means were used to create the box-plots of Figure 2.2e.

Firing rates can be highly variable for different units. Thus, when evaluating the sparsity of the observation matrix, we measured the extent to which individual units were specifically active in a few states by calculating the Gini coefficients of the rows of the observation matrix. As with transitions, we calculated mean sparsity across units for each surrogate dataset, and we then averaged over all surrogate datasets to obtain a per-session average, used in Figure 2.2f.

2.4.7 Model connectivity and sequences

To measure the degree of sequential connectivity within the graph corresponding to the transition matrix—with nodes and edges representing the states and transitions, respectively—we developed an algorithm for measuring the length of the longest path that can be taken through the graph. This method is analogous to the ‘depth-first search’ algorithm for traversing the graph’s tree structure without backtracking. First, we made an adjacency matrix for a corresponding unweighted directed graph by binarizing the transition matrix using a threshold of 0.2 on the transition probabilities. Starting from each node, we then found the longest path that ended at either a previously visited node or a terminal node (a node without any outgoing edges). To

compare models trained on actual versus surrogate datasets, we adjusted the thresholds to match the average degree (defined as the average number of edges per node) between the models, thus ruling out possible effects due to differences in the number of graph edges. We carried out this analysis on the same set of models that were generated for analyzing sparsity. To compare across sessions, we calculated the median maximum path length for each session ($n=18$) and used the per-session medians.

2.4.8 Latent state place fields

To calculate the latent state place fields, we first identified bouts of running by identifying periods when animals were running (speed >10 cm/s). We then binned the spiking during each of these bouts in 100 ms bins. Using the forward-backward algorithm (Rabiner, 1989) and the HMM model parameters learned from PBEs, we decoded each bout into a sequence of latent state probability distributions, $Pr(q_t|O_t)$. Using the track positions corresponding to each time bin, we then found the average state distribution for each position bin, x_p , and normalized to yield a distribution for each state, $Pr(x_p|q_t = i)$.

2.4.9 Decoding position from latent state sequences

We used the lsPF to decode the animal's position after determining the probability of different latent state trajectories during bouts of running. With five-fold cross validation, we estimated lsPF in a training dataset, then used the HMM model to decode latent state trajectory distributions from ensemble neural activity in the test data. The product of lsPF and decoded latent state distribution at time t is the joint distribution $Pr(x_p, q_t|O_t)$. We decode position as the mean of the marginal distribution $Pr(x_p|O_t)$.

2.4.10 Bayesian replay detection

We followed a frequently used Bayesian decoding approach to detect replay in our 1D data (Davidson et al., 2009). For each 20 ms time bin t within a PBE, given a vector comprised of spike counts from N units, $O_t = (o_{1t} o_{2t} \dots o_{Nt})$ in that bin, the posterior probability distribution over the binned track positions was calculated using Bayes' rule:

$$Pr(x_p|O_t) = \frac{Pr(O_t|x_p)Pr(x_p)}{\sum_{q=1}^P Pr(O_t|x_q)Pr(x_q)}$$

where x_p is the center of p -th linearized position bin (of P total bins). We assumed Poisson firing statistics, thus the prior probability, $Pr(O_t|x_p)$, for the firing of each unit n is equal to

$$Pr(O_t|x_p) = \prod_{n=1}^N Pr(o_{nt}|x_p) \propto \prod_{n=1}^N (\tau \lambda_{np})^{o_{nt}} \exp(-\tau \lambda_{np})$$

where τ is the duration of time bin (100 ms during estimation, 20 ms during decoding), and λ_{np} characterizes the mean firing rate of the n -th unit in the p -th position bin. We assumed a uniform prior distribution $Pr(x_p)$ over the position bins.

For each PBE, the estimated posterior probability distribution was used to detect replay as follows. Many (35,000) lines with different slopes and intercepts were sampled randomly following the approach in (Davidson et al., 2009). The Bayesian replay score for a given event was the maximum score obtained from all candidate lines, where the score for a particular line was defined as the mean probability mass under the line, within a bandwidth (of 3 cm). For time bins during which the sampled line fell outside of the extent of the track, the median probability

mass of the corresponding time bin was used, and for time bins during which no spikes were observed, we used the median probability mass across all on-track time bins. To evaluate the significance of this score, for each event we generated 5000 surrogates of the posterior probability distribution by cycling the columns (i.e., for each time bin, circularly permuting the distribution over positions by a random amount) and calculated the replay score for each surrogate. The Monte Carlo p-value for each event was obtained from the number of shuffled events with replay scores higher than the raw data. The threshold for significance was varied as described in the text. For the open field, we used previously reported criteria (Pfeiffer & Foster, 2013) to identify replay events from PBEs.

2.4.11 Replay detection via PBE model congruence

To identify replay as model congruence, for each PBEs, we used the forward-backward algorithm to calculate the sequence likelihood $Pr(O_1, \dots, O_K)$, as defined earlier. Using five-fold cross validation, the parameters of a HMM were learned from training PBE. The sequence score was then calculated for each event in the test data. To evaluate the significance of this score, for each event we generated 5000 surrogate scores using a computationally-efficient scheme. Specifically, for each surrogate, we randomly shuffle the rows of the transition matrix, excepting the diagonal. By maintaining the diagonal (i.e., transitions that begin and end in the same state) and leaving the observation model unchanged, this shuffle specifically selects against PBEs in which the latent states do not evolve in temporal sequences. The Monte Carlo pp-value for each event was calculated as the fraction of shuffled events with HMM sequence scores higher than the raw data. The threshold for significance was varied as described in the text. Note that while we describe this as HMM-congruence, we have maintained the diagonal of the transition matrix,

which specifically selects against PBEs which might be model-congruent by maintaining a single state over many time bins. In reality, there are other dimensions of the HMM that we could assess congruence against, for example the observation model, the initial state distribution, or combinations of these and the transition matrix. In comparing against Bayesian decoding, our current definition seemed most appropriate for sequence detection, but we can imagine future studies expanding on our approach.

2.4.12 Human scoring and detection comparison

We organized a group of human scorers to visually evaluate whether individual PBEs should be described as replay. More specifically, scorers were only presented with Bayesian decoded probability distributions such as those in Figure 2.4a, but without access to the spike raster or any additional information. The scorers included six graduate students (including one of the authors) and two undergraduates, all of whom were generally familiar with the concept of hippocampal replay. We built an automatic presentation system which would display each event in random order, and record one of six possible scores: ‘excellent’ (highly sequential with no jumps and covering most of the track), ‘good’ (highly sequential with few or no jumps), ‘flat’ (decoded position stayed mostly in the same place, i.e. no temporal dynamics), ‘uncertain’ (some semblance of structure, but not enough to fall into any of the previous categories) or ‘noise’ (no apparent structure, or nonsensical trajectories such as teleportation). An event was then designated as replay if it was labeled as ‘excellent’ or ‘good’ by a majority of scorers (ties were labeled as non-replay).

To calculate an ROC curve for replay detection algorithms, we used our shuffle statistics for each event to create a vector which related the significance threshold (e.g., 99%) to the label

supplied by the algorithm (i.e., significant replay or not). Then, as a function of threshold, the sensitivity (fraction of true positives identified) and selectivity (fraction of true negatives identified) were averaged over events to yield an ROC curve. To evaluate whether the AUC differed between Bayesian and model-congruence techniques we used a bootstrap approach. To generate a null hypothesis, we combined the event/threshold vectors from both groups, and then sampled two random groups (A and B) with replacement from the pooled data. The AUC for these two random groups of events were measured, and a distribution for the difference between the randomly chosen AUC was calculated. The two-sided p-value we report is the fraction of differences in random AUC which are more extreme than the actual difference.

2.4.13 HMM model quality across sessions

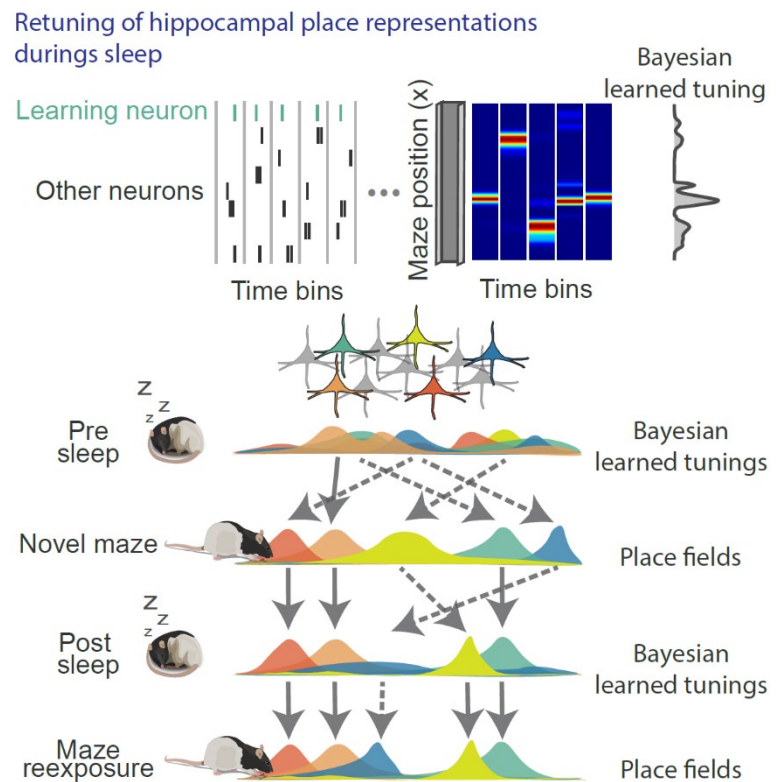
In order to understand the extent to which an HMM trained on PBEs from a given session contained sequentially-structured temporal dynamics, we calculated the ‘session quality’ (equivalently model quality) as follows. Again, using five-fold cross validation, we learn an HMM on the training subset of PBEs, and score (using the forward-backward algorithm, as before), the remaining subset of test PBEs. Then, we also score a pooled time-swap surrogate of the test PBE and we repeat this pooled time-swap scoring $n=2500$ times. Finally, we obtain a z score for each PBE by comparing the score from the actual test PBE to the distribution of pooled time-swap scores of the corresponding PBE. The session quality is then defined as the average of these z scores, over all events in a session. This measure of session quality was then used to detect the presence of putative remote replay events or other extra-spatial structure in PBEs, since a high session quality after removing local Bayesian significant events is highly suggestive of remaining (unexplained) sequential structure.

3 Retuning of hippocampal place representations during sleep

Summary

Hippocampal representations that underlie spatial memory undergo continuous refinement following initial formation. Understanding the role of sleep in this process has been challenged by the inaccessibility of place fields when animals are not actively exploring a maze. Here, we introduced a novel Bayesian learning approach to discover and dynamically track the spatial tuning of individual neurons during offline states based on the spike-triggered average decoded position of other neurons in ensemble recordings from freely moving rats. Using these dynamic tunings within hippocampal sharp-

wave ripples, we found stable spatial representations that were strongly aligned with maze place fields in slow-wave sleep following initial maze exploration. These representations were explained by a combination of factors that included pre-configured structure of firing rates in sleep before exposure to the environment, and representations



that emerged during theta oscillations and awake sharp-wave ripples on the maze. Remarkably, the ripple representations during sleep predicted the future place fields of neurons upon re-exposure to the maze, even when those fields deviated from previous place preferences, indicating that sleep drives representational drift across maze exposures. In contrast, we observed tunings with poor alignment to maze place-fields in most other time periods, including in sleep and rest before maze exposure, during rapid eye movement sleep, and during the late stages of slow-wave sleep. Overall, our population decoding approach allowed us to infer and characterize the retuning of place fields during offline periods, revealing the rapid emergence of representations following novel exploration and the active role of sleep in the representational dynamics of the hippocampus.

Highlights

- Tuning curves measured using Bayesian learning allow us to track spatial representations of hippocampal neurons during offline states.
- Stable spatial representations emerge in sharp-wave ripples, including in non-replay events, during sleep following maze exposure but not before.
- Spatial representations in sleep reflect the firing patterns of pre-configured ensembles, as well as representations that emerge during theta oscillations and awake sharp-wave ripples on the maze.
- The representations evident during early hours of sleep predict the future place fields of neurons upon subsequent exposure to the maze, even when those fields deviate from those formed on initial exposure.

3.1 Introduction

Memories are continuously refined after they form. Different stages of sleep are considered to play important roles in the transformations that memories undergo, but much about these offline processes remains unknown. Memories that involve the hippocampus are particularly affected by sleep, which alters molecular signaling, excitability and synaptic connectivity of hippocampal neurons (Havekes & Abel, 2017; Klinzing et al., 2019). Memories are considered to be represented by the activity of ensembles of neurons that form with experience. In the rat hippocampus, these ensembles are tuned to locations within a maze environment (J. O'Keefe & Dostrovsky, 1971). Indeed, an animal's position can be decoded from the spike trains recorded from a population of neurons (**Figure 3.1a**) (Zhang et al., 1998). These spatial representations, however, do not remain stationary following their initial formation. In many cases the place fields (PFs) of hippocampal neurons develop and shift during traversals of an environment (Dong et al., 2021; Frank et al., 2004), remap upon exposure to different arenas (Alme et al., 2014), and reset or remap even with repeated exposure to the same place (Mankin et al., 2012; Ziv et al., 2013).

We conjectured that modifications of spatial representations may take place during sleep when connections between some neurons are strengthened while those between other neurons are weakened (Cirelli & Tononi, 2022; Klinzing et al., 2019). Consistent with this conjecture, cells that become active in a new environment continue to reactivate for hours during sharp-wave ripples in sleep (Giri et al., 2019), suggesting that offline processes during sleep involve the spatial representations of hippocampal neurons. Moreover, the collective hippocampal map of space shows changes following sleep (Grosmark et al., 2021) and some cells express immediate early genes during that period which can mark them for sleep-dependent processing (Pettit et al.,

2022). However, while spatial representations are readily measured from the spiking activities of neurons when animals explore a maze environment, access to these non-stationary representations is lost when animals are removed from the maze, making it challenging to evaluate how spatial representations are shaped over time.

To evaluate and track the spatial preferences of a neuron across online and offline periods, we developed a novel method based on the principle of Bayesian learning (Wiskott, 2013) (**Figure 3.1b**). Under the assumption of conditional independence of Poisson spike counts from hippocampal neurons conditioned on location, we derived the Bayesian learned tuning (LT) of a neuron as its spike-triggered average of the posterior probability distribution of position based on the simultaneous spiking patterns of all other neurons in the recorded ensemble, including for time periods when animals are remote from the maze locations for which position was specified. In this formalism, a neuron's internally generated preference for a location is revealed through its consistent coactivity with other neurons in the ensemble associated with that position.

These Bayesian learned tunings allowed us to track, for the first time, the place-preferences of neurons as they evolved in long-duration (up to 12 h) hippocampal unit recordings, enabling us to identify those periods and events in which the firing activities of neurons were consistent or inconsistent with place fields on the maze, and to characterize the plastic offline changes in tuning relative to the broader ensemble, leading to multiple novel insights. We found that in sleep following exposure to a maze, hippocampal neurons rapidly reconfigured to provide spatial representations that aligned with the place fields on the maze. These representations remained stable for hours, despite the scarcity of sequential replay trajectories through the maze environment. Remarkably, the sleep representations during sharp-wave ripples predicted the place fields of neurons upon repeat exposure to the same maze, including for neurons whose

fields deviated from their previous place preferences. Overall, the Bayesian learning approach allowed us to characterize the offline dynamics of place fields and revealed the important role of sleep in the tuning the spatial representations in the hippocampus.

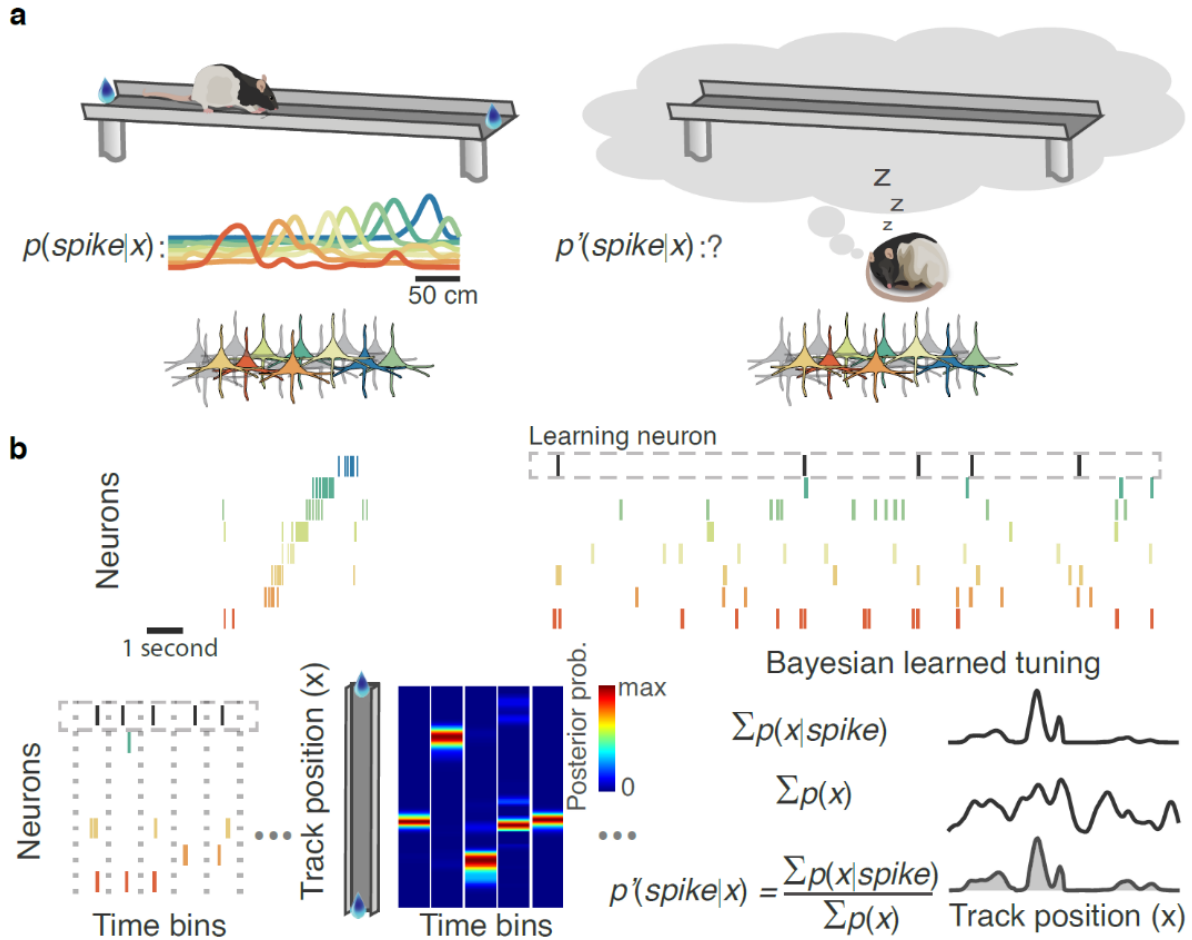


Figure 3.1: Schematic for Bayesian learning of hippocampal spatial tunings during offline states. (a) Hippocampal place cells show tuning to specific locations (place fields) on a linear track maze. When animals sleep or rest outside of the maze, the spiking of these neurons is no longer driven by maze location but may represent an internally generated simulation of x or another location. **(b)** We employed Bayesian learning to learn each neuron's tuning $p'(\text{spike}|x)$ for internally generated cognitive space, x , using the place fields of all other neurons recorded on the maze, under the assumption of conditional independence among Poisson spiking neurons conditioned on space. Top left, sample spike raster during an example traversal. Top right, spiking patterns of the same cells during a brief window in sleep. Bottom, left to right, population activity extracted for time bins in which the "learning" neuron spikes. Next, posterior probability distributions are constructed using the spikes and track tunings of the other neurons during these time-bins. The Bayesian learned tuning $p'(\text{spike}|x)$ is set to the summation of the posterior distributions over these time-bins ($\sum p(x|\text{spike})$), normalized by the overall likelihood of each track location ($\sum p(x)$) obtained across the entire offline period.

3.2 Results

3.2.1 POST but not PRE ripples align with place fields on the maze.

We first examined how tuning curves are impacted by an animal's experience on a maze by characterizing the representations of neurons from spike trains recorded from the rat hippocampus in experiments where rest and sleep in a home cage both preceded (PRE) and followed (POST) exposure to a novel track (MAZE), where rats ran for water reward. To examine spatial tunings in each brain state separately, we first separated unit and local field potential data recorded from hippocampal region CA1 into different states using general criterion (see Methods) for rapid eye movement sleep (REM, sleep featuring prominent theta), ripple periods during rest and sleep (150-250 Hz band power accompanied by high firing rates), slow-wave sleep (SWS) periods exclusive of ripples, and active wake (with prominent theta). We calculated place fields and the learned tunings for each epoch for all units with peak spatial firing rates > 1 Hz on the maze (**Figure 3.2a-c**). We limited the initial analysis to the first 4h of POST, during which we expect greater similarity with maze firing patterns (Giri et al., 2019). Learned tunings showed a wide distribution of fidelity to place fields from PRE to POST depending on brain state. Population vector (PV) correlations between spatial bins in place fields and learned tunings (**Figure 3.2b**) and LT-PF Pearson correlation coefficients (**Figure 3.2c**) demonstrated that the highest fidelities to place fields were observed in tunings during theta and ripples on the maze, as expected (Diba & Buzsáki, 2007; Dragoi & Buzsáki, 2006). However, among offline periods only tunings exhibited during POST ripples and non-ripple slow-wave sleep showed significant correlations with unit place fields in MAZE, and notably not those during PRE ripples. Surprisingly, we also failed to find representations consistent with the maze during REM sleep, when vivid dream episodes are frequently experienced (Siclari et al., 2017). This may

reflect that the bulk of REM sleep corresponds only weakly to previous experience; indeed, replays during REM sleep have thus far only been reported in over-trained rats upon repeated exposure to the same maze (Louie & Wilson, 2001). Thus, we find that only during POST slow-wave sleep do place fields maintain internal tunings consistent with their place fields on the maze. Learned tunings reveal the stability of spatial representations in sleep.

We next tracked the learned tunings of neurons over time and examined the stability and consistency of their place preferences within different epochs. We calculated LTs in 15 min windows sliding in 5 min steps during each session, from PRE through MAZE and POST. Sample unit tunings from a recording session are shown in **Figure 3.3a**. These examples show stable LTs for multiple successive time windows during POST, and in some instances, also during PRE. To quantify the LT stability for each unit, we used Pearson correlation coefficients to assess the consistency of the learned tunings across time windows within and between behavioral epochs (**Figure 3.3b**). High off-diagonal values in the correlation matrices within an epoch indicated that the LT remained stable during that epoch. For the example units in **Figure 3.3c** we compared the median LT stability values from each epoch against shuffle distributions generated by randomizing the unit identities of the LTs at each time window. This z-scored LT stability was > 0 in both PRE and POST in this session (**Figure 3.3d**) and for data pooled across all sessions (**Figure 3.3e**), but it was significantly higher in POST compared to PRE, revealing that POST sleep representations were more stable than those in PRE. When we measured the LT stability across time windows from PRE to POST epochs, to examine their consistency from before and after the novel maze exposure when place fields first form, the PRE w/ POST LT stability was significantly > 0 in the example session (median = 0.58, $p = 0.02$) as well as in the pooled data (median = 0.66, $p < 10^{-11}$, Wilcoxon signed rank test (WSRT)) but this was

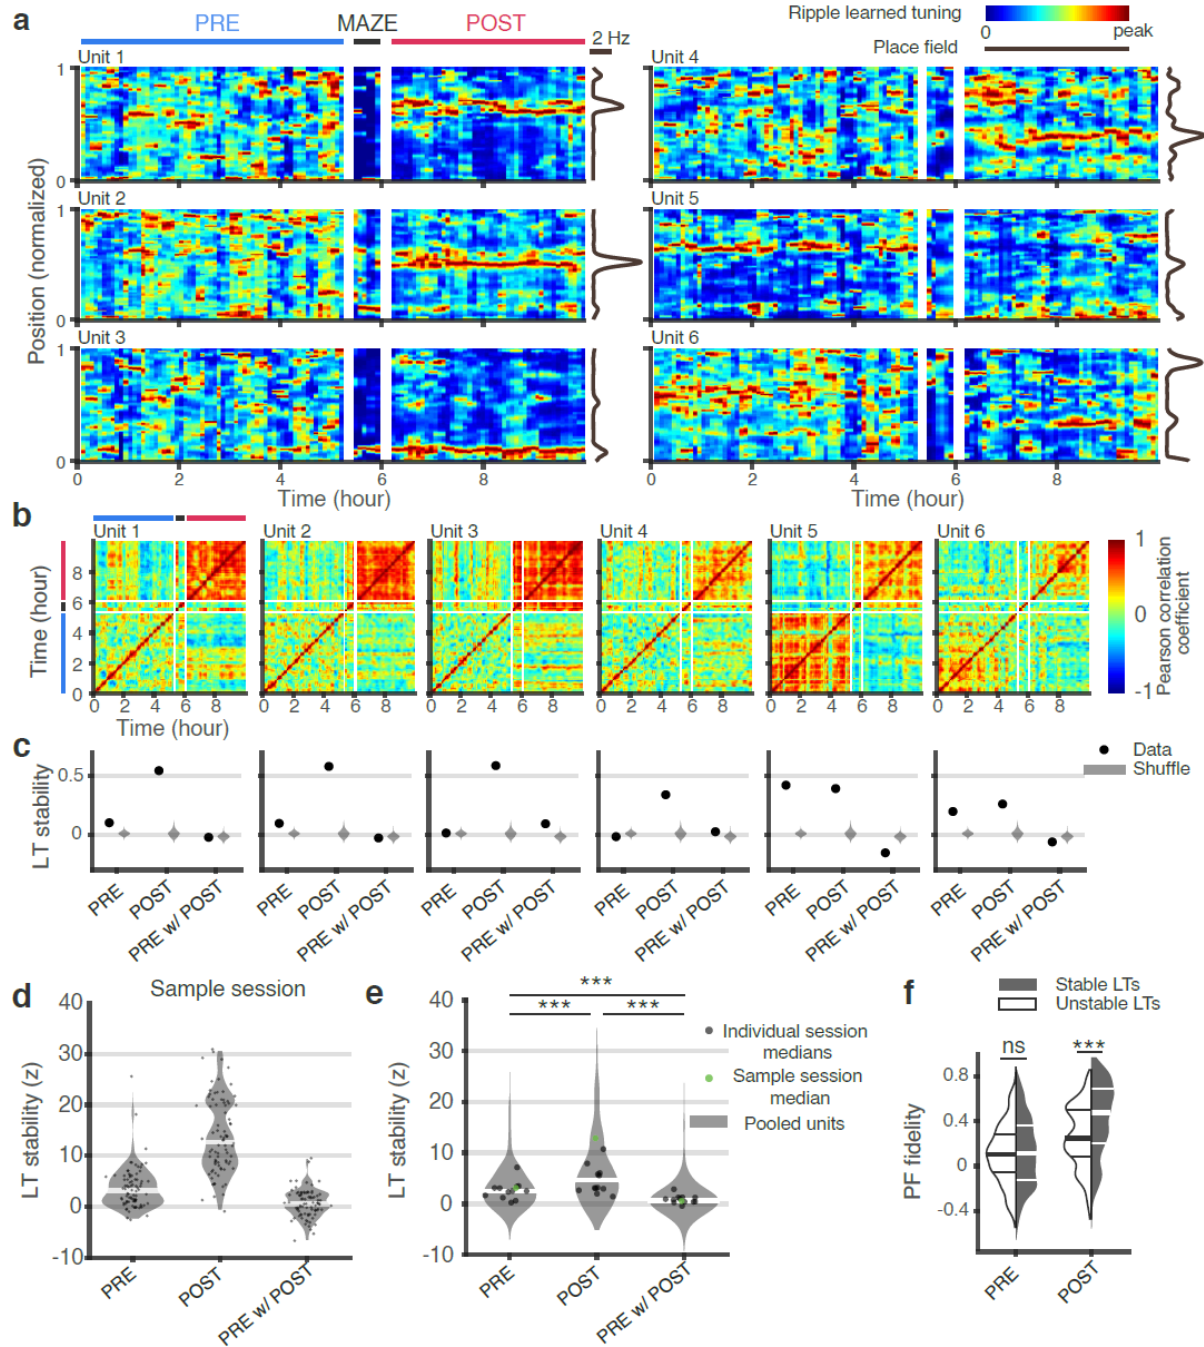


Figure 3.3: Learned tunings during ripples are stable for hours in POST. (a) Heat maps of ripple LTs for sample units in sliding 15 min windows throughout a sample session from PRE through MAZE to POST (maze PFs in gray on right) show generally stable LTs during POST. Interestingly, units 5 and 6 also show stable tunings during PRE ripples, but those tunings do not align well with their maze place fields. (b) The matrix of correlation coefficients between LTs from different time windows for the units in (a) illustrate periods of stability. (c) Stability of the LTs (black) for the units in (a) in PRE and POST, defined as the median of the correlation coefficient between LTs from non-overlapping 15-min windows. Violin plots (gray) show the chance distributions obtained from non-identical units randomly scrambled across windows (1000x). While LTs of units 5 and 6 were stable within PRE and POST, they were not consistent across these epochs. (d and e) Unit LT stabilities z-scored against unit ID shuffles were significantly > 0 for the sample session (d) (PRE: median = 3.05, $p < 10^{-4}$; POST: median = 12.59, $p < 10^{-4}$;

PRE-POST: median = 0.58, $p = 0.02$, WSRT) with individual units shown as dots, and (e) all sessions pooled together (PRE: median = 2.50, $p < 10^{-4}$; POST: median 4.68, $p < 10^{-4}$; across PRE w/ POST median = 0.66, $p < 10^{-4}$, WSRT). However, LT stability in POST was higher than for both PRE or PRE w/ POST ($p < 10^{-4}$, WSRT). (f) Distributions of PF fidelity ($r(\text{LT}, \text{PF})$) for units with stable ($z > 2$) versus unstable ($z < 2$) LTs showed no difference in PRE but were higher for stable units in POST ($p < 10^{-4}$, MWUT). *** $P < 0.001$; ns, not significant.

significantly lower than the stabilities observed within PRE and POST (PRE vs PRE w/ POST: $z = 13.2$, $p < 10^{-39}$; POST vs PRE w/ POST: $z = 18.4$, $p < 10^{-75}$, WSRT), signaling that only a small minority of units maintained the same consistent spatial tuning from before to after maze exposure.

A subset of units showed remarkably stable learned tunings during PRE which compelled us to consider whether the LTs of those units might show higher fidelity with maze PFs. To test this conjecture, we divided units into “stable” and “unstable” by whether their z-scored LT stability was $>$ or < 2 (PRE: 379 stable vs 304 unstable; POST: 491 stable vs 192 unstable), respectively, in both PRE and POST. In POST, units with both stable and unstable LTs showed significant PF fidelity ($p < 10^{-4}$, comparison against unit identity shuffles). However, the PF fidelity of units with stable LTs was significantly higher compared to units with unstable LTs in POST.

Importantly, in PRE there was no significant difference between PF fidelities of stable and unstable units, and neither of these subsets showed significantly greater PF fidelity compared to a surrogate distribution obtained by shuffling unit identities (stable LTs: $p = 0.58$; unstable LTs: $p = 0.54$). These findings demonstrate that although some units in PRE display stable learned spatial tunings, these tunings do not typically anticipate the future place fields of these neurons but rather show a broad distribution of alignments with the maze place preferences.

While the stability and fidelity of spatial tunings were significantly greater in POST, these features did not last indefinitely. In data that featured multiple hours of POST, we observed decreases in both the fidelity and stability of Bayesian learned tunings over the course of sleep

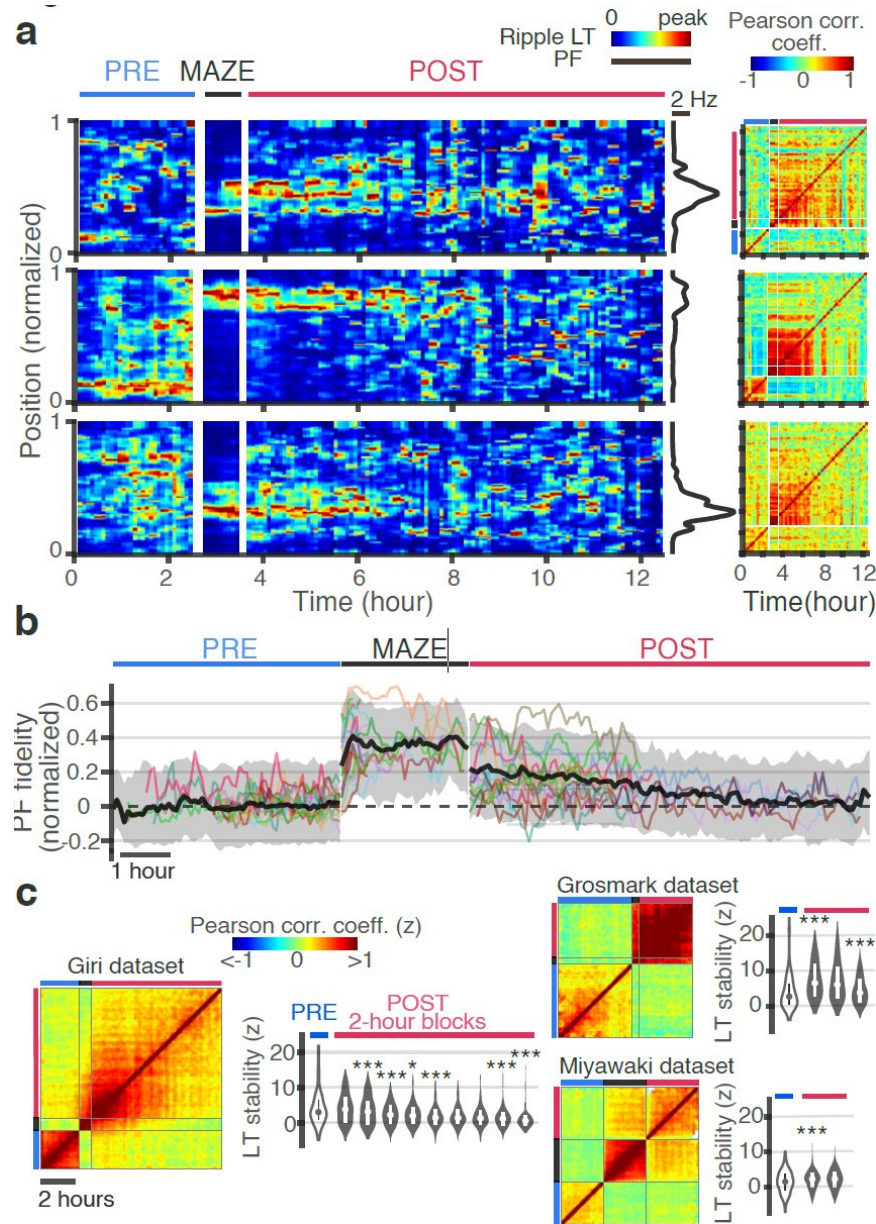


Figure 3.4: Spatial representations decay over the course of sleep. (a) Heat maps of ripple LTs for sample units in sliding 15 min windows throughout a sample long duration session show gradual decreases in LT stability over time. A matrix of correlation coefficients between LTs from different time windows is provided on the right for each unit. (b) PF fidelity (correlation coefficient between LTs and PFs) shows a gradual decrease over time in POST. The color traces show median values across units within each individual session. The black trace and gray shade depict the median and interquartile range of the pooled data. PRE and MAZE epochs of differing durations were aligned to the onset of MAZE while POST epochs were aligned to the end of MAZE. (c) Left panels show LT stability correlation matrices averaged over all recorded units, shown separately for each dataset. Here, the matrix for each unit was z-scored against unit-ID shuffles prior to averaging. Right panels show the distribution of z-scored LT stability in overlapping 2-hour blocks during POST, separately for each dataset. The asterisks above each block represent the p-value of difference in LT stability compared with the previous block. * $p < 0.05$, *** $p < 0.001$.

(**Figure 3.4a-c**). The similarity of sleep representations to maze place-fields decreased progressively over POST, eventually reaching levels similar to PRE. The stability of spatial tunings also decreased over this period, indicating that these representations become unreliable in later periods of sleep. The dissolution of representational alignment with the maze over the course of sleep may reflect an additional important aspect of sleep, distinct from that of reactivation and replay (Miyawaki & Diba, 2016; Norimoto et al., 2018).

3.2.2 Sleep dynamics predict representational consistency and drift

Recent studies report that place fields drift and frequently remap upon repeat exposures to the same environment (Grosmark et al., 2021; Kinsky et al., 2018; Mankin et al., 2012; Ziv et al., 2013) though it is unclear when and how these representational changes emerge. Given that the tunings learned during POST ripples display a diversity of place-field fidelities, some aligned but others misaligned with maze PFs, we asked how these POST ripple representations relate to the future spatial tunings of the cells. In three recording sessions, we re-exposed rats back to the maze environment after ~9h of POST rest and sleep (**Figure 3.5a**). We labelled these epochs “reMAZE” and compared the place fields across maze exposures with the ripple LTs from the intervening POST period (**Figure 3.5b-d**). POST ripple LT’s showed significant correlations with place fields from both maze exposures, indicating a continuity of representations across these periods. However, PFs were not identical between MAZE and reMAZE (**Figure 3.5b**), indicating that neuronal representations drift or remap in the rat hippocampus (Mankin et al., 2012). We hypothesized that representational remapping emerges during POST and could account for the deviations in PFs observed between repeated exposures to the maze. Consistent with this hypothesis, in instances where we saw reMAZE PFs congruent with MAZE PFs (top panel, **Figure 3.5e**), the POST LTs for those cells showed strong fidelity with the maze period.

However, in instances where reMAZE PFs deviated from the MAZE PFs (bottom panel, **Figure 3.5e**), the POST LTs for those units predicted the PFs observed during maze reexposure.

Likewise, we observed a significant correlation between PF fidelities in POST and the reMAZE-MAZE similarity (**Figure 3.5f**). To better examine whether ripple representations during POST can presage representational changes across maze exposures, we performed a multiple regression analysis to test the extent to which reMAZE PFs are explained by MAZE PFs and LTs from PRE or POST (first 4 h). We also included the average LTs (over PRE and POST) to control for the general deviations of LTs that were not specific to any unit (**Figure 3.5g**). This regression demonstrated a significant beta coefficient for MAZE PFs, as expected, indicating that there is significant continuity in place-fields across maze exposures. However, it also revealed that POST LTs, but not PRE LTs, impact the PF locations in maze reexposure, as reflected in the significant beta coefficient. Remarkably, when we repeated this analysis for last 4 hours of POST prior to reMAZE, we found no significant contribution from the late POST LTs (**Figure 3.5h**), indicating that our observations do not simply arise from temporal proximity between POST sleep and the maze reexposure, but rather reflect rapid changes in representations that are manifested in the initial hours of sleep. Overall, these results demonstrate the critical role of POST sleep in stabilizing and reconfiguring the spatial representations of hippocampal neurons across exposures to an environment.

3.2.3 Representations during high score preplays do not align with maze place fields

The sequential firing patterns observed during ripples in sleep show a diversity of faithfulness to trajectories through the maze, as captured in the distribution of replay (and preplay) scores (**Figure 3.6**) in the datasets we analyzed using a commonly used weighted-correlation measure, advocated by several studies (Farooq et al., 2019; Grosmark & Buzsáki, 2016; Silva et al., 2015). Here, each event was scored as a percentile compared to its own surrogate distributions generated using a within-ripple-event time-bin shuffle (Davidson et al., 2009) (see Methods). These distributions varied from a uniform distribution expected from chance in all datasets (expected median replay score = 50), not only for MAZE, and POST epochs, but also for PRE, consistent with previous reports (Dragoi & Tonegawa, 2011; Farooq et al., 2019; Grosmark & Buzsáki, 2016). Replay scores during MAZE showed the greatest deviation from chance. While we expected more replay than preplay based on previous reports (Farooq et al., 2019; Grosmark & Buzsáki, 2016; Silva et al., 2015), POST and PRE replay scores were only marginally different in one out of the three datasets we examined (*Grosmark dataset*: PRE: median = 53.8, POST: median = 57.2, PRE vs POST $p = 3.45 \times 10^{-12}$, Mann Whitney U Test (MWUT); *Miyawaki dataset*: PRE: median = 50.8, POST: median = 51.2, PRE vs POST $p = 0.47$, MWUT; *Giri dataset*: PRE: median = 56.6, POST: median = 55.6, PRE vs POST $p = 0.42$, MWUT).

Since replays are considered to simulate trajectories through the maze, we asked whether tunings learned from higher score ripple events in either PRE or POST might show greater fidelity to the maze PFs. We therefore calculated LTs from four subsets of ripples events with replay scores of different percentiles (**Figure 3.6b,c**). We called the tunings learned from the lowest and highest quartiles “low score” and “high score” LTs, respectively. Remarkably, both the low score and high score LTs from MAZE and POST showed strong fidelity to maze PFs, despite the absence

of sequential trajectories in low score events. In contrast, neither high nor low score LTs in PRE showed LTs consistent with the maze PFs (**Figure 3.6b**). LTs from all quartiles of replay scores showed significant PF fidelity in MAZE and POST but not in PRE (**Figure 3.6c**), with somewhat stronger PF fidelity in higher score quartiles (PRE: $\chi^2 = 7.2$, $p = 0.07$; MAZE: $\chi^2 = 143$, $p < 10^{-30}$; POST: $\chi^2 = 150.7$, $p < 10^{-31}$, Friedman test). Likewise, the spatial population vector correlations of low and high score LTs showed a strong correlation with maze PFs for both MAZE and POST epochs, but not for PRE. Overall, these results delineate that even during ripple events in POST with low replay scores, which are typically discarded as non-replays by most measures, the representational structures of neuronal spike trains remain congruent with the place fields on the maze. The firing patterns underlying these events could be detected as “reactivation” using pairwise or ensemble measures (Giri et al., 2019; Kudrimoti et al., 1999; Tingley & Peyrache, 2020), but rather than providing a sequential sweep through space that would be necessary to score high for replay, such ripple events may provide non-continuous, low momentum, or random trajectories through the maze (Krause & Drugowitsch, n.d.; Stella et al., 2019). In contrast, however, even during ripple events in PRE that appear to show sequential structure, the neurons cannot be said to represent the same locations as they do on the maze.

To better understand the dichotomy between PRE and POST ripples, we examined the decoded posterior positions and unit rasters of individual ripple events with high replay scores. In high score ripple events in PRE (e.g. examples shown in **Figure 3.6d**), even though these events appeared to show sequential trajectories leading to high scores, we could not distinguish this sequential structure in the spike trains of units sorted according to their preferred maze locations (such structure was evident in some high score events in POST, e.g. the second two panels). Inspection of individual bins revealed that units which were co-active in time-bins during PRE

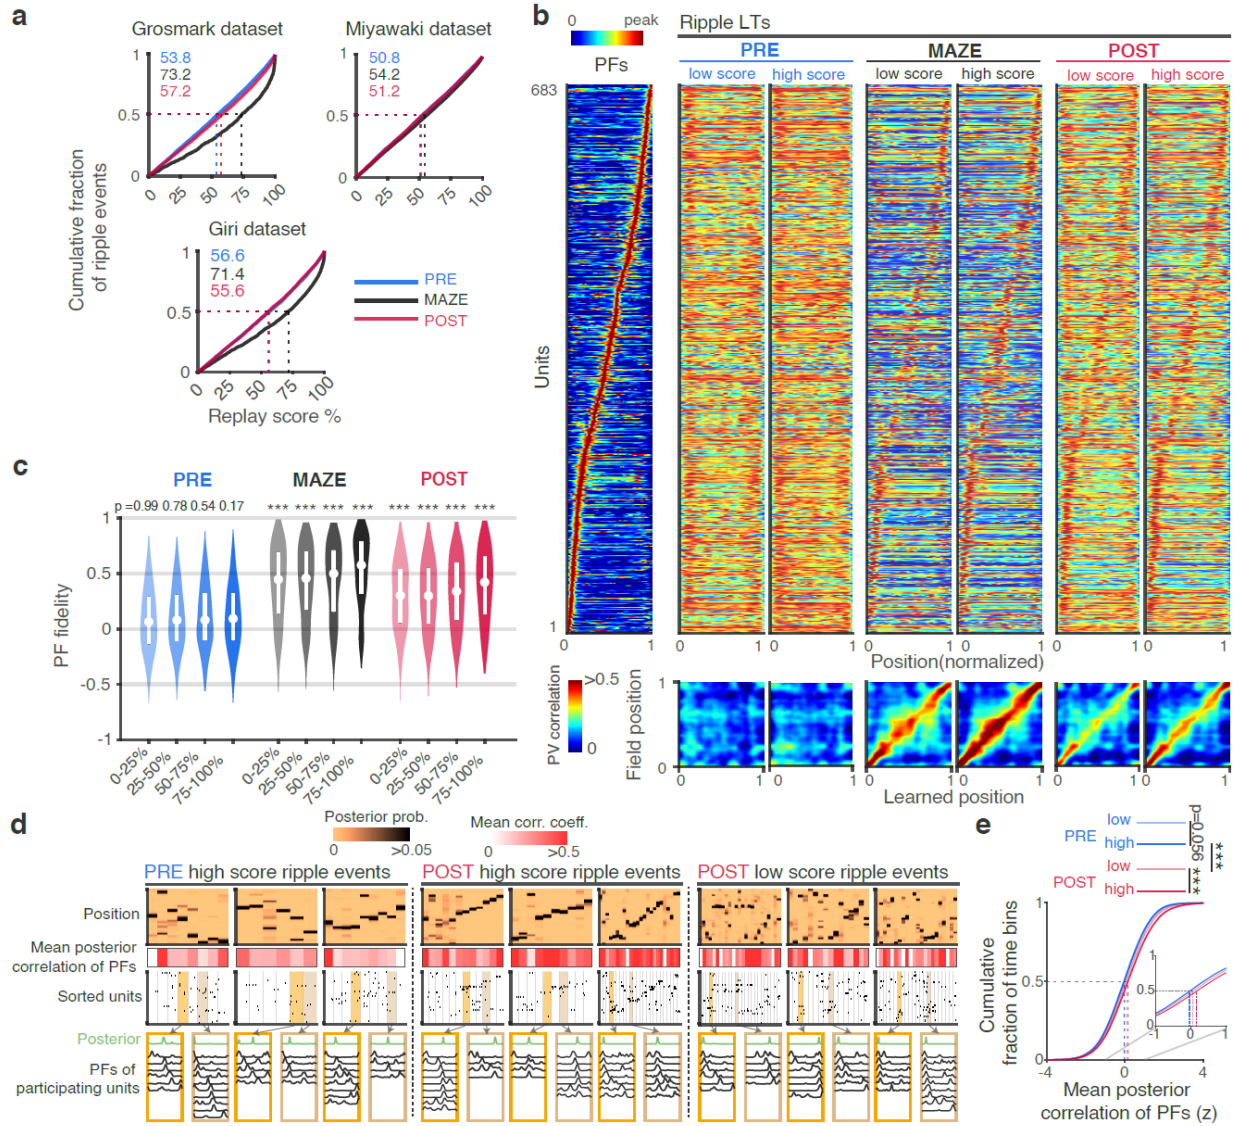


Figure 3.6: Place field fidelities do not strictly correlate with replay score. (a) Distribution of replay scores in the different datasets calculating as percentile against time shuffled bins. Median scores for different epochs are shown with dashed lines (chance median score = 50; see **Methods**). (b) Ripple events were divided into quartiles according to replay score. Top panels show the place fields and sets of LTs calculated based on low and high quartile replay score events within PRE, MAZE, and POST. Bottom panels show population vector (PV) correlations between position bins in the PFs versus different sets of LTs. (c) Distribution of PF fidelity for each ripple subset. Median PF fidelities were significant against surrogate distributions (from unit identity shuffles) in all subsets during MAZE and POST but not during PRE. (d) Place fields of participating units in replays show differing amounts of overlap with the decoded posteriors. Example events with high replay scores in PRE and POST, and low replay scores in POST showing posterior probability matrices and corresponding spike rasters of units sorted by place field order. The middle row depicts the mean correlation of the participating units' place fields with the decoding posterior in each time bin. The bottom panels show the place fields and decoded positions of participating units for example time bins. (e) Mean posterior correlation of PFs and decoded positions show increased place-field overlap in both low and high score replays in POST compared to PRE. Low and high replay score events in PRE did not differ significantly. $***p < 0.001$.

ripple events possessed highly divergent place fields with relatively low mean correlations with the collective posterior probability distributions, contrary to expectations from a unified population code. In contrast, the place fields of units in POST high and low score events showed a greater resemblance (and stronger mean correlation) with the posteriors in their respective bins, indicative of a coherent population representation. Mean posterior correlations of active PFs in bins during all ripple events in PRE and POST similarly showed no significant difference in between low and high score events in PRE, but they were significantly higher in POST relative to PRE (similar results were observed even when we restricted analysis to the subset of events that featured low jump distances (i.e., more continuous trajectories) between time bins (not shown)). The shuffling methods employed to score replays in our and other studies invariably involve assumptions that are violated in real data. These results highlight the importance of verification, as we propose via Bayesian learned tunings, to ensure that decoded positions are in fact consistent with representations of place fields on the maze.

3.2.4 Sleep representations are driven by awake ripples and theta oscillations.

Our findings thus far indicate that the neuronal firing patterns during POST ripples reflect and retune the place-field representations on the maze. We next investigated the factors that conspire to establish these patterns. One recent study (Drieu et al., 2018) reported that, more so than place field activity, the spike patterns of neurons during waking theta oscillations provide the necessary conditions for establishing the firing patterns observed during POST sleep. Another study, however, indicated that waking ripples are a primary mechanistic candidate for generating stable representations (Roux et al., 2017). Adding further complication, several studies have indicated that PRE and POST ripples share overlapping activity structure (Giri et al., 2019; Liu et al., 2019), suggesting limits on the flexibility of sleep representations. To better understand the

respective contributions of these different factors on the representations in POST sleep, we performed a multiple regression to test the extent to which POST LTs are explained by: PRE LTs, MAZE PFs, LTs of MAZE theta periods, and LTs of MAZE ripples (**Figure 3.7a**). Remarkably, we found that the beta coefficients for all of these regressors were significant. While the beta coefficient for MAZE theta LTs was significant, indicating that waking theta oscillations are important for the formation of ensemble representations, in support of the previous report (Drieu et al., 2018), MAZE ripple LTs had the largest beta coefficient, indicating that firing patterns during waking ripples on the maze have the most lasting impact on POST sleep activity. Surprisingly, the second largest beta coefficients were observed for PRE ripple LTs, indicating that next to MAZE patterns, patterns configured in PRE also provide an important determinant of POST sleep activity (Grosmark & Buzsáki, 2016; Liu et al., 2019). Consistent with this, we observed a significant correlation between the PF fidelity in PRE and the PF fidelity in POST (**Figure 3.7b**).

These observations suggest that despite the absence of maze tuning in PRE sleep, some cells maintain similar representations between PRE and POST. Sleep similarity, which measures the consistency of LTs across PRE and POST by assessing the correlation between PRE LTs and POST LTs, was significantly correlated with PF fidelity in PRE (**Figure 3.7c**); thus, PRE LT's that aligned with maze PFs, presumably by chance, maintained those LTs in POST. On the other hand, sleep similarity showed only a weak negative correlation with the PF fidelity in POST. To better understand the difference between PRE and POST LTs, we separated units into “PRE-tuned” cells (PRE PF fidelity > 0), and “PRE-untuned” cells (PRE PF fidelity < 0). PRE-tuned cells showed generally high POST PF fidelity along with high sleep similarity (**Figure 3.7e**). In contrast, PRE-untuned cells showed a significant negative correlation between sleep similarity

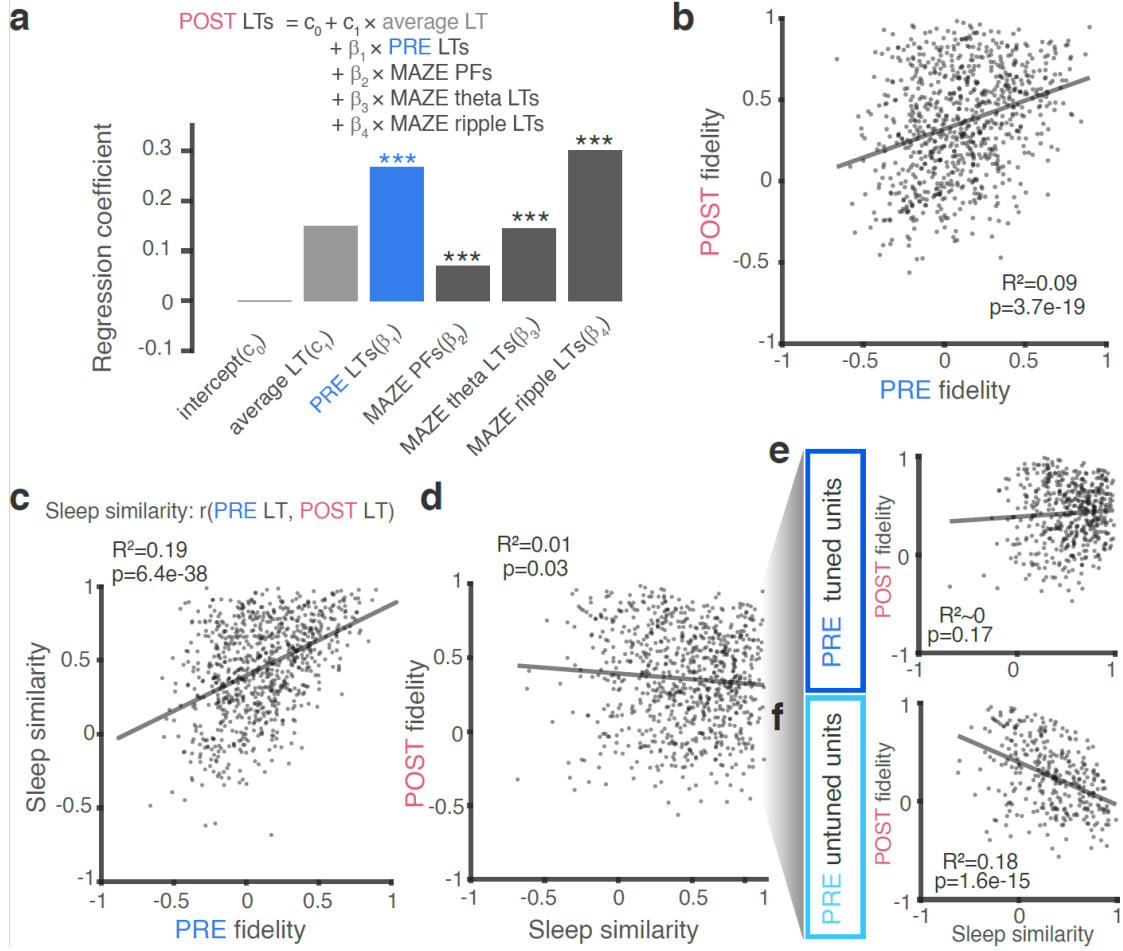


Figure 3.7: A diversity of pre-existing representations impacts the tunings in POST sleep. (a) Multiple regression analysis for estimating the dependence of POST LTs on PRE LTs, MAZE PFs, MAZE theta LTs, and MAZE ripple LTs shows that POST LTs were most significantly impacted by PRE LTs and MAZE ripple LTs (RMSE = 0.75, $R^2 = 0.43$, $p < 10^{-4}$, $c_0 = -0.0002$, $c_1 = 0.15$, $\beta_1 = 0.27$, $p < 10^{-4}$, $\beta_2 = 0.07$, $p < 10^{-4}$, $\beta_3 = 0.15$, $p < 10^{-4}$, $\beta_4 = 0.30$, $p < 10^{-4}$, p-values against surrogate distributions from 10000x unit-identity shuffles). (b) PF fidelity (correlation with MAZE PF) was significantly correlated between PRE and POST LTs. (c) Sleep similarity (correlation between PRE and POST LTs) was correlated with PRE PF fidelity, indicating that high fidelity PRE LTs are preserved in POST. (d) An overall weak negative correlation between sleep similarity and POST PF fidelity. When we split units into PRE-tuned (PRE PF fidelity > 0) and PRE-untuned PFs (PRE PF fidelity < 0), (e) there was little correlation between sleep similarity and POST fidelity for PRE-tuned cells. (f) For PRE-untuned cells, a negative correlation between POST PF fidelity and sleep similarity indicates a continuum of flexible retuning to maze PF.

and POST fidelity (**Figure 3.7f**); those with high sleep similarity were poorly tuned in POST,

while those that reconfigured from PRE to POST, showed better fidelity to maze PFs. These

analyses therefore reveal the contribution of PRE sleep to maze representations and POST

activities; cells whose representations are already aligned with maze place fields in PRE

maintain those same representations in POST, but other neurons display a broad range of flexible

reconfiguration that is inversely proportional to their rigidity (Grosmark & Buzsáki, 2016) across PRE and POST.

3.3 Discussion

The observations made possible by the Bayesian learning of spatial tuning have important implications for our understanding of how learning and sleep impact the place-field representations of hippocampal neurons. First, we found that spatial representations emerge rapidly upon exposure to a novel environment, but not beforehand. While ripple events during pre-exposure occasionally scored highly for replays, spatial representations were not coherent among active neurons during these periods, as cells with very divergent place fields often fire within the same time bins (**Figure 3.6d**). These observations suggest that continuous patterns in the decoded posteriors of spike trains could emerge spuriously. Consistent with this notion, it has been noted that the measures and shuffles used to quantify replays inevitably introduce unsupported assumptions about the nature of spontaneous activity (Farooq et al., 2019; Silva et al., 2015; Van Der Meer et al., 2020). We propose that only for those periods and events in which there is strong correspondence between the Bayesian learned tunings and neurons' place-fields, can the collective ensemble activity be considered to provide trajectories through internally generated representations of space (Diba, 2021; Van Der Meer et al., 2020).

Among the brain states we examined, sharp-wave ripples in early sleep offered the representations that best aligned with the place fields on the maze. These early sleep representations were in turn influenced by a confluence of factors, including carryover from PRE sleep (in both PRE-tuned and PRE-untuned units) (Farooq et al., 2019). Most notably, however, our analysis revealed a role in awake activity patterns during theta oscillations, and more

prominently, those during sharp-wave ripples in generating the ensemble coordination that underlies spatial representations during sleep. These observations are consistent with the hypothesis that an initial cognitive map of space is first laid down during theta oscillations (Dragoi & Buzsáki, 2006; Drieu et al., 2018; Monaco et al., 2014), then stabilized and continuously updated by awake replays based on the animal's (rewarded and/or aversive) experiences on the maze (Sen Cheng & Frank, 2008; Diba, 2021; Gupta et al., 2010; Mattar & Daw, 2018; Roux et al., 2017). Once ensembles are established, they reactive during early part of sleep (Giri et al., 2019; Wilson & McNaughton, 1994). However, sleep representations were not always exact mirror images of the maze place-fields, and our Bayesian learning approach allowed us to measure those deviations for individual neurons. Remarkably, we found that these early-sleep ripple representations proved predictive of place fields on re-exposure to the maze. Based on these observations, we propose that representational drift in fact arises rapidly from retunings that take place during early sleep replays rather than noisy deviations that develop spontaneously over time. Furthermore, we conjecture that hippocampal replay during sleep does not play a passive role in simply recapitulating the patterns already seen during learning but represents a key optimization process generating and integrating new spatial tunings within the recently formed spatial maps.

Overall, representations remained stable and consistent with the maze for hours of sleep in POST, despite the absence of strong sequential replay trajectories during ripples in POST sleep. Reconciling observations based on studies that measure neuronal reactivation using pairwise or ensemble measures with those that focus on trajectory replays has until now represented a challenge to the field (Tingley & Peyrache, 2020). Our study consolidates these views by demonstrating that faithful representations, which are consistent with pairwise and ensembles

measures of reactivation, persist for hour-long durations. However, the trajectories produced by these cell ensembles do not necessarily provide continuous high-momentum sweeps through the maze environment (Krause & Drugowitsch, 2021; Stella et al., 2019), as we found high fidelity spatial tunings even among low replay score ripple events in POST. Instead, trajectories simulated by the hippocampus during sleep may explore pathways that were not directly experienced during waking but can serve to better consolidate a cognitive map of space. Additionally, we found increasing instability and drift in the spatial representations of neurons over the course of sleep, indicating that late sleep, like PRE, features more randomized activity patterns (Genzel et al., 2014; Miyawaki & Diba, 2016). It is also worth noting that we did not find alignment between maze place fields and learned spatial tunings during REM sleep. It may be that using a different behavioral setup such as with frequently repeated maze exposures (Louie & Wilson, 2001) or salient fear memories (Boyce et al., 2016), or limiting analyses to specific phases of theta (Poe et al., 2000), we might have uncovered evidence more consistent with dream-like replays of maze place-fields (Hobson, 2009). On the other hand, it is also worth noting that the bulk of dreams do not reprise awake experiences (Vertes, 2004). The randomization of representations, as we see during REM and late stages of slow-wave sleep, may reflect an important function of sleep, driving activity patterns from highly-correlated ensembles to those with greater independence (Colgin et al., 2009; Norimoto et al., 2018), which may be important for resetting the brain in preparation for new experiences (Cirelli & Tononi, 2022). In sum, Bayesian learning provides a powerful means of tracking the stability and plasticity of representational tuning curves of neurons over time, and can be readily extended to investigate the dynamics of internally-generated representations in other systems during both sleep and

awake states, including within rehearsal, rumination, or episodic simulation (Schacter et al., 2008).

3.4 Methods

In this work, we used two previously published datasets along with new recordings. The existing datasets were also used in our previous publication (Giri et al., 2019) and described in detail there. In brief, these were recordings of unit and local field potential from the rat hippocampus CA1 region, with PRE rest and sleep, exposure to a novel MAZE, and POST rest and sleep: the Miyawaki dataset (3 rats, 5 sessions; PRE, MAZE, POST, each ~ 3 hours) (Giri et al., 2019; Miyawaki & Diba, 2016) and the Grosmark dataset (4 rats, 5 sessions; PRE, and POST, each ~ 4 hours and MAZE, ~ 45 minutes) (Grosmark & Buzsáki, 2016).

3.4.1 Behavioral task and data acquisition

For the new recordings, we trained four water-deprived rats to alternate between two water wells in a previously habituated home box. Water rewards during the alternation were delivered via water pumps interfaced with custom-built Arduino hardware. After the animals learned the alternation task, they were surgically implanted by 128 channel silicon probes (8 shanks, Diagnostic Biochips, Glen Burnie, MD) either unilaterally (one rat) or bilaterally (three rats) over the dorsal hippocampal CA1 subregions (AP:-3.36 mm , ML: \pm 2.2 mm). Following recovery from surgery, the probes were gradually lowered over a week to the CA1 pyramidal layer, which was identified by sharp wave-ripple polarity reversals and frequent neuronal firing. After ensuring recording stability, the animals were exposed to novel linear tracks during one (three rats) or two (one rat) behavioral sessions (in total five sessions from the four rats). During each session, the implanted animal was first placed in the home box (PRE, ~ 3 hours) with ad

libitum sleep (during the dark cycle). Then, the animal was transferred to a novel linear track with two water wells that were mounted on platforms at either end of the track (MAZE, ~ 1 hour). After running on the linear track for multiple laps for water rewards, the animal was returned to the home box (aligned with the start of the light cycle) for another ~10 hours of ad libitum sleep (POST). In four of these sessions, following POST the rats were re-exposed to the same linear track for another ~ 1h of running for reward (reMAZE).

Wideband extracellular signals were recorded at 30 kHz using an OpenEphys board (Siegle et al., 2017) or an Intan RHD recording controller during each session. The wideband activity was high-pass filtered with a cut-off frequency of 500 Hz and thresholded at five standard deviations above the mean to extract putative spikes. The extracted spikes were first sorted automatically using SpykingCircus (Yger et al., 2018), followed by a manual passthrough using Phy (Rossant et al., 2016)(<https://github.com/cortex-lab/phy/>). Only units with less than 1% of total number of spikes in their refractory period (based on the units' autocorrelograms) were included in further analysis. Putative neurons were classified into pyramidal and interneurons based on peak waveform shape, firing rate, and interspike intervals (Bartho et al., 2004; Petersen et al., 2021). For analysis of local field potentials (LFP, 0.5-600 Hz), signals were filtered and downsampled to 1250 Hz.

The animal's position was tracked using an Optitrack infrared camera system (NaturalPoint Inc, Corvallis, OR) with infrared-reflective markers mounted on a plastic rigid body that was secured to the recording headstage. 3D position data was extracted online using Motive software (Optitrack), sampled at either 60 Hz or 120 Hz, and later interpolated for aligning with the ephys data. Although, we attempted to track the animal's position during each entire session, including

in the home cage, the cage limited visual access from our fixed cameras. Additionally, in one session the position data for reMAZE was lost during the recording.

3.4.2 Place field calculation

To calculate place fields, we first linearized the position by projecting each two-dimensional track position onto a line that best fit the average trajectories taken by the animal over all traversals within each session. The entire span of the linearized position was divided into 2 cm position bins and the spatial tuning curve of each unit was calculated as occupancy-normalized spike counts across the linearized position bins.

3.4.3 LFP analysis and brain state detection

We estimated a broadband slow wave metric using the irregular-resampling auto-spectral analysis (IRASA) approach (Wen & Liu, 2016), following code generously shared by Dan Levenstein and the Buzsaki lab (<https://github.com/buzsakilab/buzcode>). This procedure allows estimation of the slope of the power spectrum ($1/f^\beta$, f : frequency) which we used to estimate slow-wave activity. The slow-wave metric for each session followed a bimodal distribution with a dip that provided a threshold to distinguish NREM (non “rapid eye movement” sleep) from other periods. A time-frequency map of the LFP was also calculated in sliding 1s windows, step size of 0.25 s, using the Chronux toolbox (Bokil et al., 2010). To identify high theta periods, such as during active waking or REM sleep (Buzsaki, 2002; Miyawaki & Diba, 2016), the theta/non-theta ratio was estimated at each time point as the ratio of power in theta (4-9 Hz in home cage and 6-11 Hz on the linear track) to a summation of power in delta frequency band (1-4 Hz) and the frequency gap between the first and second harmonics of theta (10-12 Hz during home cage awake and REM epochs and 11–15 Hz during MAZE). To calculate the ripple power,

multichannel LFP signals were filtered in the range of 150-250 Hz. The envelope of the ripple LFP was calculated using the Hilbert transform, z-scored and averaged across the channels. Only channels with the highest ripple power from each electrode shank were used in the averaging.

3.4.4 Detection of ripple events

For each recording session, multi-unit firing rates (MUA) were calculated by binning the spikes across all recorded single units and multi-units in 1 ms time bins. Smoothed MUA was obtained by convolving the MUA with a Gaussian kernel with $\sigma = 10 \text{ ms}$ and z-scoring against the distribution of firing rates over the entire session. Ripple events were first marked by increased MUA firing, periods when the smoothed MUA crossed $2z$ and the boundaries were extended to the nearest zero-crossing time points. The ripple events that satisfied the following criteria were considered for further analysis: a) duration between 40 and 600 milliseconds, b) occurrence during NREM or quiet waking period, c) concurrent speed of the animal below 10 cm/sec (when available), and d) concurrent ripple power in the LFP higher than 1 s.d. above the mean. To detect the ripple events during the quiet waking periods, we required either a theta-delta ratio < 1 or ripple power > 3 s.d. of the mean at the time of candidate event. All ripple events were subsequently divided into 20 millisecond time bins. The onsets and offsets of the events were adjusted to first time bins with at least two pyramidal units firing. We split ripples with silent periods $> 40 \text{ ms}$ into two or more events.

3.4.5 Bayesian learned tunings

Bayesian learning of spatial tunings or learned tunings (LTs) can be conceptualized as using offline spiking activity to update our estimate of each unit i 's spatial tuning $p(s_i|x)$, which is the probability of spiking conditioned on the network's internal estimate of position, x .

Conceptually, outside of the maze, the best initial estimate available for a cell's place preferences is its place field on the maze. Then, we use Bayesian learning (Wiskott, 2013) to update this tuning based on information available in the spike trains during the epoch of interest. We note that this approach still relies on the place-fields of neurons as measured on the maze. However, it provides a degree of separation in that a given neuron's Bayesian LT does not depend directly on its own maze place-field, but rather on the coherent firing of that neuron with the other neurons in the ensemble.

Effectively, the unit i 's LT is learned from the posterior probability distributions $p(x|s_{\forall j \neq i})$ determined from the other units' spikes for the time windows in which unit i spikes. The posterior probability distribution, $p(x|s_{\forall j \neq i})$ within individual time bins is calculated based on other units, excluding the spikes and track tuning of unit i itself, under the assumption of conditional independence among hippocampal neurons conditioned on maze position and running direction.

Since majority of the sessions (16 out of 17) consisted of two running directions on the track, we first calculated the posterior joint probability of position and travel direction and then marginalized the joint probability distribution over travel direction (Davidson et al., 2009):

$$p(x, d|s_{\forall j \neq i}) \propto p(s_1, s_2, \dots, s_{i-1}, s_{i+1}, \dots, s_n | x, d) \quad (1)$$

in which n is the total number of units and d signifies the travel direction. With the assumption of independent Poisson-distributed firings of individual units conditioned on maze position and direction, equation (1) is equal to:

$$p(x, d|s_{\forall j \neq i}) \propto \prod_{j \neq i} (f_j(x, d) \tau)^{s_j} e^{-f_j(x, d) \tau} \quad (2)$$

In equation (2), $f_j(x, d)$ characterizes the mean firing rate of unit j at position bin x and direction d and τ is the bin duration used for decoding, which was chosen = 20 ms in our analyses. By marginalizing the left hand side of equation (2) over direction d :

$$p(x|s_{\forall j \neq i}) \propto \sum_d p(x, d|s_{\forall j \neq i}) \quad (3)$$

Then, the Bayesian learned tuning for unit i was calculated as:

$$p'(s_i|x) = \sum_k k * p(s_i = k|x) \quad (4)$$

in which

$$p(s_i = k|x) = p(x|s_{\forall j \neq i} \& s_i = k) * p(s_i = k)/p(x) \quad (5)$$

Where $p(x|s_{\forall j \neq i} \& s_i = k)$ is the average posterior probability of track position, x , over time bins in which the unit i fired k spikes, calculated using only units $\neq i$. $p(s_i = k)$ is the overall proportion of time bins where unit i fired k spikes, and $p(x)$ is average posterior over all time bins.

The equation (4) for calculating the LTs was simplified to the unit i 's spike-triggered average of the posterior probability distributions:

$$p'(s_i|x) = \frac{\sum_t k_t * p_t(x|s_{\forall j \neq i})}{\sum_t p_t(x|s_{\forall j \neq i})} \quad (6)$$

in which t is the time bin and k_t is the number of spikes that unit i fired at t .

Additional requirements for potential unit clustering confounds: To avoid potential confounds from spike misclassification of units detected on the same shank (Quirk & Wilson, 1999), we placed additional inclusion requirements for LT calculations. We determined the L-ratios (Schmitzer-Torbert et al., 2005) between the unit i and each other unit recorded on the same shank, yielding the cumulative probability of the other units' spikes belonging to the unit i . Since the range of L-ratio depends on the number of included channels, to provide a consistent threshold for all datasets, the L-ratio for each pair was calculated using the four channels that featured the highest spike amplitude difference between each pair of units. Only units with L-ratio $> 10^{-3}$ (**Figure 3.8**) were used to calculate LTs for each cell.

3.4.6 Fidelity of the learned tunings across epochs

To quantify the degree to which tuning curves, $p(x|s)$, in LTs or PFs relate across epochs, we used a simple Pearson correlation coefficient of the tuning curves across position bins between the LTs/PFs. We obtained consistent results with a measure based on the Kullback-Leibler divergence (not shown). The median for each epoch were compared against a surrogate distribution of such median values obtained by shuffling the unit identities of the PFs separately within each session. In other words, we tested against the null hypothesis that learned tunings in each session may have trivial correlations with PFs. For each epoch we obtained p-values based on the number of such surrogate median values that were \geq those in the original data.

3.4.7 Learned tuning's dynamics

We further evaluated the dynamics of LTs across time in non-overlapping 15 min windows (for illustration purposes only in Figure 3.3, we used overlapping 15 min windows with a 5 min step size). A unit's LT stability was defined as the median Pearson correlation coefficient between

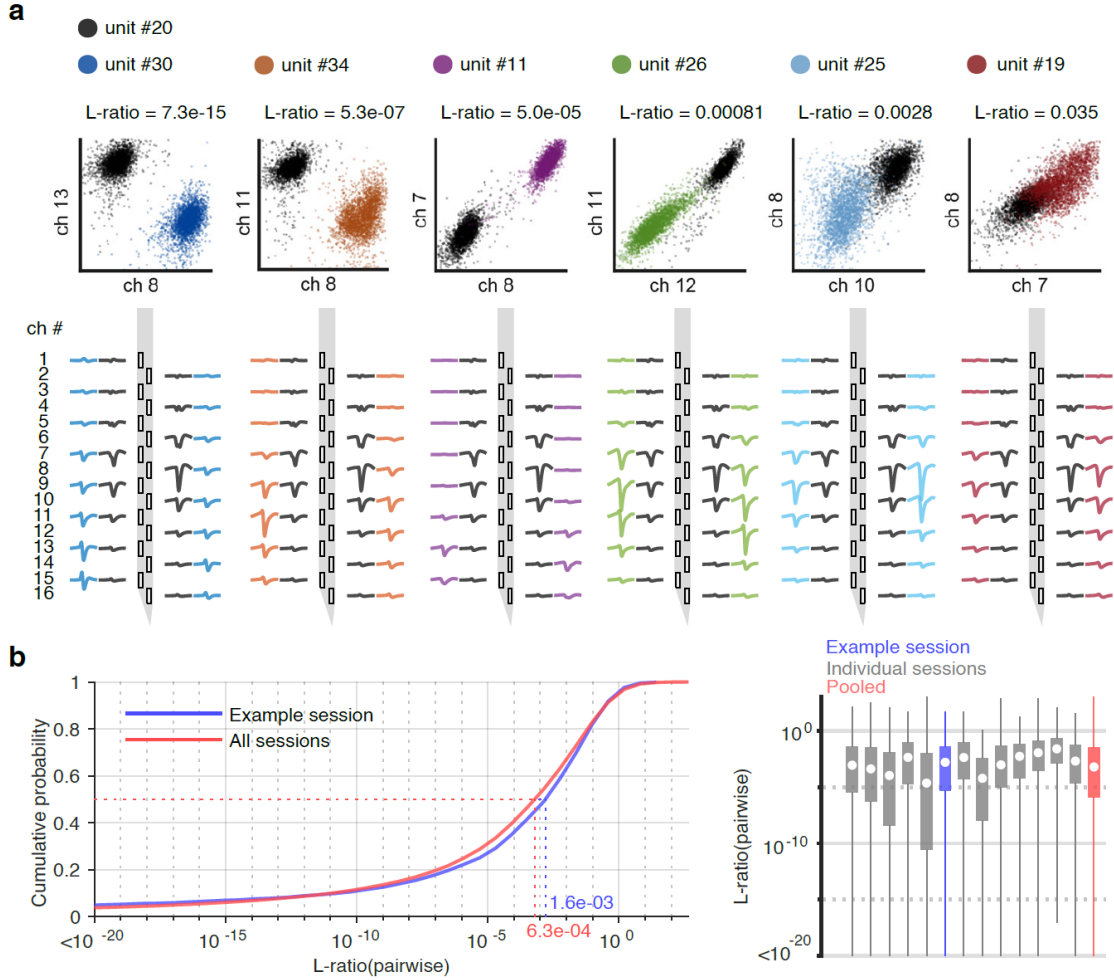


Figure 3.8: A L-ratio measure to quantify the degree of overlap in the spike feature space between pairs of units. (a) Each scatterplot on the top shows the spikes of the example unit #20 (black dots) and another unit (colored dots) recorded on the same shank in an example recording session from the Gari dataset. The axes in each scatterplot correspond to the spike amplitude on two channels with maximal distinction between the pairs. There is a range of overlap with unit #20; for example, unit #30 on the leftmost inset showed almost no overlap, while unit 19 on the rightmost inset significantly overlapped. The overlap was quantified as an L-ratio on the top of each inset. The L-ratio between unit #20 and any other unit was obtained by calculating the probability of spikes from the second unit belonging to unit 20. The insets on the second row show the comparison between the mean spike amplitudes for each unit pair as plotted in the scatterplots. (b), The cumulative distribution of L-ratio for the example session and pooled data across all sessions (left), and for each individual session (right).

that unit's LTs in all different pairs of time windows within a given epoch. Thus, units that had stable and consistent LTs across an epoch yield higher correlations in these comparisons than those with unstable LTs. These unit LT stability values were z-scored against a null distribution of median correlation coefficients based on randomizing the LTs' unit identities within each 15 min time window (1000 unit ID shuffles). Normalized stability correlation matrices in **Figure**

3.4c were calculated by z-scoring each correlation coefficient against a surrogate distribution based on shuffling the LTs' unit identities. To investigate the changes in POST LT stability over time in **Figure 3.4c**, we calculated LT stability within overlapping 2-hour blocks with a step size of one hour.

3.4.8 Ripple event replay scores

The posterior probability matrix (P) for each ripple event was calculated based on previously published methods. Replays were scored using the absolute weighted correlation between decoded position (x) and time bin (t) (Silva et al., 2015):

$$corr(t, x; p) = \frac{cov(t, x; P)}{\sqrt{cov(t, t; P) cov(x, x; P)}} \quad (7)$$

$$cov(t, x; P) = \frac{\sum_i \sum_j P_{ij} (x_j - m(x; P)) (t_i - m(t; P))}{\sum_i \sum_j P_{ij}} \quad (8)$$

$$m(x; P) = \frac{\sum_i \sum_j P_{ij} x_j}{\sum_i \sum_j P_{ij}} \quad m(t; p) = \frac{\sum_i \sum_j P_{ij} t_i}{\sum_i \sum_j P_{ij}} \quad (9)$$

in which i and j index time bin and position bin, respectively.

Each replay score was further quantified as a percentile relative to surrogate distributions obtained by shuffling the data according to the commonly used within-event time swap, in which time bins are randomized within each ripple event (Davidson et al., 2009). We preferred this method over the circular spatial bin shuffle (or column cycle shuffle, also described by (Davidson et al., 2009), as it preserves the distribution of peak locations across time bins within each event (see also related discussion in (Farooq et al., 2019)). Each ripple event was assigned

to one of four quartiles based on the percentile score of the corresponding replay relative to shuffles

3.4.9 Place fields' overlap with decoded posterior

A Pearson correlation coefficient was calculated between the PF of each unit firing (participating) in a time-bin and the posterior probability distribution for that bin based on the firings of all units. The mean posterior correlation of PFs was calculated over all participating units. Since this mean posterior correlation might be inflated when there is a low number of participating units, for each time bin with firing unit count n we generated surrogate distribution of mean posterior correlation by randomly selecting n units. Then, the mean posterior correlation in the original data was z-scored against the corresponding surrogate distribution for n randomly participating units.

3.4.10 Multiple regression analyses

To examine the extent to which a spatial tuning curve (LT or PF) within a given epoch was impacted by the tuning curves in other epochs, we performed multiple regression analyses. We modeled POST LTs and reMAZE PFs using the following equations:

$$POST\ LTs = c_0 + c_1 * average\ LT + \beta_1 * PRE\ LTs + \beta_2 * MAZE\ PFs + \beta_3 * MAZE\ theta\ LTs + \beta_4 * MAZE\ ripple\ LTs \quad (10)$$

$$reMAZE\ PFs = c_0 + c_1 * average\ LT + \beta_1 * PRE\ LTs + \beta_2 * MAZE\ PFs + \beta_3 * POST\ LTs \quad (11)$$

The dependent variables and regressors were calculated over all position bins from all units. The average LT in the analyses were calculated by averaging all unit LTs over PRE and POST. c s and β s are the regression coefficients.

In order to test the statistical significance of the regression R^2 values and each regression β coefficient, we compared these against distributions of surrogates ($n = 10,000$) which were calculated by randomizing the unit identities of the dependent variable's tuning curves. For each coefficient and R^2 values, we obtained a p-value based on the number of surrogates that were \geq those in the original data.

4 Conclusion and future direction

We developed two approaches to analyze and interpret neural activities during offline epochs, including slow wave sleep after an experience. The communications between the hippocampus and cortex underlying memory processing mainly occur during these periods, especially hippocampal SWRs. Therefore, to understand the offline memory processing, interpreting the spatiotemporal dynamics of neural activity during the SWRs is essential. It cannot be assumed that the activity patterns during sleep, especially the spontaneous activity during sleep prior to the experience, represent wake neural activity. Therefore, we developed approaches relaxing the assumption that individual neuron and ensemble activities feature the same tunings formed during the behavior.

Hidden Markov models to uncover temporal structures in SWRs. We developed a data-driven method solely based on modeling the SWR-associated PBEs in terms of transitions between a set of distinct states. The only assumption is that temporal structures remain consistent over many PBEs. We validated this approach using quiet wake PBEs during the animal’s exploration when sequence replay is strongly expressed. During these events, the states mainly represent the animal’s position, providing a ground truth to assess the model’s performance. We found that SWR population firings correspond to a clear trajectory in a state space, with many states representing the animal’s distinct positions. We were able to detect individual replay events using a cross-validation approach with a performance comparable to a standard Bayesian decoding method. Furthermore, we found that this method can extract temporal structure pertaining to remote experiences of the animal in addition to one corresponding to the local environment.

As a future direction, we can use this hidden Markov model-based approach to study temporal structures during sleep after a novel experience and compare it with pre-maze sleep. Results of fitting HMMs to neuronal population firing patterns within each behavioral epoch (PRE, MAZE, or POST) for an example session are as follows. First, we investigated whether there is a consistent temporal structure across PBEs within each epoch. The PBEs within all behavioral epochs, including PRE, showed higher than chance congruence with the HMMs trained on the same epoch (cross-validation, **Figure 4.1a**). Second, we asked whether POST's structure replays MAZE and whether there is a carryover from PRE during MAZE and POST. We investigated these by training HMMs on each epoch and calculating the congruence score of PBEs in other epochs. The POST PBEs showed high congruence with the MAZE model, confirming POST replay of MAZE temporal structure (**Figure 4.1b**, middle). There was no significant preplay during MAZE, as evident by close to the chance level congruence of MAZE PBEs with the PRE model (**Figure 4.1b**, left). On the other hand, there was a carryover from PRE in POST (**Figure 4.1b**, left), suggesting that PRE and POST share somewhat similar temporal structures, as we earlier showed in **Figure 3.7**.

Does sleep allocate a distinct group of SWRs during POST to process and protect newly encoded information from possible interference due to the pre-existing processing/prior experiences or otherwise use overlapping SWRs to integrate the recent information with the pre-existing structure? To investigate this, we asked if distinct or overlapping groups of POST PBEs are consistent with either MAZE or PRE models. Our results showed that although many POST PBEs showed perfect congruence with both PRE and MAZE models, supporting the integration hypothesis, many others differentiated between the two (**Figure 4.1c**). We need to analyze more sessions to test each hypothesis.

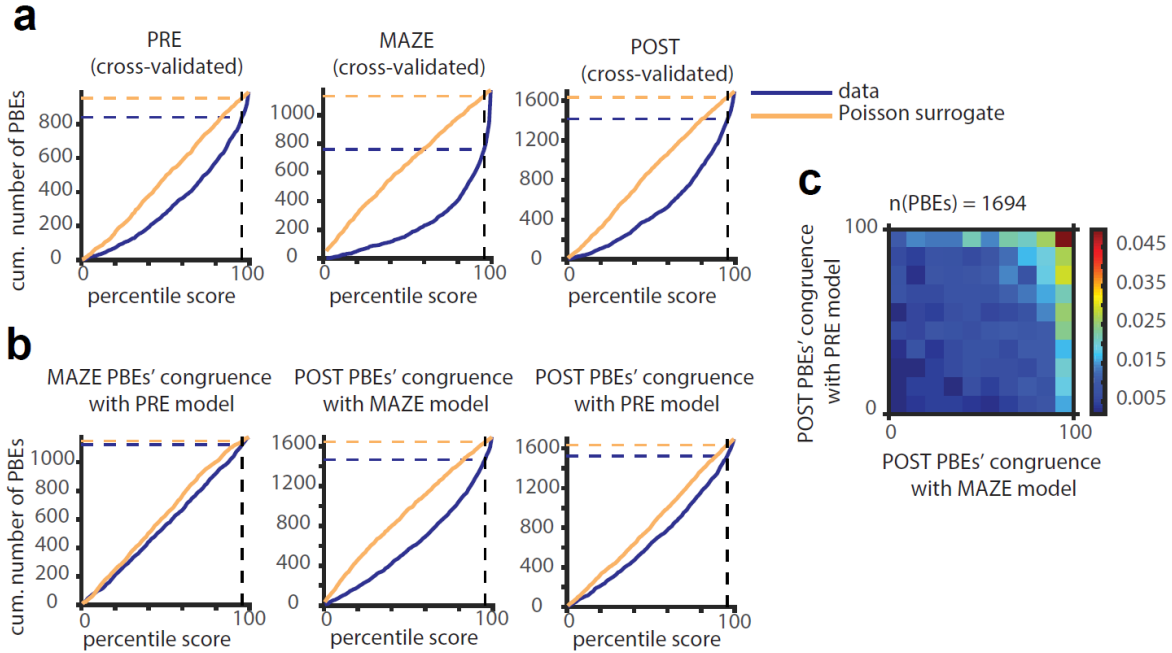


Figure 4.1: Sequential content during PBEs within each behavioral epoch. (a) The cumulative number of PBEs' congruence with HMMs trained on non-overlapping set of PBEs within the same epoch (5-fold cross-validation). The distribution from Poisson surrogate serves as a chance level. **(b)** The same as (a) but for congruence of PBEs within an epoch with a model trained on another epoch. **(c)** Joint histogram of POST PBEs' congruence with HMMs trained on PRE or MAZE.

We further calculated the latent state place fields (IsPFs) by decoding the states within each 100 ms of populations firing during animal's running on the maze and calculating each state's tendency to be decoded at each maze location (**Figure 4.2**). We found that for an example session, IsPFs during POST but not PRE represented specific maze locations. This finding suggests that location-depicting neuronal assemblies (as distinct states in HMMs) are formed upon the novel experience and are distinct from the pre-existing neuronal assemblies.

Bayesian learning of spatial tunings. We used principles of Bayesian learning to calculate spatial tunings during sleep via updating the maze place fields using ensemble firings during sleep. We conjectured that since the space is represented by ensembles formed during exploration, possible modifications of these ensembles during sleep would reflect in changes in individual neurons' spatial tunings. We found that during SWRs in early post-maze sleep/rest, learned tunings

showed diverse stability and consistency with the maze place fields. Furthermore, the learned tunings and their deviation from the place fields predicted remapping of the place fields when the animal explored the same environment after a ~9 hours delay from the first exposure. These findings suggested that sleep after a novel experience actively stabilizes or retunes the place fields, impacting the spatial representation upon re-exposure. On the other hand, the pre-maze learned tunings were inconsistent with the upcoming maze place fields, suggesting that ensembles form de novo upon experiencing a novel environment.




Figure 4.2: The states during POST but not PRE represented maze locations. Left, IsPFs of example HMM states in PRE. Right, the same as left panel but for POST.

The findings of this project have implications for evaluating and revising methods to measure extended replays during sleep. Given that the spatial tunings during the pre-sleep are not predictive of upcoming maze place fields, the preplay reported in previous studies might be spurious. A replay score for each SWR event equals the degree to which the sequence of decoded positions in successive snapshots of population activity (20 millisecond time bins) within the event resembles the animal's trajectories on the maze. To evaluate the statistical significance of the replay scores, they are compared against distributions generated using null hypotheses. Under these null hypotheses, the incidence of replay within an event is equally likely as surrogates with removed temporal structure but otherwise intact features like the event's firing

characteristics. Different methods have been used interchangeably to calculate the replay scores and generate the null distributions (**Figure 4.3**). Although most of these methods perform well in detecting replay during intermittent quiet wake periods (when the replay is less noisy), there was a striking lack of correlation between these methods in scoring sleep SWRs. Therefore, some of these methods often result in false detection of sleep replay (in POST) and preplay (in PRE) contents. We think this problem is due to the wrong presumption that the place fields remain consistent during sleep. As we showed in **Figure 3.6d**, place field fidelity of the learned tunings relates to the degree of overlap between the maze place fields of the cofiring neurons. We observed a lack of significant overlap during many SWRs, especially those during PRE. Therefore, we conjectured that non-replay events would not stand out compared to surrogate events in which neurons with divergent place fields cofire. We suggest that the unit-identity shuffle would generate the optimal null distribution to test the statistical significance of replay scores for individual SWRs.

To find the most reliable replay scoring approach, we need to characterize false detection rate corresponding to each of these approaches. Since there is no ground truth replay score, we plan to conduct a survey in which we ask human subjects to score a set of SWRs, one at a time, in terms of how likely they believe that each event looks a “true” replay.

We present the subjects with images (or heatmaps) of the decoded position, like the one in **Figure 4.3a**, but without access to the spike rasters or any additional information. The subjects are asked whether they perceive a line that goes from one corner to another corner of the image (a trajectory). They will provide a score for each SWR event in a range from zero to nine, based on the quality of the line that they perceive. A zero is equal to absence of a clear line and a nine is equal to a continuous line spanning the whole image. Events with partly continuous line

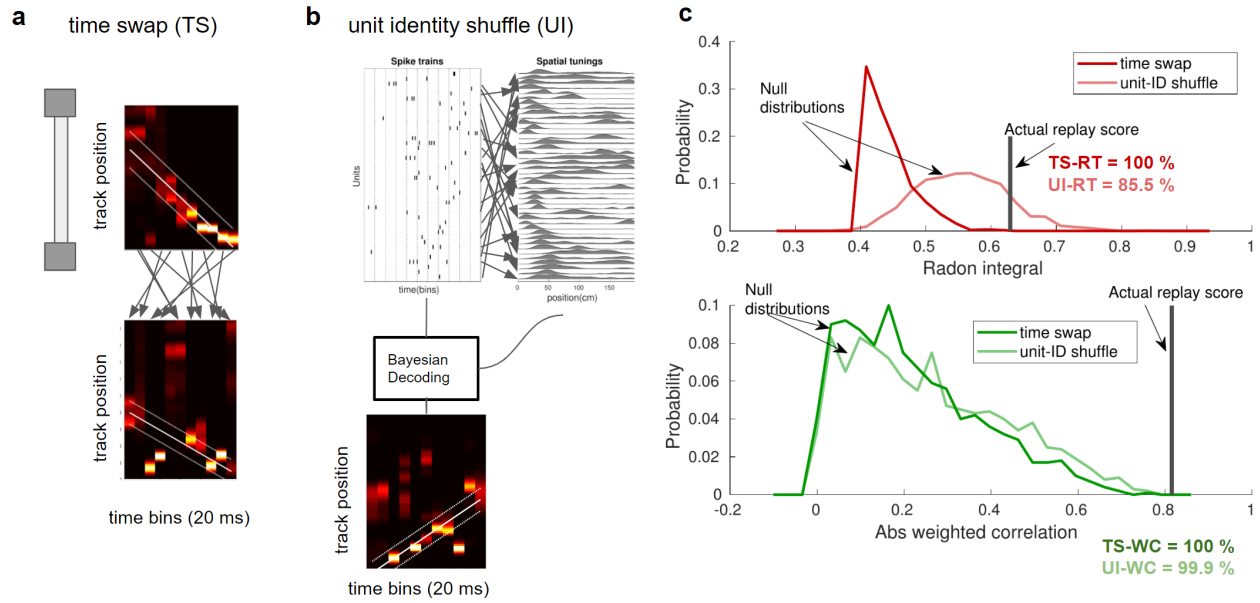


Figure 4.3. Examples of methods to calculate replay scores and their statistical significance. (a, b) Two methods for generating null distributions: **(a)** Time-swap (TS) PBEs in which time bins within the example PBE are randomly swapped. **(b)** Unit identity (UI) shuffle in which the place fields are randomly reassigned to neurons, which changes the decoded positions given intact PBE firings. Replay scores in (a) and (b) are obtained using two methods: (1) Radon transform (RT), in which a line is fitted to the peak positions across successive time bins and the summation of the decoding probabilities along the best fit line is calculated (Davidson et al., 2009); (2) weighted correlation (WC), which is the absolute correlation between the decoded positions and time within a PBE (Grosmark & Buzsáki, 2016). The peak decoded positions are less aligned with the best fit line (white line with a deviation tolerance band) for both shuffling methods. The replay scores for the actual PBE (top) are compared against 1000 time-swap PBEs. **(c)** The replay scores (based on RT or WC) for the example PBE compared to null distributions from TS or UI methods. The actual replay scores are calculated as the percentiles of each null distribution.

segments are given a score in between. The scores for each subject are collected and compared with the scores that previously were calculated using the automatic approaches. An approach is evaluated as reliable if it results in replay scores that significantly correlate with scores provided by the human subjects. Preliminary results showed that combination of Radon transform as a replay scoring method and unit-identity shuffle for generating the null distribution performs better than the other methods (**Figure 4.4**). We need data from more subjects to confirm the results.

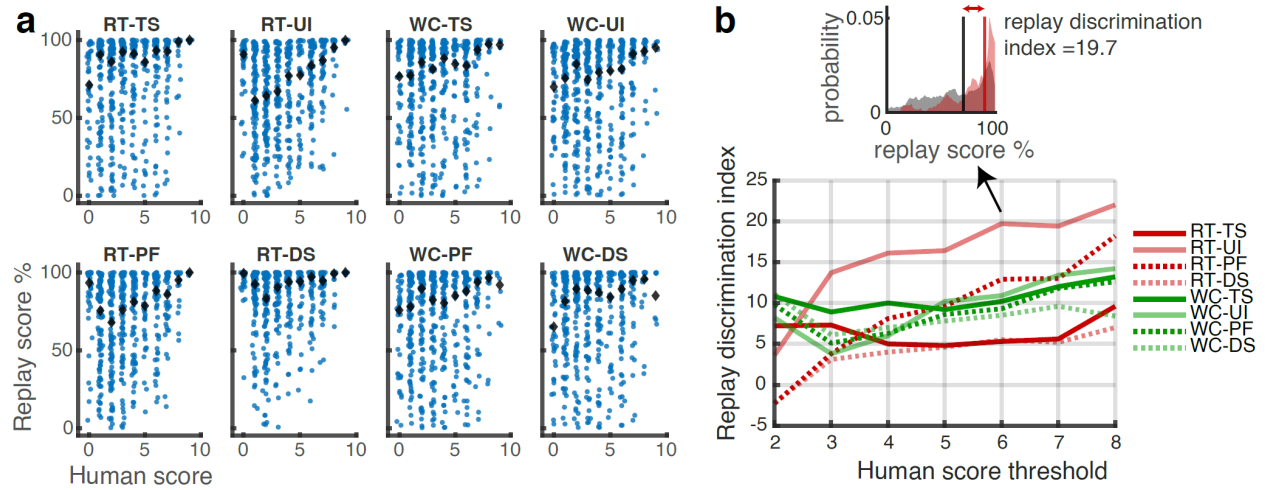


Figure 4.4: Agreement of each algorithmic replay score method with human scores. (a) The scatter plots of algorithmic replay scores versus human scores for eight methods. RT: Radon transform, WC: weighted correlation, TS: time swap, UI: unit-ID shuffle, PF: place field random circular shift, DS: decoded position random circular shift. The blue dots show individual SWRs and the black diamonds show the median replay scores corresponding to each human score. **(b)** Difference in median replay scores between SWRs with human scores above or below a threshold, considered as putative replay and non-replay events, respectively. The results are shown for different thresholds in the range of 2 to 8. The top inset shows distribution of RT-UI replay scores for replays (red, human score > 6) and non-replays (gray, human score < 6). The RT-UI method has a higher replay discrimination index compared to other methods, regardless of the threshold of human score to determine replay vs non-replay events.

References

- Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M. B. (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences of the United States of America*, 111(52), 18428–18435. <https://doi.org/10.1073/pnas.1421056111>
- Ambrose, R. E., Pfeiffer, B. E., & Foster, D. J. (2016). Reverse Replay of Hippocampal Place Cells Is Uniquely Modulated by Changing Reward. *Neuron*, 91(5), 1124–1136. <https://doi.org/10.1016/j.neuron.2016.07.047>
- Aronov, D., Nevers, R., & Tank, D. W. (2017). Mapping of a non-spatial dimension by the hippocampal-entorhinal circuit. *Nature*, 543(7647), 719–722. <https://doi.org/10.1038/nature21692>
- Bartho, P., Hirase, H., Zugaro, M., Harris, K. D., Zu-, M., & Harris, K. D. (2004). Characterization of Neocortical Principal Cells and Interneurons by Network Interactions and Extracellular Features. *J Neurophysiol*, 92, 600–608.
- Bellesi, M., Pfister-genskow, M., Maret, S., Keles, S., Tononi, G., & Cirelli, C. (2013). *Effects of Sleep and Wake on Oligodendrocytes and Their Precursors*. 33(36), 14288–14300. <https://doi.org/10.1523/JNEUROSCI.5102-12.2013>
- Bokil, H., Andrews, P., Kulkarni, J. E., Mehta, S., & Mitra, P. P. (2010). Chronux : A platform for analyzing neural signals. *Journal of Neuroscience Methods*, 192(1), 146–151. <https://doi.org/10.1016/j.jneumeth.2010.06.020>
- Boyce, R., Glasgow, S. D., Williams, S., & Adamantidis, A. (2016). Sleep research: Causal evidence for the role of REM sleep theta rhythm in contextual memory consolidation. *Science*, 352(6287), 812–816. <https://doi.org/10.1126/science.aad5252>
- Buhry, L., Azizi, A. H., & Cheng, S. (2011). *Reactivation , Replay , and Preplay : How It Might All Fit Together. 2011*. <https://doi.org/10.1155/2011/203462>
- Buzsaki, G. (1989). Two-stage model of memory trace formation: a role for “noisy” brain states. In *Neuroscie* (Vol. 31, Issue 3, pp. 551–570). Neuroscience.
- Buzsáki, G., & Tingley, D. (2018). Space and Time: The Hippocampus as a Sequence Generator. *Trends in Cognitive Sciences*, 22(10), 853–869. <https://doi.org/10.1016/j.tics.2018.07.006>
- Buzsaki, Gyorgy. (1986). Hippocampal Sharp Waves: Their Origin and Significance. *Brain Research*, 398(2), 242–253.
- Buzsaki, Gyorgy. (2015). Hippocampal sharp wave-ripple A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10), 1073–1188.
- Buzsaki, György. (1998). Memory consolidation during sleep: a neurophysiological perspective. *Journal of Sleep Research*, 7(S1), 17–23. <https://doi.org/10.1046/j.1365-2869.7.s1.3.x>
- Carr, M. F., Jadhav, S. P., & Frank, L. M. (2011). Hippocampal replay in the awake state: A potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, 14(2), 147–153. <https://doi.org/10.1038/nn.2732>
- Chen, Z., Gomperts, S. N., & Wilson, M. A. (2014). *Neural Representation of Spatial Topology in the Rodent Hippocampus*. 39, 1–39. <https://doi.org/10.1162/NECO>
- Chen, Z., Kloosterman, F., Brown, E. N., & Wilson, M. A. (2012). *Uncovering spatial topology represented by rat hippocampal population neuronal codes*. 227–255. <https://doi.org/10.1007/s10827-012-0384-x>
- Cheng, S., & Frank, L. M. (2011). THE STRUCTURE OF NETWORKS THAT PRODUCE THE TRANSFORMATION FROM GRID CELLS TO PLACE CELLS. *NSC*, 197, 293–306. <https://doi.org/10.1016/j.neuroscience.2011.09.002>
- Cheng, Sen, & Frank, L. M. (2008). New Experiences Enhance Coordinated Neural Activity in the Hippocampus. *Neuron*, 57(2), 303–313. <https://doi.org/10.1016/j.neuron.2007.11.035>

- Cirelli, C., & Tononi, G. (2022). The why and how of sleep-dependent synaptic down-selection. *Seminars in Cell and Developmental Biology*, 125(January 2021), 91–100. <https://doi.org/10.1016/j.semcdb.2021.02.007>
- Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M. B., & Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, 462(7271), 353–357. <https://doi.org/10.1038/nature08573>
- Dabaghian, Y., Brandt, V. L., & Frank, L. M. (2014). *Reconceiving the hippocampal map as a topological template*. 1–17. <https://doi.org/10.7554/eLife.03476>
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal Replay of Extended Experience. *Neuron*, 63(4), 497–507. <https://doi.org/10.1016/j.neuron.2009.07.027>
- Diba, K. (2021). Hippocampal sharp-wave ripples in cognitive map maintenance versus episodic simulation. *Neuron*, 109(19), 3071–3074. <https://doi.org/10.1016/j.neuron.2021.09.006>
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, 10(10), 1241–1242. <https://doi.org/10.1038/nn1961>
- Diekelmann, S., & Born, J. (2010). *The memory function of sleep*. 11(2), 114–126. <https://doi.org/10.1038/nrn2762>
- Diekelmann, S., Büchel, C., Born, J., & Rasch, B. (2011). Labile or stable: Opposing consequences for memory when reactivated during waking and sleep. *Nature Neuroscience*, 14(3), 381–386. <https://doi.org/10.1038/nn.2744>
- Dong, C., Madar, A. D., & Sheffield, M. E. J. (2021). Distinct place cell dynamics in CA1 and CA3 encode experience in new environments. *Nature Communications*, 12(1), 1–13. <https://doi.org/10.1038/s41467-021-23260-3>
- Dragoi, G., & Buzsáki, G. (2006). Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron*, 50(1), 145–157. <https://doi.org/10.1016/j.neuron.2006.02.023>
- Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330), 397–401. <https://doi.org/10.1038/nature09633>
- Drieu, C., Todorova, R., & Zugaro, M. (2018). Nested sequences of hippocampal assemblies during behavior support subsequent sleep replay. *Science*, 362(6415), 675–679. <https://doi.org/10.1126/science.aat2952>
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, 13(8), 995–1002. <https://doi.org/10.1038/nn.2599>
- Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial.pdf. *Hippocampus*, 1(20), 1–10.
- Eichenbaum, H. (2014). Time cells in the hippocampus : a new dimension for mapping memories. *Nature Publishing Group*, 15(October). <https://doi.org/10.1038/nrn3827>
- Eschenko, O., Ramadan, W., Mölle, M., Born, J., & Sara, S. J. (2008). Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*, 15, 222–228. <https://doi.org/10.1101/lm.726008.15>
- Farooq, U., Sibille, J., Liu, K., & Dragoi, G. (2019). Strengthened Temporal Coordination within Pre-existing Sequential Cell Assemblies Supports Trajectory Replay. *Neuron*, 103(4), 719–733.e7. <https://doi.org/10.1016/j.neuron.2019.05.040>
- Feng, T., Silva, D., & Foster, D. J. (2015). Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession. *Journal of Neuroscience*, 35(12), 4890–4902. <https://doi.org/10.1523/JNEUROSCI.2614-14.2015>
- Fernández-Ruiz, A., Oliva, A., de Oliveira, E. F., Rocha-Almeida, F., Tingley, D., & Buzsáki, G. (2019). Long-

- duration hippocampal sharp wave ripples improve memory. *Science*, 364(6445), 1082–1086. <https://doi.org/10.1126/science.aax0758>
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084), 680–683. <https://doi.org/10.1038/nature04587>
- Frank, L. M., Stanley, G. B., & Brown, E. N. (2004). Hippocampal plasticity across multiple days of exposure to novel environments. *Journal of Neuroscience*, 24(35), 7681–7689. <https://doi.org/10.1523/JNEUROSCI.1958-04.2004>
- G, B. (2002). Theta Oscillations in the Hippocampus. *Neuron*, 33, 1–16.
- Genzel, L., Kroes, M. C. W., Dresler, M., & Battaglia, F. P. (2014). Light sleep versus slow wave sleep in memory consolidation: A question of global versus local processes? *Trends in Neurosciences*, 37(1), 10–19. <https://doi.org/10.1016/j.tins.2013.10.002>
- Gillespie, A. K., Maya, D. A. A., Denovellis, E. L., Roumis, D. K., Eden, U. T., & Frank, L. M. (2021). Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice Hippocampal replay reflects specific past experiences rather than a plan for subsequent choice. *Neuron*, 109(19), 3149–3163.e6. <https://doi.org/10.1016/j.neuron.2021.07.029>
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10), 1222–1223. <https://doi.org/10.1038/nn.2384>
- Giri, B., Miyawaki, H., Mizuseki, K., Cheng, S., & Diba, K. (2019). Hippocampal reactivation extends for several hours following novel experience. *Journal of Neuroscience*, 39(5), 866–875. <https://doi.org/10.1523/JNEUROSCI.1950-18.2018>
- Gridchyn, I., Schoenenberger, P., Neill, J. O., Csicsvari, J., Gridchyn, I., Schoenenberger, P., Neill, J. O., & Csicsvari, J. (2020). Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. *Neuron*, 106(2), 291–300.e6. <https://doi.org/10.1016/j.neuron.2020.01.021>
- Grosmark, A. D., & Buzsáki, G. (2016). *Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences*. *Science*. <https://doi.org/10.1126/science.aad1935>
- Grosmark, A. D., Sparks, F. T., Davis, M. J., & Losonczy, A. (2021). Reactivation predicts the consolidation of unbiased long-term cognitive maps. *Nature Neuroscience*, 24(11), 1574–1585. <https://doi.org/10.1038/s41593-021-00920-7>
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal Replay Is Not a Simple Function of Experience. *Neuron*, 65(5), 695–705. <https://doi.org/10.1016/j.neuron.2010.01.034>
- Guzmán-marín, R., Suntsova, N., Stewart, D. R., Gong, H., Szymusiak, R., & McGinty, D. (2003). Sleep deprivation reduces proliferation of cells in the dentate gyrus of the hippocampus in rats. *Journal of Physiology*, 563–571. <https://doi.org/10.1113/jphysiol.2003.041665>
- Havekes, R., & Abel, T. (2017). The tired hippocampus: the molecular impact of sleep deprivation on hippocampal function. *Current Opinion in Neurobiology*, 44, 13–19. <https://doi.org/10.1016/j.conb.2017.02.005>
- Hobson, J. A. (2009). REM sleep and dreaming: Towards a theory of protoconsciousness. *Nature Reviews Neuroscience*, 10(11), 803–814. <https://doi.org/10.1038/nrn2716>
- Hurley, N., Member, S., Rickard, S., & Member, S. (2009). *Comparing Measures of Sparsity*. 55(10), 4723–4741.
- Inostroza, M., & Born, J. (2013). Sleep for preserving and transforming episodic memory. *Annual Review of Neuroscience*, 36, 79–102. <https://doi.org/10.1146/annurev-neuro-062012-170429>
- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087), 1454–1458. <https://doi.org/10.1126/science.1217230>

- Jenkins, J. G., & Dallenbach, K. (1924). Obliviscence during Sleep and. *The American Journal of Psychology*, 35(4), 605–612.
- Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*, 10(1), 100–107. <https://doi.org/10.1038/nn1825>
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913–918. <https://doi.org/10.1038/nn.2344>
- Karni, A., Tanne, D., Rubenstein, B. S., Askenasy, J. J. M., & Sagi, D. (1994). Dependence on REM sleep of overnight improvement of a perceptual skill. *Science*, 265(5172), 679–682. <https://doi.org/10.1126/science.8036518>
- Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., Liu, D. F., & Frank, L. M. (2020). Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell*, 180(3), 552–567.e25. <https://doi.org/10.1016/j.cell.2020.01.014>
- Kemere, C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). *Rapid and Continuous Modulation of Hippocampal Network State during Exploration of New Places*. 8(9). <https://doi.org/10.1371/journal.pone.0073114>
- Khodagholy, D., Gelinas, J. N., & Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science*, 358(6361), 369–372. <https://doi.org/10.1126/science.aan6203>
- Kinsky, N. R., Sullivan, D. W., Mau, W., Hasselmo, M. E., & Eichenbaum, H. B. (2018). Hippocampal Place Fields Maintain a Coherent and Flexible Map across Long Timescales. *Current Biology*, 28(22), 3578–3588.e6. <https://doi.org/10.1016/j.cub.2018.09.037>
- Klinzing, J. G., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, 22(October). <https://doi.org/10.1038/s41593-019-0467-3>
- Kovacs, K. A., O'Neill, J., Schoenenberger, P., Penttonen, M., Guerrero, D. K. R., & Csicsvari, J. (2016). Optogenetically Blocking Sharp Wave Ripple Events in Sleep Does Not Interfere with the Formation of Stable Spatial Representation in the CA1 Area of the Hippocampus. *PLOS One*, 1–19. <https://doi.org/10.1371/journal.pone.0164675>
- Krause, E. L., & Drugowitsch, J. (n.d.). *A large majority of awake hippocampal sharp-wave ripples feature spatial 1 trajectories with momentum*. 2. <https://doi.org/10.1101/2021.05.13.444067>
- Krause, E. L., & Drugowitsch, J. (2021). A large majority of awake hippocampal sharp-wave ripples feature spatial 1 trajectories with momentum. 2. *BioRxiv*, 2021.05.13.444067. <https://doi.org/10.1101/2021.05.13.444067>
- Krishnan, G. P., Chauvette, S., Shamie, I., Soltani, S., Timofeev, I., Cash, S. S., Halgren, E., & Bazhenov, M. (2016). Cellular and neurochemical basis of sleep stages in the thalamocortical network. *ELife*, 1–29. <https://doi.org/10.7554/eLife.18607>
- Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience*, 19(10), 4090–4101. <https://doi.org/10.1523/jneurosci.19-10-04090.1999>
- Larkin, M. C., Lykken, C., Tye, L. D., Wickelgren Jeanette Graham, & Frank, L. M. (2014). Hippocampal output area CA1 broadcasts a generalized novelty signal during an object-place. *HIPPOCAMPUS*, 24(7), 773–783.
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6), 1183–1194. [https://doi.org/10.1016/S0896-6273\(02\)01096-6](https://doi.org/10.1016/S0896-6273(02)01096-6)
- Liu, K., Sibille, J., & Dragoi, G. (2019). Preconfigured patterns are the primary driver of offline multi-neuronal sequence replay. *Hippocampus*, 275–283.
- Logothetis, N. K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H. C., Besserve, M., & Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. *Nature*,

491(7425), 547–553. <https://doi.org/10.1038/nature11618>

- Louie, K., & Wilson, M. A. (2001). *Temporally Structured Replay of Awake Hippocampal Ensemble Activity during Rapid Eye Movement Sleep*. 29, 145–156.
- Magee, J. C., & Johnston, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science*, 275(5297), 209–213. <https://doi.org/10.1126/science.275.5297.209>
- Mainingret, N., Girardeau, G., Todorova, R., Goutierre, M., & Zugaro, M. (2016). Hippocampo-cortical coupling mediates memory consolidation during sleep. *Nature Neuroscience*, 19(7), 959–964. <https://doi.org/10.1038/nn.4304>
- Mankin, E. A., Sparks, F. T., Slayyeh, B., Sutherland, R. J., Leutgeb, S., & Leutgeb, J. K. (2012). Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, 109(47), 19462–19467. <https://doi.org/10.1073/pnas.1214107109>
- Marshall, L., Helgadóttir, H., Mölle, M., & Born, J. (2006). Boosting slow oscillations during sleep potentiates memory. *Nature*, 444(7119), 610–613. <https://doi.org/10.1038/nature05278>
- Mattar, M. G., & Daw, N. D. (2018). Prioritized memory access explains planning and hippocampal replay. *Nature Neuroscience*, 21(11), 1609–1617. <https://doi.org/10.1038/s41593-018-0232-z>
- Miyawaki, H., & Diba, K. (2016). Regulation of Hippocampal Firing by Network Oscillations during Sleep. *Current Biology*, 26(7), 893–902. <https://doi.org/10.1016/j.cub.2016.02.024>
- Monaco, J. D., Rao, G., Roth, E. D., & Knierim, J. J. (2014). Attentive scanning behavior drives one-trial potentiation of hippocampal place fields. *Nature Neuroscience*, 17(5), 725–731. <https://doi.org/10.1038/nn.3687>
- Morris, R. G. M., Anderson, E., Lynch, G. S., & Baudry, M. (1986). Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, 319, 774–776.
- Nádasdy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*, 19(21), 9497–9507. <https://doi.org/10.1523/jneurosci.19-21-09497.1999>
- Norimoto, H., Makino, K., Gao, M., Shikano, Y., Okamoto, K., Ishikawa, T., Sasaki, T., Hioki, H., Fujisawa, S., & Yuji Ikegaya, †. (2018). Hippocampal ripples down-regulate synapses. *Science*, 359, 1524–1527. <http://science.sciencemag.org/>
- O’Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map . Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*, 34(1), 171–175. <http://www.ncbi.nlm.nih.gov/pubmed/5124915>
- O’Keefe, J., & Speakman, A. (1987). Single unit activity in the rat hippocampus during a spatial memory task. *Experimental Brain Research*, 68(1), 1–27. <https://doi.org/10.1007/BF00255230>
- O’Keefe, John. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*, 51(1), 78–109. [https://doi.org/10.1016/0014-4886\(76\)90055-8](https://doi.org/10.1016/0014-4886(76)90055-8)
- Pastalkova, E., Itskov, V., Amarasingham, A., & Buzsaki, G. (2008). Internally Generated Cell Assembly Sequences in the Rat Hippocampus. *Science*, 321(September), 1322–1328.
- Peigneux, P., Laureys, S., Fuchs, S., Collette, F., Perrin, F., Reggers, J., Degueldre, C., Phillips, C., Fiore, G. D., Aerts, J., Luxen, A., & Maquet, P. (2004). Are Spatial Memories Strengthened in the Human Hippocampus during Slow Wave Sleep? *Neuron*, 44(3), 535–545. <https://doi.org/10.1038/nn0311-272>
- Petersen, P. C., Siegle, J. H., Steinmetz, N. A., & Mahallati, S. (2021). NeuroResource CellExplorer : A framework for visualizing and characterizing single neurons II II NeuroResource CellExplorer : A framework for visualizing and characterizing single neurons. *Neuron*, 109(22), 3594–3608. <https://doi.org/10.1016/j.neuron.2021.09.002>

- Pettit, N. L., Yap, E.-L., Greenberg, M. E., & Harvey, C. D. (2022). Fos ensembles encode and shape stable spatial maps in the hippocampus. *Nature*, 609(September). <https://doi.org/10.1038/s41586-022-05113-1>
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447), 74–79. <https://doi.org/10.1038/nature12112>
- Pfeiffer, B. E., & Foster, D. J. (2015). *Autoassociative dynamics in the generation of sequences of hippocampal place cells*. 349(6244), 180–184. <https://doi.org/10.5061/dryad.gf774.SUPPLEMENTARY>
- Poe, G. R., Nitz, D. A., McNaughton, B. L., & Barnes, C. A. (2000). Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. *Brain Research*, 855(1), 176–180. [https://doi.org/10.1016/S0006-8993\(99\)02310-0](https://doi.org/10.1016/S0006-8993(99)02310-0)
- Prince, T. M., Wimmer, M., Choi, J., Havekes, R., Aton, S., & Abel, T. (2014). Sleep deprivation during a specific 3-hour time window post-training impairs hippocampal synaptic plasticity and memory. *Neurobiology of Learning and Memory*, 109, 122–130. <https://doi.org/10.1016/j.nlm.2013.11.021>
- Quirk, M. C., & Wilson, M. A. (1999). Interaction between spike waveform classification and temporal sequence detection. *Journal of Neuroscience Methods*, 94(1), 41–52. [https://doi.org/10.1016/S0165-0270\(99\)00124-7](https://doi.org/10.1016/S0165-0270(99)00124-7)
- Rabiner. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceeding of the IEEE*, 77(2), 257–286. <https://doi.org/10.25300/MISQ/2017/41.3.08>
- Rasch, B., & Born, J. (2022). *ABOUT SLEEP ' S ROLE IN MEMORY*. 681–766. <https://doi.org/10.1152/physrev.00032.2012>
- Rossant, C., Kadir, S. N., Goodman, D. F. M., Schulman, J., Hunter, M. L. D., Saleem, A. B., Grosmark, A., Belluscio, M., Denfield, G. H., Ecker, A. S., Tolias, A. S., Solomon, S., Buzsáki, G., Carandini, M., & Harris, K. D. (2016). Spike sorting for large , dense electrode arrays. *Nature Neuroscience*, 19(4). <https://doi.org/10.1038/nn.4268>
- Roux, L., Hu, B., Eichler, R., Stark, E., & Buzsáki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nature Neuroscience*, 20(6), 845–853. <https://doi.org/10.1038/nn.4543>
- Sadowski, J. H. L. P., Jones, M. W., & Mellor, J. R. (2016). Sharp-Wave Ripples Orchestrate the Induction of Synaptic Plasticity during Reactivation of Place Cell Firing Patterns in the Hippocampus. *Cell Reports*, 14(8), 1916–1929. <https://doi.org/10.1016/j.celrep.2016.01.061>
- Schacter, D. L., Addis, D. R., & Buckner, R. L. (2008). Episodic simulation of future events: Concepts, data, and applications. *Annals of the New York Academy of Sciences*, 1124, 39–60. <https://doi.org/10.1196/annals.1440.001>
- Schmitzer-Torbert, N., Jackson, J., Henze, D., Harris, K., & Redish, A. D. (2005). Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience*, 131(1), 1–11. <https://doi.org/10.1016/j.neuroscience.2004.09.066>
- Scoville, W. B., & Milner, B. (1957). *Loss of recent memory after bilateral hippocampal lesions*.
- Shin, J. D., Tang, W., & Jadhav, S. P. (2019). Dynamics of Awake Hippocampal-Prefrontal Replay for Spatial Learning and Memory-Guided Decision Making. *Neuron*, 104(6), 1110–1125.e7. <https://doi.org/10.1016/j.neuron.2019.09.012>
- Siclari, F., Baird, B., Perogamvros, L., Bernardi, G., Larocque, J. J., Riedner, B., Boly, M., Postle, B. R., & Tononi, G. (2017). *The neural correlates of dreaming*. 20(6). <https://doi.org/10.1038/nn.4545>
- Siegle, J. H., Lopez, A. C., Patel, Y. A., Abramov, K., Ohayon, S., & Voigts, J. (2017). Open Ephys : an open-source , plugin-based platform for multichannel electrophysiology. *Journal of Neural Engineering*, 14.
- Silva, D., Feng, T., & Foster, D. J. (2015). Trajectory events across hippocampal place cells require previous experience. *Nature Neuroscience*, 18(12), 1772–1779. <https://doi.org/10.1038/nn.4151>

- Snyder, F., Hobson, J. A., Morrison, D. F., & Goldfrank, F. (1964). Changes in respiration , heart rate , and systolic blood pressure in human sleep. *Journal of Applied Physiology*, 19(3), 417–422.
- Stella, F., Baracska, P., O'Neill, J., & Csicsvari, J. (2019). Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. *Neuron*, 102(2), 450–461.e7. <https://doi.org/10.1016/j.neuron.2019.01.052>
- Thompson, L. T., & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*, 9(7), 2382–2390. <https://doi.org/10.1523/jneurosci.09-07-02382.1989>
- Tingley, D., & Peyrache, A. (2020). On the methods for reactivation and replay analysis. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1799). <https://doi.org/10.1098/rstb.2019.0231>
- Tononi, G., & Cirelli, C. (2014). Perspective Sleep and the Price of Plasticity : From Synaptic and Cellular Homeostasis to Memory Consolidation and Integration. *Neuron*, 81(1), 12–34. <https://doi.org/10.1016/j.neuron.2013.12.025>
- van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal Offline Reactivation Consolidates Recently Formed Cell Assembly Patterns during Sharp Wave-Ripples. *Neuron*, 92(5), 968–974. <https://doi.org/10.1016/j.neuron.2016.10.020>
- Van Der Meer, M. A. A., Kemere, C., & Diba, K. (2020). Progress and issues in second-order analysis of hippocampal replay. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1799). <https://doi.org/10.1098/rstb.2019.0238>
- Vertes, R. P. (2004). Memory consolidation in sleep: Dream or reality. *Neuron*, 44(1), 135–148. <https://doi.org/10.1016/j.neuron.2004.08.034>
- Wang, Y., Romani, S., Lustig, B., Leonardo, A., & Pastalkova, E. (2015). Theta sequences are essential for internally generated hippocampal firing fields. *Nature Neuroscience*, 18(2), 282–288. <https://doi.org/10.1038/nn.3904>
- Wen, H., & Liu, Z. (2016). Broadband Electrophysiological Dynamics Contribute to Global Resting-State fMRI Signal. *The Journal of Neuroscience*, 36(22), 6030–6040. <https://doi.org/10.1523/JNEUROSCI.0187-16.2016>
- Wikenheiser, A. M., & Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, 18(2), 289–294. <https://doi.org/10.1038/nn.3909>
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, 261(5124), 1055–1058. <https://doi.org/10.1126/science.8351520>
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172), 676–679. <https://doi.org/10.1126/science.8036517>
- Wiskott, L. (2013). *Lecture Notes on Bayesian Theory and Graphical Models*. 1–34.
- Xu, H., Baracska, P., Neill, J. O., Csicsvari, J., Xu, H., Baracska, P., Neill, J. O., & Csicsvari, J. (2019). Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the Radial Eight-Arm Maze Article Assembly Responses of Hippocampal CA1 Place Cells Predict Learned Behavior in Goal-Directed Spatial Tasks on the. *Neuron*, 101(1), 119–132.e4. <https://doi.org/10.1016/j.neuron.2018.11.015>
- Yamamoto, J., & Tonegawa, S. (2017). Direct Medial Entorhinal Cortex Input to Hippocampal CA1 Is Crucial for Extended Quiet Awake Replay. *Neuron*, 96(1), 217–227.e4. <https://doi.org/10.1016/j.neuron.2017.09.017>
- Yger, P., Spampinato, G. L. B., Esposito, E., Lefebvre, B., Vision, I. De, Umrs, I., & Um, U. (2018). A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *ELife*, 1–23.
- Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology*, 79(2), 1017–1044. <https://doi.org/10.1152/jn.1998.79.2.1017>

Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., Gamal, A. El, & Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, *16*(3), 264–266.
<https://doi.org/10.1038/nn.3329>