Stability and change in the memory system during rest

by

Lars Bollmann

August, 2024

A thesis submitted to the
Graduate School
of the
Institute of Science and Technology Austria
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Committee in charge:

Carrie Bernecky, Chair
Jozsef Csicsvari
Gašper Tkačik
Cristina Savin

The thesis of Lars Bollmann, titled *Stability and change in the memory system during rest*, is approved by:

Supervisor: Jozsef Csicsvari, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Gašper Tkačik, ISTA, Klosterneuburg, Austria

Signature: _____

Committee Member: Cristina Savin, New York University, New York, USA

Signature: _____

Defense Chair: Carrie Bernecky, ISTA, Klosterneuburg, Austria

Signature: _____

Signed page is on file

ISTA Thesis, ISSN: 2663-337X

I hereby declare that this thesis is my own work and that it does not contain other people's work without this being so stated; this thesis does not contain my previous work without this being stated, and the bibliography contains all the literature that I used in writing the dissertation.

I accept full responsibility for the content and factual accuracy of this work, including the data and their analysis and presentation, and the text and citation of other work.

I declare that this is a true copy of my thesis, including any final revisions, as approved by my thesis committee, and that this thesis has not been submitted for a higher degree to any other university or institution.

I certify that any republication of materials presented in this thesis has been approved by the relevant publishers and co-authors.

Signature: _____

Lars Bollmann
August, 2024

Signed page is on file

Abstract

Acquiring, retaining, and retrieving information over a wide range of timescales are crucial functions of the brain. The successful processing of memories affects many aspects of our lives and enables us and many other organisms to operate in a complex environment and to interact with it. In this context, the hippocampus and functionally connected brain areas, such as the prefrontal cortex, are central and have been subject to intensive research in the past decades. Storage of memories is believed to rely on distributed neural activity within these neural circuits. Additionally, neural memory traces of recent experience are reinstated during periods of rest or sleep. These reactivations are thought to play an outstanding role in the consolidation of memories and potentially facilitate the transfer of information from the hippocampus to cortical areas for long-term storage and integration into existing knowledge.

However, there is growing evidence that memory-related neural representations in the hippocampus are not as stable as initially thought and that they change even in the absence of learning. It has been suggested that these changes reflect the accumulation of experience, but the influence of interspersed consolidation periods has not been considered. Previous studies have analyzed consolidation periods by detecting activity that strongly resembled neural activity during the acquisition of memory. Besides being often limited to only non-rapid eye movement (NREM) sleep, the used approaches were not capable of tracking changes in neural representations over extended temporal periods. More fluid representations do not only challenge our understanding of how information is stored, but they also affect the transfer of information between brain areas during the consolidation process.

For this thesis, I investigated the evolution of memory-related activity during sleep periods expected to be involved in consolidation in the hippocampus and between the hippocampus and prefrontal cortex. I found that reactivated activity in the hippocampus gradually transformed during prolonged periods of sleep and inactivity. In the beginning, neural activity strongly resembled acquisition activity, whereas, with the progression of time, it became more similar to the subsequent recall activity. NREM periods drove this process, while rapid-eye movement (REM) periods showed a resetting effect. This reactivation drift was due to firing rate changes of a subset of cells and mirrored the representational changes from the acquisition to the recall. A stable subset of cells withstood the drift and maintained their activity. Therefore, my results indicate that memory-related representations undergo spontaneous modifications during consolidation periods and that these changes are predictive of representational drift.

Furthermore, I found that the amount of change in the neural activity during subsequent sleep periods was biased by prior behavioral performance. Observed changes in the hippocampus and the prefrontal cortex were synchronized and increased after poor performance, highlighting a potential role in the exchange of information. Low-variance

periods with distinct, more stable activity from a subset of cells significantly contributed to the heightened synchrony between both areas. Hence, interleaved phases of more stable neural activity could facilitate the information transfer between brain areas.

In conclusion, my investigations underline the fluidity of memory-related representations and assign a prominent role to sleep reactivation periods in their evolution. In addition, I identified a potential mechanism of stable activity phases that might facilitate the synchronization across hippocampal-prefrontal activity despite ongoing changes. Reconciling and integrating findings from both spontaneous and behaviorally-related representational changes in functionally related brain areas will help to broaden our understanding of how knowledge is stored, maintained, updated, and transferred between brain areas.

Acknowledgements

First of all, I want to thank Federico Stella for his support and guidance throughout my Ph.D. studies. The amount and quality of my work would not have been possible without his continuous feedback, advice, and input. I especially want to highlight the positive impact our regular meetings had on my motivation and commitment.

I am also grateful for the opportunity to do my research in the Csicsvari group. I want to thank Jozsef Csicsvari, who allowed me to join his group despite having no formal training in either Data Science or Neuroscience.

I want to extend my gratitude to Gašper Tkačik and Wiktor Mlynarski, who were always there for me in challenging times. Both helped me shape my research project, guided me in finding appropriate analyses for questions I had in mind, and explained various analysis approaches to me.

I also want to thank the other members of the Csicsvari group, especially Michele Nardin, Andrea Cumpelik, Heloisa Chiossi, and Predrag Zivadinovic, who were always there to exchange research-related ideas and made my whole Ph.D experience much more pleasant.

In addition, I am grateful that Cristina Savin agreed to be a member of my Ph.D. committee without reservations and for dealing with my concerns and questions.

Finally, I want to thank Anitha Andiappan for her unconditional support, patience, and cheerful attitude that helped me overcome difficult times.

About the Author

Lars Bollmann holds a Master's degree in Biomedical Engineering and completed his studies at ETH Zürich, Graz University of Technology, and the University of Cambridge. Before joining the Cscicsvari group at ISTA in 2019, he was working as an R&D engineer at Ottobock in the field of high-tech prosthetics and active implants.

Lars has done research covering a variety of topics including soft tissue biomechanics at Graz University of Technology, human-robot interaction at the University of British Columbia, medical robotics at CNRS in Montpellier, and cell mechanics at the University of Cambridge & Osaka University. He published his work on microglia mechanics in Frontiers in Cellular Neuroscience in 2015 and gave several talks about his research. His main interest lies in topics at the intersection between biology and technology.

During his Ph.D. studies, Lars presented his research at the Computational Neuroscience meeting (2021), the Bernstein conference (2021, 2022, 2023), the ANA meeting (2021), and the FENS forum (2022).

List of Collaborators and Publications

- Peter Baracskay:
  - recording of neural data for Chapter 2
  - pre-processing data (clustering and sleep scoring) for Chapter 2
  - writing first part of the Method section (until SWRs detection) for Chapter 2
- Karola Käfer
  - recording of neural data for Chapter 3
  - pre-processing data (clustering) for Chapter 3
- Federico Stella
  - input regarding data analysis, figures and writing for Chapter 1, Chapter 2, Chapter 3 and Chapter 4
- Jozsef Cscicsvari
  - input regarding data analysis, figures and writing for Chapter 1, Chapter 2 and Chapter 3

Table of Contents

Abstract vii

Acknowledgements ix

About the Author x

List of Collaborators and Publications xi

Table of Contents xiii

List of Figures xiv

1 Introduction 1
- 1.1 Hippocampus and memory . . . 1
- 1.2 Medial prefrontal cortex . . . 4
- 1.3 Rest periods and memory consolidation . . . 5
- 1.4 Aims of the Thesis . . . 7

2 Sleep stages antagonistically modulate reactivation drift 9
- 2.1 Abstract . . . 9
- 2.2 Introduction . . . 10
- 2.3 Results . . . 10
- 2.4 Discussion . . . 22
- 2.5 Supplementary Figures . . . 24
- 2.6 Methods . . . 33

3 Prior behavioral performance determines stability and synchrony in the memory system during rest 43
- 3.1 Abstract . . . 43
- 3.2 Introduction . . . 44
- 3.3 Results . . . 44
- 3.4 Discussion . . . 61
- 3.5 Supplementary Figures . . . 65
- 3.6 Methods . . . 74

4 Conclusion & outlook 79

Bibliography 83

List of Figures

2.1 Stable wireless tetrode recordings in CA1 during extended-time learning paradigm . . . 11
2.2 Rest neural activity continuously transitions from being similar to acquisition neural activity towards greater similarity with recall neural activity. . . . 13
2.3 NREM and REM sleep stages show opposing effects on the observed drift. 15
2.4 Cells that modulate their firing from acquisition to recall drive memory drift, while persistent cells reduce it. . . . 17
2.5 Memory drift is driven by firing rate modulations of a subset of cells. . . . 18
2.6 Persistent cells maintain their spatial map and population activity from acquisition to recall. . . . 19
2.7 Similar remapping for persistent and decreasing cells during acquisition. . 20
2.8 Persistent cell activity during rest is more strongly influenced by the interleaved acquisition session than decreasing cell activity. . . . 21
S2.1 Stable recordings . . . 24
S2.2 pHMM model fit and its quality . . . 25
S2.3 Spatial selectivity of decoded pHMM states. . . . 26
S2.4 Drift during sleep was also observed using Bayesian decoding. . . . 27
S2.5 Likelihoods and reactivation dynamics for REM and NREM . . . 28
S2.6 Decoded content in REM and NREM sleep. . . . 29
S2.7 Differences in firing rates between persistent, decreasing and increasing cells. 30
S2.8 Changes in firing rate probabilities between REM and NREM sleep epochs. 31
S2.9 Remapping of differet subsets and likelihoods for decoding pre-probe . . . 32

3.1 Offline activity drift in the hippocampus and prefrontal cortex. . . . 47
3.2 Synchrony of short-term fluctuations in neural activity between HPC and PFC. 49
3.3 Different regimes of low frequency fluctuations. . . . 50
3.4 Correlated low-variance periods in PFC and HPC drive subsets of cells. . . 52
3.5 Co-activity between HPC and PFC during sharp wave ripples and low-variance periods. . . . 55
3.6 Features of clustered cells. . . . 58
3.7 Population activity during low-variance periods is constrained and changes slowly. . . . 60
S3.1 Measuring behavioral performance during rule switch paradigm. . . . 65
S3.2 Short-term fluctuations . . . 66
S3.3 Different regimes and LFP . . . 67
S3.4 Low-variance periods . . . 68
S3.5 Low-variance periods - Interneurons . . . 69
S3.6 Communication channel . . . 70
S3.7 Features of clustered cells . . . 71
S3.8 Features of clustered cells & Overall modulation during rest . . . 72

S3.9 Low-variance state activity 73

CHAPTER 1

Introduction

The ability to acquire, update, and retain knowledge about the external world and one's own experiences is a crucial function of the brain. Accessing memories allows an organism to successfully deal with recurring obstacles using previously acquired strategies and solutions. On the other hand, while facing novel scenarios, the recall of previous approaches and the familiarity with the constraints of the environments facilitate the inference of new strategies[1,2]. Memory also plays a critical role when it comes to the avoidance of adverse experiences and is central to social interactions[3,4]. Taken together, memory processes significantly increase the chance of survival and, therefore, can be found across a wide range of species[5].

This thesis focuses on spatial memory, a specific type of memory that is concerned with information about the locations of objects such as food and which facilitates spatial navigation[6]. A subcortical brain area, the hippocampus, has been identified as a prime location involved in the processing of these memories[6–8].

The first research chapter characterizes the evolution of neural activity in the rat hippocampus from the acquisition to the recall of a spatial memory. In this regard, I focused on neural activity during long rest periods, which are of great importance for the stabilization of memories[9,10].

In the second research chapter, I investigated changes in neural activity in rats during rest following a spatial rule-switch task in two relevant brain areas. The successful execution of this task was not only dependent on the hippocampus but also relied on the prefrontal cortex[11]. The prefrontal cortex is a cortical brain area involved in working memory, decision-making, and memory consolidation[12–14]. By analyzing synchronous neural activity in both brain areas during rest periods after the task, I tried to identify features related to prior task performance and interregional coordination[15].

1.1 Hippocampus and memory

The hippocampus is part of the medial temporal lobe and is located many synapses away from sensory receptors and motor neurons. Scientific investigations spanning over a century have established it as playing a major role in the acquisition and consolidation of different types of memories.

Our modern understanding of hippocampal involvement in memory functions can be traced back to a study from the 1950s, which was the first one to provide experimental evidence implicating the medial temporal lobes in memory processes[16]. Following the surgical removal of the temporal lobes to stop seizures, the patient Henry Molaison (referred to as patient H.M.) showed very specific memory impairments. Although he was able to recall episodic memories (personal experiences) from before the intervention, he was unable to acquire new ones. Furthermore, he was not able to acquire or recall spatial memories. On the other hand, neither his intellectual nor perceptual abilities nor his capacity to acquire or recall procedural memories (motor skills) was affected by the resective surgery. In the following years, further human case studies and a growing number of animal studies confirmed the outstanding role of the hippocampus and adjacent, anatomically related brain regions for the processing of memories[17], more specifically for declarative memories. Nowadays, the identification of the hippocampus as the initial storage site for memories has been widely accepted and has been further integrated into several broader theories of memory processing in the brain. All of these theories imply the existence of a two-stage memory consolidation, that is, the process leading to the transfer of hippocampal memory traces to cortical areas for long-term storage and integration[18].

The hippocampus receives inputs from various brain regions representing multiple sensory modalities[19]. In this context, the entorhinal cortex, an adjacent brain region, acts as a gateway and comprises many direct projections to and from the hippocampus. The multitude of inputs is combined in the hippocampus to form more abstract neural representations[20–22]. In the classical view, input signals reach the hippocampal subregion dentate gyrus first. Resulting patterns of neural activity are believed to encode information in a distinct, non-overlapping way to avoid interference of similar neural activity (pattern separation)[23,24]. The neural signal is then relayed to the subregions CA3 via the mossy fiber pathway[25]. The neurons in the CA3 show strong recurrent connectivity and have been hypothesized to perform pattern completion to allow the generalization and the association in the presence of incomplete inputs[26,24]. Neural signals then travel to the hippocampal subregion CA1 via the Schaffer collateral pathway[25]. The CA1 is the major output site of the hippocampus and has many projections to the subiculum, the entorhinal cortex, and the prefrontal cortex[25]. Neural signals originating in the CA1 have been shown to be relayed to cortical areas[27–29] and different CA1 cells are believed to transmit different types of information in parallel[30].

Learning and the acquisition of new memories are thought to rely on changes in synaptic strength and on the intrinsic excitability of involved neurons[31–34]. Modifications of functional connectivity between cells can be achieved through the means of Hebbian learning[35]. The experimentally observed long-term potentiation (LTP) of synaptic connectivity has been hypothesized as a potential mechanism to store information in biological neural networks such as the hippocampus[36–39,22,40] In this context, glutamatergic NMDA receptors facilitate certain types of synaptic plasticity and are relevant for the storage of different types of memories[41–44]. During the early phase of a Hebbian LTP, simultaneous pre- and post-synaptic activity triggers a signaling cascade, which eventually can lead to the insertion of AMPA receptors in the postsynaptic membrane and a strengthening of the connection (glutamatergic transmission) between both cells[45–47]. The late phase of LTP facilitates structural changes in the synapse, including the enlargement and addition of new dendritic spines (functional contact sites between cells)[48,49].

Dopamine release in the hippocampus, which can be induced by novelty or surprise,

has been shown to promote LTP and therefore has a modulatory effect on memory processes[50–54]. There are indications that LTP positively correlates with memory performance, whereas long-term depression (LTD) leads to memory impairment[55]. Although an increasing number of studies support the close relationship between NMDA receptor-dependent plasticity, LTP, and memory, there remains doubt about a causal relationship[56,39,57,58].

According to the famous theory of Donald Hebb, the modification of synaptic connectivity through experience and learning leads to groups of cells that tend to fire together to associate and ultimately form cell assemblies[59,60]. The collective activity of these assemblies is believed to code for memory-related information and is, therefore, crucial for the acquisition and retrieval of a memory[61–66]. But how are neurons assigned to cell assemblies? It has been shown that the assembly recruitment during learning depends on the excitability of neurons in the circuit. That means that cells that can be efficiently activated during the acquisition phase have a higher chance of becoming members of the associated assembly[67–70]. The excitability of single neurons is hypothesized to depend on intrinsic membrane properties and is potentially driven by the synchrony of synaptic inputs[71,68]. In this context, inhibitory neurons play a major role in regulating the activity of cells and are possibly involved in shaping cell assemblies during the acquisition process[72,73].

In addition to their role in memory processes, many principal neurons in all three hippocampal subregions of rodents are modulated by the current spatial position of the animal[74,75]. Hence, each of these cells (called "place cells") is preferentially active in a distinct part of the environment (its "place field"). It has been a major endeavor to reconcile the concepts of memory processing and spatial coding in the hippocampus[76]. One theory is that the neural representation in the hippocampus is a substrate for navigating in both the physical and the mental space[6]. Along these lines, hippocampal neural activity has also been shown to encode non-spatial dimensions such as frequencies of a tone or time[77,78].

How can we relate the concepts of memory, plasticity, and the emergence of assemblies to the spatial firing of hippocampal cells? First of all, there is substantial evidence that the spatial representations in the hippocampus are distinct for different environments and that these representations quickly emerge during the first couple of minutes of exposure[76]. The reorganization of the spatial representation, which is accompanied by the formation of new place fields and the modification of single cell activity (firing rates), is referred to as remapping[76]. Remapping between environments could therefore facilitate distinct experiences and minimize the risk of interference[76]. In agreement with this hypothesis, even small changes in the environment and alterations in non-spatial cues can cause changes in the hippocampal spatial representation[79–82].

Observations that hippocampal spatial maps and representations of functionally connected areas are modified during learning support the existence of memory encoding in the hippocampal circuitry[83–85,23,86–88]. In this context, assemblies are formed by binding together cells with overlapping place fields and the recruitment of place cells can be influenced by changing their excitability during behavior[89,90]. To recall previously acquired information, patterns of activity representing these memory items are thought to be reinstated. Therefore, it has been proposed that memories are stored in stable activity states (attractors) in the network and that similar inputs will lead to discrete or at least distinct representations in the hippocampus[91–94]. Experimental evidence that the activation of hippocampal neurons that were active during learning can trigger the recall

of the acquired memory was first shown by Liu et al[95]. However, the relationship between memory and the stability of spatial hippocampal maps remains inconclusive.

Early studies suggested that the spatial firing of hippocampal neurons in rats within an environment is stable across re-exposures over several weeks[96]. The reinstatement of a stable spatial representation was hypothesized to constitute the successful retrieval of components of the animal's memory[97]. In support of this hypothesis, Wang et al. showed that place cells remapped during the learning of an aversive stimulus and that the resulting spatial representation exhibited long-term stability when the stimulus was presented again over the course of 5 days[98]. Learning-induced changes in the hippocampal spatial map and maintained spatial firing during subsequent recall have also been reported for hippocampal-dependent navigational tasks[99]. In this context, the stability of spatial representation was promoted by the attention and engagement of the animal[97]. Nonetheless, there is a growing body of evidence that hippocampal spatial representations do change across exposures to the same environment or task in the absence of learning[100–102]. This phenomenon, called representational drift, has been mainly observed in mice and may be less prominent in rats and differs for different hippocampal sublayers[103–106]. While spatial representations in the CA1-CA3 continuously change, the spatial code of granule cells in the dentate gyrus seems to be more stable[104].

The observed instability of hippocampal maps could be a passive, noise-driven process based on the turnover dynamics of hippocampal synapses[56,107]. A more exciting possibility is that the evolution of spatial representations in the hippocampus reflects to some degree the incorporation of new experiences and associations into existing knowledge[108,107]. Along these lines, some studies propose that changes in the spatial firing across multiple exposures could potentially timestamp events and therefore provide different neural codes to distinguish between overlapping but temporally distinct experiences[109,100,110–113]. Additionally, the study by Hayashi reports that continuous change in place-cell activity is NMDA receptor-dependent and might therefore describe an active rather than a passive process[100]. Further evidence comes from two recent studies which claim that experience and not the passage of time is the driving force of representational drift[114,115]. In support of their finding, Blair et al. describe a stronger reorganization of the spatial firing in the hippocampus following an aversive event within a familiar environment[116]. Results by Geva et al. put forward the idea that there might be different mechanisms coding for time (changes in activity rates) and different experiences (changes in spatial tuning)[114,117].

Given the observed representational drift, is it still possible to retain important behaviorally relevant information from the hippocampal circuitry? A set of experiments suggests exactly that. Kinsky et al. showed that hippocampal neurons with behaviorally relevant coding exhibit heightened stability across multiple days[118]. Another study found that contextual coding in the hippocampus was maintained, albeit in the presence of representational drift in terms of spatial coding[119]. Along these lines, it has been suggested that the drift in neuronal activity could be constrained to directions unrelated to behaviorally relevant encoding[120,121,107]. Changes in neural activity could therefore occur as long as the correct output patterns would be generated for the respective inputs[122]. Alternatively, plasticity based on a local learning rule in downstream populations could potentially compensate for representational drift and, therefore, correct for changes in hippocampal output patterns[123].

As a last thought, findings of representational drift outside the hippocampus indicate that changes in representations are not constrained to memory and might therefore describe

a more general phenomenon with potentially different implications for different brain areas[124–126,123,127].

1.2 Medial prefrontal cortex

The medial prefrontal cortex (mPFC) has been assigned a variety of functions related to decision-making and adaptive behavior[12]. For example, the mPFC is involved in reward-guided learning and cognitive control[128,129]. Lesions of the mPFC resulted in reduced behavioral flexibility and impaired the ability to shift between different strategies[130,80].

In addition, several studies uncovered the mPFC's role in retrieving recent and remote memories - a prerequisite for adaptive behavior[131,132]. The mPFC is also involved in processing short-term memories and is important for the consolidation of memories[13,14]. Disrupting mPFC activity immediately after learning impaired the memory performance for a range of tasks and the consolidation of a fear memory[133–137]. The critical temporal window for consolidation in the mPFC seems to be 1 to 2 hours after learning because memory impairment does not occur for disruptions beyond this window[12]. With regard to memory processes, the mPFC's functional interaction with the hippocampus is crucial. A multitude of direct and indirect pathways between both areas have been identified[138]. These allow for the bi-directional communication between the hippocampus and the mPFC, a functional connection that has been shown to support many cognitive processes[139,140,130,141–143].

1.3 Rest periods and memory consolidation

There is ample evidence that sleep and rest periods greatly enhance memory performance[9]. One of the first studies investigating the influence of sleep on remembering nonsense syllables dates back to the year 1924. Jenkins and Dallenbach found that forgetting happened at a much slower rate when subjects slept during the intervening period instead of being awake[144]. Shortly after, two main brain states during sleep were discovered: slow-wave sleep (SWS or NREM) with slow oscillations and rapid eye movement (REM, also referred to as paradoxical sleep) sleep resembling awake neural activity[10]. Early studies focused on REM sleep because of its association with dreaming and found that REM sleep-deprived rats showed deficits in memory and learning[145–147]. Further confirmation of REM sleep's role in memory consolidation came from experiments involving the presentation of a learning-associated cue while the animal was asleep. Animals that received the cue during bouts of REM sleep were much better at remembering the task afterwards[148]. But how are memories, encoded during behavior, strengthened during REM sleep? A study by Ravassard et al. provided the first evidence that long-term potentiation during REM sleep might be crucial for the stabilization of certain memories[149]. Additionally, REM sleep is characterized by strong theta (6-10 Hz) oscillations in the hippocampus and these oscillations potentially facilitate memory consolidation[150–152].

However, the role of REM sleep in the consolidation process is still controversial and evidence from human studies underlined the great importance of NREM sleep for the consolidation and retention of hippocampal-dependent memories[153,154]. Many NREM-centered investigations of memory consolidation were inspired by the two-stage model of memory consolidation brought forward by Gyorgy Buzsaki[155,18]. According to this model, information from the neocortex is transferred to the CA3, triggering heterosynaptic

potentiation during behavior. In subsequent sleep and rest, long-term synaptic modifications are induced in the same subsets of CA3 neurons and their targets in CA1 through high synchrony activity accompanied by high-frequency oscillations (sharp wave ripples, SWRs) during NREM phases. This implied that activity observed during behavior should be reactivated in a "time-compressed manner" to allow the strengthening of synapses (long-term potentiation) during SWRs in sleep. Many studies in the following years confirmed the compressed, sequential replay of awake hippocampal activity during SWRs in subsequent sleep and rest, providing support for the two-stage model of memory[156–158]. Along these lines, it was also shown that reactivation of hippocampal activity and the number of SWRs was correlated with the successful acquisition and recall of a task-related memory[84,159,160]. Reactivated activity in NREM sleep has been detected as recurring single-cell activity, correlated firing between pairs of cells, and preserved temporal order of firing strongly resembling the neural activity from the encoding phase[161,162,157,158].

Due to the advancement in technology, recent animal studies could directly relate reactivations during SWRs to memory performance: disruption of hippocampal ripples during post-training consolidation periods caused an impairment of the associated memory[163,164]. On the other hand, artificially prolonging SWRs in the hippocampus facilitated subsequent memory recall[165]. As further proof of the outstanding role of reactivations during SWRs, Grydchin et al. selectively impaired a distinct spatial memory by inhibiting SWRs reactivations of the associated experience through optogenetics[166].

However, the observation that many memories are only hippocampus-dependent during the encoding and consolidation period led to the realization that memory processes must involve brain areas beyond the hippocampus[167,161,15,16]. In the system consolidation view, the replay of hippocampal activity during SWRs in slow-wave sleep triggers the gradual transformation and integration of memory-relevant representations in neocortical networks[161]. The observation of simultaneous widespread cortical activation in correspondence with SWRs additionally boosted scientists' interest in NREM sleep's role in memory consolidation[168–170]. It shall be noted though, that neocortical areas are involved in the acquisition and encoding phase as well[171,172]. In the special case of episodic memories, components of the experience are believed to be encoded in different brain areas, representing certain aspects of the experience, whereas the hippocampus binds them together[161,15,173]. By reinstating the respective hippocampal representation during consolidation periods, related representations distributed in the neocortex would be simultaneously reactivated[9]. This mechanism could potentially strengthen representations in the neocortex and allow the integration of newly acquired information into pre-existing memories[161,174]. The recall of most (non-spatial) memories stored in the neocortex would then become hippocampal-independent after the consolidation period[161,16]. A set of experiments has shown the replay of behaviorally relevant activity in the prefrontal cortex during subsequent sleep and rest, underpinning the system consolidation theory[175,176]. Further evidence comes from studies in which mPFC activity was disrupted immediately after learning and caused severe memory impairments[133,12,136,14].

But how is the activity between the hippocampus and neocortical areas synchronized during the consolidation period? There is evidence that hippocampal SWRs coincide with certain intervals of cortical spindles (10-18Hz) and that these "spindle-ripple events" facilitate the hippocampus-to-neocortical transmission of reactivated memory information[177,178,161,179–184]. Additionally, hippocampal ripple activity is modulated by slow oscillations (<1.5Hz) generated in the neocortex, leading to a suppression of ripples

during the down-state and an up-regulation during the up-state[185,179]. Through these mechanisms, the cortex could potentially trigger hippocampal reactivations selectively, which then, in turn, would reinforce cortical reactivations of relevant neural patterns[12,183].

Another phenomenon related to memory consolidation in the brain during sleep is synaptic rescaling[161]. Memory representations for long-term storage need to be selectively strengthened, whereas representations of the initial encoding period are believed to be weakened. This can be achieved by synaptic upscaling and downscaling, respectively[186]. Overall, the number of synapses in the cortex, functional synaptic strength, hippocampal and cortical firing rates, and cortical excitability decrease throughout sleep in agreement with the synaptic homeostasis hypothesis[187–189,161,190–193,106]. However, locally the number of synapses, the area of synaptic contact, and the firing rates of subgroups of cells do increase, a sign of the selective strengthening of distinct representations[194–197]. In this regard, REM and NREM are believed to play different roles in the synaptic up- and downscaling; contrasting results from different studies make a clear conclusion difficult though[161]. Nevertheless, there seems to be a tendency of REM sleep to be more strongly involved in the global synaptic downscaling, while maintaining newly formed synapses[161,196]. NREM sleep, on the other hand, might be primarily involved in strengthening newly formed connections and promoting the formation of new ones[161,197].

Although most studies have focused on either NREM or REM sleep, there have been early attempts to combine memory processes occurring during both sleep stages into one joint framework[198,10]: the sequential hypothesis states that NREM and REM sleep participate in the offline consolidation process in a sequential manner. In this way, the selection of relevant information could be performed during one sleep stage, whereas the processing and integration into existing knowledge would be done in the other[198,108,10,2].

1.4 Aims of the Thesis

Sleep's role in memory consolidation has attracted much attention in the past 100 years[9,161]. The observation that sleep significantly enhances memory performance and the prediction and detection of reactivated activity during consolidation periods have led to a wide range of follow-up studies. Advances in technology enabled scientists to directly manipulate the consolidation process and therefore impair or reinforce the recall of distinct memories.

However, the specific roles of NREM and REM sleep in the consolidation process are still topics of ongoing research. So are mechanisms enabling the transfer and transformation of newly formed memories into existing knowledge - a process that involves multiple brain areas. By examining the neural activity during rest and sleep from the perspectives of reactivations, consolidation, and its contribution to awake representational drift, my Ph.D. thesis aims to answer the following questions:

1. How stable is the activity during rest and is there a connection with (a) changes in neural activity across behavioral episodes and, (b) prior behavioral performance?

2. How do long-duration rest episodes affect neural activity and do different sleep stages play a role?

3. If there are systematic changes, is this a population-wide effect, or do different cells contribute to a different extent?

4. Are systematic changes in neural activity during rest from different, interacting populations (HPC & PFC) synchronized?

5. If systematic changes are synchronized between the hippocampus and PFC, is this coherence dependent on prior performance, and if yes, what drives the synchrony?

CHAPTER 2

Sleep stages antagonistically modulate reactivation drift

2.1 Abstract

Acquiring, consolidating, and recalling spatial memories are critical brain functions for survival. In this context, a major role has been assigned to the reactivation of awake hippocampal activity during rest periods, which was shown to significantly contribute to memory performance. Thus far, studies have only considered reactivations strongly resembling prior neural activity ignoring the possibility that reactivations could change as part of the consolidation process. I addressed this option by tracking hippocampal reactivations during an extended rest period (∼20h) following a spatial-memory acquisition paradigm.

I found that reactivations during long rest periods underwent a gradual transformation. At the beginning of rest, reactivations strongly mirrored the neural activity from the acquisition of the memory, as previously shown. However, with increasing time during rest, reactivated activity progressively resembled neural activity observed during the recall of the memory. This drift in reactivation activity occurred mainly during non-rapid eye movement (NREM) periods and was driven by firing rate changes of a plastic subset of cells. Conversely, rapid eye movement (REM) sleep showed a resetting effect and drove the neural activity during rest toward prior observed activity. A persistent subset of cells maintained their firing rates from the acquisition to the recall and therefore withstood the reactivation drift. These persistent cells showed constant spatial coding before and after sleep and, during memory acquisition, showed similar remapping characteristics to the supposedly more plastic cells.

My findings unveil the active role of rest in reshaping memory representations. This process was mainly driven by firing rate changes in a subset of cells, while a minority of cells maintained their firing rates. The modifications during rest could potentially contribute to representational drift across days previously described in the literature. Furthermore, the offline reactivation drift and distinct effects of different sleep stages suggest that the gradual changes in memory representations could be part of a brain-wide consolidation process.

2.2 Introduction

The Hippocampus has been characterized as the initial storage site for spatial memories in the brain. Hippocampal neuronal activity changes during the acquisition of a memory and is replayed in subsequent periods of rest[199,76,181]; these reactivations of neural activity patterns play a critical role in the consolidation of recently established memories and have been linked to memory performance[9,84,188,161,200]. For example, prolonging the sharp-wave ripples (SWRs) that usually accompany hippocampal reactivations facilitates subsequent memory recall[201]. On the other hand, disruption of SWRs caused impaired memory performance[164,166]. Although some studies showed that reactivation events are observed during extended periods of rest[202], the specific reactivated activity patterns have mostly been studied during short rest periods (<1h). Activity associated with recent experience is preferentially reinstated during the initial phases of rest, but this time interval might only capture the initial phase of memory consolidation[203]. In addition, the study of reactivations has been restricted to the identification of activity strongly resembling previously observed neural activity.

However, there is some evidence that reactivations in the visual cortex are predictive of future stimulus responses and therefore reactivations might incorporate additional functionality rather than simply replaying constant activity patterns[204]. It is unknown whether hippocampal reactivations show similar dynamics. In support of less static hippocampal representations, recent studies have shown that the hippocampal coding of the same environment changes through repeated experiences[114,115,102]. The representational drift has been linked to the updating or processing of memories. This mechanism could potentially be supported by changes in reactivation content during offline periods that are predictive of changes occurring between experiences.

In this study, I tracked hippocampal reactivation dynamics during an extended rest period following a spatial memory paradigm to assess changes and reactivation relationship to the acquisition and recall of the memory. I found that hippocampal activity during rest gradually changes from acquisition-like to recall-like. This process was driven by non-rapid eye movement (NREM) sleep, whereas rapid eye movement (REM) sleep showed a resetting effect. I identified firing rate changes of a plastic subset of neurons as the main cause of the memory drift. A smaller subset of cells maintained its firing rates and spatial coding from the acquisition to the recall acting as a stable backbone. Therefore, offline reactivations of memory-relevant activity might directly contribute to changes in neural representations over days, a phenomenon that might support updating and processing of memories.

2.3 Results

Stable recordings during acquisition, long rest, and successful recall of a spatial memory

Principal neurons in the dorsal CA1 of three long-Evans rats were recorded using bilaterally implanted 32-channel tetrode microdrives (see Methods). The food-deprived animals learned a set of three reward locations on the cheeseboard and successfully recalled this spatial memory after an extended rest period (Fig. 2.1a). The excess path (deviation from the optimal path between reward locations) on the first trial of recall was not significantly

longer than the excess path of the third trial during acquisition (Fig. 2.1b). Hence, animals were able to immediately recall the direct paths between reward locations after the extended rest period.

To confirm the stability of the tetrode recordings throughout the experiment, I analyzed the temporal stability of waveforms of identified spikes from all neurons (Fig. 2.1c-f). This allowed me to assess whether spikes from different neurons were well separated and if they showed temporal consistency in terms of their waveform shape.

First, I confirmed that for each cell the first twelve waveform based clustering features remained stable over time (Figure 2.1d). I then investigated whether waveforms of different neurons were well separated for the duration of the experiment. To do so, I compared the distance between the cell's mean waveform from the nth hour and the first hour with the distance between the cell's mean waveform from the nth hour and the other cells' mean waveforms from the first hour (see Methods). The distance to each cell's mean waveform from the first hour was significantly smaller yielding a ratio much smaller than one and no significant differences for different time points during the experiment (Fig. 2.1e). The same was true when I compared the distance to each cell's mean waveform with the distance of other cells' mean waveform from the nth hour with the cell's mean waveform from the first hour (Fig. S2.1a). Next, I compared the distance between the mean waveform of the first hour with the mean waveform of the last hour for each cell with the across cells' waveform distance during the first hour (Fig. 2.1f). The across-cell distance during the first hour was significantly larger than the distance between mean waveforms from the first and the last hour of the experiment for each cell ($p < 0.001$). Therefore, single-cell waveforms were much more similar to each other than compared to other cells' waveforms. Taken together, my analysis confirmed the separation of single units and the stability of our recordings based on spike waveform-based features for the entire duration of the experiment.

Sleep neural activity undergoes transformation from patterns observed during acquisition to patterns occurring during recall

After confirming the stability of the recordings, I investigated reactivations during the extended rest period separating the acquisition and the recall of the spatial memory. First, I fit two Poisson hidden Markov models (pHMMs) to the acquisition and the recall neural data, respectively. This allowed me to capture stereotypical neural patterns in both behavioral episodes. By relating the activity during rest to these representations I was able to identify acquisition-like and recall-like neural patterns. The optimal number of hidden states for both pHMMs was derived using the cross-validated log-likelihood (Fig. S2.2a). The optimal number was comparable for the acquisition and the recall model (Fig. S2.2b).

Next, I verified that the pHMMs accurately reproduced basic statistics of the neural data (Fig. S2.2c-d). By aligning hidden state activations with the location tracking data of the animal, I was able to further analyze single states. Despite not providing location data to the model explicitly, most hidden states were spatially constrained with a large proportion of states showing activations around reward locations (Fig. S2.3a-b). Furthermore, I was able to decode the animal's location using pHMM state activations (see Methods) with errors comparable to standard Bayesian decoding (Fig. S2.2c-f). In summary, the pHMMs accurately captured the neural activity during acquisition and recall. In addition,

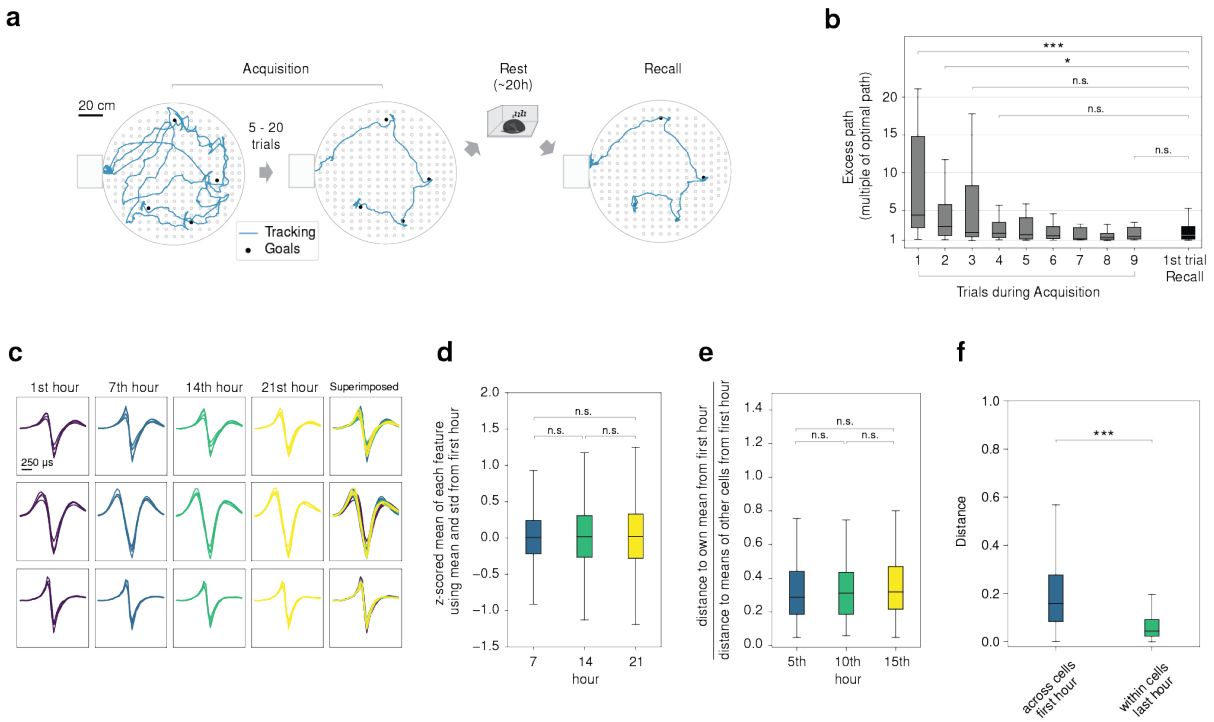

Figure 2.1: **Stable wireless tetrode recordings in CA1 during extended-time learning paradigm.** (a) Structure of analyzed behavioral sessions and tracking data for one example session: learning a novel set of goal locations (acquisition), long rest period (20h), and re-exposure to the same goal arrangement for memory assessment (recall). (b) Excess path measured as a multiple of the optimal path between goals during the first nine trials of memory acquisition and the first trial of recall. The excess paths of the first trial during recall and trials three to nine during acquisition are not significantly different (data from all sessions, * $p < 0.05$, *** $p < 0.001$, n.s. $p > 0.05$, one-sided Mann-Whitney U test, corrected for multiple comparisons). (c) Waveforms over the 25 hours recording period for three exemplary principal cells. (d) Distribution of first 12 clustering features z-scored for different time intervals for all sessions. The mean of each feature is computed in the corresponding time window and z-scored using the mean and std of the same feature during the first hour ($p > 0.4$ for all comparisons, two-sided Mann-Whitney U test). (e) Distance with the cell's mean from the first hour versus distance with other cells' means from the first hour for one example session. For each time interval the distance (1-Pearson R) between the first 12 clustering features during the interval (5th, 10th, 15th hour), and the mean of the first 12 clustering features from the first hour is calculated (no sign. Difference between 7th and 21st hour for all sessions, $p>0.15$, two-sided Mann-Whitney U test). (f) Waveform distance per cell for one example session. For each cell, the mean of the first 12 clustering features during the first hour is computed. Then, the across-cell distance is computed using the mean features of the cell of the first hour and the other cells' features from the first hour. The within distance is the distance (1-Pearson R) of the cell's mean from the first hour with its waveforms during the last hour of the experiment. There was a s´ignificant difference for all sessions ($p < 0.001$, two-sided Mann-Whitney U test).

hidden states represented spatial and non-spatial components of the neural activity during behavior.

To evaluate the similarity of the neural activity during rest with either acquisition or

recall, I first looked at the probability of decoding either an acquisition or a recall state (Fig. 2.2a-c). Initially, acquisition states were much more likely to be decoded than recall states. By the end of sleep recall states had a higher probability of being decoded ($p < 10e-8$, one-sided t-test).

Next, I computed the drift score to analyze the rest neural activity on finer timescales: for each neural activity bin (12 spikes, see Methods), the normalized difference between maximum log-likelihoods given the acquisition pHMM and the recall pHMM was calculated (Fig. 2.2a and Methods). Confirming the changes in decoding probabilities, the drift score was initially below zero, corresponding to greater similarity of the rest neural activity with the acquisition, and gradually moved to values greater than zero at the end of rest (Fig. 2.2d-e). However, the drift did not occur uniformly throughout sleep. The first half of sleep contributed more strongly to the drift than the second half (Fig. S2.4a). I confirmed my findings using a standard Bayesian decoding approach (see Methods and Fig. S2.4b-c).

To check whether the drift was driven by an overall decrease in decoding quality or by the actual transformation of the neural activity during rest from acquisition-like to recall-like, I performed additional analysis steps. Although the number of significant SWR reactivations decreased throughout rest due to fewer SWRs during later stages of rest (Fig. S2.4d), the fraction of significant reactivations stayed constant (Fig. S2.4e-g). Maximum log-likelihoods of the acquisition pHMM generally decreased throughout rest, whereas maximum log-likelihoods of the recall pHMM increased (Fig. S2.4h-i). When I projected the rest neural activity onto the vector between the decoded acquisition and the decoded recall state (Fig. S2.4j and Methods), I found that the relative distance to the decoded acquisition states continuously increased ($p=0.0017$, two-sided t-test). On the other hand, the relative distance to the decoded recall state was reduced ($p=0.012$, two-sided t-test).

In summary, the neural activity during rest gradually changed from being very similar to the acquisition of the memory at the beginning of rest towards being more similar to the neural activity during recall at the end of rest.

Non-rapid eye movement sleep accelerates changes, while rapid eye movement sleep has a resetting effect

Sleep stages differ in their physiological properties and are thought to serve different functions in the memory consolidation process. I separated the data based on the theta/delta ratio of the local field potential (LFP) into non-rapid eye movement sleep (NREM) and rapid eye movement sleep (REM) to assess the contributions of each to the overall drift in the rest neural activity. Surprisingly, REM and NREM sleep epochs had opposing effects on the drift (Fig. 2.3a-c). While NREM sleep showed a positive contribution in the transition from acquisition-like to recall-like patterns of activity, REM sleep had a negative, oppositely directed effect, akin to resetting the memory representation to its previous state (Fig. 2.3b).

The same applied to the cumulative contribution of NREM and REM epochs to the drift (Fig. 2.3c and Fig. S2.5a). When I compared the overall change in the drift score throughout rest with the cumulative change of NREM and REM epochs, I found a difference of two orders of magnitude (Fig. 2.3d). Short-term changes in drift score occurring during NREM and REM epochs were not correlated with the longer-term

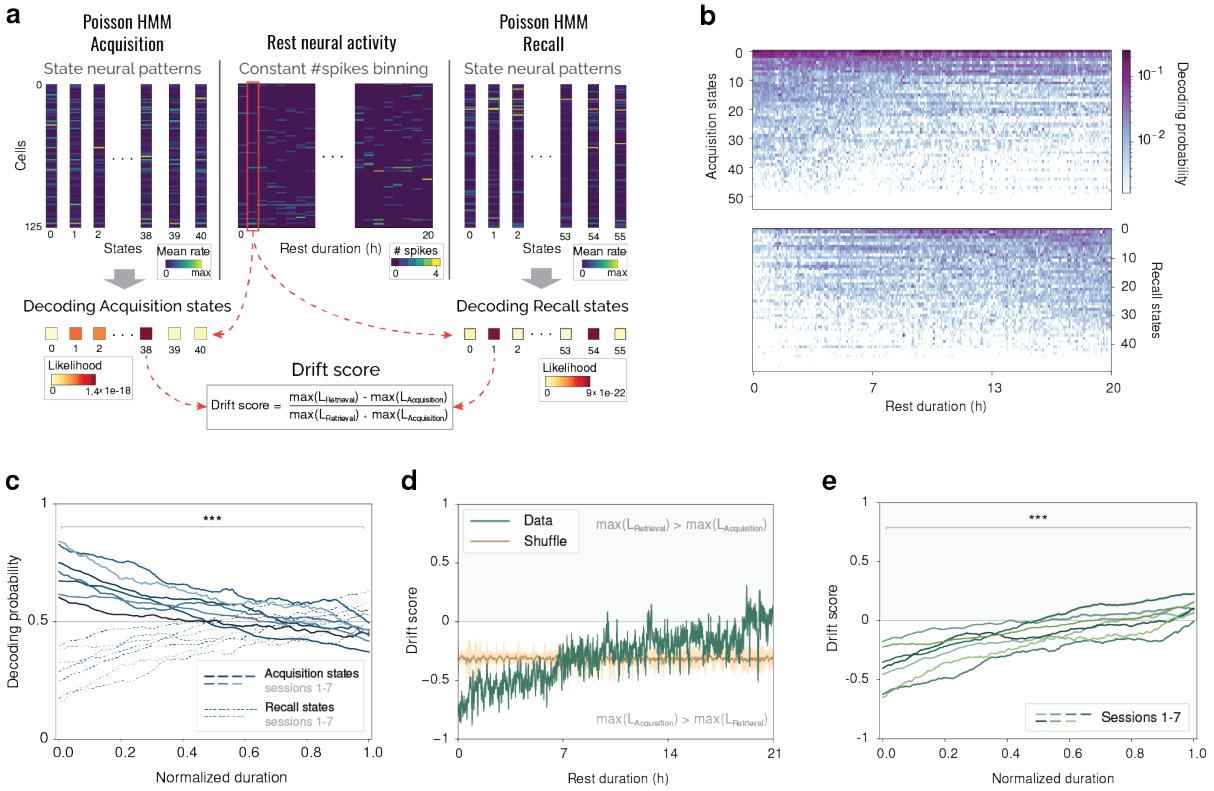

Figure 2.2: **Rest neural activity continuously transitions from being similar to acquisition neural activity towards greater similarity with recall neural activity.** (a) Decoding neural activity during rest. Two separate Poisson Hidden Markov models are fit to the awake neural data from acquisition and recall. Rest data is binned using a constant number of 12 spikes. The data likelihood for each bin is computed using acquisition and recall states from the Poisson HMM models. The Drift score is calculated using the maximum likelihoods from acquisition and recall. (b) Reactivation probabilities for acquisition and recall pHMM states during sleep for one example session. (c) Summed reactivation probabilities for acquisition and recall pHMM states respectively for all sessions. The summed reactivation probabilities at t=0 differ significantly from the corresponding values at t=1 ($p < 10e\text{-}8$, one-sided t-test). (d) Drift score (normalized difference between max. recall and max. acquisition likelihoods) computed using rest activity as a function of time for one example session (green). Time bin shuffle (n=500) mean and standard deviation (yellow). (e) Drift scores as a function of normalized sleep duration for all sessions. The Drift scores are below zero at t=0 ($p < 0.001$, one-sided t-test), above zero for t=1 ($p < 0.01$, one-sided t-test), and significantly different between t=0 and t=1 ($p < 10e\text{-}5$, one-sided t-test).

changes (Fig. S2.5b). Hence, the drift score fluctuated strongly on shorter timescales without a clear direction, whereas the long-term change was directed and much smaller in magnitude.

Next, I tested if the opposing contributions of REM and NREM sleep were an overall effect or whether there was a direct link to interactions between subsequent sleep epochs. The contributions of neighboring epochs to the drift were anti-correlated while the contributions of non-neighboring epochs were not correlated (Fig. 2.3e-h). This effect was even stronger using a fixed succession of NREM epochs being followed by REM epochs (Fig. 2.3e). Sleep epochs of the same type (separated by a sleep epoch of the opposite type) were only

weakly correlated (Fig. 2.3g-h). Taken together, these findings suggest that there is an immediate correspondence between subsequent NREM and REM epochs. The sleep neural activity during NREM epochs undergoes modifications that make it less acquisition-like and more recall-like. During subsequent REM epochs, these modifications are partially reset making the sleep neural activity again more similar to the activity observed during the acquisition.

In support of this finding, I found that the normalized likelihoods per decoded pHMM state were greater during REM epochs than during NREM epochs (Fig. S2.5c-d), implying a more "pure" reactivation of acquisition states in REM sleep. Additionally, the temporal dynamics of the decoded activity resembled the awake activity much better during REM epochs (Fig. S2.5e-j).

Next, I explored if differences in reactivated content could partially explain why the activity during REM epochs was a better reflection of the activity during the acquisition. E.g. the possibility that certain locations or awake activity patterns were only reactivated during REM and not during NREM yielding cleaner reactivations for the former. First, I focused on the decoded spatial locations during sleep using the Bayesian decoding approach (Fig. S2.6a-b). Decoded positions during REM and NREM were highly correlated. Secondly, I computed the number of times a pHMM state was decoded during either REM or NREM epochs (Fig. S2.6c). I found that pHMM state reactivations in REM and NREM were highly correlated. Both results (Fig. S2.6d) suggest that the content being reactivated during REM and NREM strongly overlaps. Therefore, differences in reactivated content cannot explain the differential effect of NREM and REM sleep stages on the drift.

Different subsets of cells contribute distinctively to reactivation drift

I then asked if the drift was a population-wide phenomenon or whether different cells contributed distinctively. To do so, I separated cells into three groups based on the difference in their firing rate distributions between acquisition and recall: an increasing, a decreasing, and a persistent subset (Fig. 2.4a-b and Methods). Not surprisingly, the difference in mean firing rates was confirmed when I looked at the acquisition and recall states: persistent cells showed similar mean firing rates, increasing cells increased their firing rates from acquisition to recall states and decreasing cells reduced their firing rates (Fig. S2.7a).

Interestingly, more than half of the recorded cells were from the decreasing group, and only a smaller portion of cells had stable firing rates over the entire experiment (Fig. S2.7b). Cells with persistent activity, on average, exhibited lower firing rates than the changing firing rate groups but maintained a comparatively higher rate during sleep (Fig. S2.7c-e and Fig. 2.4c). I did not find differences in SWR firing gain or waveform stability between the different subsets (Fig. S2.7f-g), nor did cells belonging to the same subset cluster around one tetrode (Fig. S2.7h).

Next, I computed the drift score using the identified subsets of cells: using only persistent cells yielded a stable drift score around zero - implying similar activity patterns during acquisition and recall and consistent reactivations during rest (Fig. 2.4d-e). However, using either decreasing or increasing cells for the decoding resulted in drift score changes similar to the ones calculated using all cells. Therefore, unstable, rate-changing cell groups were the primary driver of the observed reactivation drift.

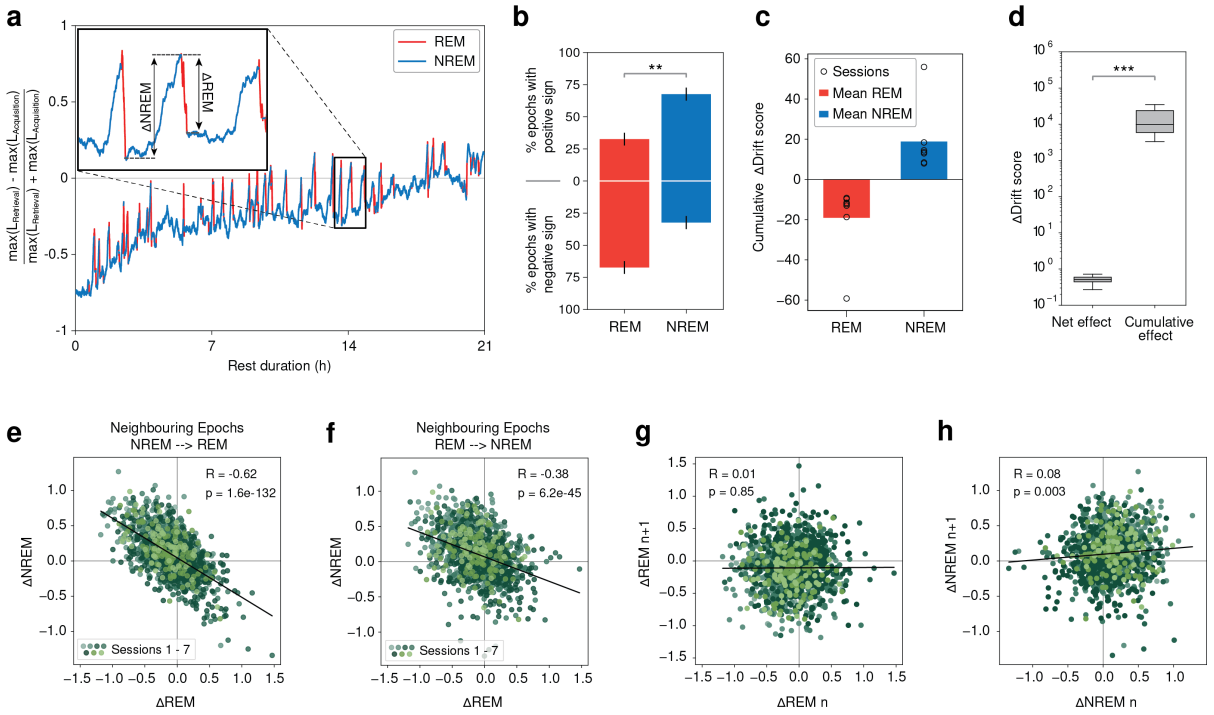

Figure 2.3: **NREM and REM sleep stages show opposing effects on the observed drift.** (a) Drift score as a function of sleep duration (REM periods in red, NREM periods in blue) for one example session. Computation of the ΔDrift score for the respective sleep stage (inset). (b) Contributions of REM and NREM epochs to drift. For each epoch, the delta score was computed as the difference in the drift score between end and beginning. Percentage of epochs with positive and negative deltas are depicted for REM (red) and NREM (blue). Data from all sessions (mean±SEM, ** $p < 0.001$, two-sided Mann-Whitney U test). (c) Summed ΔDrift score for REM and NREM epochs for all sessions ($p < 0.01$, two-sided Mann-Whitney U test). (d) The net effect of change in drift score (difference in drift score between the beginning and the end of rest) and the cumulative effect of change in drift score (sum of absolute ΔDrift score for REM and NREM periods). Data from all sessions ($p < 0.01$, two-sided Mann-Whitney U test). (e)-(f) ΔDrift score values of neighboring NREM and REM epochs considering the order of sleep periods. (e) NREM epoch followed by a REM epoch, R = -0.62, p = 1.6e-132. (f) REM epoch followed by an NREM epoch: R = -0.38, p = 6.2e-45. (g)-(h) ΔDrift score values of subsequent sleep epochs of the same type. (g) R = 0.01, p = 0.85, REM. (h) R = 0.08, p = 0.003, NREM.

Firing rate modulations of an unstable cell population correlate with changes in Drift score

The observed reactivation drift could have been driven by changes in the coordinated firing between cells, changes in firing rates of single cells, or a combination of both. Since I identified rate-changing cells as strong contributors to the reactivation drift, I hypothesized that rate changes are the main cause of the drift. First, I investigated the role of sleep-related changes in the structure of pairwise activity correlations after having controlled for firing rate changes (Fig. S2.8a-b and Methods). Consistently with a rate-dominated drift, I could find only minor rearrangements in the correlation values of persistent cells, showing that cell co-activations were only mildly affected, besides the

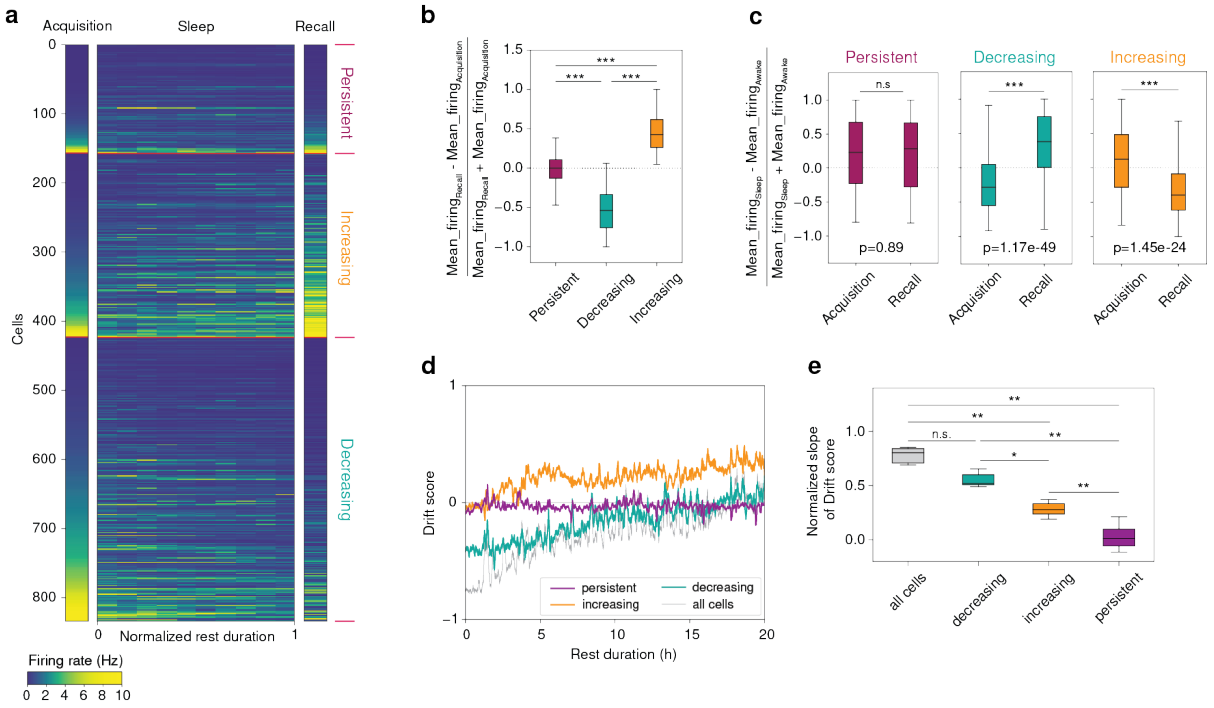

Figure 2.4: **Cells that modulate their firing from acquisition to recall drive memory drift, while persistent cells reduce it.** (a) Average firing rates of persistent, increasing, and decreasing cells during acquisition, sleep, and recall (all sessions). Cells were classified based on the difference between their firing rate distributions during Acquisition and Recall. (b) Normalized firing rate change from Acquisition to Recall for persistent, decreasing, and increasing cells (all $p < 0.001$, two-sided Mann-Whitney U test, corrected for multiple comparisons). (c) Mean firing rate ratio for decreasing, increasing, and persistent cells using mean firing from acquisition, sleep, and recall (n.s. $p > 0.05$, ** $p < 0.001$, two-sided Mann-Whitney U test). (d) Drift score using either all cells (gray), only persistent cells (violet), only increasing cells (orange), or only decreasing cells (turquoise) for one example session. (e) Normalized slope fit to Drift score computed using different subsets of cells (** $p < 0.01$, * $p < 0.05$, n.s.: $p > 0.05$, two-sided Mann-Whitney U test, corrected for multiple comparisons).

overall changes in firing rate in different cell groups. The pairwise activity correlations using all cells did show small but significant alterations during sleep, not surprisingly, given the use of a mixture of cells moving in and out of the reactivated assemblies and spanning a wide interval of average firing rate changes (Fig. S2.8b).

I then correlated firing probability changes and the drift score changes for each NREM and REM epoch to directly measure the contribution of firing rate modulations (Fig. 2.5a). For the increasing cell group, I found a positive correlation (Fig. 2.5b-c). This result suggests that increases in firing probability for this subset during NREM go hand-in-hand with increases in the drift score. At the same time, a stronger decrease in firing probability resulted in a stronger "reset" of the drift score during REM epochs. On the other hand, decreasing cells showed a negative correlation (Fig. 2.5d-e). Therefore, a decrease in firing probability for decreasing cells causes an increase in the drift score during NREM. Increasing their firing probability in REM resulted in the reduction of the drift score. Persistent cells showed weaker, but significant correlations with drift score changes (Fig. 2.5f-g). Similarly to the drift score changes, neighboring REM and NREM

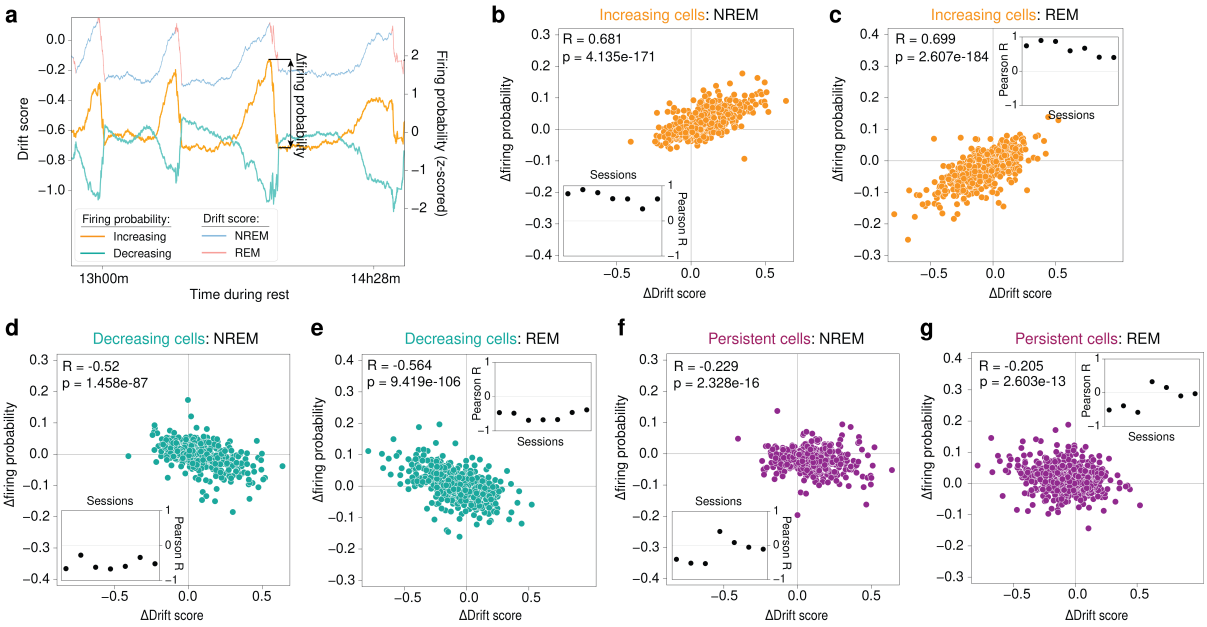

Figure 2.5: **Memory drift is driven by firing rate modulations of a subset of cells.** (a) Illustration of drift score and firing probability fluctuations of decreasing and increasing cells for one example session. (b-g) Correlation between ΔDrift score and the change in firing probability for increasing cells during NREM and REM periods. Inset: correlation values per session(increasing cells: NREM (b) R = 0.68, p = 4.13e-171, REM (c) R=0.70, p = 2.61e-184; decreasing cells: NREM (d) R = -0.52, p = 1.46e-87, REM (e) R = -0.56, p = 9.42e-106); persistent cells: NREM (f) R = -0.23, p = 2.33e-16, REM (g) R = -0.21, p = 2.60e-13).

periods exhibited correlated firing rate changes for both the decreasing and increasing groups with a stronger effect for NREM periods being followed by REM periods (Fig. S2.8c-f). Firing rate changes between sleep epochs of the same type were only weakly correlated (Fig. S2.8g-j).

In summary, firing rate modulations of unstable cells that occurred majorly during NREM periods were the main driver of the observed drift. Since cells were classified using the acquisition and recall periods, observed offline firing modulations are directly linked to how the representation of the spatial memory changes from acquisition to recall.

Persistent cell population maintains their population activity and spatial coding from acquisition to recall

The identified persistent subset of cells, which did not participate in the reactivation drift, showed consistent firing rates between the acquisition and recall of the spatial memory. Hence, I asked whether this subset of cells also maintained their spatial coding across the behavioral episodes. First, I looked at the spatial coding properties of the identified subsets. Persistent cells had lower spatial information than the decreasing cells during the acquisition and lower spatial information than increasing cells during the recall (Fig. 2.6a). Therefore, the persistent subset did not exhibit superior spatial coding ability compared to the more active unstable subsets during the respective behavioral session.

However, when I computed the population vector correlations between acquisition and

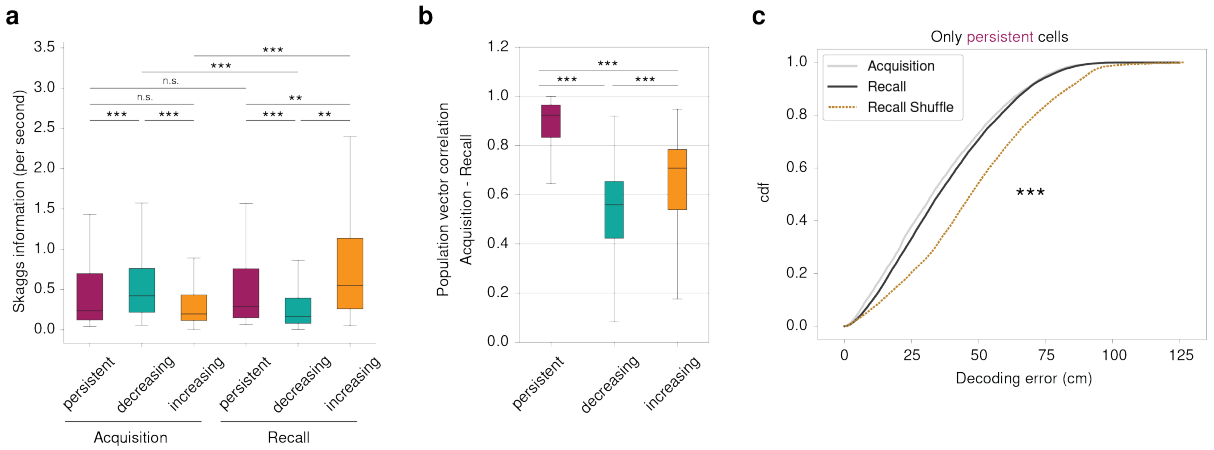

Figure 2.6: **Persistent cells maintain their spatial map and population activity from acquisition to recall.** (a) Skaggs spatial information for persistent, decreasing, and increasing cells during acquisition and recall (*** $p < 0.001$, ** $p < 0.01$, n.s. $p > 0.05$, two-sided Mann-Whitney U test, corrected for multiple comparisons). Only cells with a mean firing rate above 1 Hz were considered. (b) Mean population vector correlations between acquisition and recall for persistent, increasing, and decreasing cells (*** $p < 0.001$, two-sided Mann-Whitney U test, corrected for multiple comparisons). (c) Decoding error using Bayesian decoder for persistent cells trained on 50% of acquisition data. Error on held-out acquisition data (light gray) and error on recall data (dark gray). Spatial bin shuffle in brown. Both acquisition and recall have significantly lower decoding errors than shuffle (one-sided Mann-Whitney U test, both $p < 0.001$).

recall, I found that persistent cells showed a significantly greater overlap in terms of population-wide spatial coding (Fig. 2.6b). In support of this finding, a Bayesian decoder for animal position trained using only persistent cell activity from acquisition performed significantly above chance during recall (Fig. 2.6c). Taken together, persistent cells not only "bridged" the acquisition and recall of the memory by maintaining their firing rates, but also by having a consistent spatial map between these two episodes.

Similar amount of remapping for persistent and decreasing cells during the acquisition of the spatial memory

Since the persistent subset seemed to play a linking role between acquisition and recall, they might also play a distinctive role in the actual acquisition of the memory. Therefore, I investigated the changes in spatial coding of cells during the acquisition, first independently, and then with respect to preceding (pre-probe) and subsequent (post-probe) behavioral episodes (Fig. 2.7a). The coding for goal locations is potentially crucial for the cheeseboard task. Hence, I computed the distances between the peak firing location per cell and the closest goal location (Fig. 2.7b). Throughout acquisition, these distances decreased for persistent and decreasing cells without major differences between the two subsets (for distances during recall see Fig. S2.9a-b). Next, I looked at population vector correlations across behavioral episodes. First, I divided the experimental paradigm into blocks and computed mean population vector correlations between these blocks for persistent and decreasing cells (Fig. 2.7d). When I divided the population vector correlations between acquisition and recall by the values from acquisition vs. pre-probe (see red rectangles), the ratio increased continuously and leveled well above a value of 1 for persistent cells

(Fig. 2.7e). This result implies that learning-induced remapping of persistent cells during the acquisition carried over to the recall of the memory. Decreasing cells and increasing cells (Fig. S2.9c-d) did show an increase in population vector correlations with the recall during acquisition, but the values were below those of persistent cells and did not go above one. Therefore, the spatial map of decreasing cells at the end of acquisition was not more similar to the spatial map of the recall than to the spatial map of the pre-probe.

In conclusion, persistent and decreasing cell activity was modulated during the acquisition of the memory with persistent cells showing a greater connection with the recall activity for the population vector correlation analysis at the end of acquisition.

I then asked whether persistent cell activity was only consistent between the acquisition and recall or whether this subset maintained similar activity throughout all behavioral episodes of the experiment (Fig. 2.7a). To answer this question, I binned activity from all behavioral episodes (see Methods) and trained separate support-vector machines (SVMs) to evaluate how well behavioral episodes could be distinguished based on persistent cell activity (Fig. 2.7f-g). Interestingly, I found that decoding accuracy was the lowest, while still above chance, for the discrimination of acquisition and recall ($p > 0.01$, MWU test). This implies that, although persistent cell activity did change from the acquisition to the recall, the persistent cell activity during these two episodes was more similar to each other than to the persistent cell activity of any other episode.

Persistent cell activity during rest is more strongly modulated by the interleaved acquisition session than decreasing cell activity

After relating changes in neural activity during acquisition to preceding and subsequent behavioral episodes, I assessed its relation with rest activity before and after. Previous studies have shown that behaviorally relevant neural activity is replayed during subsequent rest periods[181]. Other studies demonstrated that certain neural activity patterns observed during behavior can be pre-played during prior rest periods[205].

I explored these two possibilities by first fitting pHMMs to the neural activity of decreasing and persistent cells during the acquisition, separately (Fig. 2.8a and Methods). Increasing cells were excluded from this analysis due to their very low firing rates during the initial behavioral episodes. Then, I computed the data likelihoods in the rest before the acquisition (pre-rest) and in the first hour of the rest after the acquisition (rest) for the decoded neural states given the two pHMMs. By computing a normalized ratio of these likelihoods in pre-rest and rest, I identified neural states that were either more likely to be active during rest before (ratio = +1) or rest after (ratio = -1) the acquisition (Fig. 2.8b). States with a ratio around zero were not differentially expressed between both rest episodes. Normalized ratios of neural states from the persistent cell pHMM had more ratios closer to 1 and -1 as compared to neural states from the decreasing cell pHMM ($p < 0.01$, MWU test using absolute ratios, Fig. 2.8c), suggesting that behavioral neural activity patterns for persistent cells were more selectively active in either the rest before or the rest after. In support of this finding, persistent cell activity during acquisition, as captured by the pHMM model, was significantly different from pre-probe activity (Fig. S2.9e). For increasing and decreasing cells, I did not observe a difference (Fig. S2.9f-g).

Decreasing and persistent cells had a similar number of states that were either replayed (normalized ratio > 0.5) or pre-played (normalized ratio < 0.5, S9h-i).

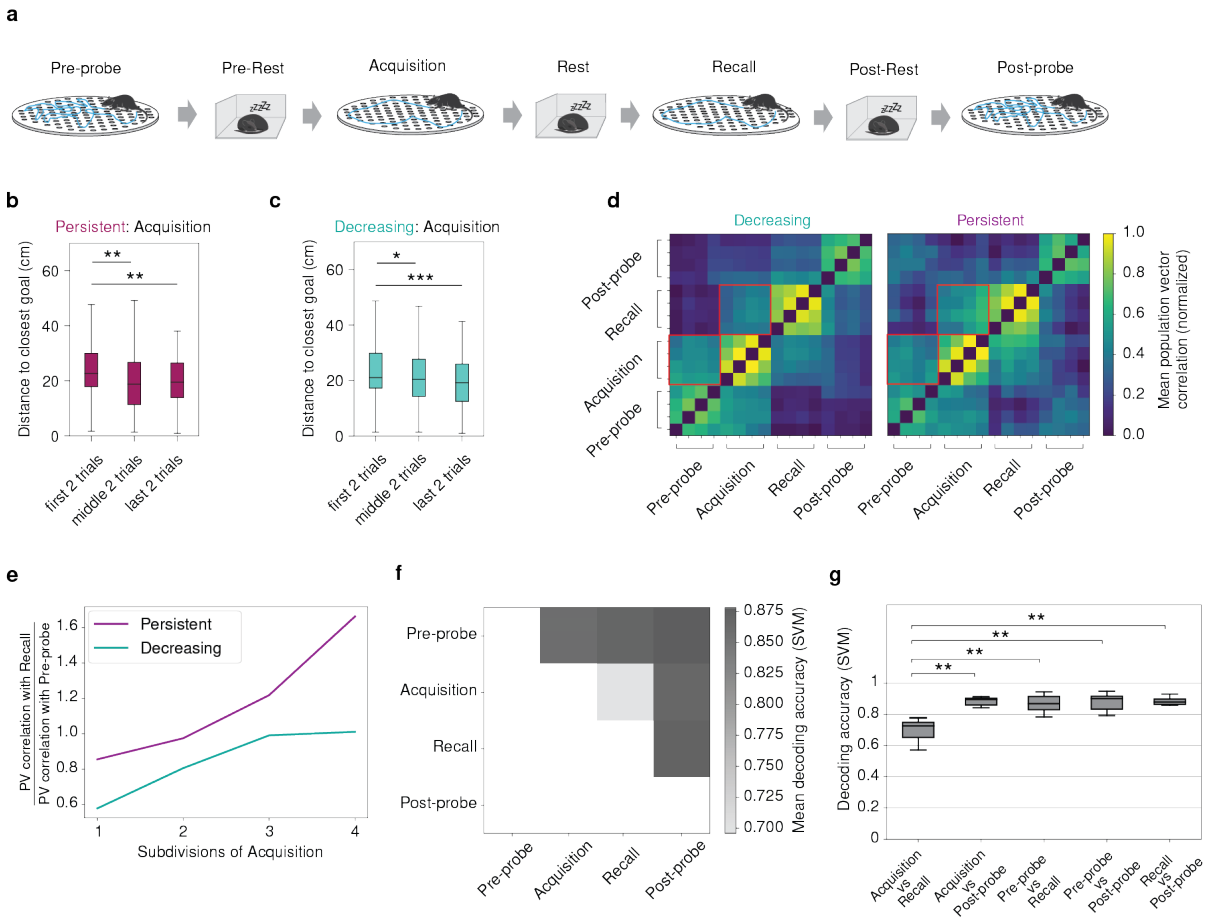

Figure 2.7: **Similar remapping for persistent and decreasing cells during acquisition.** (a) Schematic showing all experimental episodes. (b)-(c) Distances between place field peak firing locations and goal locations during acquisition for persistent (b) and decreasing (c) cells (data from all sessions, *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, two-sided Mann-Whitney U test, corrected for multiple comparisons). (d) Spatial population vector correlations between behavioral episodes. Spatial bins that were not visited in all behavioral episodes and the first trial of learning were excluded. All diagonal elements were set to zero and the remaining values normalized to lie within 0 and 1. (e) Ratio between mean population vector correlations (normalized) from pre-probe and acquisition and recall and acquisition for decreasing and persistent cells. (f) Mean decoding accuracy of SVM decoder trained to distinguish between behavioral episodes. Data from all sessions. (f Decoding accuracy of SVM decoder trained to distinguish between behavioral episodes (data from all sessions). Decoding accuracy is significantly lower for acquisition vs. recall ($p < 0.01$, two-sided Mann-Whitney U test). (g) Decoding accuracy of SVM decoder trained to distinguish between behavioral episodes for all sessions (** $p < 0.001$, two-sided Mann-Whitney U test).

Finally, I investigated whether the entire spectrum of the acquisition activity was reactivated during the first hour of the rest after the acquisition. To do so, I computed the fraction of pHMM states that were never decoded during this interval. I found that for persistent cells only a small part of the acquisition states were reactivated (Fig. 2.8d). On the contrary, most of the decreasing cell pHMM states were decoded during the first hour of rest.

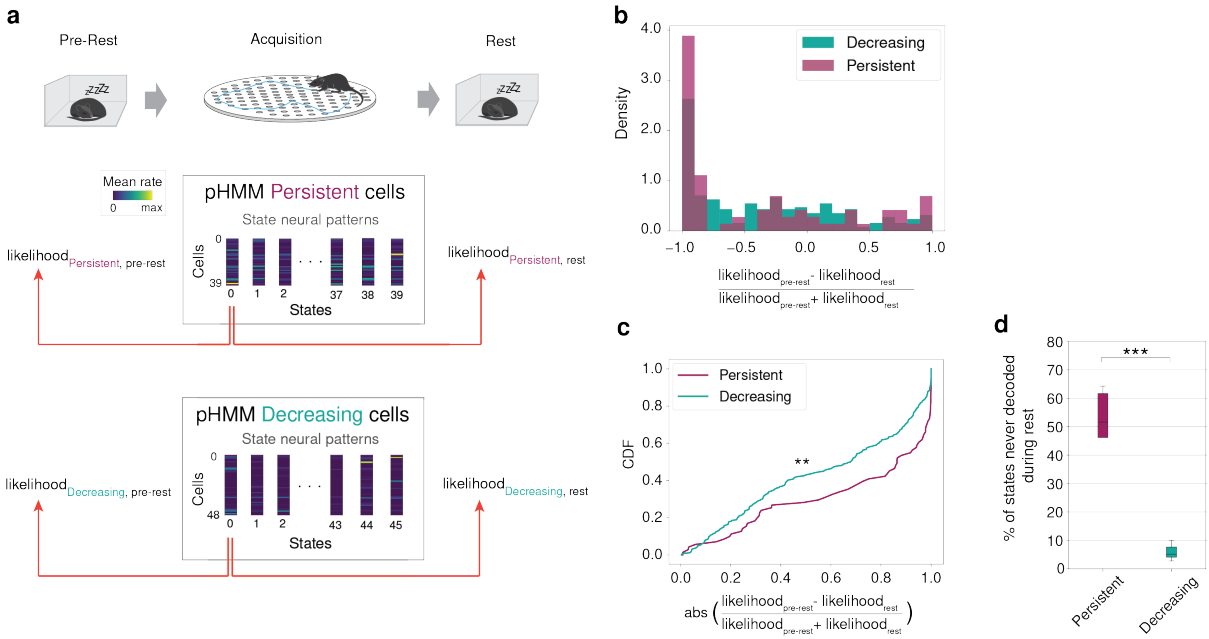

Figure 2.8: **Persistent cell activity during rest is more strongly influenced by the interleaved acquisition session than decreasing cell activity.** (a) Schematic showing persistent and decreasing cell pHMM models during acquisition and likelihood computation in pre-sleep and sleep (only the first hour of the long rest was used to compute the likelihood in sleep.). (b) Difference in expression between pre-rest and first part of long rest for decoded persistent cell pHMM states and decreasing cell pHMM states (see Methods). A normalized likelihood ratio < -0.5 defines pre-play states, whereas a normalized likelihood ratio > 0.5 corresponds to pre-play states. (c) Absolute differential expression of decoded persistent cell pHMM states and decoded decreasing cell pHMM states (** $p < 0.01$, two-sided Mann-Whitney U test). (d) Number of acquisition states that were never decoded during the first hour of long rest for decreasing and persistent cells (*** $p < 0.001$, Mann-Whitney U test).

In summary, persistent cell activity during rest was more strongly modulated by the interleaved acquisition session than for decreasing cells. However, the number of replayed and pre-played states for persistent and decreasing cells was similar. Furthermore, only a fraction of persistent cell activity patterns were reactivated during subsequent rest, whereas decreasing cell activity was reactivated close to its entirety.

2.4 Discussion

I showed how reactivated activity in the hippocampus during rest underwent a gradual transformation. Initially, reactivations showed a great overlap with neural activity observed during the acquisition of the spatial memory. Towards the end of rest, reactivations were more recall-like and therefore resembled future activity that only occurred after the rest epoch. This memory drift was mainly driven by the firing rate changes of an unstable subset of cells. The modulation of these cells arose during NREM epochs, whereas REM sleep showed a resetting effect. The stable subset maintained firing rates and its spatial coding from the acquisition to the recall. In this way, stable cells provided a link between multiple experiences involving the same set of goal locations. While remapping properties of unstable and persistent subsets during the acquisition were comparable, the

relationship with the rest before and after was distinct for both subsets. The interleaved acquisition session led to greater differences between the rest activity before and after for the persistent cells. However, only a fragment of persistent cell activity was actually reactivated during the rest after the acquisition.

The discrepancy between my results and previous studies which did not identify temporal changes in hippocampal reactivations during rest might be due to the following two reasons. Firstly, previous studies used data from much shorter rest periods (< 6 hours). Secondly, the most common approach to quantifying reactivations is to compare rest activity with prior neural activity during behavior and shuffled neural activity data.

During the initial part of rest, my results showed acquisition-like reactivation activity and therefore support previous studies[188]. I also found that the number of significant reactivations decreased over time due to a decrease in SWR occurrences, confirming previous results (Fig. S2.4d)[202]. However, since I compared each activity bin in rest with prior (acquisition) and subsequent (recall) neural activity I was able to identify a gradual drift - a finding that could not have resulted from the standard reactivation analysis. This gradual transformation was most strongly pronounced within the first $\sim$10h leading to more recall-like activity towards the end of rest.

The occurrence of neural activity patterns from future behavioral episodes during rest has been characterized before[204]. Unlike in our study, pre-play was identified during rest before the exploration of a novel environment. In my case, the environment and the goal locations associated with the spatial memory were kept constant between acquisition and recall. Hence, it is natural to relate my findings to representational drift: recent studies have shown that repeated experiences in an environment lead to changes in hippocampal spatial representations[114,115,102]. These changes were hypothesized to only occur when the animal actively engaged with a memory as part of an updating process. Yet, my results suggest that this process is at least partially initiated during interleaved rest periods. Since a large portion of the identified pHMM states were spatially constrained, a transition from acquisition states towards recall states during rest could translate to changes in hippocampal spatial representations between both behavioral episodes[206]. In addition, I found that changes in firing rates during rest were consistent with firing rate changes between acquisition and recall - further confirming the strong interrelatedness between offline and online modulations of the hippocampal neural activity.

The memory drift did not occur homogeneously across the whole population. Low-firing, persistent cells acted as a "backbone" and bridged the acquisition and recall of the spatial memory by maintaining their firing rates and their spatial coding. A previous study has also identified low-firing cells that code for novel features of an experience in the CA1 neural population[207,104]. In support of their findings, I show that persistent cell activity is significantly modulated during the learning of goal locations and that this modulation causes significant differences between the rest before and after[208]. However, I did not find a superior SWR gain for my persistent subset compared to the unstable subset (Fig. S2.7f).

During rest, decreasing and increasing cells gradually left or joined the neural assemblies anchored by persistent cell activity. Since this process showed differential contributions by NREM and REM sleep, it could be partially explained by the different functions and brain states that are associated with these sleep stages.

My results suggest that the reshaping of neural activity is more strongly driven during

NREM sleep and consists of increasing and decreasing activity of distinct subsets. Since SWRs happening during NREM have been associated with the transfer of information to neocortical areas[170], I propose that modulations of hippocampal activity are driven by the global network state. One possibility is that newly acquired information is integrated into existing knowledge to promote generalization[209]. Another explanation could be that memory representations are transformed in a way to make them sparser while maintaining the information content[210]. Alternatively, modulations could be due to homeostatic effects such as synaptic downscaling[211]. However, the downscaling of synapses has been associated with REM sleep mainly.

A resetting effect of REM periods on firing rate increases happening during NREM sleep and the symbiotic relationship between both sleep stages has recently received increasing support[188,212]. I found that REM reactivations are more similar to the awake activity and therefore resemble a more "pure" form of the memory. Therefore, my findings support the idea that both sleep stages perform distinctive functions: NREM sleep edits the memory-related activity, whereas REM sleep has a stabilizing role[213,214]. In this way, recent memories are reactivated in a pure form during REM sleep, which is followed by more global computations during NREM phases supporting the 2-stage model for memory consolidation[18].

Future experiments will be needed to identify how the observed modulations during sleep affect previously reported sequential replay[215,216] in the hippocampus, something my pHMM approach is potentially capable of. In addition, simultaneously recorded sleep activity of the hippocampus and the neocortex are essential to better understand the connection between REM-NREM modulations of memory-related activity and the 2-stage memory consolidation process.

In summary, my results suggest that rest periods contribute to changes in neural representations across days and that these modifications are functionally driven by repeated REM-NREM cycles to potentially support the updating and consolidation of spatial memories.

2.5 Supplementary Figures

a

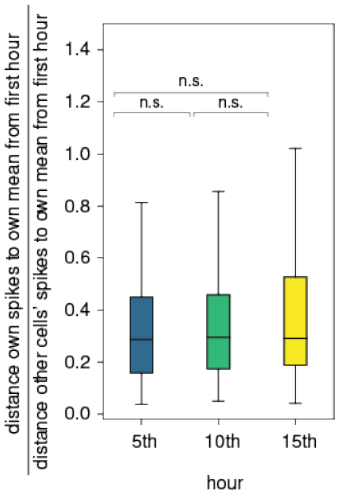

Figure S2.1: (a) Distance (1 - Pearson R) between the mean (first 12 clustering features) during the corresponding interval (5th, 10th, 15th hour) and the cell's mean from the first hour divided by the distance between the other cells' means during the corresponding interval and the cell's mean from the first hour for one example session (for all sessions $p>0.13$, two-sided Mann-Whitney U test).

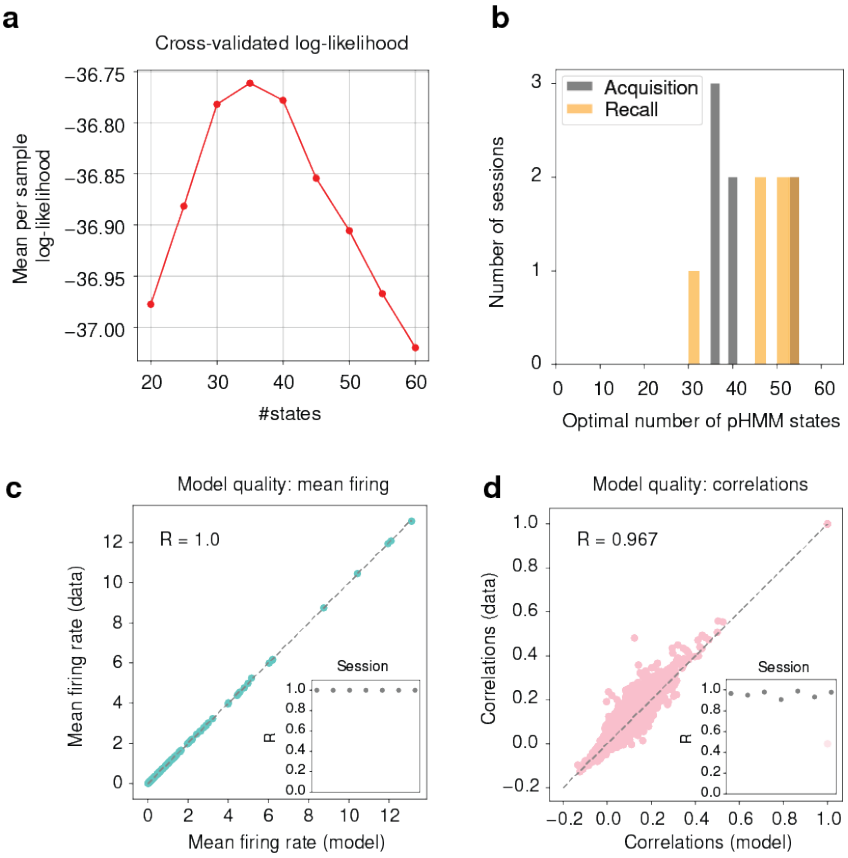

Figure S2.2: **pHMM model fit and its quality.** (a) Cross-validated log-likelihood to identify optimal number of hidden states (example session). (b) Optimal number of states for acquisition and recall for all sessions. (c) Mean firing rates of real data versus mean firing rates of data sampled from the pHMM for one example session (R = 1.0, p =2.3375e-262). Inset: Pearson R values for all sessions (all R > 0.99088). (d) Correlation of instantaneous firing rate of cell pairs of real data versus sampled data from the pHMM for one example session (R = 0.967, p = 0.0). Inset: Pearson R values for all sessions (all R > 0.85709).

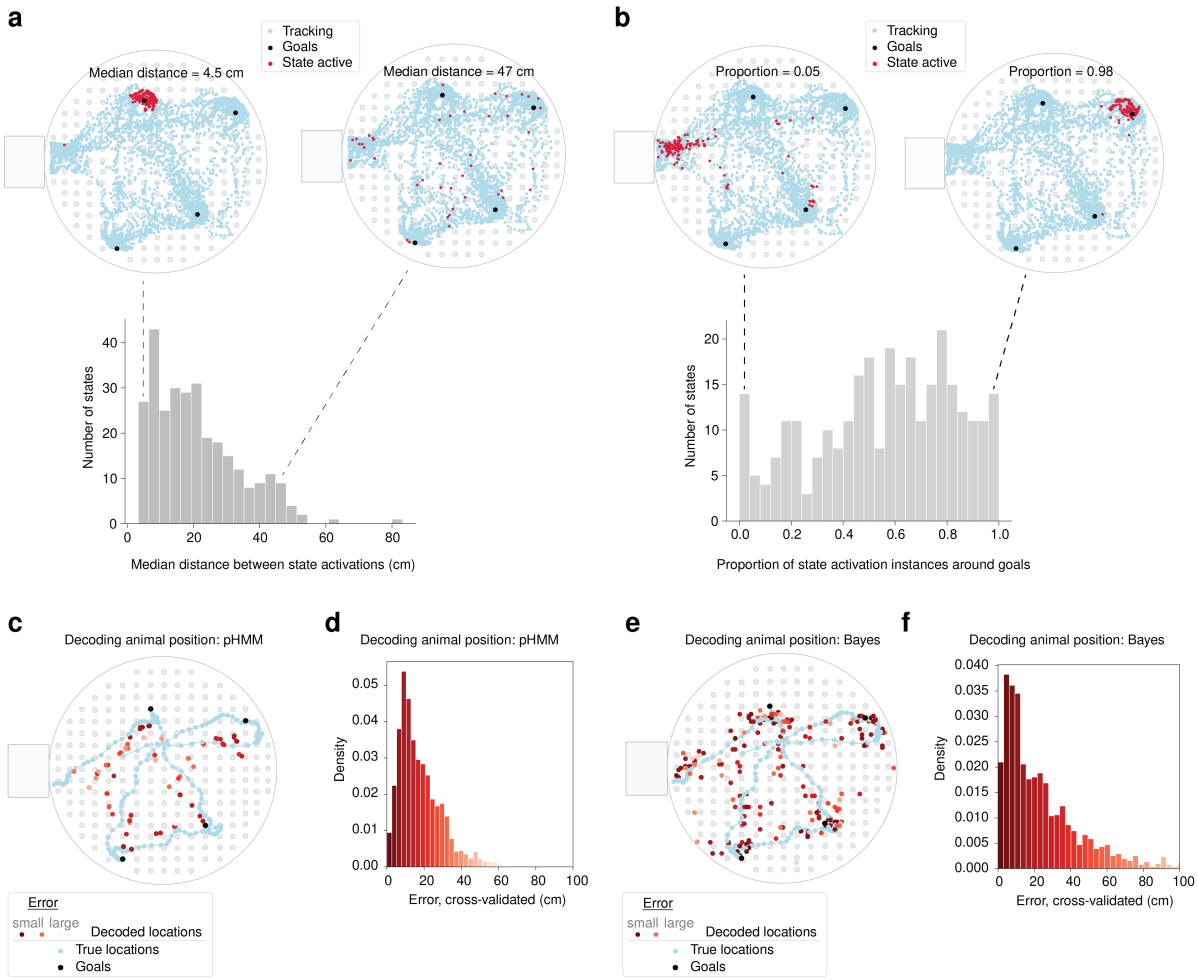

Figure S2.3: **spatial selectivity of decoded pHMM states.** (a, b) Top: examples showing the cheese board locations where a given pHMM state was decoded during acquisition. Bottom: (a) distribution of median distances of decoded locations for each pHMM state for all sessions. (b) Distribution of goal selectivity of pHMM states measured as the proportion of instances when a pHMM state was active near (<10cm) a goal for all sessions. (c) True and decoded locations using the pHMM states for decoding (see Methods) for one example trial. (d) The distribution of cross-validated decoding errors using the pHMM decoding approach for all sessions. (e-f) Same as (c-d) but using Bayesian decoding.

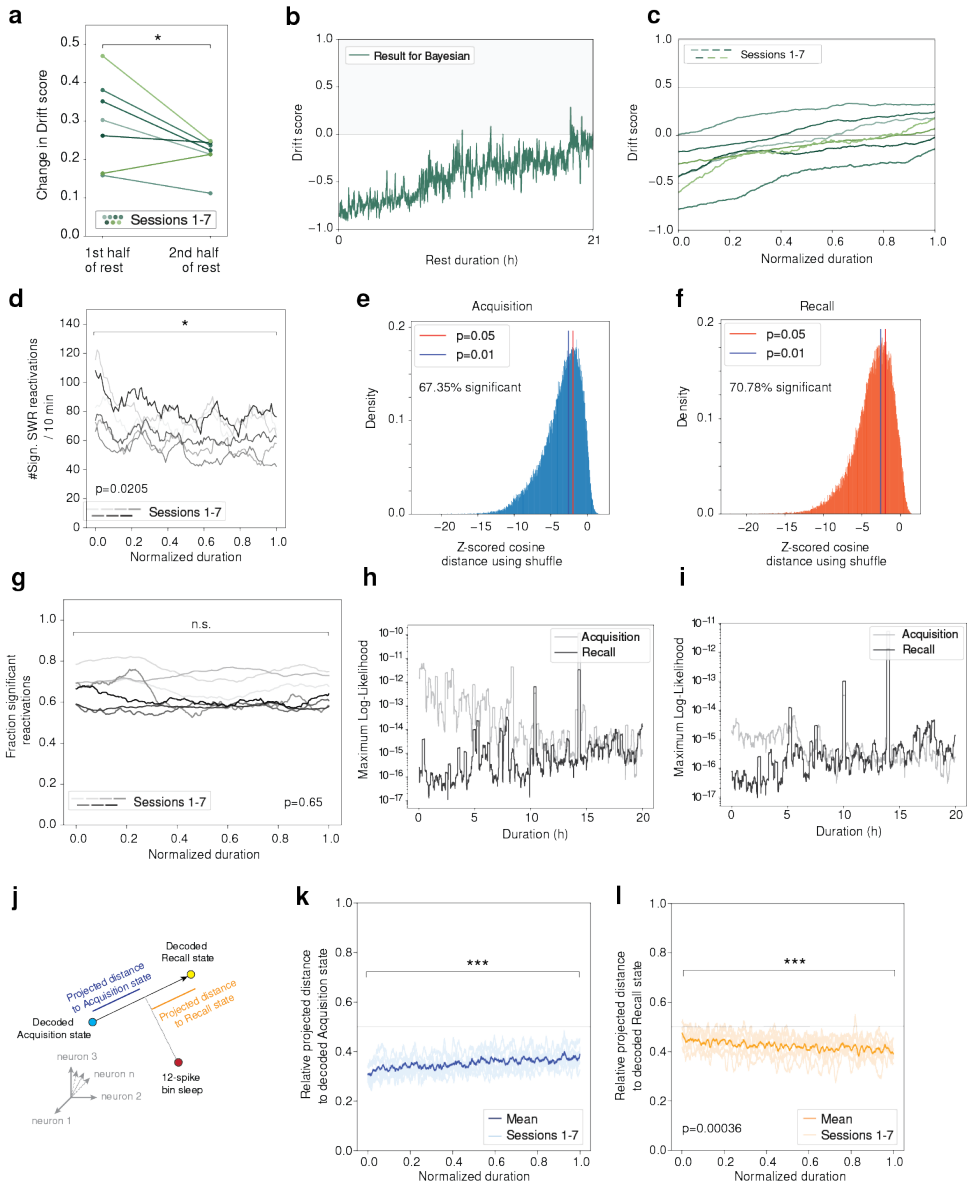

Figure S2.4: (a) Drift score changes more strongly during the first half of sleep than during the second half of sleep (alternative hypothesis: ratio first half/second half is greater than 1, $p < 0.05$, one-sided t-test). (b) Drift score using Bayesian decoding (one example session). (c) Drift score using Bayesian decoding for all sessions. (d) Number of significant sharp wave ripple (SWR) reactivations during sleep. (e) Decoding quality of Acquisition states for one example session (see Methods). (f) Decoding quality of Recall states for one example session (see Methods). (g) Fraction of significant reactivations throughout sleep (beginning vs. end, p=0.1129, two-sided T-test). Result for all sessions. (h) Maximum likelihoods of sleep neural data for acquisition pHMM and for recall pHMM (one example session). (i) Likelihoods like in (i) for another example session. (j) Schematic showing the projection of the sleep neural activity (12-spike bin) onto the vector between the decoded Acquisition state and the decoded Recall state. The relative projected distance to the decoded state is computed by dividing the length of the projection by the length of the vector between decoded Acquisition and decoded Recall state. (k) Relative projected distance to decoded Acquisition state throughout sleep (beginning vs. end, p=0.0017, two-sided T-test). Results for all sessions are shown. (l) Relative projected distance to decoded Recall state throughout sleep (beginning vs. end, p=0.012, two-sided T-test). Results for all sessions.

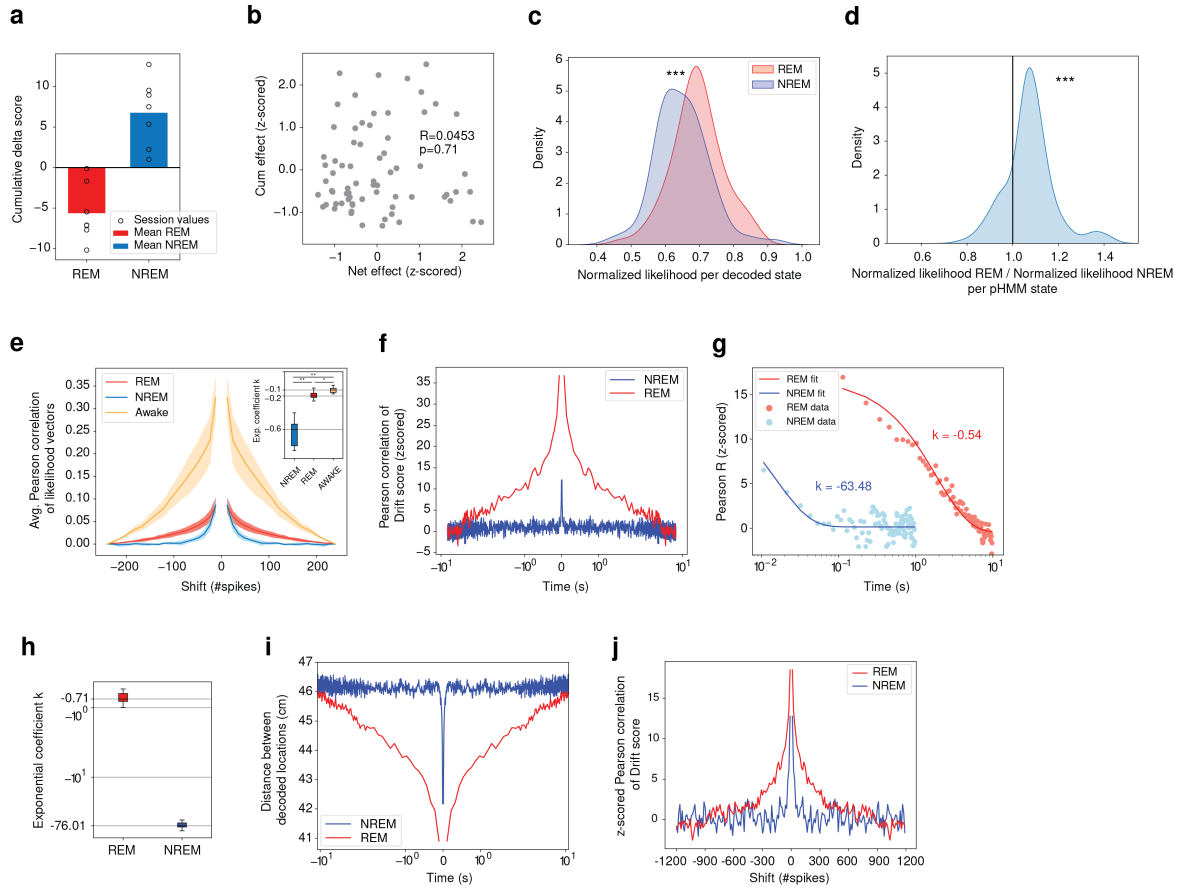

Figure S2.5: (a) Cumulative delta score for REM and NREM for all sessions using Bayesian decoding of sleep activity. (b) Correlation between cumulative effect (sum of absolute ΔDrift score values per epoch) and net effect (difference in ΔDrift score between the end and the beginning of each epoch) for all sleep epochs (data from all sessions, see Methods). (c) Normalized likelihoods per decoded pHMM state in NREM and REM. Only modes with significant differences in terms of normalized likelihood between NREM and REM were used. Data from all sessions ($p < 0.001$, two-sided Mann-Whitney U test). (d) Distribution of ratios between normalized likelihood in REM and normalized likelihood in NREM per state. Only states with significant differences in terms of normalized likelihood between NREM and REM were used. Data from all sessions ($p < 0.001$, one-sided T-test comparing distribution of ratios to a mean of 1). (e) Auto-correlation of likelihood vectors for pHMM decoding in REM (red), NREM (blue) and awake (yellow) for one example session. Exponential coefficients for REM, NREM and awake data (inset, two-sided Mann-Whitney U test). (f) Auto-correlation of the drift score in time for REM and NREM sleep (one example session). (g) Exponential fit for auto-correlation of the drift score. Exponential coefficient k for REM and NREM data from one example session. (h) Exponential coefficient k like in (g) for REM and NREM for all sessions. (i) Distance between subsequently decoded locations during sleep using Bayesian decoding. (j) Auto-correlation of the drift score as a function of shift (nr. of spikes) in NREM and REM.

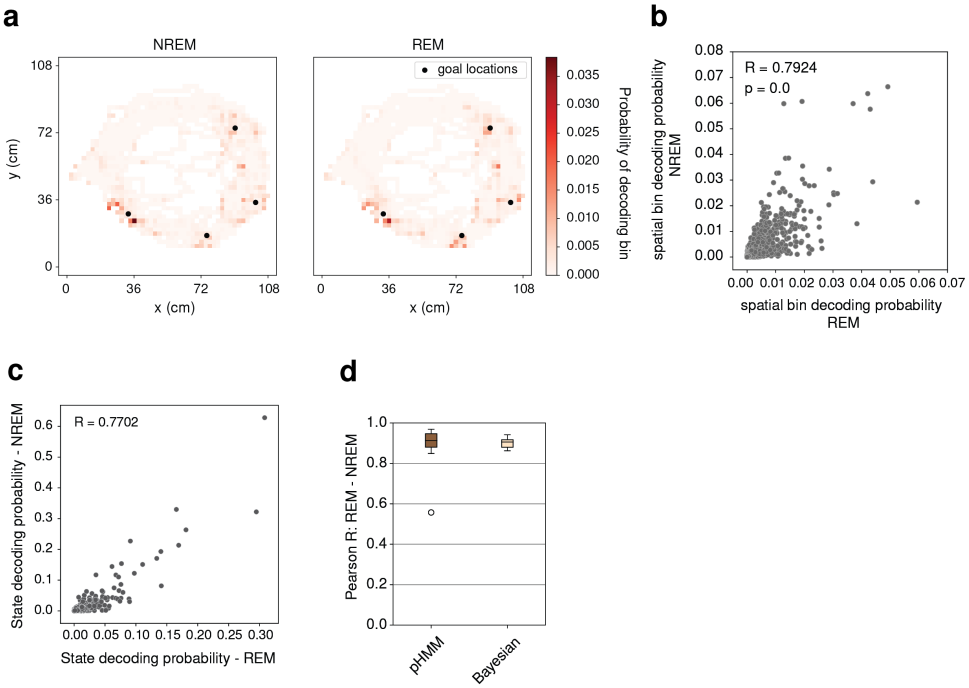

Figure S2.6: (a) Example decoding probability maps using Bayesian decoding for NREM (left) and REM (right) for one session (using acquisition map, Pearson R = 0.89, p = 0). (b) Correlation of spatial bin decoding probabilities from NREM and REM periods for all sessions. (c) Correlation between state decoding probabilities of REM and NREM for pHMM decoding and correlation between spatial bin decoding probabilities of REM and NREM for Bayesian decoding. (d) Correlation between reactivation probabilities in NREM and REM for pHMM and Bayesian decoding.

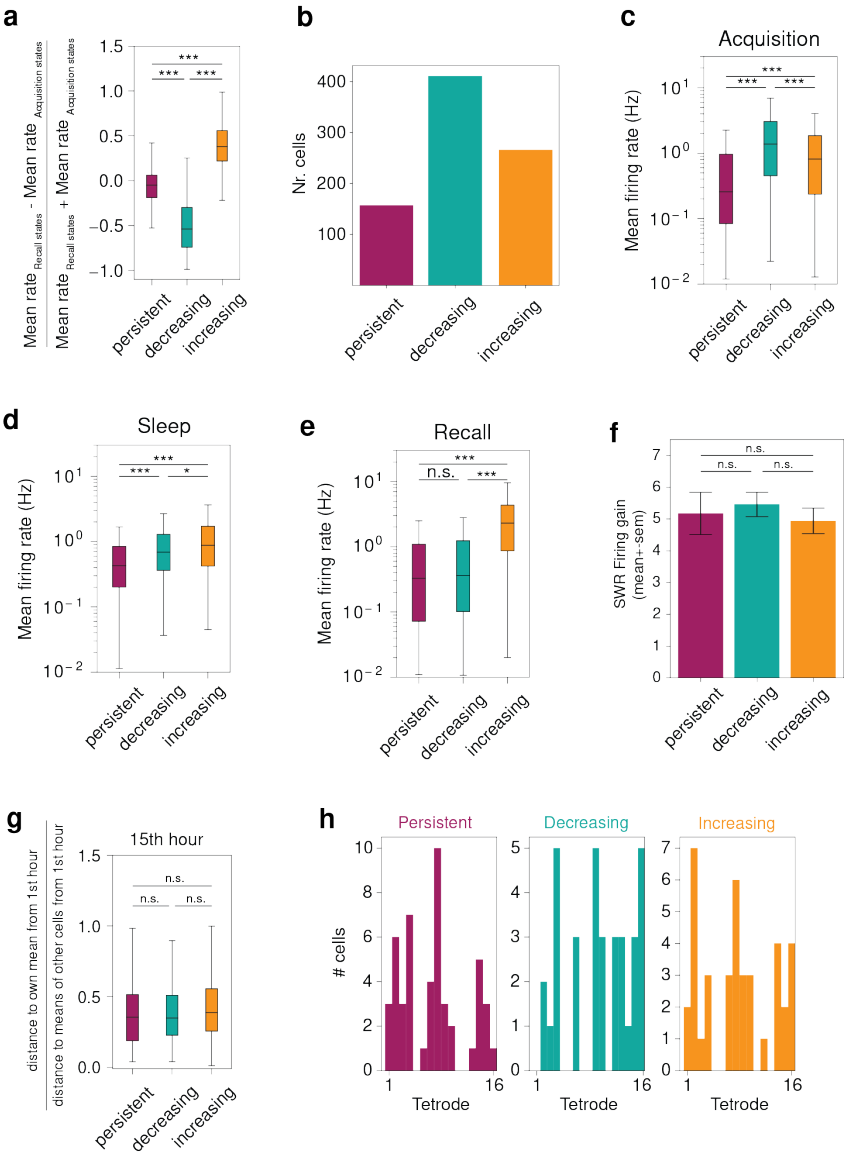

Figure S2.7: (a) Normalized difference in firing rates between acquisition and recall states for persistent, decreasing and increasing cells (*** $p < 0.001$, two-sided Mann-Whitney U test, Bonferroni correction). (b) Number of persistent, decreasing and increasing cells for all sessions. (c-e) Mean firing rates during acquisition (c), rest (d) and recall (e) sessions (*** $p < 0.001$* $p < 0.05$, n.s. $p > 0.05$, two-sided Mann-Whitney U test, Bonferroni correction). (f) SWR firing rate gain (mean±SEM, ratio of baseline to SWR peak firing rate) for persistent, decreasing and increasing cells (all $p>0.29$, two-sided Mann-Whitney U test). (g) Cluster stability of different cell groups. Distance (1 - Pearson R) between the mean cluster (first 12 clustering features) between first and fifteenth hour. Ratio of within vs. across cluster distance is shown. Data from all sessions (all $p > 0.05$, two-sided Mann-Whitney U test). (h) Number of persistent, decreasing and increasing cells detected across the different tetrodes of an example recording session.

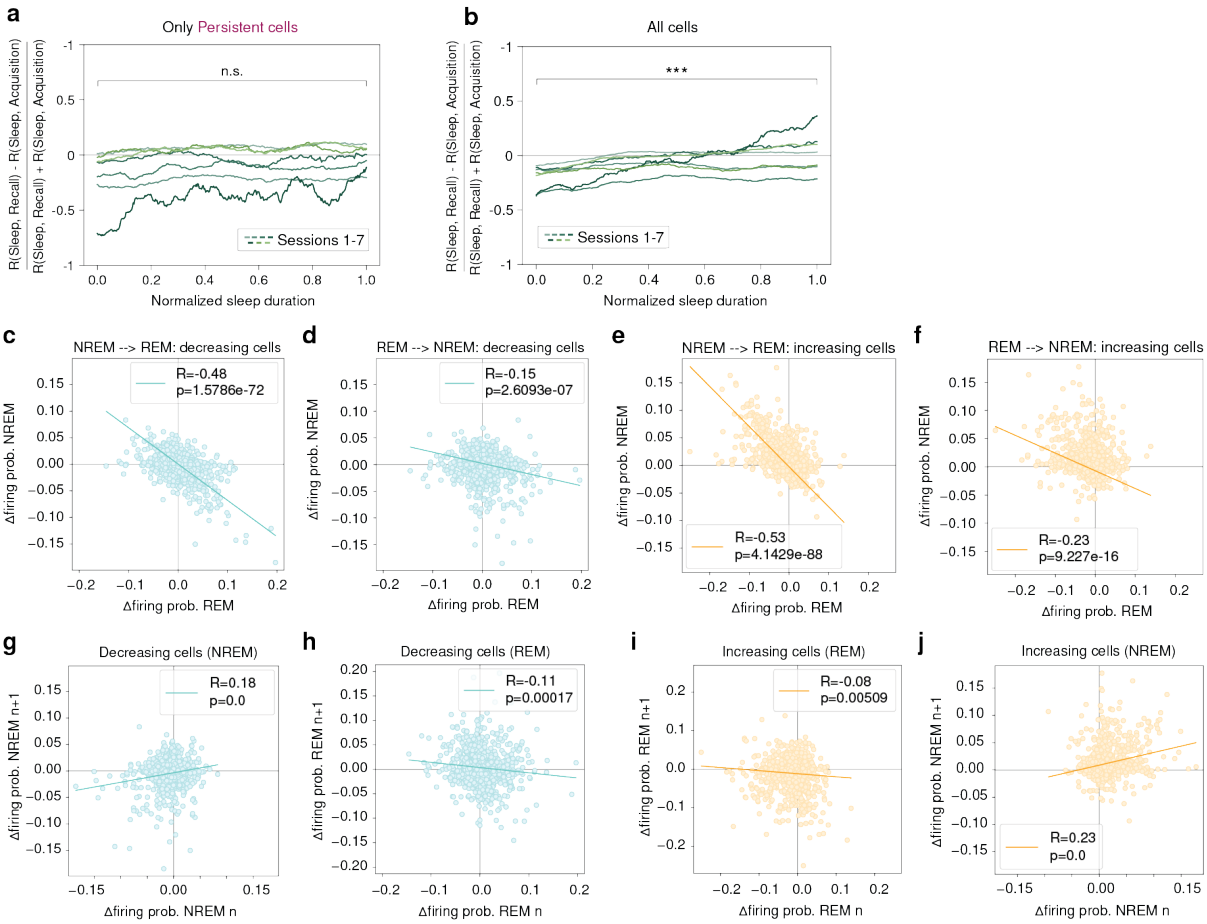

Figure S2.8: (a) Similarity of across cell firing correlations during sleep compared to correlations of acquisition and recall. Only persistent cells with equalized firing rates over time were used (beginning vs. end, p>0.1, T-test). (b) Similarity of across cell firing correlations during sleep compared to correlations of acquisition and recall. All cells with equalized firing rates over time were used (beginning vs. end, p<0.001, T-test). (c)-(f) Correlation between changes in firing probability in neighboring NREM and REM periods for (c, d) decreasing and (e, f) increasing cells considering the order of subsequent sleep type periods ($R < -0.47$, $p < 1.579e\text{-}71$ for NREM followed by REM and $R < -0.14$, $p < 2.609e\text{-}6$ for REM followed by NREM). (g)-(j) Correlation between firing probability changes in subsequent sleep epochs of the same type for (g, h) decreasing and (i, j) increasing cells ($R < 0.24$ for NREM, $R < -0.07$ for REM).

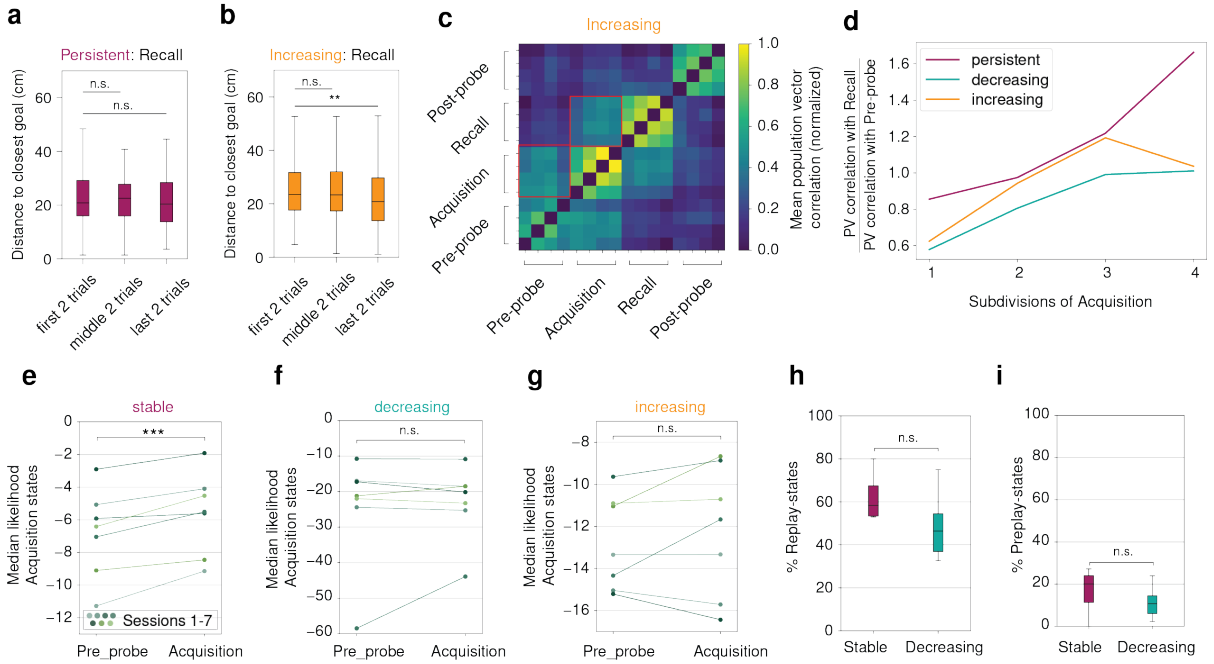

Figure S2.9: (a-b) Distances between place field peak firing locations and goal locations during Recall for (a) persistent cells and (b) increasing cells (n.s. $p > 0.05$, ** $p < 0.01$, two-sided Mann-Whitney U test). (c) Mean population vector correlation between different experimental episodes for increasing cells. The diagonal was set to zero and all remaining values normalized to lie between 0 and 1. (d) Ratio between mean population vector correlations (normalized) from pre-probe and acquisition and post-probe and acquisition for decreasing, increasing and persistent cells. (e-g) Median likelihood of pHMM states trained during acquisition using (e) only persistent, (f) only decreasing and (g) only increasing cells. Significance per session was computed by using one-sided Mann-Whitney U test (alternative: likelihoods pre-probe < likelihoods acquisition) and all sessions with $p<0.01$ were considered as showing statistical difference. For all sessions a binomial test was performed ($p=0.0078$ for persistent cells, $p=0.9375$ for decreasing cells, $p=0.22656$ for increasing cells). (h-i) Percentage of (h) replay (differential expression between pre-sleep and first part of sleep < -0.5, see Methods) and (i) preplay (differential expression between pre-sleep and first part of sleep > 0.5) pHMM states (n.s. $p > 0.05$, two-sided Mann-Whitney U test).

2.6 Methods

Animals and surgery

Animals were implanted with microdrives housing 32 (2x16) independently movable tetrodes targeting the dorsal CA1 region of the hippocampus bilaterally. Each tetrode was fabricated out of four 10 μm tungsten wires (H-Formvar insulation with Butyral bond coat California Fine Wire Company, Grover Beach, CA) that were twisted and then heated to bind them into a single bundle. The tips of the tetrodes were then gold-plated to reduce the impedance to 200-400 kΩ. During surgery, the animal was under deep anesthesia using isoflurane (0.5%-3% MAC), oxygen (1-2l/min), and an initial injection of buprenorphine (0.1mg/kg). Two rectangular craniotomies were drilled at relative to bregma (centered at AP = -3.2; ML = $\pm$1.6), the dura mater removed and the electrode bundles implanted into the superficial layers of the neocortex, after which both the exposed cortex and the electrode shanks were sealed with paraffin wax. Five to six anchoring screws were fixed onto the skull and two ground screws (M1.4) were positioned above the cerebellum. After removal of the dura, the tetrodes were initially implanted at a depth of 1-1.5 mm relative to the brain surface. Finally, the microdrive was anchored to the skull and screws with dental cement (Refobacin Bone Cement R, Biomet, IN, USA). Two hours before the end of the surgery the animal was given the analgesic Metacam (5mg/kg). After a one-week recovery period, tetrodes were gradually moved into the dorsal CA1 cell layer (stratum pyramidale). After completion of the experiments, the rats were deeply anesthetized and perfused through the heart with 0.9% saline solution followed by a 4% buffered formalin phosphate solution for the histological verification of the electrode tracks.

Data Acquisition, Training and Behavior

The animals were housed individually in a separate room under a 12h light/12h dark cycle. Following the postoperative recovery period, rats were reduced to and maintained at 85% of their age-matched preoperative weight. Water was available ad libitum. Each animal was handled and familiarized with the recording room and with the general procedures of data acquisition. Behavioral training was performed after electrode implantation during days when the electrodes were moved towards the hippocampus, but before they reached the hippocampus. Overall 9 rats were trained to perform the seek of the hidden rewards task on the cheeseboard maze[84,217] and come back to the start-box. In order to achieve this, random groups of visible food pellets (MLab rodent pellet 20 mg, TestDiet) were spread out on the surface of the cheeseboard maze while the rat was inside the start box. Then we opened the door and left the animal freely foraging the entire maze and once the animal returned to the start box we closed the start-box door. With the help of this training protocol we could shape our animals behavior to automatically explore the entire maze and return to the start-box. Despite the automatic behavior we could ensure that under experiments rats had no or limited experience in performing the cheese-board maze task during the time of the recordings. Each daily experiment consisted of a sequence of nine recording sessions in the following order: a free exploration session on a familiar environment, half hour immobility/sleep rest session in the animal own cage, free exploration on the cheese-board, an immobility/sleep rest session (own cage), a learning session (4 randomly selected invisible locations) on the cheese-board, 20 hour continuous monitoring in the cage, recall learning session with the same bait locations, an immobility/sleep rest session (own cage), post probe on the cheese-board (free exploration,

unrewarded) and another free exploration session on the same familiar environment as first. Under the learning session once animals learned the four invisible goal locations, they performed only 5 additional trials to master the task. Following the 20 hour long rest session we tested the recall performance of the animal on the cheeseboard maze in 50 further rewarded trials with the same bait locations that were learned 20 hours before.

To be able to record extracellular electric signals continuously over a long period of time we adapted to our commonly built 3D printed microdrive a new, high-fidelity 128-channel wireless recording system from TBSI (Triangle BioSystems, Durham, NC, W128) to use in our experiments. Using this experimental preparation, we recorded cell population activity continuously over 30 hours during learning, long periods of rest where reactivation takes place and memory recall tests in the end. This telemetry system has been able to amplify and transmit wide band (0.7 Hz to 9 kHz) signals above 20 kHz on 128 analog channels which were then digitized at 20 kHz.

Two small light-emitting diodes (LEDs) mounted on the preamplifier headstage were used to track the location of the animal via an overhead video camera. The animal's location was constantly monitored throughout the experiment. We detected 2 LEDs with a custom made tracking software (positrack, github.com/kevin-allen/positrack) made by Kevin Allen. Video signal has been triggered and tracked continuously with a TTL pulse sent by the camera's computer on a common analog channel.

Spike detection, sorting and stability

The spike detection and sorting procedures, clustering were performed as previously described[218,219]. Continuously recorded wide-band signals were digitally high-pass filtered (0,8-5 kHz). Action potentials were extracted by first computing power in the 800-9000 Hz range within a sliding window (12.8 ms). Action potentials with a power of > 5 SD from the baseline mean were selected and spike features were then extracted by using principal components analyses (PCA). The detected action potentials were segregated into multiple putative single units by using automatic clustering software (http://klustakwik.sourceforge.net/). These clusters were manually refined by a graphical cluster cutting program. Only units with clear refractory periods in their autocorrelation ($<20\mu$s) and well-defined cluster boundaries were used for further analysis. We further confirmed the quality of cluster separation by calculating the Mahalanobis distance between each pair of clusters[220]. To be able to analyze changes in the firing patterns of neuronal ensembles over time, we have to guarantee that our set of putative cells was sampled from clusters with stable firing over the whole recording. To ensure this we clustered together periods of waking and rest sessions and then we plotted spike features over time by plotting 2-dimensional unit PCA cluster plots across the whole recording in addition to the stability of spike waveforms. With the help of this method we could exclude those spike clusters which are overlapped during the course of recording. To further verify spike cluster stability we used the t-student stochastic neighbor embedding (t-sne) dimensionality reduction method: T-sne embeds the n-dimensional extracellular spikes (n = number of features by which each spike is decomposed) into a low dimensional space[221]. T-sne focuses on ensuring that the local structure remains intact while it ignores the global structure, therefore when we expressed T-sne features over time we could visually exclude those clusters which were unstable during the whole recording due to electrode drifting.

Putative pyramidal cells and putative interneurons in the CA1 region were discriminated by their autocorrelations, firing rate, and waveforms, as previously described[218].

Sleep classification

In recordings exploratory and immobility or sleep sessions were manually separated off-line as previously described[222,219]. For each session, the theta/delta ratio was plotted against speed so that the behavioral state could be manually identified. The theta/delta field power ratio was measured in 1,600 ms segments (800 ms steps between measurement windows) with Thomson's multitaper method[223,224]. Waking behavior included periods of locomotion and/or the presence of theta oscillations (visible in the theta/delta ratio), with no more than 2.5 s of transient immobility. Rest epochs were selected when both the speed and theta-delta ratio dropped below a pre-set threshold (speed: <4cm/s, theta/delta ratio: <2) for at least 2.4 s. During periods of active waking behavior, theta-oscillatory waves detection was performed as previously described[222,219] using the negative peaks of individual theta waves from the filtered trace of the local field potential (5–28 Hz). The band used for the detection was wider than the theta band in order to precisely detect the negative peaks of the theta waves, which otherwise would have smoothed out in using a narrow theta band. Sleep segments have been identified by longer periods of immobility (because of the lengths of our recording at least 10 min) and the clear presence of REM-theta and slow-wave field oscillations.

Sharp wave ripple detection

For the detection of SWRs, local field potentials were band-pass filtered (150–250 Hz), and a reference signal (to ensure the lack of ripple activity we left a tetrode above the hippocampus as a reference) was subtracted to eliminate the so-called common noise (muscle artifacts due to scratching, twitching etc.). The power (root mean square) of the filtered signal was calculated for each electrode and summed across electrodes designated as being in the CA1 stratum pyramidale. The threshold for SWR detection was set to 7 SD above the background mean. The SWRs detection threshold was always set in the first sleep session, but longer (at least 35 min) as it was described earlier and the same threshold was used for all other sessions[219].

Stability of clustering features over time

To assess the temporal stability of waveforms, I used the first 12 clustering features. For each feature, the mean was computed for the corresponding interval (5th, 10th, 15th hour) and z-scored using the mean and std of the first hour.

Cell separation over time

I estimated the separation of cells by using the first 12 clustering features and three different approaches. First, for each cell, I divided the distance (1 - Pearson R) between the mean value of each feature during the corresponding interval (5th, 10th, 15th hour) and the cell's mean feature values from the first hour by the distance between the mean feature values during the corresponding interval and the other cells' feature means from the first hour. Using this measure I investigated whether a cell's waveform at a later stage during the experiment is more similar to its waveform from the first hour than to first

hour waveforms from other cells. Second, for each cell, I divided the distance (1 - Pearson R) between the mean feature values during the corresponding interval and the cell's mean values from the first hour by the distance between the other cells' feature means during the corresponding interval and the cell's feature means from the first hour. This measure estimates whether other cells' waveforms at later stages during the experiment are more similar to the cell's waveform from the first hour than its waveforms at later stages. Third, I compared the across-cell distance with the within-cell distance. For the across-cell distance, the distance between the mean feature values per cell of the first hour and the other cells' feature means from the first hour was calculated. The within distance corresponds to the distance between the cell's feature means from the first hour and the same cell's feature means of the last hour of the experiment.

Clustering feature stability

For each cell I calculated the mean and std for each of the first 10 PCA clustering features using data from the first hour. I then computed the mean per clustering feature of each cell at different time intervals (7th, 15th and 21nd hour) of the experiment. In order to test whether clustering features drift away from the initial values, I z-scored the means during different time intervals using the mean and std from the first hour.

Measuring excess path during acquisition and recall

I assessed the animal's ability to learn and recall goal locations on the cheeseboard by computing the excess path: once the rat had left the start box, I measured the length of the path the animal took to reach any of the four goals (animal position within 10 cm radius around goal location). Next, I detected when the animal left the goal again and measured the path length to the next goal. This procedure was repeated for the remaining two goals. I then calculated the optimal paths as straight lines between either start location and the first goal or between subsequently visited goals. Each taken path length was then divided by the optimal path length to yield the excess path as a multiple of the optimal path.

Poisson hidden markov model (pHMM) and model fitting

I trained two separate hidden Markov models with Poisson emissions (pHMM) on the neural data obtained during the cheeseboard task before (acquisition) and after (recall) the long sleep. Only data from running periods (speed > 5cm/s) was used. The acquisition data length was matched in terms of duration to the recall data to have the same training data length for acquisition and recall. Then, the neural data was binned using temporal bins of 100 ms length. The pHMM model assumes the temporal evolution of an unobserved discrete state q_t as described in[225]. In short, the probability of observing an ensemble O_t of N independently firing neurons at time t for state i can be modeled as:

$$P(O_t \mid q_t = i) = \prod_{n=1}^{N} P(o_{n,t} \mid q_t = i) \propto \prod_{n=1}^{N} (\lambda_{n,i})^{o_{n,t}} exp(-\lambda_{n,i}) \tag{2.1}$$

where $o_{n,t}$ is the number of spikes of neuron n at time t. The firing rate is modeled according to a Poisson process with a mean number of spikes $\lambda_{n,i}$ defined by the unobserved discrete state i. The transition probabilities between M unobserved states is captured by

the M x M transition matrix A. The hyperparameter M defining the number of states of the model was determined using the cross-validated maximum likelihood (Fig. S2.2). All model parameters were inferred using the EM-algorithm.

Rate map generation for Bayesian decoding

For each session and cell one rate map for acquisition and one rate map for recall was computed.

In order to reconstruct the two-dimensional spatial distribution of each cell's firing, I used spatial bins of 5 cm size and temporal windows of 10ms. For each spatial bin, the mean number of spikes per cell was computed using all temporal windows, which corresponded to the animal being within the spatial bin. Periods of rest (speed < 5cm/s) were excluded.

Decoding sleep activity using pHMM

In order to compensate for differences in temporal dynamics between REM and NREM sleep, I binned the sleep data using bins with a constant number of 12 spikes. Since the awake pHMM models were trained on temporal bins of 100 ms I computed a scaling factor between awake and sleep neural activity to match the two. First, I calculated the mean number of spikes occurring within 100 ms time bins during awake behavior $n_{awake,100ms}$. The scaling factor γ_{phmm} is defined as:

$$\gamma_{phmm} = \frac{12}{n_{awake,100ms}} \tag{2.2}$$

The likelihood of the sleep activity at time t given the discrete pHMM state q with N neurons is computed as follows:

$$L(q_t = i) = \prod_{n=1}^{N} \frac{(\gamma_{phmm} \cdot \lambda_{i,n})^{o_{n,t}}}{o_{n,t}!} exp(-\gamma_{phmm} \cdot \lambda_{i,n}) \tag{2.3}$$

where $o_{n,t}$ is the number of spikes of neuron n at time t and $\lambda_{i,n}$ is the mean number of spikes (per 100ms) of the neuron in state I. Notice that for my decoding procedure the transition probabilities across states were not considered. To assess which state was most likely reactivated at time t during sleep the state with the maximum log-likelihood was selected.

Decoding sleep activity using Bayesian decoding

Equivalent to the pHMM decoding approach the sleep data was binned using bins with a constant number of 12 spikes. In the case of Bayesian decoding the acquisition and recall rate maps were computed using 10 ms time bins. Therefore, the scaling factor $\lambda_{bayesian}$ is computed using the mean number of spikes occurring within 10 ms time bins nawake,10ms during awake behavior:

$$\gamma_{bayesian} = \frac{12}{n_{awake,10ms}} \tag{2.4}$$

The likelihood of the sleep activity at time t given the spatial bin x is given by:

$$L(x_t = i) = \prod_{n=1}^{N} \frac{(\gamma_{bayesian} \cdot \lambda_{i,n})^{o_{n,t}}}{o_{n,t}!} exp(-\gamma_{bayesian} \cdot \lambda_{i,n}) \tag{2.5}$$

where $o_{n,t}$ is the number of spikes of neuron n at time t and $\lambda_{i,n}$ is the mean number of spike (per 10ms) of that neuron in spatial bin x. To reconstruct the spatial bin that was most likely reactivated at time t, the spatial bin with the maximum log-likelihood was selected.

Quantifying drift during sleep

For each time point t in sleep, I calculated the Drift score in the following way:

$$Driftscore = \frac{max(L_{recall}) - max(L_{acquisition})}{max(L_{recall}) + max(L_{acquisition})} \tag{2.6}$$

with L_{recall} and $L_{acquisition}$ being the maximum likelihoods across all states or spatial bins for the acquisition or recall models (pHMM or rate maps) at time t, respectively. The resulting Drift score was smoothed across time.

Spike shuffle

As a control I used the following spike shuffling procedure: two random constant spike bins were selected. Within these constant spike bins two cells with at least one spike each were randomly chosen. Then, one spike for each cell was transferred by moving it from one bin to the other. This maintains mean firing rates per cell, but shifts the spike times. The swapping of spikes was repeated 50 times the number of total spike bins. I computed means $m_{shuffle}$ and $std_{shuffle}$ of the shuffled data by repeating the above procedure 30 times. The z-scored Driftscore with regards to the shuffled ratios was computed as follows:

$$Driftscore_{zscored} = \frac{Driftscore - m_{shuffle}}{std_{shuffle}} \tag{2.7}$$

From the above expression the p-values related to the difference between similarity ratios and the shuffle were computed.

First versus second half of sleep

In order to validate whether the drift is more prominent during the first half of sleep, I computed the delta of the similarity ratio for the first and second half of sleep separately. Then, I tested whether the ratio between the delta of the first half and the delta of the second half was greater than 1 using the student T-test.

Net effect vs. cumulative effect of drift

In order to assess the amount of memory drift with respect to different timescales, I compared the net effect and cumulative effect. For the net effect, I computed the difference in drift score between the beginning and the end of the rest period. On the other side, the cumulative effect was computed by summing the absolute values of the drift score throughout the rest period.

Decoding quality

Decoding quality per population vector in sleep was assessed in the following way: I computed the cosine distance between the population vector (12 spike bin) and the decoded Acquisition or Recall state. Next, I computed the cosine distance between the population vector and 100 artificially created Recall or Acquisition states (shuffle). For each artificial state, every neuron's mean firing rate is randomly chosen from one of the pHMM states from the Acquisition or Recall model, respectively. If the distance to the decoded Acquisition or Recall state is 1.96 std below the mean of the distances to the shuffle, the decoding quality was classified as sufficient.

Projecting sleep activity onto decoded Acquisition state - decoded Recall state axis

For each sleep activity bin (12-spike bin), I determined the decoded Acquisition and the decoded Recall state using my maximum likelihood approach. Assuming Poisson firing, I generated a 12-spike bin from the decoded Acquisition and the decoded Recall state, respectively. This was done to have all vectors in the same format for the subsequent computations. I derived the vector between the decoded Acquisition and Recall state (12-spike bins) in neuron space. Then, the sleep activity bin was projected onto this vector. The length of the resulting projected vector was divided by the length of the vector between the decoded Acquisition to the decoded Recall state. This relative projected distance indicates whether the projected sleep activity bin is closer to either the decoded Acquisition or the decoded Recall state. Next, I selected all sleep activity bins where the maximum likelihood across Acquisition states was greater than the maximum likelihood across Recall states (Drift score < 0) and plotted the relative projected distance to the decoded Acquisition state. The same was done for distance to the decoded Recall state for sleep activity bins with maximum likelihood across Recall states $>$ maximum likelihood across Acquisition states (Drift score > 0).

Opposing effect on drift in NREM and REM

NREM and REM periods were identified as described above. Neighboring sleep epochs of the same type were merged to obtain a set of alternating NREM/REM epochs. For each epoch, the change in similarity ratio was computed by subtracting the first value of the epoch from the last value of the epoch.

Differences in NREM and REM reactivation quality

To assess differences in the quality of sleep reactivations in NREM and REM, I first looked at the maximum likelihoods per constant number spike bin during sleep (see Decoding sleep activity). For each bin I computed the likelihoods for all pHMM states and selected the maximum likelihood per bin. This procedure was done for NREM and REM bins separately and the distributions of maximum likelihoods compared. Next, I analyzed REM and NREM posterior probabilities per pHMM state. For each pHMM state I identified when this state was reactivated during sleep (using my maximum a-posterior decoding approach) for REM and NREM epochs separately. Then, I identified states that showed a significant difference between REM and NREM epochs. The difference for a given state was considered significant if the ratio between posterior probabilities of REM and NREM

epochs deviated more than 2 std from the mean of the REM posterior probabilities. The ratios of all significantly different states were then used to assess the quality of sleep reactivations in REM and NREM.

Temporal differences in NREM and REM reactivations

I analyzed the decay of reactivations in NREM and REM by computing the auto-correlation of the similarity ratio. In order to display the decay as a function of time instead of the number of constant number spike bins, I computed the mean duration of the 12 spike bins in sleep for REM and NREM. Then, I z-scored the resulting auto-correlation values using the mean and std of the tails. The exponential coefficient k was determined by fitting an exponential function using the standard numpy optimize library.

Similarity in NREM and REM reactivations

In order to investigate the similarity in reactivation content between REM and NREM, I used my decoding results from pHMM and Bayesian decoding. For the Bayesian decoding, I computed the probability of decoding one spatial bin using either all NREM or REM epochs. Then, the correlation of the flattened probability maps was computed. The equivalent approach for the pHMM decoding correlates the probabilities of decoding one state for all states between REM and NREM epochs.

Persistent and unstable subsets

Persistent and unstable cell subsets were identified based on changes in their firing rate distributions from acquisition to recall. For each cell the distributions of firing rates for acquisition and recall were computed separately. If the acquisition distribution was significantly greater than the recall distribution (one-sided Mann-Whitney U test, $p<0.01$), the cell was assigned to the decreasing subset. If, on the other hand, the recall distribution was significantly greater than the acquisition distribution (one-sided Mann-Whitney U test, $p<0.01$), the cell was labeled as increasing. All other cells which did not show a significant difference in their firing rate distributions from acquisition to recall made up the persistent subset.

Drift using different subsets of cells

To evaluate the effect of using only a subset of cells for my sleep decoding procedure on the observed drift I proceeded as follows. I removed all cells not contained in the subset from my constant number spike bins and computed the Drift score using the maximum likelihoods from the acquisition and recall states. Then, the Drift score for the entire sleep duration was split into four parts of equal length. For each part, a line was fit and the resulting slope calculated. The mean slope of the four parts was then compared to the equivalent value of the Drift score using the entire cell population.

Firing probability changes in NREM and REM

REM and NREM sleep epochs were identified as described above. Thereafter, for each subset of stable, increasing and decreasing cells I computed the change in firing probability from the beginning to the end of each epoch. The firing probability is defined as the

number of spikes contributed by the subset to the total number of spikes per constant spike bin. I computed the change by subtracting the firing probability of the first bin of the epoch from the value of the last bin of the epoch. Only epochs with significant changes in firing probability were considered.

Measuring spatial information

To assess the spatial information of single cells, I computed the sparsity and spatial information per second as previously described[226]. The sparsity for X spatial bins is defined as follows:

$$s = \frac{(\sum_{i=1}^{X} p_i \lambda_i)^2}{\sum_{i=1}^{X} p_i \lambda_i^2} \tag{2.8}$$

with pi being the probability of being in spatial bin i and λ_i being the mean firing rate in that bin. The spatial information per second was computed using the following equation:

$$I_{sec} = \sum_{i=1}^{X} p_i \lambda log(\frac{\lambda_i}{\lambda}) \tag{2.9}$$

where pi and λ_i are the probability of occupying and the firing rate of bin i, respectively. Parameter λ describes the mean firing rate of the cell in the environment.

Decoding positions using neural activity during behavior

Bayesian decoding I applied standard Bayesian decoding using 5 cm spatial bins. First, I computed the mean number of spikes $\lambda_{n,i}$ per 10ms for each cell n and spatial bin i. Given the spikes O_N of N neurons at time t, I computed the likelihood of being in bin i using the following equation:

$$L(O_N \mid x = i) = \prod_{n=1}^{N} \frac{\lambda_{i,n}^{o_{n,t}}}{o_{n,t}!} exp(-\lambda_{i,n}) \tag{2.10}$$

with o_n being the number of spikes (per 10ms) of neuron n at time t. The spatial bin with the highest likelihood represented the decoded location.

Decoding using pHMM

Using the entire neural data from acquisition and the trained pHMM, I inferred the most likely state sequence using the Viterbi algorithm. By matching the sequence of states with the tracking data of the animal I identified a mean spatial location for each state. Given the activity of N neurons at time t, I computed the normalized likelihood for each state of my pHMM. The decoded location was then calculated by weighing the mean location of each state with its normalized likelihood and computing the average position across all states.

Computing mean firing rates

Acquisition, sleep and recall were split into 5 min. chunks to computed mean and maximum firing rates of the different cell subsets.

Distance between peak firing and closest goal

For each cell, I computed its rate map and determined the location on the maze with the highest firing rate. Next, I calculated the distance between the location with the peak firing rate and the closest goal location.

Population vector correlations

Neurons were separated into persistent, increasing and decreasing cells according to the procedure defined above. Neural activity per subset was binned using spatial bins of 10 cm2 size yielding one population vector per spatial bin. Spatial bins which were not visited in all relevant behavioral episodes were excluded. Then, Pearson correlations between population vectors of the same spatial bin during the different behavioral episodes were computed. Computing similarity in co-firing with equalized firing rates In order to assess the similarity of co-firing during rest with either the acquisition or recall and controlling for firing rate changes I performed the following analysis steps. First, I excluded cells with very low firing rates. Then, I computed the correlation matrices for the acquisition and recall using the unmodified firing rate rasters. I divided the rest activity into chunks of 200 seconds and computed the minimum number of spikes per 200 second window for each cell. For each cell, I then removed random spikes from each 200 second window until the minimum number of spikes was reached. This procedure allowed us to remove the effect of firing rate changes throughout rest. I repeated the subsampling procedure fifty times and computed the correlation matrix between cell firing for each iteration. Then, the firing-equalized correlation matrices for each window were correlated with the acquisition and recall correlation matrices. Using the normalized difference, I obtained a temporal measure for the relative similarity of the co-firing during rest with either acquisition or recall.

Pre- & replay states analysis

I assessed the extend of activity being pre-played during the sleep before (pre-sleep) the actual experience and the extend of activity re-played during the sleep after (sleep) in the following manner. Cells were split into stable and decreasing subsets in accordance with the aforementioned procedure. For each subset I fit a separate pHMM model to the activity during acquisition. Then, for each subset I used the corresponding pHMM model to decode the activity during the pre-sleep and the first part of the long sleep. I only used the first part of the long sleep to avoid effects of the observed drift in activity. For both sleeps separately, I computed the mean likelihood per decoded pHMM state. This yielded per state, the mean likelihood for pre-sleep $L_{pre-sleep}$ and the mean for the first part of the long sleep L_{sleep}. I then computed the normalized difference between the two to evaluate if a state is differentially expressed:

$$difference_{expression} = \frac{L_{pre-sleep} - L_{sleep}}{L_{pre-sleep+L_{sleep}}} \tag{2.11}$$

I defined states with $difference_{expression}$ < -0.5 as replay states and states with $difference_{expression}$ > 0.5 as preplay states.

Support Vector Machine (SVM) to decode behavioral episodes

I trained a support vector machine to distinguish between behavioral episodes. First, the neural activity of two behavioral episodes (e.g. acquisition and recall) was binned into 2s temporal bins. Then I assigned labels to the temporal bins identifying which episode each bin belonged to. The data was then split into a training and a test set. After training the SVM, I computed the accuracy on the held-out training data.

CHAPTER 3

Prior behavioral performance determines stability and synchrony in the memory system during rest

3.1 Abstract

A major challenge to our current views on information storage in the brain comes from growing evidence indicating that neural activity in cortical and subcortical areas undergoes changes and significant reorganization even in the absence of novel experiences or learning. Such fluidity in the nature of neural activity posits challenges for our understanding of processes based on the exchange of information between multiple brain regions, such as memory acquisition and consolidation. Here, I focused on the coordination of activity between the hippocampus (HPC) and the medial prefrontal cortex (PFC) - two regions known to be relevant to memory functions. Specifically, I studied how prior performance in a rule-switching task impacts the stability and coordination of the resting activity in both areas.

Rest periods are known to be important for the consolidation of memories, potentially involving the transfer of information between the hippocampus and the prefrontal cortex. I found that during rest neural activity underwent systematic changes both in the PFC and in the HPC. Such reorganization differed in the two regions in terms of degree and relationship with previous behavior: while hippocampal reorganization increased after poor task performance, such a trend was not observed for the prefrontal cortical activity. Nevertheless, the neural activity in the PFC showed more systematic changes during rest than the hippocampal activity.

Poor behavioral performance was also followed by an increase in coherence between the short-term fluctuations in neural activity in the two regions, possibly as an effect of a heightened flow of information between HPC and PFC in response to cognitive demands. The synchrony was mostly driven by sustained periods of low-variance (∼5-20s) in the neural activity, which coincided in both populations and which were entrained by an ultraslow oscillation in the range of 0.1 Hz driving the alternation between them in slow wave sleep (SWS) periods. The majority of PFC and HPC cells were down-regulated during these brief periods, but a small subset of cells consistently increased their activity. I identified opposing firing rate modulations during SWRs, indicating that active cells during low-variance periods differ from SWR up-regulated cells. The identified subsets

of cells showed significant differences in coding task-relevant features and they might contribute to different aspects of solving and learning the rule switch task.

3.2 Introduction

Many situations of everyday life require the ability to flexibly adapt behavior based on changes in the environment. Evidence about these changes might not always be readily available and one might need to infer the accurate behavior based on the outcome of performed actions. This cognitive process has been linked to activity in the prefrontal cortex (PFC)[227,228]. Previous studies have shown that PFC is involved in decision-making, flexible behavior, and storage of knowledge[229–231]. To solve different tasks, knowledge about previous experiences and more generalized knowledge about when to change behavior is crucial[2]. The process of acquiring and consolidating relevant memories has been attributed to the hippocampal-neocortical network[232,233]. Along these lines, sleep and rest have been found to elicit the strengthening of relevant memories[9,161]. One mechanism contributing to the cooperative processing of information is the coupling of hippocampal sharp-wave ripples (SWRs) and cortical spindles[182]. However, both areas could potentially communicate through alternative mechanisms such as slow oscillatory activity or during periods of reduced cortical activity as well[234,170]. A more thorough understanding of simultaneously occurring neural activity in the hippocampus and the prefrontal cortex might therefore uncover how information is processed to facilitate flexible behavior in the future.

In addition to the open questions regarding the transfer of information between brain areas, previous studies have shown that neural representations are not as stable as initially thought[114,107,102]. Recent investigations, including the study of Chapter 1, uncovered systematic changes in neural activity during rest/sleep or without external stimuli[235]. There is also evidence that these changes actually actively contribute to changes in representations across matching sensory experiences[235]. The link between offline neural activity changes and changes in behavioral neural activity could therefore help us to understand the mechanisms that drive the fluidity of representations. A first step in this direction is to uncover the relationship between prior behavior and subsequent offline activity drift. Furthermore, the coordination of changes in the neural activity during rest from populations of the same functional network might be a critical component concerning information transfer and remains to be explored[107].

In this study, I assessed systematic changes in neural activity in two memory-relevant brain regions (HPC and PFC) during rest after the execution of a rule-switch task. I found a greater amount of change in the PFC neural activity as compared to activity from the hippocampus. For the hippocampal neural activity during rest, good performance was followed by an increase in activity reconfiguration. Slow fluctuations in the neural activity across the two populations were more synchronized after poor performance, a potential marker of active transfer of information between the two areas. The synchrony was primarily driven by the co-occurrence of low-variance periods in both populations. Active cells during these periods differed from those cells that were recruited during rest SWRs, suggesting that there might be an alternative mechanism besides SWR-spindle coupling to coordinate activity between the HPC and PFC.

3.3 Results

Offline activity drift was observed during short rest and changed with prior performance

To assess the influence of prior behavior on the offline activity drift during rest, I used data from a rule-switch paradigm with interleaved rest periods which was previously recorded in our lab[11]. Animals were trained to apply one of two pre-trained rules in order to navigate from the start location of the plus maze to the rewarded arm (Fig. 3.1a). During the first behavioral episode, the food-deprived animal had to follow the first of the two pre-trained rules (rule A) to receive a food reward. Subsequently, the animal was placed in a rest box for 40 minutes (rest 1). After the rest period, the rewarded arms were still assigned according to rule A. At a random point in time, the rule was switched from rule A to the other pre-trained rule B without indicating the change to the animal (rule switch). Therefore, the animal had to independently adapt its strategy to rule B to navigate to the correct arm to receive a reward. I measured the behavioral performance by counting the number of trials until the animal correctly executed five consecutive trials adhering to rule B (Fig. S3.1a). After the rule switch session, the animal was allowed to rest in the rest box for another 40 minutes (rest 2) before concluding the experiment with another behavioral episode with rule B.

As a first step, I wanted to quantify if the neural activity during rest periods underwent systematic changes. I binned the neural activity of each population using temporal bins (1 s) and trained a linear regression model that predicted time passed during rest based on the neuronal firing (Fig. 3.1b). I then tested the model's performance on held-out data (80/20 train-test split) from the same session and computed the coefficient of determination (report mean value across 100 train-test iterations with changing training and test splits). In this way, I was able to evaluate how tightly changes in neuronal firing were linked to time passed using a linear model.

I observed that the prefrontal cortical neural activity was a better predictor of time compared to the hippocampal neural activity and therefore showed stronger offline activity drift ($p < 0.05$, MWU test, Fig. 3.1c). I confirmed this result using a multinomial logistic regression model (see Methods and Fig. S3.1b). Moreover, the offline activity drift was more pronounced in rest 2 after the rule switch episode as compared to rest 1 ($p < 0.05$, MWU test). This might be an indication of more pronounced systematic changes in the neural activity in both regions after a more challenging behavioral episode.

After establishing that offline neural activity drift could be detected during the rest periods using the described ridge regression approach, I wondered if there was a relationship between the amount of drift and prior rule switch performance. Therefore, I correlated the difference in offline activity drift between rest 2 and rest 1 with the behavioral performance as measured by the number of trials the animal took to adjust its strategy to the changed rule (Fig. 3.1d-f). I found a negative, but not significant correlation for the PFC ($R = -0.5$, $p = 0.116$) and a significant positive correlation for the HPC ($R=0.63$, $p = 0.036$). Although not significant, there was a weak trend for PFC neural activity to undergo more systematic changes in rest 2 after good rule switch performance with respect to rest 1 (Fig. 3.1e). After poor rule switch performance, neural activity in the PFC tended to reduce the amount of change in the rest after (Fig. 3.1f). The amount of hippocampal offline activity drift increased after poor performance but showed less change between rest 1 and rest 2 after good performance.

Next, I investigated if the offline activity drift showed temporal consistency between HPC and PFC. For each population, I projected the neural activity onto the predictive axis of the corresponding ridge regression model to obtain two time series (Fig. 3.1g). Cross-correlating the time series yielded a clear peak at zero offset with a mean correlation value of R = 0.61 (SEM=0.196, Fig. 3.1h). Thus, the offline activity drift in the HPC and the offline activity drift in the PFC were temporally aligned, despite the difference in absolute magnitude.

In summary, my analysis revealed that the neural activity during short consolidation periods ($\sim$ 40 min) underwent systematic changes and that the magnitude of change was biased by the prior performance of the animal. I found that changes in the hippocampal neural activity were more pronounced after good performance. On the other hand, the neural activity in the prefrontal cortex showed the opposite trend, exhibiting more stable activity after good performance. Despite their differential relationship with the animal's rule-switch performance, changes in neural activity in both brain regions were strongly correlated. This hinted at performance providing only modulatory effects on the overall activity structure, which appeared to be dominated by the synergy between the two regions.

Synchrony of low-frequency fluctuations in the neural activity between HPC and PFC is modulated by prior behavioral performance

To further investigate how changes in neural activity were related between HPC and PFC, I performed the following analysis. First, I computed the cosine distance between subsequent population vectors (1s bins), for each brain region independently, to capture time-point to time-point changes in neural activity (Fig. 3.2a). The two resulting time series of short-term fluctuations showed a significant correlation with a peak at zero offset (mean R = 0.371, SEM R = 0.052, all 10 sessions, except for one, had an R $>$ 0.17 and p $<$ 10e-15, Fig. 3.2b).

Since the HPC-PFC circuitry is potentially involved in consolidating knowledge related to the successful execution of the rule switch task, I speculated that changes in the synchrony between both areas could be related to the prior task performance. To test this possibility, I correlated the difference in synchrony between rest 2 and rest 1 for each offset separately with the number of trials the animal took to adapt its behavior after the rule switch (Fig. 3.2c). I found a positive correlation for all offsets between -10 and +10 seconds (R $>$ 0.6, all p $<$ 0.05) indicating that the synchrony in the rest after increases with poor rule switch performance. This finding was most prominent when I only used the first half of each rest (Fig. 3.2c) but held true when using the entire rest data (S2a-b). The highest correlation between the difference in synchrony and rule switch performance was found for an offset of -2 seconds (R=0.96, p=3.35e-6, Fig. 3.2c-d), corresponding to changes in which hippocampal activity led by two seconds relative to the prefrontal cortex activity. Still, it is important to note that there was an overall increase in synchrony following poor rule switch performance that was not exclusive to a specific temporal offset between both neural populations (Fig. S3.2c).

The observed change in synchrony could potentially have been initiated during the behavioral rule switch episode itself. To exclude this option I applied the same analysis, cross-correlating the time series of cosine-distances, to the awake neural activity during

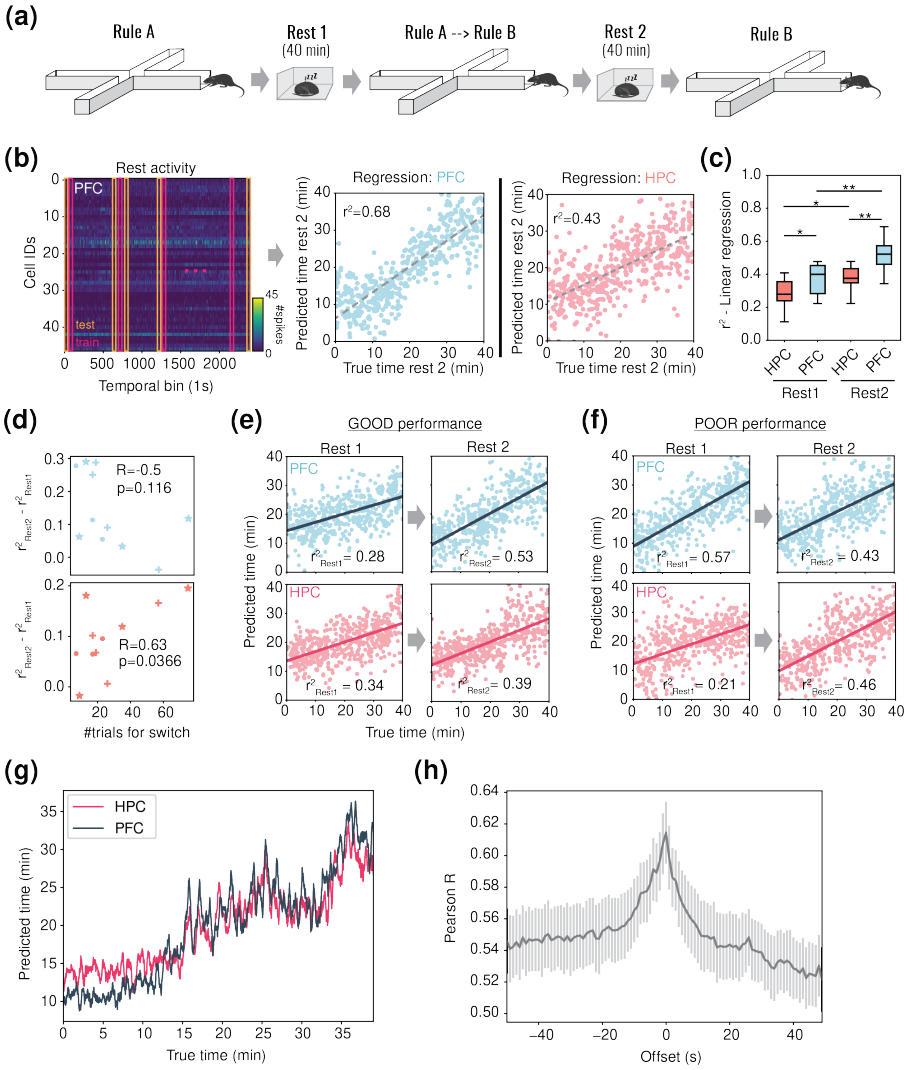

Figure 3.1: **Offline activity drift in the hippocampus and prefrontal cortex.**(a) Experimental paradigm of the rule-switch task: a rewarded arm of the plus maze according to the pre-learned rule A (first behavioral episode), rest interval of 40 min (Rest 1), rule switch task (the rule is switched to pre-learned rule B unannounced), rest interval of 40 min (Rest 2), rewarded arm of the plus maze according to the pre-learned rule B (last behavioral episode). (b) Predicting the time passed during rest to measure the amount of offline activity drift. Left: example raster from rest 2. Predicted and true time during rest 2 using PFC neural activity (middle) and HPC neural activity (right). (c) Coefficient of determination of the ridge regression to predict time for HPC and PFC during rest 1 and rest 2. (d) Difference in coefficient of determination between the rest before (rest 1) and rest after (rest 2) as a function of rule switch performance. Different symbols correspond to different animals. (e)-(f) True and predicted time during rest before and after for two example sessions with good and poor rule-switch performance. (g) Time traces of smoothed predicted time values for HPC and PFC for one example session. (h) Cross-correlation of un-smoothed predicted time traces from HPC and PFC for all sessions (mean: 0.614, standard error of the mean: 0.0196).

the rule switch task (Fig. S3.2d). Since I did not find a significant correlation (R=0.107, p=0.75) between the synchrony during behavior and the performance of the animal, I concluded that the increase in synchrony is an isolated phenomenon happening during

rest after the task.

The cosine distance between subsequent population vectors showed a strong alignment with the overall activity of the population. That means that periods of low cosine distance values overlapped with periods of low overall neural activity and periods of high cosine distance values coincided with periods of high overall activity (Fig. S3.2e). When I used the synchrony between the overall activity of HPC and PFC, and correlated the difference between rest 1 and rest 2 with the animal's performance, I didn't find a significant relationship (R=0.38, p=0.25 for an offset of zero seconds, Fig. S3.2f). Therefore, the synchrony of short-term fluctuations in neural activity, although correlated with the overall population activity, seemed to be more tightly modulated by prior performance.

The computation of cosine distances between subsequent population vectors (Fig. 3.2a) was performed using 1s temporal bins. To investigate whether the observed modulation of synchrony with respect to prior rule switch performance was primarily driven by fluctuations at this temporal resolution, I resorted to a filter analysis approach. High-pass filtering the time series of the cosine distances of HPC and PFC before cross-correlating them, revealed that the low-frequency components were the main driver of the observed behavioral modulation (Fig. 3.2e). Applying a band-pass filter uncovered the frequency range between 0 and 0.05 Hz to be crucial (Fig. 3.2f-g). Consequently, prior rule-switch performance most strongly affected the synchrony of changes in the neural activity between the HPC and PFC at these ultra-slow frequencies.

Taken together, I found that fluctuations in neural activity between the hippocampus and prefrontal cortex aligned strongly and that their synchrony was increased after poor rule switch performance - a potential sign of heightened information exchange between both areas as part of the consolidation process. Furthermore, I identified low-frequency components in these fluctuations which mainly drove the synchrony between the hippocampus and prefrontal cortex.

Different regimes of short-term fluctuations correspond to different brain states

The time series of cosine distances between subsequent population vectors showed distinct regions with stereotypical distributions of cosine distance values (Fig. 3.2a). This raised the question if these regimes corresponded to different brain states and whether the synchrony in neural activity fluctuations was modulated by prior behavior for all of them. As a first step, I clustered the hippocampal time series of cosine distances using a semi-automatic clustering approach (Fig. 3.3a, Fig. S3.3a, and Methods). Temporal windows of 60 seconds duration were assigned to one of three clusters: regime 1 showed high fluctuations in neural activity (high cosine distance values) with brief periods of stable activity (low cosine distance values). Regime 2 corresponded to generally more stable activity as indicated by lower cosine distances between subsequent population vectors. The transition cluster captured time periods of great variability in cosine distance values. I found that regime 1 was the predominant one during rest ($p < 0.01$, regime 1 duration vs. regime 2 duration for all sessions) and that its duration did not change from rest 1 to rest 2 ($p > 0.05$, MWU test, data from all sessions, Fig. S3.3b).

Next, I computed the power of standard frequency ranges in the HPC and PFC, respectively, and aligned the results with the cosine distance clusters (Fig. 3.3a and Fig. S3.3c). Regime 1 was further subdivided into periods of high cosine distances (HI) and periods

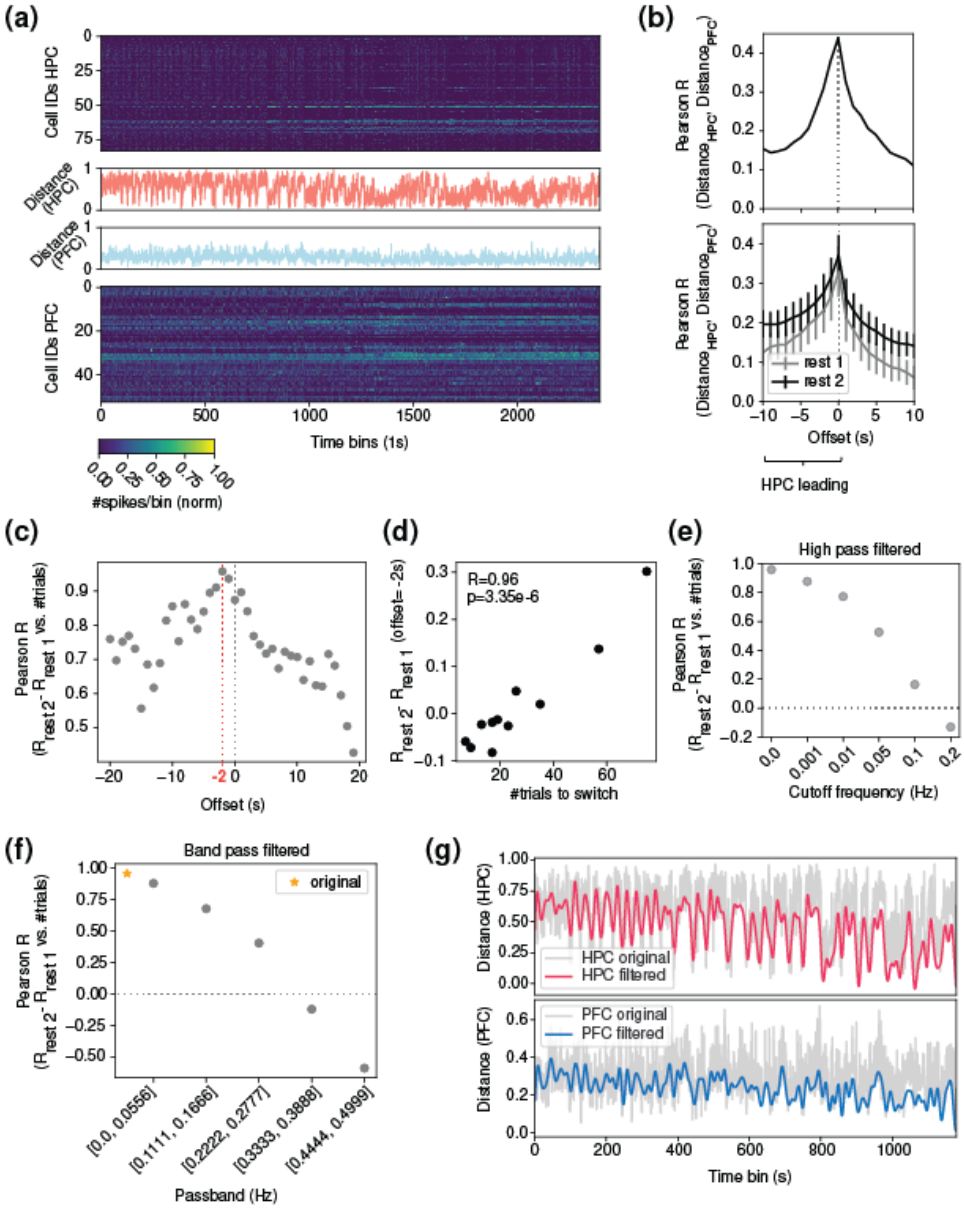

Figure 3.2: **Offline activity drift in the hippocampus and prefrontal cortex.**(a) Computing short-term fluctuations in neural activity. Raster plots of neural activity during rest from HPC (top) and PFC (bottom). Cosine distance between neighboring population vectors in HPC (salmon) and PFC (blue). (b) Cross-correlation of short-term change from HPC and PFC. Top: data from (a). Bottom: Data from all sessions for rest 1 and rest 2. (c) Correlation between the difference in cross-correlation (short-term changes) from rest 1 and rest 2 using the first half of each rest for different offsets and the number of trials the animal took for the rule-switch. (d) Number of trials the animals took to switch the rule versus the difference in correlation (short-term fluctuations) from rest 1 to rest 2 for an offset of -2s (HPC leading). (e)-(f) Correlation between the difference in cross-correlation (short-term fluctuations) from rest 1 to rest 2 (offset: -2s) and the number of trials the animal took for the rule-switch after filtering the short-term fluctuations in HPC and PFC with different filters. Results for high-pass filtered (e) and bandpass filtered (f) short-term fluctuations. (g) Bandpass filtered (passband: [1e-5 Hz, 0.55 Hz]) short-term fluctuations in HPC and PFC for one example session.

of low cosine distances (LO). I identified significant differences in delta, theta, gamma, ripple, and spindle power between the different regimes ($p < 0.001$ for all comparisons, Fig. 3.3b-c, Fig. S3.3c-e), implying that the regimes corresponded to different brain states. Given the spectral properties of the two regimes and the lower Theta-Delta ratio for regime 1 (Fig. 3.3c) it is very likely that regime 1 corresponds to NREM sleep and to an alternation of SWS with periods of low-delta power regulated by an ultraslow oscillation of around 0.1Hz. Regime 2 properties are instead highly overlapping with those expected from REM sleep.

When I looked at the modulation of synchrony of low-frequency fluctuations in the neural activity between HPC and PFC (see previous section) due to prior behavior for each regime separately, the effect remained for regime 1 (max correlation for offsets between -5s and +5s: R=0.75, p=0.0068) and the transition regime (R=0.674, p=0.023), but not for regime 2 (R=0.408, p=0.213, Fig. S3.3e-g). Therefore, the changes in synchrony in regime 1 from rest 1 to rest 2 seemed to be most sensitive to the animal's rule switch performance.

To summarize, fluctuations in neural activity could be assigned to distinct regimes during the consolidation period. This classification was based on the distribution of cosine distances between subsequent population vectors: the different regimes corresponded to periods where the neural activity fluctuated to different extents over small temporal scales. Interestingly, the identified regimes showed well-defined oscillatory profiles in the local field potential and therefore are likely to overlap with specific brain states. Increased synchrony between the fluctuations in neural activity from HPC and PFC after poor performance was most prominent during regime 1. This regime showed periods of high fluctuations interspersed with periods of high stability in neural activity and most likely corresponded to NREM sleep.

low-variance periods in PFC and HPC coincide and different cells are up- and down-regulated

Regime 1 was characterized by strong fluctuations in neural activity interspersed with periods of high stability in neural activity (Fig. 3.3a and Fig. 3.4a). These low-variance states lasted for 1-25 seconds (Fig. S3.4a), coincided with a decrease in oscillation power ($p < 0.001$, MWU test, Fig. 3.3b-c, Fig. 3.4a and Fig. S3.4b), and were a main driver for the synchrony in low fluctuations between HPC and PFC. Hence, when I removed intervals of low-variance from the data, the correlation between low fluctuations in HPC and PFC significantly decreased ($p < 0.05$, MWU test, Fig. S3.4c). Not surprisingly, the low-variance periods in HPC and PFC showed a strong temporal overlap (68% of absolute temporal differences between troughs < 5s, Fig. 3.4c-d). However, neither the number of low-variance periods nor the temporal fraction of low-variance periods in regime 1 showed a positive relationship with the animal's behavior ($p > 0.05$, Fig. S3.4d-e).

During a low-variance period, the overall population activity of excitatory cells decreased, while the activity of a subset of cells was up-regulated (Fig. 3.4a,e-g). This effect was more prominent in the HPC as compared to the PFC (p=3.2754e-98, one-sided MWU). When I looked at the activity of interneurons, the difference between the HPC and PFC was even more prominent (HPC > PFC, p = 5.347e-5, one-sided MWU, Fig. S3.5a-c). While the majority of interneurons in the HPC decreased their firing during a low-variance period (fraction of positively correlated cells with HPC cosine distance > 0.5, p = 0.00452,

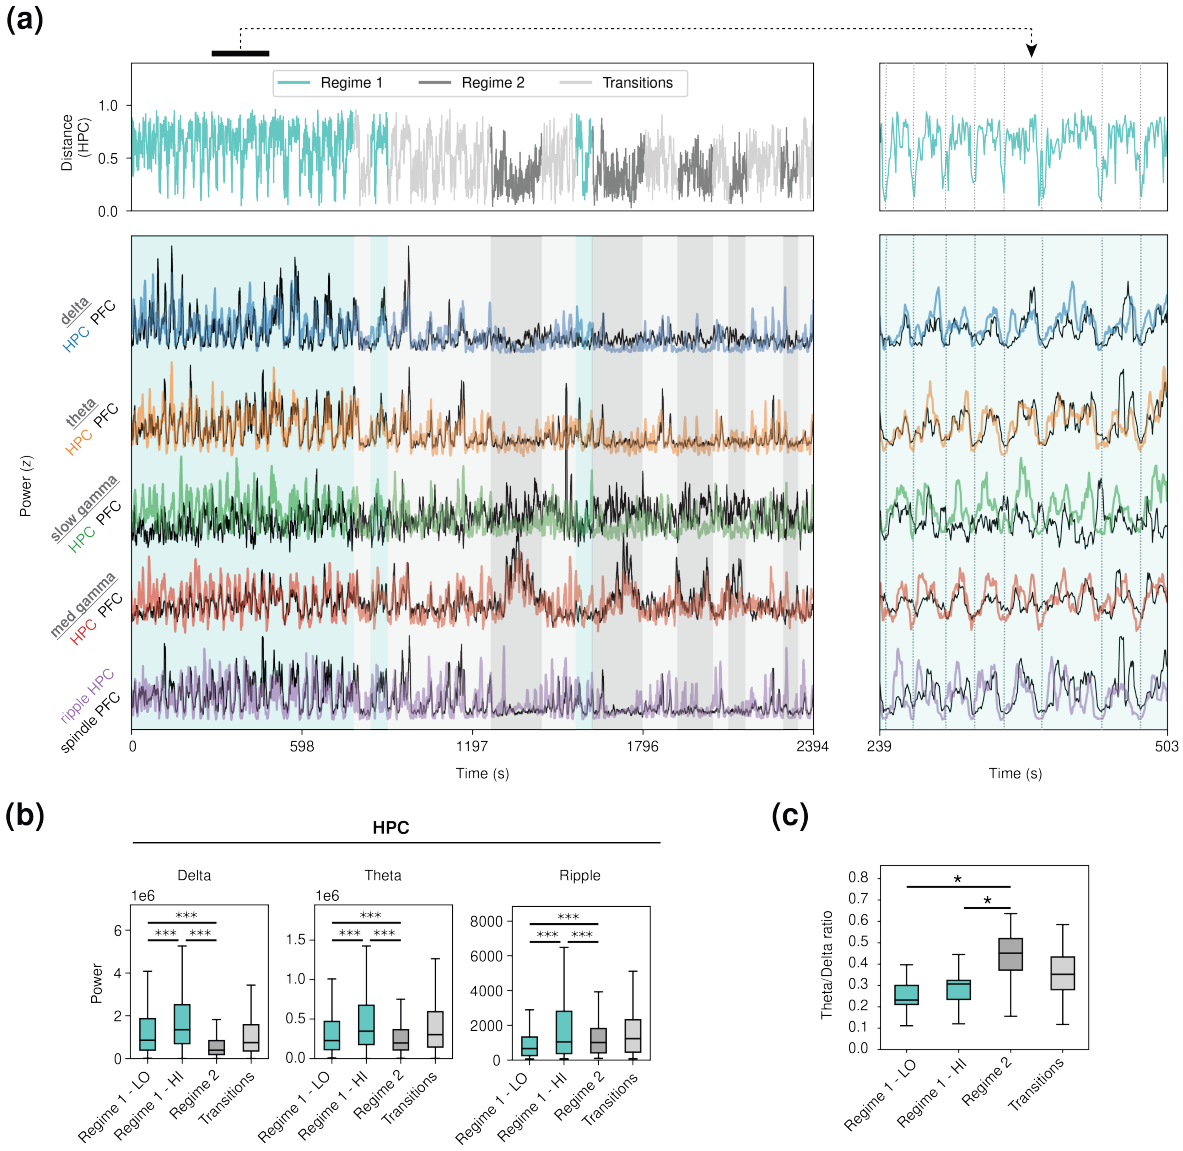

Figure 3.3: **Different regimes of low frequency fluctuations.**(a) Result of semi-automatic clustering approach applied to hippocampal neural activity using Gaussian mixtures to divide short-term change (cosine distance between subsequent population vectors) into different regimes for one example session (top left). Top right: zoomed-in snippet from regime 1 showing low-variance periods. Bottom: Delta, Theta, slow Gamma, medium Gamma, and ripple/spindle power in the hippocampus (colored) and prefrontal cortex (black). (b) Delta, Theta, and ripple/spindle power in the hippocampus for different regimes (data from all sessions, *** $p < 0.001$, Mann-Whitney U with Bonferroni-correction). (c) Hippocampal Theta-Delta ratio for different regimes.

binomial test) in a similar fashion like the excitatory cells, many PFC interneurons were up-regulated (fraction of negatively correlated cells with HPC cosine distance > 0.5, p = 8.472e-8, binomial test). This result might imply that the majority of the HPC network was shut down during a low-variance period, whereas the PFC network showed a more local inhibition potentially driven by PFC interneurons.

In conclusion, I found that low variance periods majorly contributed to the synchrony between fluctuations in the neural activity between HPC and PFC. The discovery that a certain subset of cells from both populations was up-regulated during these periods led to

the following hypothesis: the co-activity of distinct cells during low variance periods could potentially be a mechanism to synchronize the activity between the HPC and the PFC.

HPC-PFC co-activity differs between low-variance periods and SWRs

To test the possibility that both populations synchronized their activity during low-variance periods, I first wanted to establish if the same subset of cells was consistently up-regulated during these periods. I computed the normalized difference in mean firing rates for excitatory cells between low-variance and high variance periods in regime 1 for rest 1 and rest 2, separately. The firing rate modulations were strongly correlated between rest 1 and rest 2, indicating that a similar subset of cells was consistently up- or down-regulated during low-variance periods ($R > 0.66$, $p < 3.2e\text{-}118$, Fig. 3.5a-b).

HPC-PFC communication has been ascribed to the coupling of hippocampal sharp-wave ripples (SWR) with cortical spindles (10-18Hz oscillations of the local field potential in the cortex). To identify the relationship between SWR-driven activity and the activity during low-variance periods, I identified the subsets of cells that were recruited in both cases. Surprisingly, the SWR-modulation of firing rates was anti-correlated with the low-variance period-modulation ($R < -0.35$, $p < 2.2e\text{-}28$, Fig. 3.5c-d). This suggested that cells which are up-regulated during low-variance periods are down-regulated during SWRs and vice versa. Hence, certain excitatory cells from HPC and PFC were upregulated together during low-variance periods and this co-activity involved a different subset of cells than the one recruited during SWRs.

To further investigate the possibility of coherent co-activity in HPC and PFC during low-variance periods, I resorted to an approach inspired by the idea of a communication subspace between the two areas 20 (Fig. 3.5e). To uncover how the activity of one population was related to the activity of the other, I performed the following analysis: For each hippocampal excitatory cell, I fit a regression model predicting its activity (in 100ms temporal bins) using the co-occurring activity of all recorded excitatory prefrontal cortical neurons. This yielded one regression direction (i.e. a line indicating a direction defined by the regression weights of all PFC cells) for each HPC cell in the PFC activity space (Fig. 3.5e). Then, I clustered all regression directions based on the cosine distance between them to identify hippocampal cells which were modulated by similar prefrontal cortex activity (F6f, F6h and Methods). I separated the regression directions into three clusters anticipating that there could be one cluster capturing low-variance coactivity, one cluster for SWR coactivity, and one additional cluster representing the remaining coactivity. For each cluster, I computed the mean weight per PFC neuron (using all regression directions of this cluster) describing the stereotypical relationship between PFC and HPC activity (Fig. 3.5g). When I correlated the cluster-specific mean regression weight with the firing rate modulation during SWRs and low-variance periods for all PFC cells, the three clusters showed clear differences (Fig. 3.5i and Fig. 3.5k): for one cluster (cluster 1), the activity of cells with large regression weights tended to be down-regulated during low-variance periods ($R = -0.46$, $p < 5.24e\text{-}40$) and up-regulated during SWRs ($R = -0.42$, $p < 8.79e\text{-}33$). The regression weights of PFC cells for another cluster (cluster 2) showed a positive correlation with low-variance modulation ($R=0.28$, $p < 4.31e\text{-}15$) and a negative correlation with SWR modulation ($R=-0.15$, $p < 3.11e\text{-}5$). For the third cluster (cluster 3), regression weights showed a negative correlation with the firing rate modulation during low-variance periods ($R=-0.28$, $p < 3.57e\text{-}15$) and a positive correlation

with the modulation during SWRs (R=0.45, p < 5.78e-38). Therefore, the two clusters seemed to be more similar in terms of firing rate modulation (first and third cluster) than cluster 2. The greater similarity of these two clusters was further confirmed by the smaller angle between the mean predictive directions of the two ($p < 0.05$, Fig. S3.6a).

To support the argument that there is a state-dependent link between the activity of HPC and PFC cells, a potential indicator for communication between the two populations, I wanted to test if the predicted HPC cell activity showed the same modulation during low-variance periods and SWRs as the predicting PFC cell activity. Since each predictive axis corresponded to one HPC cell, I assigned each HPC cell to the corresponding cluster. This implied that the activity of HPC cells from the same cluster was modulated by similar PFC activity. Next, I compared the firing rate modulation of the HPC cells from different clusters during low-variance periods and SWRs (Fig. 3.5j and Fig. 3.5l). Coherently, I found that HPC cells belonging to a cluster characterized by a reduction in the activity of predicting PFC cells during low-variance periods, also showed down-regulated activity during these periods (low-variance periode modulation < 0, p<4.1545e-103, one-sided T-test). On the other hand, the same HPC cells increased their firing during SWRs and therefore mirrored the modulation of the predicting PFC cells (SWR modulation > 0, p<4.3545e-113, one-sided T-test). HPC cells from the cluster with an up-regulation of predicting PFC activity during low-variance periods were less likely to reduce or increase their firing rates during low-variance periods and SWRs, respectively ($p < 0.001$, Fig. 3.5l).

When I repeated the above analysis using HPC activity to predict PFC activity, the results showed the same trends (Fig. S3.6b-d). However, the differences in firing rate modulations between the clusters were smaller in magnitude.

It is important to note that the emergence of clusters with different cell activity modulation during low-variance periods and SWRs was not an explicit constraint of the prediction and clustering procedure, but occurred by itself. To exclude the possibility that the finding was mainly due to the chosen analysis approach, I performed a canonical correlation analysis (CCA) using the same data. This allowed me to identify the contribution of single cells to the correlation of neural activity between HPC and PFC. In support of the idea that distinct activity patterns from both populations co-occur during low-variance periods, the cell-specific weights defining the first pair of canonical variables showed a significant positive correlation with the firing rate modulation during low-variance periods ($R > 0.5922$, p < 1.94e-71, Fig. S3.6e-f).

In summary, the used prediction and clustering approach identified subsets of cells in one population (HPC or PFC) whose activity could be predicted using distinct activity patterns from the other population. Two of the identified clusters (cluster 1 and cluster 3) contained cells that were up-regulated during SWRs and down-regulated during low-variance periods. For these clusters, the mean regression weights of the cells from the other population were positively correlated with the SWR firing rate modulation and negatively correlated with the low-variance period modulation. Therefore, the co-occurrence of activity patterns from HPC and PFC for these clusters was mainly driven by SWRs.

On the other hand, the mean regression weights of the predicting cells from the remaining cluster (cluster 2) showed a negative correlation with the SWR firing modulation and were positively correlated with the low-variance period modulation. For this cluster, the predicted cell activity was less down-regulated during low-variance periods and showed a significantly reduced up-regulation during SWR. Consequently, I identified activity

patterns that co-occurred in HPC and PFC during low-variance periods which were different in terms of involved cells from activity patterns co-occurring during SWRs.

Taken together, my analysis revealed co-activity patterns in the hippocampus and prefrontal cortex that were either driven by low variance periods or SWRs. The observation that different subsets of cells in the HPC and PFC were co-active during either low variance periods or SWRs might indicate that these subsets serve different functions.

Difference in spatial coding and rule coding between SWR and low-variance period favoring cells

The finding that different subsets of cells from HPC and PFC were co-active either during SWRs or low-variance periods during rest raised the question of whether these cells also differed in terms of activity during the rule switch paradigm.

Mean firing rates of hippocampal cells from the three clusters showed mild differences during the rule switch session with significantly higher mean firing rates for cluster 3 (cells with intermediate down-regulation during low-variance periods and strong up-regulation during rest SWRs) in the behavioral episode after rest ($p < 0.05$, Fig. 3.6a-b). Mean firing rates stayed relatively constant between the rule switch episode and the behavioral episode of the recall after rest for all three clusters ($p > 0.05$, Fig. 3.6c). When I repeated the same analysis for PFC cells (clusters derived from predicting PFC activity using HPC activity during rest), I did not find any difference in mean firing rates across clusters nor changes between behavioral episodes ($p > 0.05$, Fig. S3.7a-c). Hence, only hippocampal cells that were intermediately down-regulated during low-variance periods and strongly up-regulated during SWRs in rest showed higher mean firing rates than the other cells and maintained this activity in the recall episode.

Since the rule switch paradigm has a spatial component and hippocampal cells are coding for space, I was curious to see whether there were differences in spatial information for the cells from different clusters (Fig. 3.6d-f). Interestingly, cells from the low-variance favoring cluster had lower spatial information during the rule switch task ($p < 0.01$, MWU test) and showed a moderate, but significant increase in spatial information in the recall session after the rest ($p < 0.001$, MWU test). On the contrary, cells from the SWRs favoring cluster showed a reduction in spatial information from the rule switch episode to the recall episode ($p < 0.001$). Using sparsity as an alternative measure for spatial coding confirmed the reduced spatial selectivity of cells from the low-variance favoring cluster ($p < 0.01$, MWU test, Fig. S3.7d). Changes in sparsity from the rule switch episode to the recall episode for the clusters only partially agreed with the observed changes in spatial information (Fig. S3.7e-f).

It has been shown that PFC cells can also exhibit some degree of spatial coding[11]. When I compared spatial information and sparsity of PFC cells from the different clusters, cells from cluster 1 (up-regulation during SWRs and down-regulation during low-variance periods) showed a higher spatial selectivity during the rule switch task ($p < 0.05$, MWU test, Fig. S3.7g-m). Interestingly, all PFC cells displayed an increase in spatial information from the rule switch episode to the recall episode ($p < 0.05$, MWU test, Fig. S3.7i). The same trend was observed when using sparsity as a measure for spatial coding (Fig. S3.7m).

In short, I identified differences in spatial information for the hippocampal cells from different clusters, which might imply that they exhibit different functionality while the

animal is engaged in different episodes of the rule switch paradigm.

Previous studies have shown that the neural activity from HPC and PFC can be used to decode the rule the animal is currently following[176]. To test the possibility of cells from different clusters coding for the rule to different extents, I performed the following analysis: for each population, I split the rule switch data into two parts. The first part consisted of the neural activity while the animal was performing the task using rule A. The second part captured the neural activity after the animal successfully switched its strategy to adhere to rule B. Then, I trained a support vector machine (SVM) to separate the activity of rule A and rule B. I confirmed the decodability by cross-validation (mean accuracy > 0.636 for all models). When I compared the magnitude of SVM coefficients for the hippocampal cells from the three clusters, cells from cluster 2 (low-variance favoring cluster) and cluster 3 (cells with intermediate down-regulation during low-variance periods and strong up-regulation during rest SWRs) contributed more to the classification of rule A and rule B activity ($p < 0.01$, MWU test, Fig. 3.6g). However, the hippocampal cells from different clusters did not show a distinct up- or down-regulation of activity between both rules as indicated by similar SVM coefficient distributions ($p > 0.05$, MWU test, S8a). If I used PFC activity instead, there was no significant difference between the clusters (derived from predicting PFC activity using HPC activity during rest) neither for the SVM coefficients nor the absolute SVM coefficients ($p > 0.05$, MWU test, Fig. S3.8b-c).

The SVM rule decoding analysis was purely based on firing rate changes of cells between both rules (rate remapping). Nonetheless, there remained the possibility that cells coded for the rule by changing their spatially defined activity between rules. For each cell, I generated one rate map (firing rates per spatial bin) using the data from rule A and one rate map using the data from rule B. Then, I computed the correlation between these rate maps to evaluate their similarity. Hippocampal cells from cluster 3 (cells with intermediate down-regulation during low-variance periods and strong up-regulation during rest SWRs) showed a slightly, but significantly higher correlation between rule A and rule B ratemaps ($p < 0.05$, Fig. 3.6h) indicating that these cells maintain their spatial tuning across both rules more strongly. The same analysis using the PFC clusters yielded no significant differences ($p > 0.05$, MWU test, Fig. S3.8d).

The animal needed to recall rule B during the behavioral episode after the rest and hippocampal-cortical activity was potentially involved in this process. Hence, I wondered how consistent the spatial tuning of hippocampal cells was between the rule B part of the rule switch episode and the rule B recall episode after rest. Again, hippocampal cells from cluster 3 showed a more consistent spatial tuning between these episodes as indicated by higher correlation values between the corresponding rate maps ($p < 0.01$, MWU test, Fig. 3.6i). PFC cells from different clusters did not show a significant difference in this regard ($p > 0.05$, MWU test, Fig. S3.8e). Finally, I wanted to test whether the preferential firing rate modulation of cells during either low-variance periods or SWRs was somehow related to the overall change in firing rates across rest. Interestingly, the overall change in firing rates across sleep (slope of smoothed firing rate series) was correlated with the low-variance modulation for HPC and PFC cells ($p < 0.001$, MWU test, Fig. S3.8f-g) and anti-correlated with the SWR modulation ($p < 0.001$, MWU test, Fig. S3.8h-i). Hence, cells that were up-regulated during low-variance periods also increased their firing rates throughout rest. On the other hand, cells that were recruited during SWRs tended to decrease their firing rates across sleep. This finding might provide a link to the offline

activity drift analysis from the first part of this chapter.

In summary, hippocampal cell clusters that were computed using their relationship with PFC activity during rest, did show differences in mean firing rates, spatial coding, and rule coding while the animal was performing different episodes of the rule switch paradigm. This indicates that the differential up-regulation of subsets of cells during low variance periods or SWRs might indeed modulate cells based on their functionality during behavior.

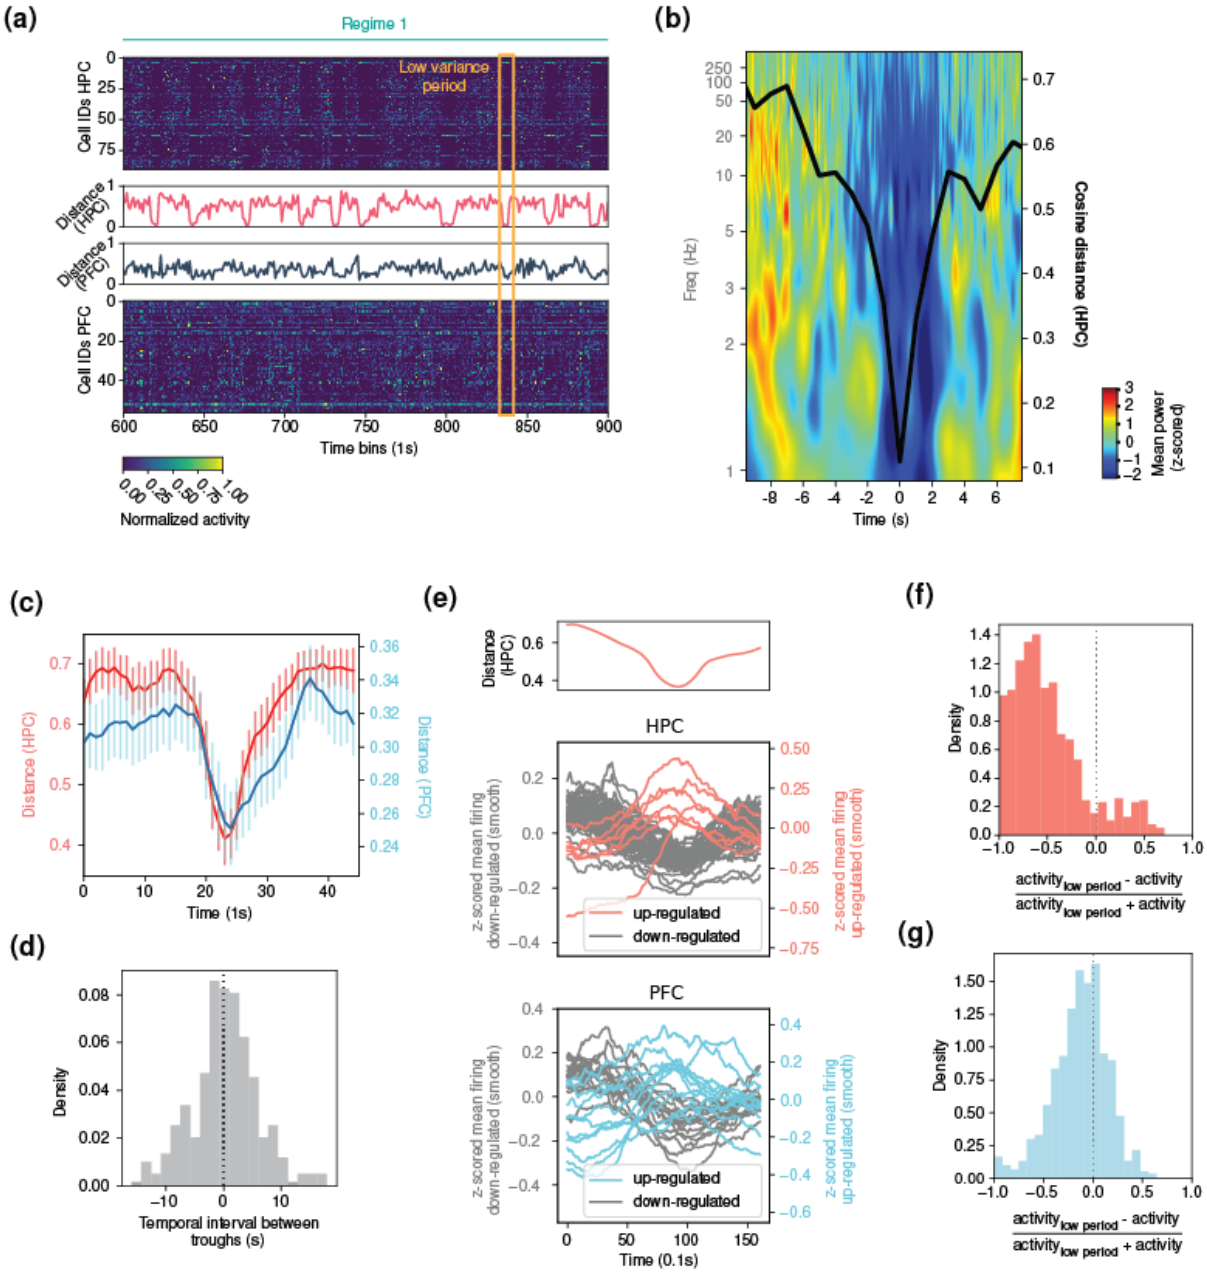

Figure 3.4: **Correlated low-variance periods in PFC and HPC drive subsets of cells.**(a) low-variance periods in regime 1 for one example session. (b) LFP power around low-variance periods for one example session (mean across all low-variance periods, centered at the minimum mean cosine distance value across periods). (c) Short-term changes in neural activity as measured by the cosine distance between subsequent population vectors around low-variance periods in HPC and PFC for one example session (smoothed mean+-std across all low-variance periods). (d) Temporal interval between troughs of the short-term changes in neural activity for all low-variance periods from all sessions. A negative interval corresponds to a trough that occurs first in the HPC with respect to the PFC. A value of zero indicates perfect temporal overlap between the troughs from HPC and PFC. (e) Event-triggered mean of the short-term changes (cosine distance) in HPC (top). Event-triggered z-scored mean firing rates during low-variance periods of up- and down-regulated cells in the HPC (middle) and PFC (bottom) for one example session. (f)-(g) Normalized difference in firing rates within and outside low-variance periods for hippocampal (f) and prefrontal cortical excitatory cells (g).

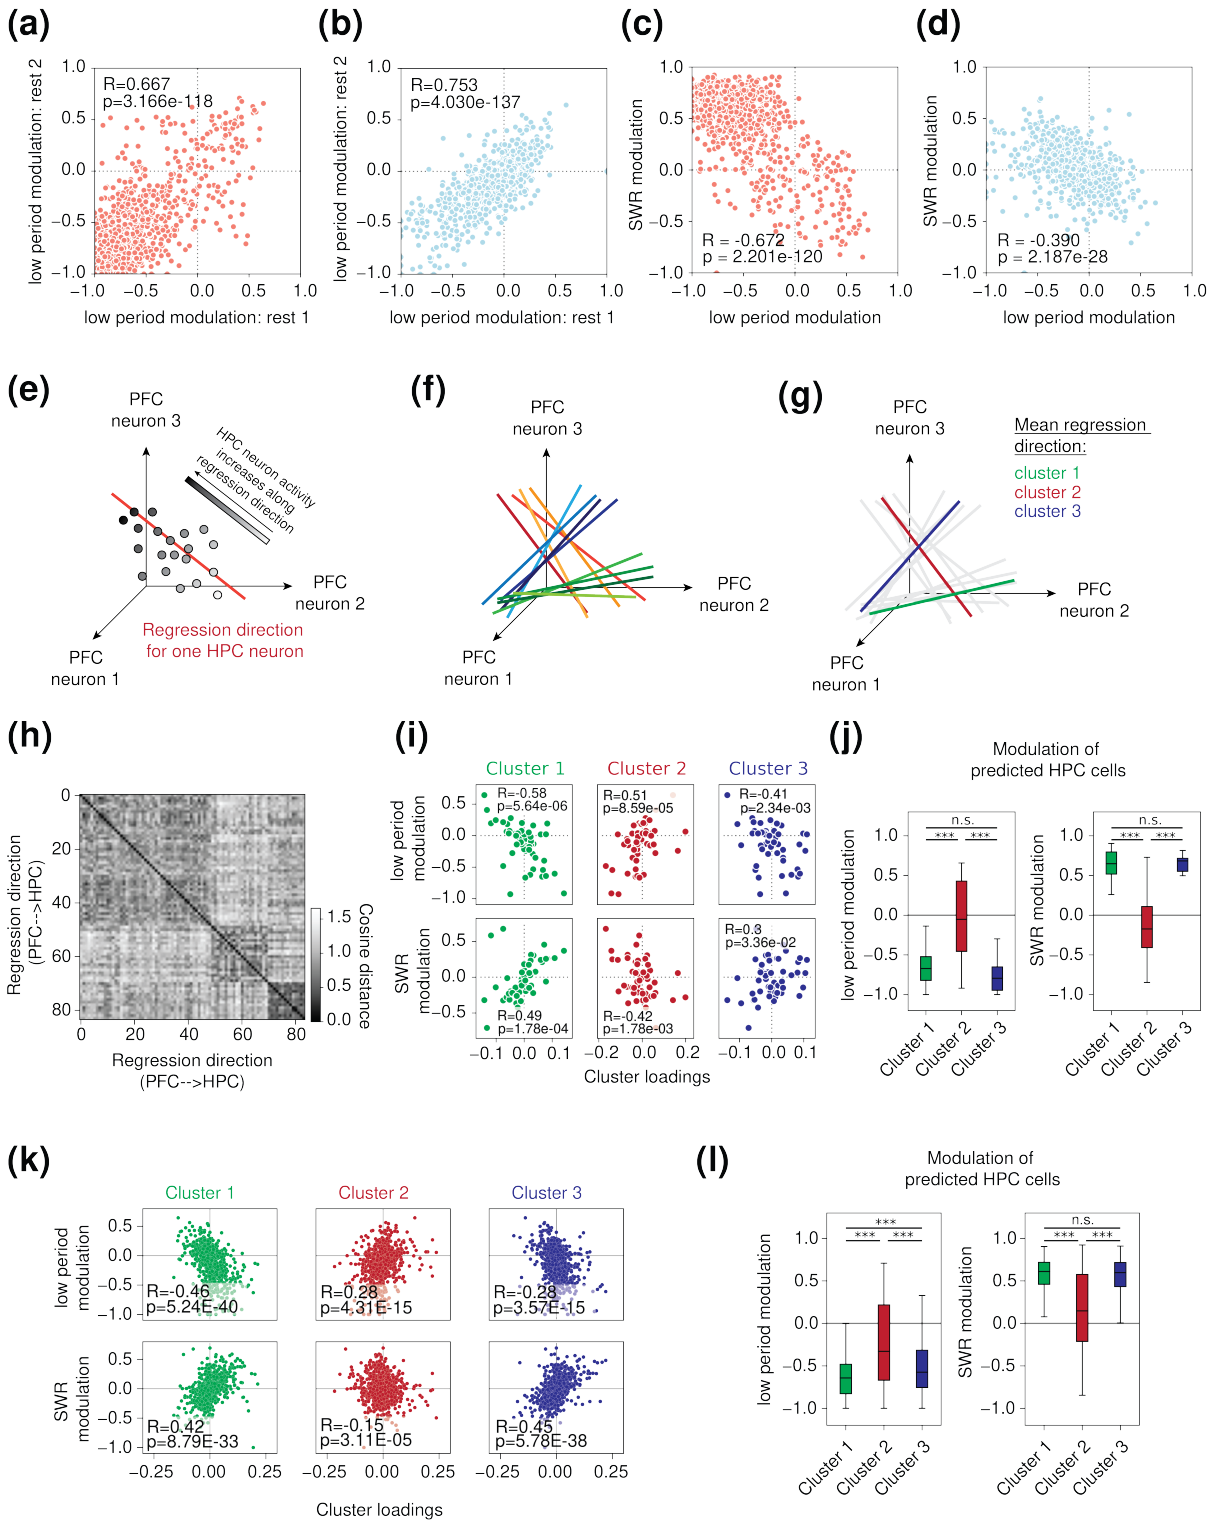

Figure 3.5: **Co-activity between HPC and PFC during sharp wave ripples and low-variance periods.**(a)-(b) Firing rate modulation of excitatory hippocampal (a) and prefrontal cortical (b) cells during low-variance periods in rest 1 and rest 2. (c)-(d) Firing rate modulation of hippocampal (c) and prefrontal cortical excitatory cells (d) during low-variance periods and SWRs. (e)-(g) Schematic of prediction and clustering approach for PFC-HPC co-activity. Visualization of regression directions for one HPC neuron (e, adapted from 20), several HPC neurons (f), and clustered regression directions (g).

Figure 3.5: continued from previous page

(h)-(j) Communication channel results for one example session. (h) Clustering regression directions for all hippocampal neurons (activity for each HPC neuron is predicted by all PFC neurons) using hierarchical clustering for one example session. Each entry corresponds to the cosine distance between the regression directions on the x- and y-axis. (i) Firing rate modulation during low-variance periods and SWRs for PFC neurons (regressors) with the corresponding loadings for the three clusters from (h). (j) Firing rate modulations of predicted HPC neurons during low-variance periods and SWRs for the three different clusters (*** $p < 0.001$, MWU, Bonferroni-correction). (k)-(l) Communication channel results for all sessions. (k) Correlation between firing rate modulation and cluster loadings for PFC neurons (regressors). (l) Firing rate modulations of predicted HPC cells for the three different clusters (*** $p < 0.001$, MWU, Bonferroni-correction).

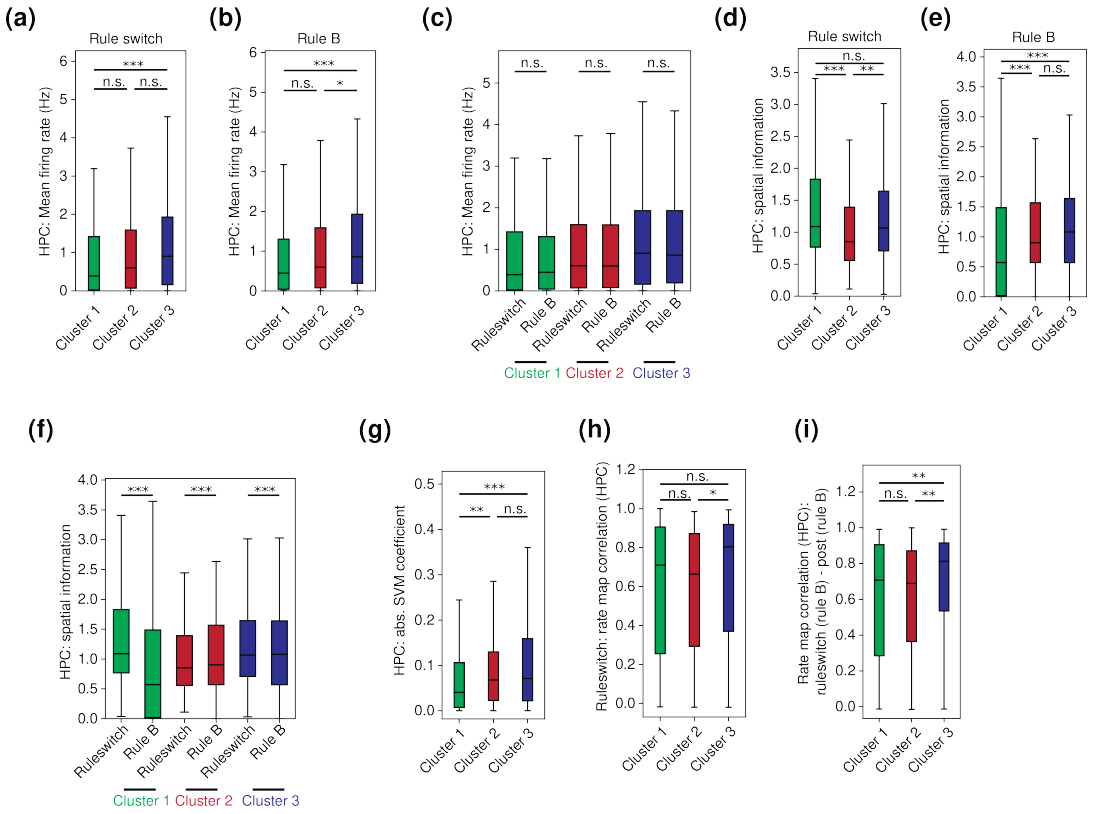

Figure 3.6: **Features of clustered cells.** (a)-(b) Mean firing rates of hippocampal cells from different clusters (see previous section, Cluster 1 & 3: SWRs favoring cells, Cluster 2: low-variance periods favoring cells) during the rule switch task (a) and behavioral episode after rest (rule B, panel b). (c) Comparison of mean firing rates between the rule switch task and the behavioral episode after rest for hippocampal cells. (d)-(e) Spatial information of hippocampal cells during the rule switch task (d) and the behavioral episode after rest (e) for the three different clusters. (f) Changes in spatial information for hippocampal cells from the rule switch task to the rule B recall episode after rest for the three different clusters. (g) Rule coding: magnitude of support vector machine coefficients of hippocampal cells used to classify rule A vs. rule B activity during the rules switch. Each coefficient was assigned to a cluster based on the corresponding hippocampal cell's cluster assignment. (h) Rate map correlations between rule A and rule B during the rule switch task for hippocampal cells. (i) Rate map correlations between rule B of the rule switch task and rule B during the recall episode after rest for hippocampal cells. For all panels: *** $p < 0.001$, ** $p < 0.01$,* $p < 0.05$,n.s. $p > 0.05$, Mann-Whitney U Test with Bonferroni-correction.

Population activity during low-variance periods

After establishing that a subsets of cells from HPC and PFC consistently up-regulated their activity during low-variance periods simultaneously - a potential means of communication - I wanted to provide a better description of the population activity occurring during these periods. Furthermore, I attempted to find a link between low-variance period activity from rest 2 and the other episodes of the experiment.

To start with, I identified low-variance periods during rest 2 using hippocampal neural activity (see Methods). I selected population vectors from HPC and PFC coinciding with the identified time intervals and computed the cosine distance between all pairs within each population. Then, I compared the resulting distribution of distances with the distribution of cosine distances between population vectors of the remaining activity of regime 1 (Fig. 3.7a-b). Not surprisingly, for both populations, low-variance activity appeared to be much more constrained, as indicated by lower cosine distances, than the activity outside of these periods ($p < 10e\text{-}50$, MWU test).

Since low cosine distance values between subsequent population vectors did not only occur during low-variance periods but were also prominent in regime 2 (see Fig. 3.3a), I wondered if the activity during low-variance periods and the activity of regime 2 would be similar to some extent. Indeed, when I computed the pairwise distances between population vectors from the low-variance periods and the population vectors from regime 2, the resulting values were significantly lower than the pairwise distance between low-variance activity and the remaining activity of regime 1 ($p < 10e\text{-}50$, MWU test, Fig. 3.7c-d). Hence, population activity from low-variance periods was constrained and showed greater similarity with activity from regime 2.

To further investigate the stability constraints and evolution of the neural activity during low-variance periods, I calculated the pairwise cosine distances between population vectors considering the number of population vectors that separated the pair (Fig. 3.7e-f). If population vectors occurred in close proximity, they were much more similar than population vectors that were separated by a large number of other population vectors ($p < 0.001$, MWU test, Fig. 3.7f). This result applied to both HPC and PFC low-variance neural activity and implied that the neural activity evolved over time and did not stay constant across low-variance periods. However, changes across a few population vectors in low-variance neural activity were significantly smaller as compared to changes across a few population vectors in the remaining neural activity of regime 1 ($p < 0.001$, MWU test, Fig. 3.7g and Fig. S3.9a). Computing the distance between population vectors as a function of time confirmed the aforementioned findings (Fig. 3.7b).

For the hippocampus, the similarity in neural activity of regime 1 outside low-variance periods was very low (high cosine distance values) and did not show a dependency on the number of population vectors in between ($p > 0.05$, MWU test, Fig. 3.7g). Hence, the corresponding neural activity was neither correlated across longer or shorter timescales. On the other hand, changes in PFC activity from low-variance periods and the remaining regime 1 intervals increased with an increasing number of population vectors in between, implying a stronger correlation of activity on shorter timescales.

Finally, I wanted to investigate whether low-variance period activity patterns were exclusive to rest 2 or also occurred during other episodes of the rule switch paradigm. For each brain region, I used all population vectors from the low-variance periods to compute one mean low-variance population vector as a template. The cosine distances between

the neural activity of rest 1, the rule switch task, rest 2 and the template were then used to evaluate the similarity in neural activity (Fig. S3.9c-e). Not surprisingly, the activity from rest 2 showed the greatest similarity (lowest cosine distance) with the template ($p < 0.001$, MWU test) followed by rest 1 activity. The neural activity during the rule switch task was least similar to the template, indicating that the neural activity during behavior did not resemble low-variance activity from rest 2. I obtained similar results when I correlated the time traces of cosine distances with the templates from HPC and PFC: when the activity in rest 2 was similar to the low-variance activity template in HPC, the PFC activity was also more similar to the corresponding template from the PFC (Fig. S3.9d). Applying the same measure to rest 1 revealed lower, but not significantly different correlation values as compared to rest 2. For the awake neural activity during the rule switch task, there was a reduced co-occurrence of low-variance period-like activity from HPC and PFC as indicated by very low correlation values.

To summarize, neural activity during low-variance periods showed much smaller changes as a function of time than the activity outside these periods. This finding held especially true for activity from the HPC. low-variance period neural patterns and their co-activity in HPC and PFC were mostly observed during rest and only weakly resembled neural activity during behavior.

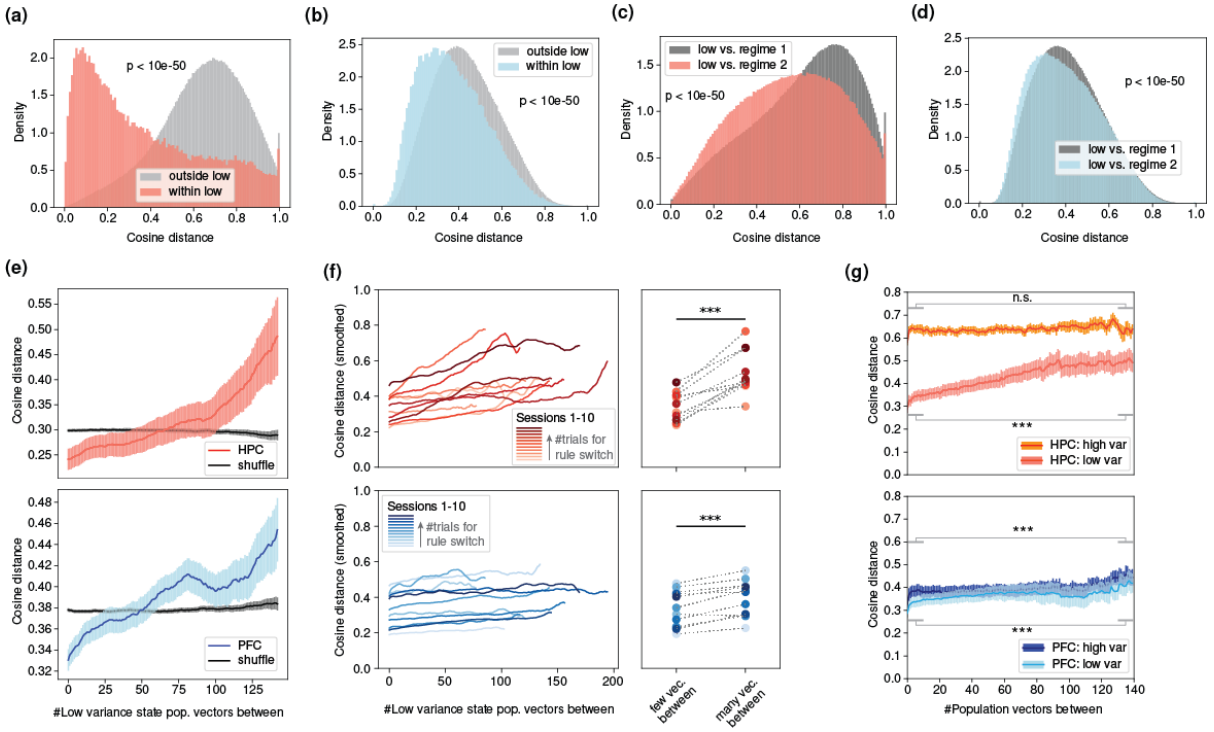

Figure 3.7: **population activity during low variance periods is constrained and changes slowly.** (a)-(b) Distribution of cosine distances between population vectors within low variance periods (salmon and light blue) and within high variance periods (gray) for regime 1. Results for hippocampal (a) and prefrontal cortical activity (b). (c)-(d) Cosine distances between low variance activity and high variance activity of regime 1 (gray) and between low variance activity (regime 1) and activity from regime 2. Results for hippocampal (c) and prefrontal cortical activity (d). (e) Cosine distance between population vectors of low variance periods as a function of the number of population vectors in between for one example session. Top: hippocampal activity (salmon) and results from population vector ID shuffle (gray). Bottom: prefrontal cortical activity (light blue) and shuffle (gray). Mean and sem for data and shuffle. (f) Cosine distance between population vectors of low variance periods as a function of the number of population vectors in between for all sessions. Top panel: hippocampal activity. Bottom: Prefrontal cortical activity. Right panels: Mean cosine distances with few population vectors in between (first 10 data points from left panel) versus mean cosine distances with many population vectors in between (last 10 data points from left panel). *** $p < 0.001$, T-test. (g) Cosine distances between population vectors of low variance periods and between vectors of high variance periods as a function of the number of population vectors in between for all sessions. Top: hippocampal activity. Bottom: prefrontal cortical activity. Distributions of distances using data from the first 10 entries between low variance and high variance periods are significantly different for hippocampal and prefrontal cortical activity ($p < 0.001$, Mann-Whitney U). *** $p < 0.001$, n.s. $p > 0.05$ Mann-Whitney U.

3.4 Discussion

My analyses revealed that offline activity in the hippocampus showed stronger systematic changes after poor rule switch task performance. On the contrary, the PFC neural activity showed a non-significant tendency to change more after good rule switch performance. While activity in the two areas presented such a dissociation over long time scales

(minutes and hours), short-time fluctuations in the order of seconds and tens of seconds were temporally aligned between the hippocampus and prefrontal cortex. Also, this phenomenon presented a relationship with behavior. In fact, poor rule switch performance was followed by an increase in synchrony of ultra-slow fluctuations (∼0.05 Hz) in neural activity between the hippocampus and the prefrontal cortex during rest. Short-term changes in neural activity, which were driving these fluctuations, were clustered into different regimes based on the distribution of similarity between subsequent population vectors. Taking into consideration the power spectrum profile associated with each of these regimes, I estimated that each regime corresponded to a different brain state. Strong delta and ripple components pointed to an identification of regime 1 with NREM sleep and SWS, while the high theta/delta ratio and enhanced medium gamma power appeared to indicate that regime 2 coincided with REM sleep.

Prior rule switch performance significantly modulated the synchrony between HPC and PFC for regime 1, which was characterized by strong fluctuations in neural activity interspersed with brief (1-25s) periods of low-variance in neural activity. These low-variance periods in HPC and PFC coincided and significantly contributed to the synchrony between both populations. During these periods, different subsets of cells were consistently up- or down-regulated and this modulation was anti-correlated with the firing rate modulation during sharp-wave ripples (SWRs). An across-area neural activity prediction and clustering approach revealed the existence of co-active subsets of cells from HPC and PFC. Cells of one subset tended to be consistently more active during low-variance periods and were down-regulated during SWRs. The hippocampal cells of this subset had less spatial information during the rule switch task and were involved in coding for the rule during behavior. Finally, I showed that the population activity during the identified low-variance periods, in particular in the hippocampus, was constrained and its evolution in time was much slower than comparable neural activity recorded during other sleep periods.

Offline activity drift

The offline activity drift I described in the first part of this chapter stands in contrast to the reactivation drift analysis I performed for Chapter 2. Here, I did not relate the neural activity during rest to the neural activity before or after. This can explain why I was able to detect systematic changes in neural activity also during shorter rests (40 min). A complete assessment of the properties of offline activity drift and reactivation drift would require a more thorough analysis and specifically designed datasets.

Although I did not explicitly detect sleep stages, the ability to regress time using rest neural activity could be partially explained by the occurrence of different sleep stages at different times during rest. Sleep stages have been shown to modulate the activity of cells (previous chapter and[188,236]) and their temporal distribution could therefore strongly influence the ability to regress time. The dataset from Chapter 2 could potentially be used to distinguish systematic changes in neural activity happening during REM and NREM sleep separately.

The reduction in offline activity drift in the HPC and the increase in systematic changes in PFC neural activity after good rule switch performance could be interpreted as follows. One possibility is that after the task structure is sufficiently learned, the hippocampus replays the relevant neural activity[237,176] (hence a reduction in offline activity drift) whereas the prefrontal cortical activity is modified to encode this knowledge[238,239]. Changes in the occurrence of reactivations as a function of e.g. novelty have been identified in the

hippocampus before[202]. An alternative interpretation would be that the prefrontal cortical neural activity during the rule switch task had already captured the necessary information. Along these lines, the activity during rest was pushed to another location in neural state space to avoid interference with this information[240].

The observed decrease in changes in the PFC neural activity and the increase in offline activity drift in the HPC after poor rule switch performance could imply that the activity in the PFC is not updated because the task structure is not sufficiently learned. Alternatively, relevant activity could have been reactivated in the PFC to encode this knowledge leading to a reduction in offline activity drift[175,176]. At the same time, less stability in the hippocampal activity during rest might have been due to fewer reactivations - a possibility that could readily be tested using the current data set.

It is also important to note, that for the investigated experimental paradigm nothing new was learned (because the animal had learned both rules before the actual experiment). Still, the animal needed to associate the absence of a reward with the rule change.

Different regimes of instantaneous change

Although I did not detect sleep stages explicitly, I identified differences in oscillations of the LFP which are often used to identify sleep stages[222,224]. The higher Theta-Delta ratio (Fig. S3.3e) and the reduced ripple power (Fig. 3.3b) for regime 2 are strong indicators that this regime corresponds to REM sleep. Regime 1, on the other hand, showed a lower Theta-Delta ratio and a significantly higher ripple power. Therefore, regime 1 might majorly overlap with NREM sleep.

HPC-PFC synchrony

The two-stage model of system consolidation has established the idea that newly acquired information is initially stored in the hippocampus and is then transferred to other cortical areas during subsequent rest or sleep[18,161]. Although SWRs-spindle coupling has been identified as a primary mechanism for this transfer of information, recent studies have considered the possibility of slower frequency oscillation contribution[234,241]. The results presented in this chapter revealed differences in the coupling of hippocampal and cortical activity in an ultra-slow frequency range with a dependency on prior behavior. The finding that cells either up-regulated their activity during SWRs or low-variance periods, might indicate that these two communication channels are also transmitting different types of information or serve different functions. What remains to be investigated is whether cells that are modulated during SWS (down-states correspond to time intervals < 1s)[241] are also the cells that increase their firing during low-variance periods (time intervals 2-20s).

low-variance period co-activity versus SWRs co-activity

I used three clusters to separate HPC and PFC cell subsets that tended to either increase their firing during low-variance periods or SWRs. This was an arbitrary choice and might have impacted the other analyses using these clusters. A more insightful approach would be to first identify the optimal number of clusters needed to separate cells based on their low-variance period or SWR modulation. One possible criterion could be the angle between the mean predictive axis of clusters.

The sub-optimal separation of cells into clusters using the neural activity during rest might have diminished differences between these clusters in terms of spatial coding and rule coding during the actual behavior. A more principled analysis of the different cell subsets could potentially reveal functional differences of cells that were up-regulated during low-variance periods more clearly.

Nevertheless, I found that hippocampal cells which were not as strongly down-regulated during low-variance periods exhibited less spatial information and coded for the rule during behavior. Therefore, these cells might be involved in spatial coding to a lesser extent and might have served a different function. On the other hand, cells that showed an intermediate modulation during low-variance periods and a strong up-regulation during SWRs tended to have higher mean firing rates during behavior, more spatial information, were strongly coding for the rule, and exhibited higher rate map stability (Fig. 3.6). This implies that these cells played a prominent role in spatial and rule coding during the rule switch task[11].

Assuming that the co-activity of certain cell subsets from HPC and PFC implies some form of synchronization between the two areas that might be related to inter-area communication[242,243]. Along these lines, the directionality of this communication is a crucial feature. For the consolidation process, communication could be happening in both directions: information transfer from the cortex has been linked to the transformation of working memory into long-term memories[244]. At the same time, it has been shown that memory traces which are initially located in the hippocampus are then transferred to more cortical areas[230]. In this study, I did not assess the directionality but tried to predict the activity of one population using the activity of the other using a simple regression model. Since I used relatively large time bins (100 ms) and did not consider the influence of activity not happening simultaneously, I cannot make any statements about which population activity might have influenced the activity of the other. However, answering the question about directionality might be crucial to also identify the function of HPC-PFC subsets which tend to be active during low-variance periods, and the function of HPC-PFC subsets active during SWRs.

low-variance periods and population activity

low-variance periods majorly contributed to the observed synchrony of low-frequency fluctuations in the neural activity between HPC and PFC (Fig. S3.4d) which was increased after poor performance. Interestingly, I found that the actual number of low-variance periods showed the opposite effect: poor performance was followed by a reduction in the number of low-variance periods (Fig. S3.4b-c). This implies that not the total number of low-variance periods is the driver of synchrony, but rather the temporal alignment between low-variance periods in the HPC and the PFC.

Good performance was followed by an increase in low-variance periods and the cells that were up-regulated during these periods also tended to change their firing across the subsequent rest more (Fig. S3.8f-g). This observation could be the link between low-variance activity and the increased offline activity drift in the PFC as described in the first part of this chapter. Relating the two concepts of across-population synchrony and offline activity drift remains to be established by future research.

One intriguing hypothesis regarding the drift in neural activity and communication between areas is that activity that is relevant for across-area communication is more stable

(Fig. 3.7g). In this way, the effect of neural activity drift on the transfer to downstream areas would be minimized.

The concept of brain areas linking their activity using slow-global neural ensembles, whereas local ensembles are fast, has been brought forward before[245]. My findings support the idea that neural ensembles that occur during low-variance periods change slowly (Fig. 3.7) and are potentially involved in across-area communication. The finding that the activity of these neural ensembles is more constrained, might relate to the idea of a communication subspace[243], positing that only a subset of the activity dimensionality is actually relevant for cross-area communication[246]. Along these lines, fast-changing activity (e.g. non low-variance activity of regime 1) could correspond to more local activity. I found low-variance period activity to be more similar to the activity occurring in regime 2 (which likely corresponds to REM sleep). There is the possibility that the activity from those two time intervals, which do not occur in close temporal proximity, is functionally related. It remains to be explored how interleaved phases of REM sleep influence the changes in the activity during low-variance periods in regime 1 (potentially NREM sleep). There is evidence that REM and NREM sleep serve different functions ([214], Chapter 2), and establishing a link between the two could improve our understanding of the involved processes.

I attempted to uncover the relevance of low-variance period activity for behavior but only found that the low-variance period activity is significantly more similar to neural activity during rest than neural activity occurring during the rule switch task. Using one mean population vector to describe the low-variance period activity is a very crude approach, especially considering that low-variance period activity does change (Fig. 3.7e-f). I regard the link between identified low-variance activity patterns and the activity occurring during prior or subsequent behavior as a crucial step to understanding the relevance of low-variance periods. A more thorough decoding analysis, for example using hidden Markov models as used for the first chapter, could provide a better description of the relationship between low-variance activity and neural activity during behavior.

The presented results suggest that low-variance periods in neural activity during sleep might be involved in the synchronization of activity between the hippocampus and the prefrontal cortex. A phenomenon that is influenced by prior rule switch performance of the animal. Additionally, there is a discrepancy between involved cells, which seem to be also driving offline activity drift, and cells that are recruited during sharp wave ripples. Therefore, the consolidation of task-relevant knowledge in the cortical memory system might involve more processes apart from the coupling of SWRs and spindles. How these processes potentially trigger systematic changes in neural activity during rest and how these changes might influence neural activity during behavior remains to be explored.

3.5 Supplementary Figures

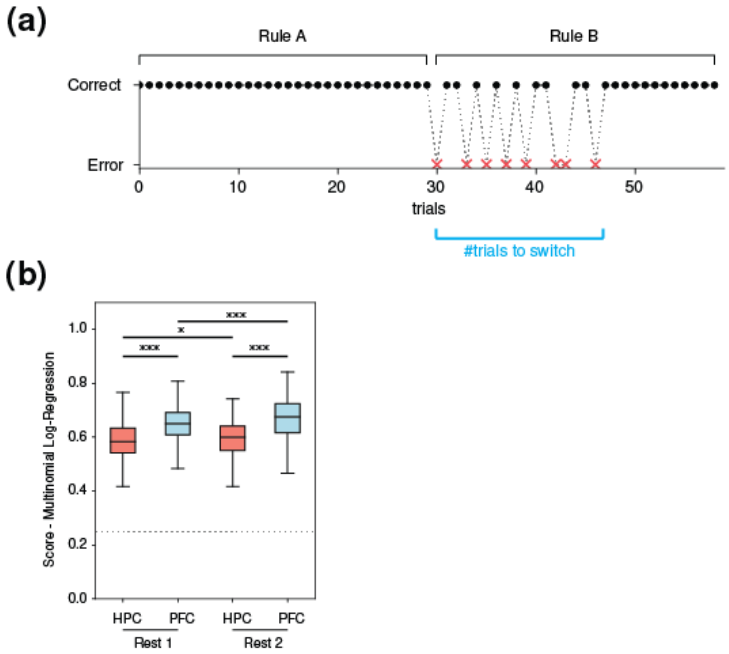

Figure S3.1: (a) Measuring behavioral performance during rule switch paradigm. (b) Offline activity drift measured using multinomial logistic regression (* $p < 0.05$, *** $p < 0.001$, Mann-Whitney test, Bonferroni-correction).

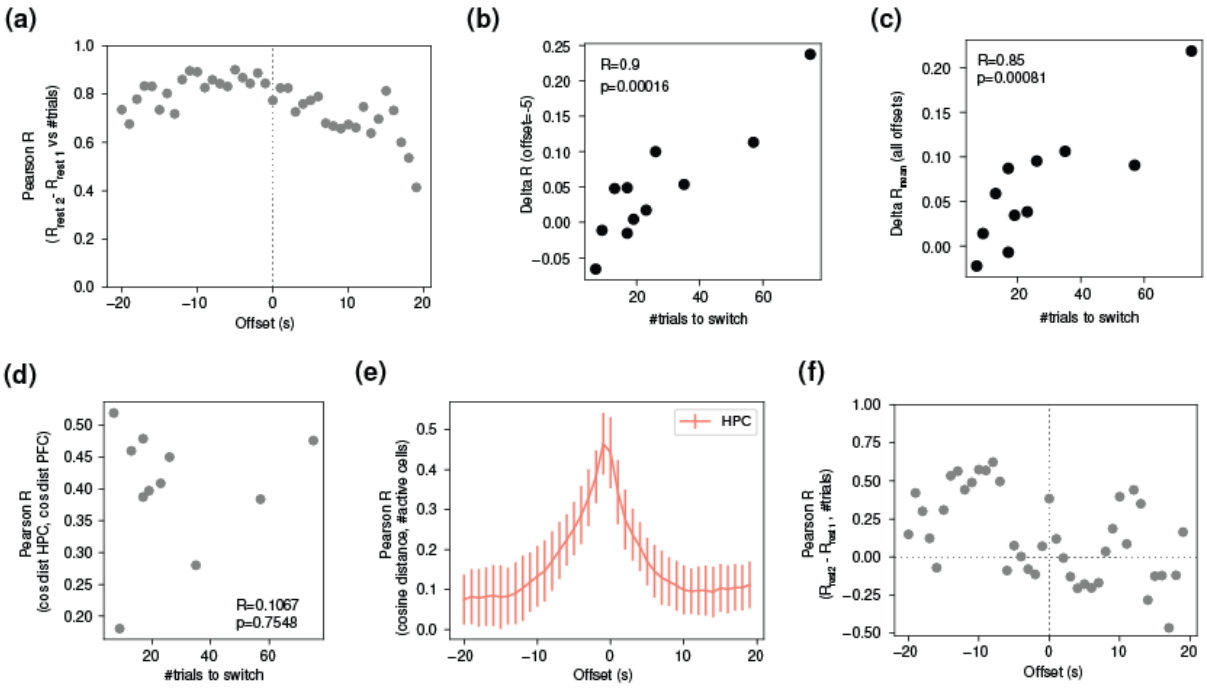

Figure S3.2: (a) Correlation between the difference in cross-correlation (short-term fluctuations) from rest 1 and rest 2 using the entire rest duration for different offsets and the number of trials the animal took for the rule-switch. (b) Number of trials the animals took to switch the rule versus the difference in correlation (short-term fluctuations) from rest 1 to rest 2 using the entire rest for an offset of -5s (HPC leading). (c) Correlation between the difference of the mean correlation (short-term fluctuations HPC and PFC, all offsets from -20 to 20 seconds) from rest 1 to rest 2 using the entire rest and the number of trials the animal took for the rule switch. (d) Synchrony of short-term fluctuations (cosine distance) between HPC and PFC during the rule switch task (offset = 0s) versus the number of trials the animal took for the rule switch. (e) Cross-correlation between the cosine distance and the number of active cells per time bin for the hippocampus (all sessions). (f) Correlation between the difference in cross-correlation (number of active cells per time bin) between rest 1 and rest 2 and the number of trials for the rule switch.

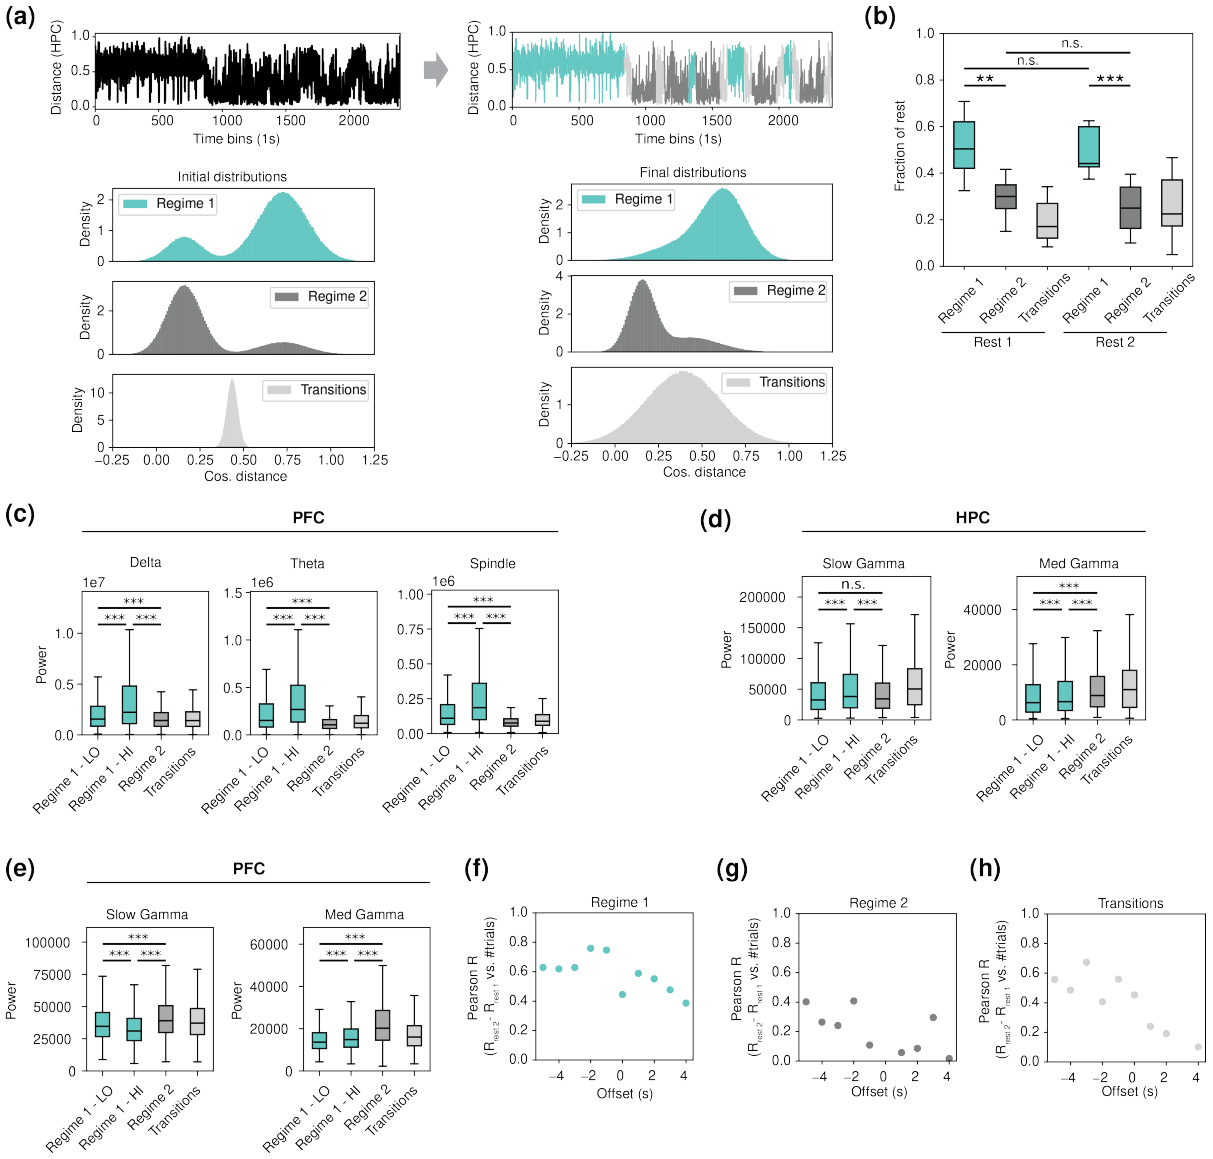

Figure S3.3: (a) Clustering algorithm to distinguish different regimes of short-term changes in neural activity. Left: unclustered time trace of short-term changes (cosine distance between subsequent population vectors) in the HPC and initial distributions of cosine distance for regime 1, regime 2 and transitions. Right: clustered time trace and distributions of cosine distance for the different regimes after clustering. (b) Duration of regime 1, regime 2 and transition as a fraction of the entire rest 1 and rest 2. (c) Delta, Theta and ripple / spindle power in the PFC for different regimes (data from all sessions, *** $p < 0.001$, Mann-Whitney U with Bonferroni-correction). (d)-(e) Power of slow and medium gamma for different identified regimes in the hippocampus (d) and the prefrontal cortex (e). (f)-(h) Correlation between the difference in cross-correlation from rest 1 and rest 2 for different offsets and the number of trials the animal took for the rule-switch using different regimes.

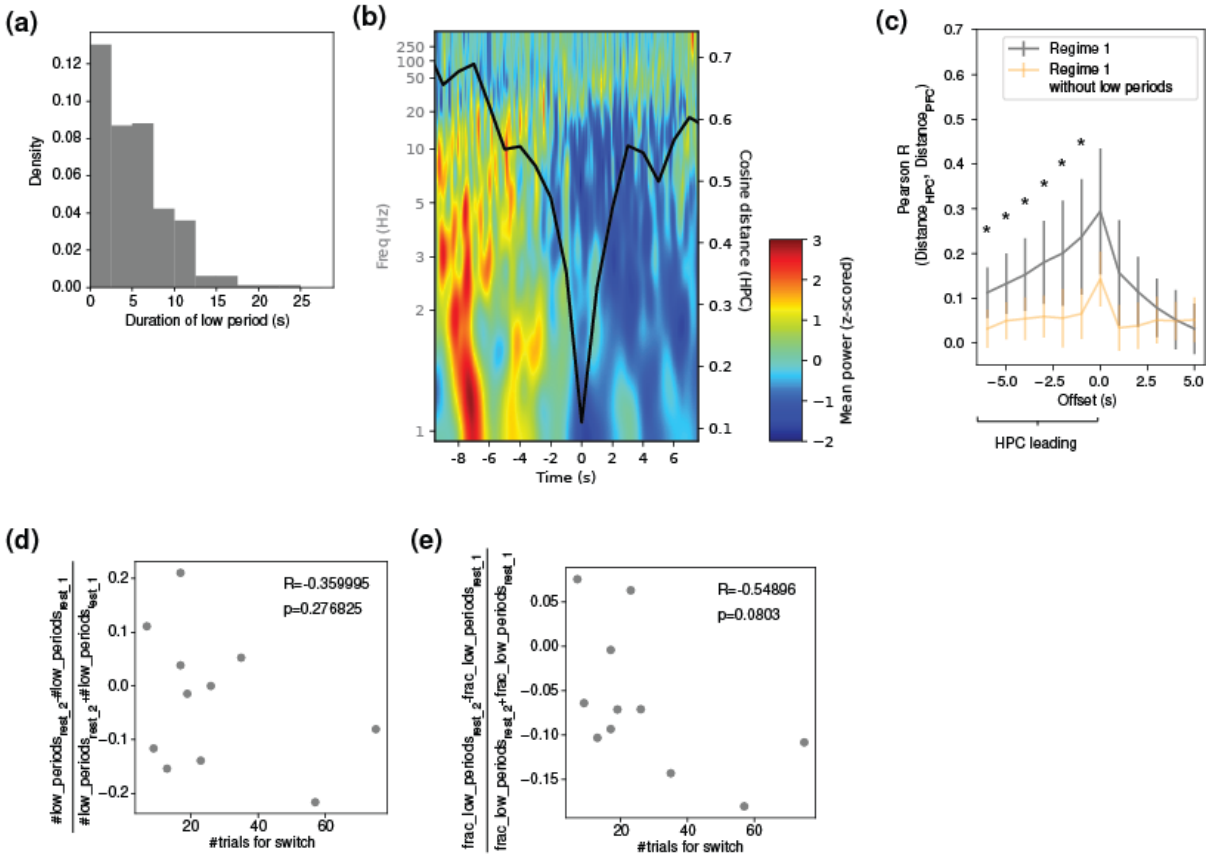

Figure S3.4: (a) Duration of low-variance periods from all sessions. (b) LFP power around low-variance periods for one example session. (c) Influence of low-variance periods on the synchrony of short-term change between HPC and PFC for regime 1 (result from all sessions, * $p < 0.05$, T-test, Bonferroni-correction for multiple comparisons). (d) Difference in number of low-variance periods in the hippocampus between rest 2 and rest 1 versus number of trials the animal needed to switch the rule. (e) Difference in fraction of low-variance periods in regime 1 between rest 2 and rest 1 in the hippocampus versus number of trials for the rule switch.

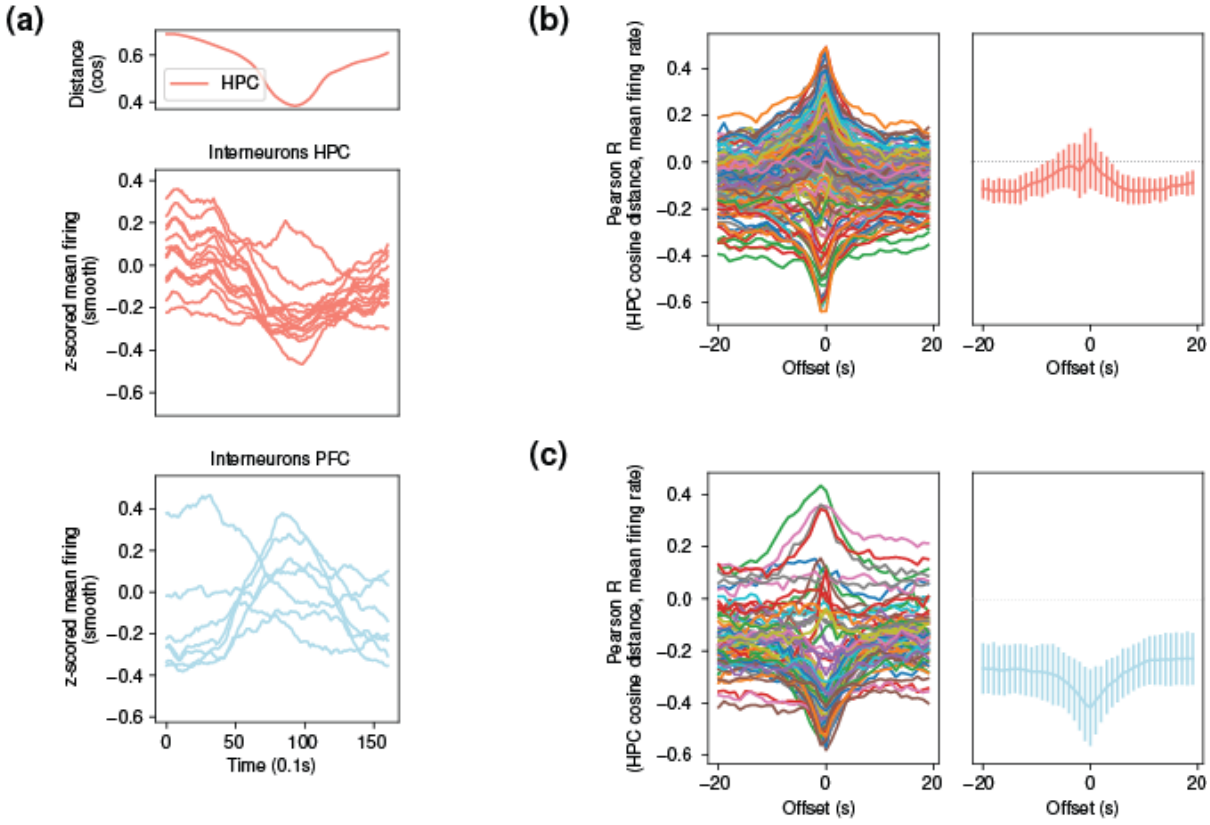

Figure S3.5: (a) Firing rates of hippocampal and cortical interneurons around low-variance periods. (b)-(c) Cross-correlogram of hippocampal (b) and cortical (c) interneuron firing and short-term fluctuations (cosine distance between subsequent population vectors of hippocampal pyramidal cell activity).

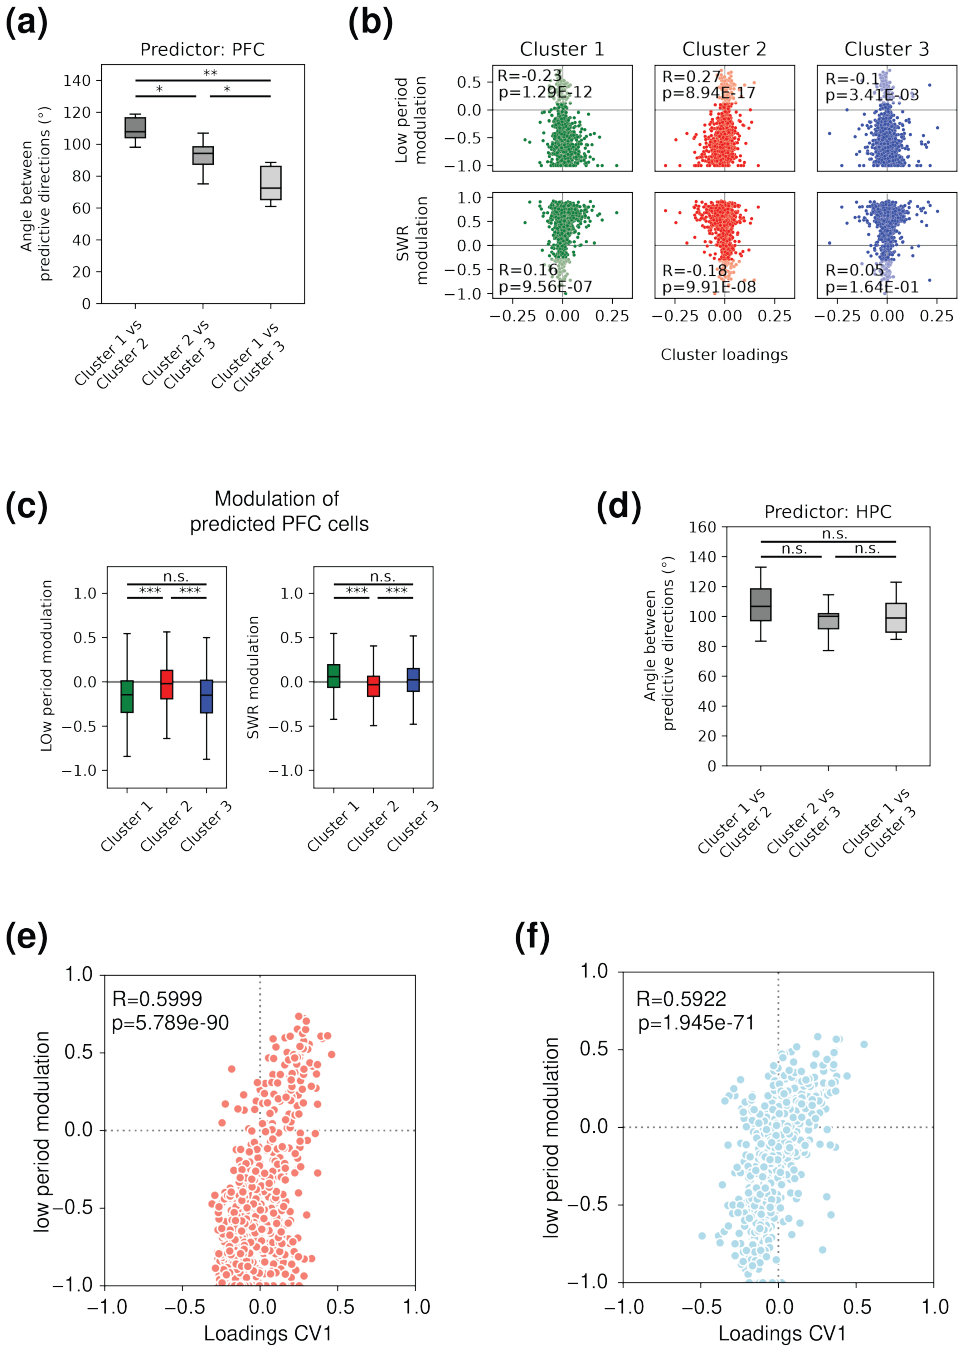

Figure S3.6: (a) Angles between mean predictive directions for the different clusters using PFC neurons' firing rates as regressors (* $p < 0.05$, ** $p < 0.01$, MWU, Bonferroni-correction). (b)-(d) Communication channel results for all sessions using HPC neuronal activity to predict PFC neuronal activity. (b) Firing rate modulations of HPC neurons (regressors) during low-variance periods and SWRs with the corresponding cluster loadings for the three identified clusters. (c) Firing rate modulations of predicted PFC neurons during low-variance periods and SWRs for the three clusters. (d) Angles between mean predictive directions of the different clusters using HPC neuronal activity as regressors (n.s. $p > 0.05$, MWU, Bonferroni-correction).(e)-(f) Correlation between low-variance period modulation and canonical variate 1 (CV1) for HPC (e) and PFC (f).

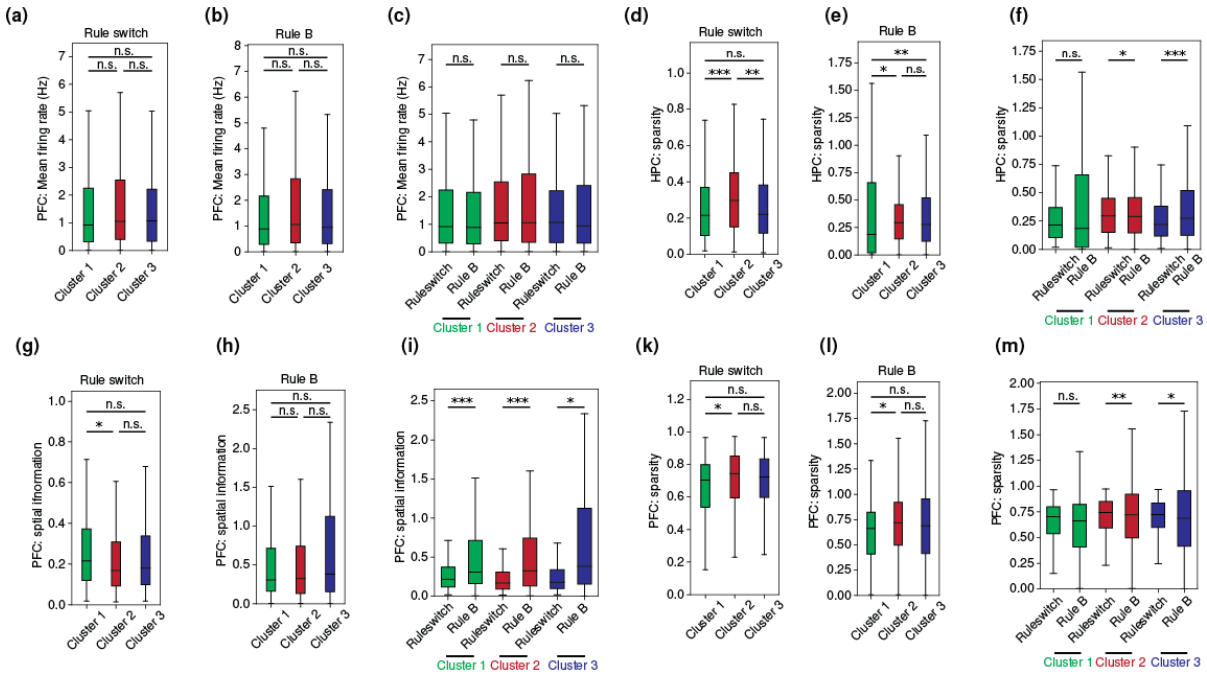

Figure S3.7: (a)-(b) Mean firing rates of prefrontal cells from different clusters (see previous section, Cluster 1 & 3: SWRs favoring cells, Cluster 2: low-variance periods favoring cells) during the rule switch task (a) and behavioral episode after rest (rule B, panel b). (c) Comparison of mean firing rates between the rule switch task and the behavioral episode after rest for prefrontal cells. (d)-(f) Sparsity of hippocampal cells during the rule switch task (d) and the behavioral episode after rest (e) for the three different clusters. (g)-(i) Spatial information of cortical cells during the rule switch task (d) and the behavioral episode after rest (e) for the three different clusters. (k)-(m) Sparsity of cortical cells during the rule switch task (d) and the behavioral episode after rest (e) for the three different clusters. For all panels: *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, n.s. $p > 0.05$, Mann-Whitney U Test with Bonferroni-correction.

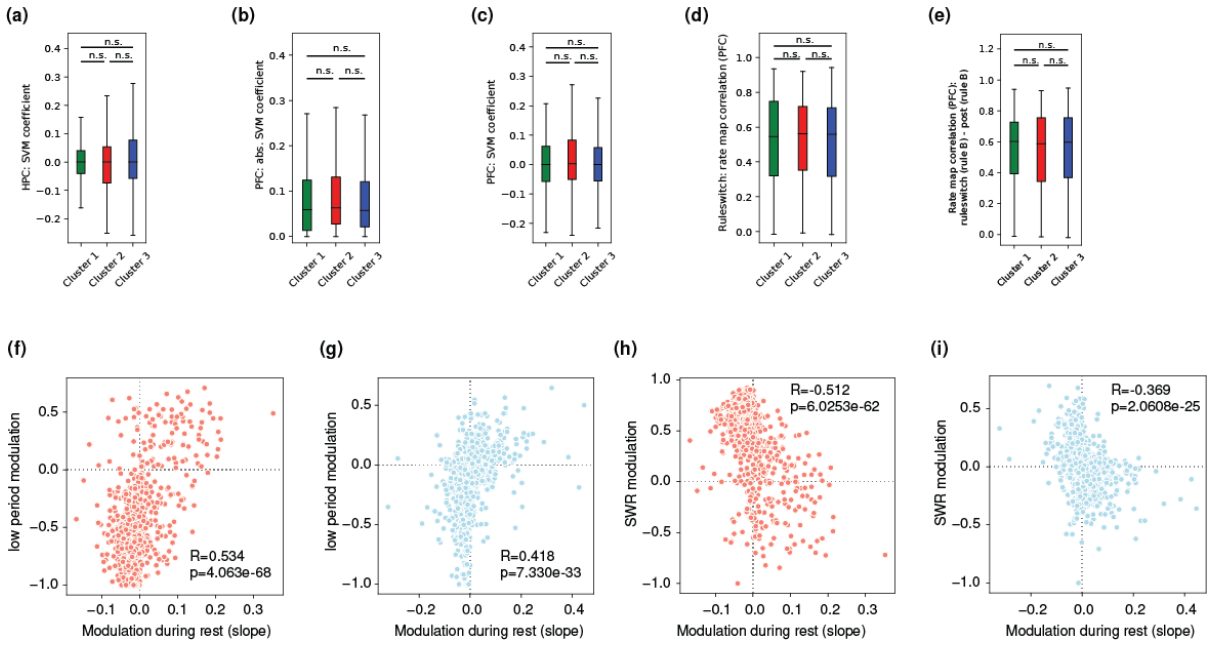

Figure S3.8: (a) Coefficients of the support vector machine (SVM) to classify rule A and rule B activity during the rule switch for hippocampal cells. (b)-(c) Magnitude of coefficients (b) and coefficients (c) of SVM to classify rule A and rule B activity during rule switch for cortical neurons. (d) Rate map correlations between rule A and rule B during the rule switch task for cortical cells.(e) Rate map correlations between rule B of the rule switch task and rule B during the recall episode after rest for cortical cells. (f)-(i) low-variance period modulation versus overall modulation (slope of fitted linear curve to firing rates throughout rest) for HPC (a) and PFC (b). Data from all sessions. For all panels: *** $p < 0.001$, ** $p < 0.01$,* $p < 0.05$,n.s. $p > 0.05$, Mann-Whitney U Test with Bonferroni-correction.

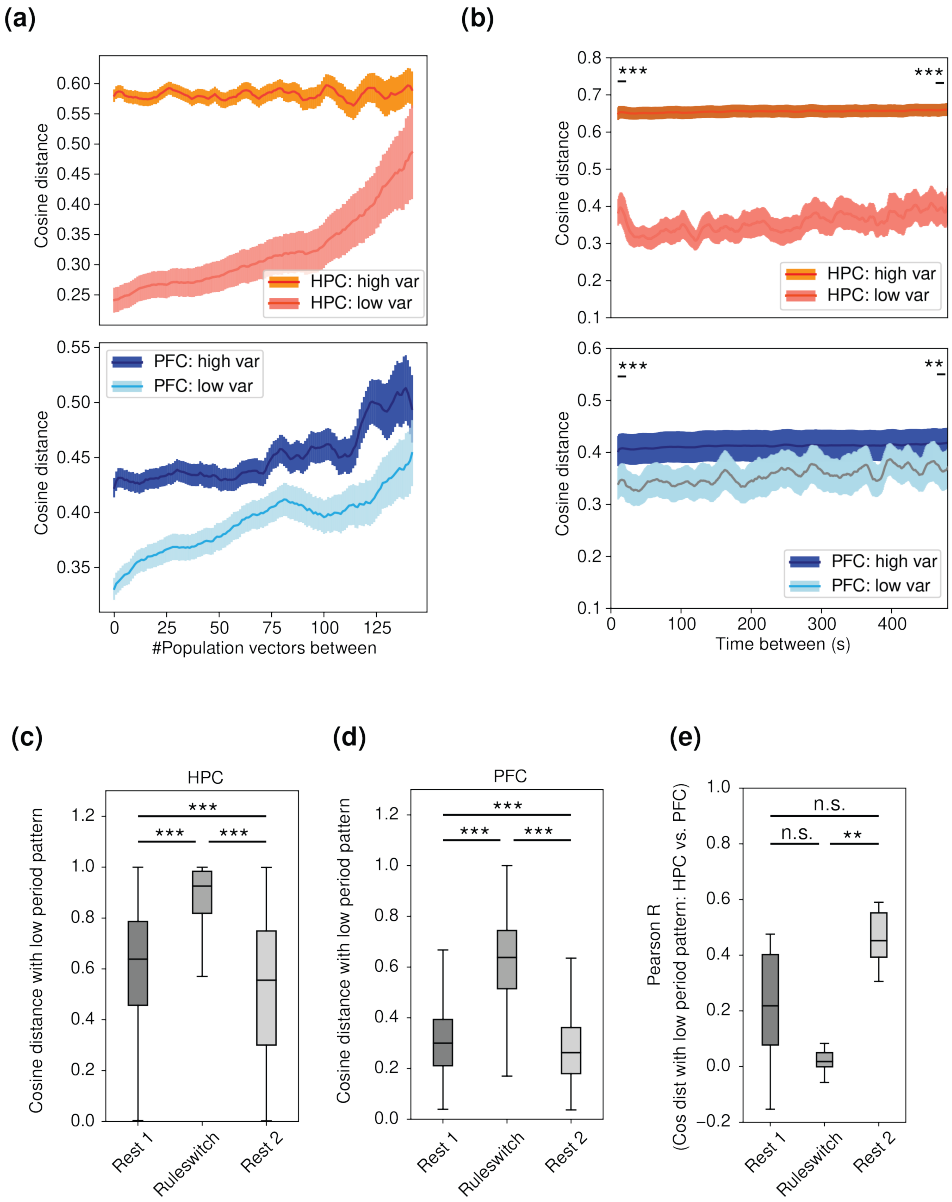

Figure S3.9: (a) Cosine distances between population vectors of low-variance periods and between vectors of high variance periods as a function of the number of population vectors in between for one example session. Top: hippocampal activity. Bottom: prefrontal cortical activity. (b) Cosine distance between population vectors as a function of the temporal gap between them for all sessions. Top: hippocampal activity. Bottom: prefrontal cortical activity. Statistics were done on the first 10 and last 10 values from all sessions, comparing low and high variance state cosine distances. (c)-(d) Cosine distance between activity from rest 1, rule switch, rest 2 and the mean activity during low-variance periods for hippocampal (b) and prefrontal (c) cells. The awake neural activity (rule switch) was binned using 100ms windows, whereas for the rest activity (rest 1 and rest 2) a temporal window of 1s was used to compute the population vectors. (d) Co-occurrence of low-variance period-like activity in the hippocampus and prefrontal cortex for different experimental episodes: across-area correlation of the cosine distance time series between the neural activity and the mean low-variance activity. *** $p < 0.001$, ** $p < 0.01$, n.s. $p > 0.05$ Mann-Whitney U with Bonferroni correction.

3.6 Methods

Data & experimental setup

For details on the surgery, training and data acquisition see[11]. In short, four Long-Evans rats were implanted with 32-tetrode microdrives targeting the right dorsal hippocampus (CA1 sub-region, HPC) and left medial prefrontal cortex (prelimbic area, PFC). An average of 63 PFC and 78 HPC principal cells were recorded per session.

Animals were placed in one of two start arms (north/south arm) and learned to retrieve food in one of two goal arms (east/west arm) based on a spatial or a light-response rule. For the spatial rule, the animal only received a reward in either the east or west arm, irrespective of the starting arm. During the light-response rule, the animal was only rewarded in the light-on arm. Importantly, the light in either the east or west arm was also switched on (not necessarily indicating the rewarded arm) while the animal was supposed to follow the spatial rule. Therefore, for a successful execution of the task using the spatial rule, animals had to ignore the conflicting light stimulus. On each day, the animal first had to recall the rule of the previous day (rule A) and was allowed to rest (rest 1) before performing the rule switch task (rule A $\rightarrow$ rule B). During the rule switch task, the animal needed to first recall and apply rule A until reaching a performance criterion (12/15 successful trials for the spatial rule, 24/30 successful trials for the light-response rule). Then, the rule was switched unannounced by the experimenter. The animal had to figure out the change in rule based on whether it received a reward and therefore navigated to the correct goal arm. The rule switch task was followed by a rest of 40 minutes (rest 2) before the recall of rule B was tested (rule B).

Measuring behavioral performance

In order to assess the behavioral performance of the animal during the rule switch task, I counted the number of trials it took the animal after the unannounced rule switch to perform at least five successive, successful trials using the new rule (rule B). See also Fig. S3.1.

Regressing time during rest

I used the ability to regress time during rest using an ordinary least square model (scikit learn) based on the neural activity as a proxy for offline neural activity drift. First, I binned the neural activity using temporal bins of 100ms duration. Next, I fitted the model to 60% of the data and validated its performance on the remaining 20% of the data. The resulting coefficient of determination for the test data was then used to characterize the amount of the offline neural activity drift.

Multinomial logistic regression

As a control, I used a multinomial logistic regression to confirm that there are systematic changes in the neural activity during sleep. I binned the sleep neural activity using 100ms temporal bins and split the data into 4 chunks of equal length. All the bins within one chunk were assigned the same ID and a multinomial logistic regression trained to predict the ID based on the activity of the temporal bin. I then tested the model on held out data (10% of the total data) to obtain the mean accuracy of the model.

Computing synchrony between short-term fluctuations between PFC and HPC

To compute short-term fluctuations in the neural activity of each brain area, I computed the cosine distance between subsequent population vectors (temporal binning with 1s bins). This resulted in two time traces, one for each population. Then, I calculated the cross-correlation of these two time traces to assess their temporal alignment.

Frequency filtering of short-term fluctuations

To analyze the influence of different frequency components on the synchrony between short-term fluctuations of both areas, I first filtered the time traces from each population using a highpass or bandpass Butterworth filter. Then, the filtered time traces were cross-correlated.

Clustering regimes of short-term fluctuations

I obtained regimes with similar distributions of short-term fluctuations using the following analysis. By visual inspection, I identified one regime which was dominated by high cosine values with interspersed low cosine value intervals (regime 1). Another regime showed mostly low cosine values with some high cosine values (regime 2). The remaining data resembled an intermediate distribution of cosine values (transition). Therefore, I decided to represent the cosine distance distributions of regime 1 and regime 2 using a mixture of two Gaussians each. On the other hand, the distribution of cosine distances of the transition regime was captured by one Gaussian only. I defined the means of the initial Gaussians using the 20th, 50th and 90th percentile of all cosine distance values. For regime 1, the weights were set to 0.2 (variance: 0.01) for the Gaussian with the smaller mean (corresponding to the 20th percentile) and to 0.8 (variance: 0.02) for the Gaussian with the larger mean (corresponding to the 90th percentile). For regime 2 the values were assigned in reversed order corresponding to predominantly small cosine distance values. For the initial Gaussian for the transitions, I set the mean the 50th percentile and the variance to 0.001. A window of 60s duration (without overlap) was then moved along the time trace of cosine distances. For each window the likelihood of the cosine distance distribution given the three models (regime 1, regime 2, transitions) was computed and all cosine distance values of the window assigned to the model with the maximum likelihood. After processing all windows, the models were updated using the assigned data. These steps were repeated until model assignments per window did not change across iterations resulting in one model assignment per window (regime 1, regime 2, transitions). Finally, I visually verified the correct assignment of labels and re-assigned windows if there was a mismatch.

LFP analysis

To obtain the power in different frequency bands, I applied a fourier transform (spectrogram function from scipy.signal package) to the local field potential (LFP) data from one tetrode in the hippocampus and one tetrode in the prefrontal cortex. The following bands were used:

- delta : 1-4 Hz

- theta : 8 - 12 Hz
- spindle : 10 - 16 Hz
- low gamma : 25 - 40 Hz
- med gamma : 60 - 80 Hz
- ripple : 150-250Hz

Detecting low-variance periods

To identify low-variance periods, I first selected only cosine distance values from regime 1 using hippocampal activity if not indicated otherwise. The obtained time trace was then z-scored. Time points with values less than 2 standard deviations below the mean were selected and extended by 1s in both directions to yield intervals with low-variance in the population activity.

Computing firing rate modulation during low-variance periods and SWRs

Differences in firing rates of cells between low-variance periods and the other activity during regime 1 was calculated the following way:

$$modulation_{low_{p}eriod} = \frac{m_{withinLow} - m_{outsideLow}}{m_{withinLow} + m_{outsideLow}} \tag{3.1}$$

where m is the the mean firing rate during the corresponding time interval.

To compute the modulation during SWRs, low-variance periods were excluded to calculate the baseline activity firing meanoutside SWR

$$modulation_{SWR} = \frac{m_{withinSWR} - m_{outsideSWR}}{m_{withinSWR} + m_{outside_{S}WR}} \tag{3.2}$$

Prediction and hierarchical clustering analysis

The neural activity during rest was binned using 100ms temporal windows for each population separately. Then the activity of each neuron from one population was regressed using the activity of all the cells from the other population. The resulting regression weights for each predictive dimension were normalized. The cosine distance between all obtained regression dimensions (one per predicted cell activity) was computed and used to assign predictive dimensions to one of three clusters (using Agglomerative clustering from sklearn). Per cluster, the mean predictive dimension was then calculated using all predictive dimensions belonging to that particular cluster.

Applying canonical correlation analysis to rest data

I applied a standard canonical correlation analysis (CCA, sklearn package) to the time binned rest data (1s time bins) to identify each cell's contribution to the correlation between the neural activity from HPC and PFC.

Measuring spatial information

To assess the spatial information of single cells, I computed the sparsity and spatial information per second as previously described[226]. The sparsity for X spatial bins is defined as follows:

$$s = \frac{(\sum_{i=1}^{X} p_i \lambda_i)^2}{\sum_{i=1}^{X} p_i \lambda_i^2} \tag{3.3}$$

with pi being the probability of being in spatial bin i and λ_i being the mean firing rate in that bin. The spatial information per second was computed using the following equation:

$$I_{sec} = \sum_{i=1}^{X} p_i \lambda log(\frac{\lambda_i}{\lambda}) \tag{3.4}$$

where pi and λ_i are the probability of occupying and the firing rate of bin i, respectively. Parameter λ describes the mean firing rate of the cell in the environment.

Rule decoding

The neural activity from either HPC or PFC during the rule switch task was binned using 100ms temporal bins. Bins which occurred before the experimenter switched the rule were assigned one label (rule A), bins occurring after the animal successfully switched its behavior were assigned another label (rule B). A support vector machine (sklearn package) was trained to predict the label of each temporal bin and tested on held-out data. The coefficient for each cell of the trained model was used as a proxy for its contribution to rule coding.

Rate map correlations

To assess the similarity of per cell rate maps between different episodes, I first computed the number of spikes occurring within each spatial bin (spatial bin size = 5 cm) and divided the resulting matrix by the occupancy matrix. Gaussian smoothing was applied to acquire one rate map per cell and episode. The similarity between the rate maps of the same cell between different episodes was estimated using the Pearson correlation between the flattened rate maps.

Computing modulation during rest

The overall modulation of firing rates per cell throughout rest was computed in the following manner. First, the neural activity was binned using 1s temporal bins. Per cell the firing rate data was normalized, smoothed and a line was fit to the resulting data. The coefficient describing the slope was used to describe the overall modulation of the cell's firing rate throughout rest.

Cosine distance as a function of number of population vectors

The neural activity was binned using 1s temporal bins. Afterwards, bins were assigned to either low-variance periods (see Detecting low-variance periods) or high variance periods

(remaining activity of regime 1). The temporal order of bins was maintained. Then, the cosine distance between all pairs of bins belonging to one group (low or high variance) was computed yielding a varying number of cosine distances per gap between the population vectors of each pair. The mean per gap (number of population vectors in between) was then calculated using all corresponding cosine distances.

Correlation of similarity with low-variance period neural patterns

To start with, a mean population vector across all low-variance activity temporal bins (1s, rest 2) was calculated and used as a template. The similarity of activity with this template was evaluated using the cosine distance. For rest 1, the neural activity was binned using 1s time bins and for the neural activity during the rule switch temporal bins of 100ms were used.

CHAPTER 4

Conclusion & outlook

My analyses revealed that neural activity during consolidation periods undergoes distinct changes. In the case of a long period of rest in between exposures to the same configuration of goal locations, the evolution of the neural activity resembled changes observed from the acquisition to the recall of the spatial memory. Consequently, cells that increased or decreased their activity from the acquisition to the recall, showed the same trends during the consolidation period. This phenomenon is potentially related to the stabilization and integration of newly acquired information into existing knowledge[161,174]. Furthermore, my findings challenge the view that changes in representations across exposures are purely due to the accumulation of experiences or adaptations during behavior[114,115]. In fact, the predictive power of reactivated activity for future neural activity was recently shown in the visual cortex as well[235]. I observed that the activity at the end of the consolidation period was similar but not identical to the subsequently observed neural activity during recall. Therefore, it seems likely that the circuit is primed to the new pattern configuration during rest, and that novel experience only causes additional modifications on top of those already emerged during sleep[67–70].

The question remains if the observed continuous transformation of neural activity during consolidation periods was actually related to memory processes or indeed was the effect of physiological adjustments intrinsic to the passage of time. In this regard, it would be useful to run control experiments with repeated exposure to a familiar environment. The analysis of interleaved long rest periods could confirm or reject the active role of memory processes in driving changes in neural activity and provide a direct link to representational drift. Along these lines, it could be informative to investigate how the reactivated activity of two distinct experiences (e.g. from different environments or contexts) evolves throughout the consolidation period. Are spontaneous changes in the corresponding neural representations related or do we observe that the representations become increasingly orthogonal[76,210]?

Another exciting endeavor might be the identification of differences or commonalities between representational changes due to learning, consolidation periods, or simply the passage of time[107]. The analysis of time scales and affected neurons could be a good starting point to evaluate if e.g. changes triggered by learning are simply continued in subsequent consolidation periods. Furthermore, one could ask how trial-to-trial variability is related to changes in neural representations. In this context, it is important to identify what kind of representation might be relevant for a given task. Studies investigating

representational drift in the hippocampus have been focused on changes in spatial coding, despite the observed coding for non-spatial dimensions[119]. Considering the coding for different features, it could be informative to see how continuous changes affect the different representations and whether task-relevant ones tend to be maintained[118].

To perform the corresponding experiments, reliable long-term neural recordings during behavior and offline periods for extended time intervals are necessary. So far, many studies investigating representational drift relied on calcium imaging with head-fixed animals and therefore were not able to collect neural data during consolidation periods. The use of alternative recording technologies such as neural probes, grin lenses, and sophisticated implants, might enable scientists to record the required neural data over several weeks[247]. Recent developments in recording technology make it possible to record from a large number of neurons from several brain regions in parallel[248]. These advancements, in conjunction with the ability to record from the same neurons over long durations of time, could lead to major discoveries related to fluid representations and their brain-wide distribution and processing. My findings from Chapter 3 indicate that changes in neural activity between HPC and PFC are synchronized and that certain brain states support inter-area communication. Using these findings as a starting point, one could ask if the functional coupling between both areas is maintained despite changes in neural activity. Alternatively, the inter-area functional connectivity could also change. This question could potentially be tackled by applying approaches related to a communication subspace[243].

Longer, simultaneous recordings of the HPC and the PFC would also allow relating changes in neural activity during rest to reactivations (as done in Chapter 2). In this way, the co-evolution of memory-related neural representations in multiple, functionally-connected areas could be investigated. Along these lines, the influence of SWRs, low variance periods, and REM sleep on changes in neural activity occurring in both areas in parallel could be analyzed further.

As a starting point, my analysis from Chapter 2 and Chapter 3 uncovered a range of phenomena that either promoted changes in neural activity during consolidation periods or had a stabilizing effect. I showed that SWRs occurring during NREM sleep caused the neural activity to change from acquisition-like to recall-like, whereas REM periods showed a more stable acquisition-like activity. Interestingly, the analysis from Chapter 3 revealed that there were periods of stable activity during NREM sleep outside of SWRs. The activity during these low-variance periods strongly resembled REM-like neural activity and might therefore provide a functional link between NREM and REM sleep[198,10].

The differential contribution of cell subsets to the stability and changes in neural activity provides another exciting direction for future studies. In this context, I showed in Chapter 2 that changes in neural activity were driven by a plastic subset of cells, whereas a persistent subset of cells maintained their activity from the acquisition to the recall of a spatial memory. On the other hand, the analysis of Chapter 3 demonstrated that different cells were either recruited during (change-promoting) SWRs or (stable activity) low-variance periods. The next step would be to investigate the overlap between the described subsets to get a coherent picture of involved cells and their relevance for memory-related behavior.

Given the future existence of multi-area, long-term recordings, appropriate analysis approaches to answer the aforementioned questions need to be developed as well. This might include latent-state models[249], time-series analysis for high-dimensional data to assess its non-stationary characteristics (e.g. dynamical systems), and models that

incorporate temporal dynamics. Since representational drift and continuous changes in neural activity have been identified in other brain areas as well[124,126,127], these approaches have the potential to be extensively applied across a large number of datasets. In this context, the pHMM analysis from Chapter 2 might be one solution. Since the model does not assume the coding for a specific feature or dimension, it could be easily extended to brain areas beyond the hippocampus.

As a last thought, theoretical advances, clarifying if fluid representations are advantageous for neural computations or just a constraint of the biological system will help to inspire future experiments[107,123]. Investigating information encoding, storage, integration and transfer from the perspective of stability and change in neural representations will certainly broaden our understanding of how the brain performs these fundamental operations.

Bibliography

[1] Daniel L Schacter and Donna Rose Addis. The cognitive neuroscience of constructive memory: Remembering the past and imagining the future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481):773–786, May 2007. ISSN 0962-8436. doi: 10.1098/rstb.2007.2087.

[2] Dorothy Tse, Rosamund F. Langston, Masaki Kakeyama, Ingrid Bethus, Patrick A. Spooner, Emma R. Wood, Menno P. Witter, and Richard G. M. Morris. Schemas and Memory Consolidation. *Science*, 316(5821):76–82, April 2007. doi: 10.1126/science.1135935.

[3] S. Maren, G. Aharonov, and M.S. Fanselow. Neurotoxic lesions of the dorsal hippocampus and Pavlovian fear conditioning in rats. *Behav.Brain Res.*, 88(0166-4328):261–274, November 1997.

[4] Matthew Schafer and Daniela Schiller. Navigating Social Space. *Neuron*, 100(2): 476–489, October 2018. ISSN 08966273. doi: 10.1016/j.neuron.2018.10.006.

[5] D.F. Sherry and D.L. Schacter. The Evolution of Multiple Memory Systems. *Psychological Review*, 94(4):439–454, 1987.

[6] György Buzsáki and Edvard I. Moser. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16(2):130–138, February 2013. ISSN 1546-1726. doi: 10.1038/nn.3304.

[7] Praveen K. Pilly and Stephen Grossberg. How Do Spatial Learning and Memory Occur in the Brain? Coordinated Learning of Entorhinal Grid Cells and Hippocampal Place Cells. *Journal of Cognitive Neuroscience*, 24(5):1031–1054, 2012. ISSN 0898-929X. doi: 10.1162/jocn_a_00200.

[8] Yael Shrager, Peter J. Bayley, Bruno Bontempi, Ramona O. Hopkins, and Larry R. Squire. Spatial memory and the human hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*, 104(8):2961–2966, February 2007. ISSN 0027-8424. doi: 10.1073/pnas.0611233104.

[9] Susanne Diekelmann and Jan Born. The memory function of sleep. *Nature Reviews Neuroscience*, 11(2):114–126, February 2010. ISSN 1471-0048. doi: 10.1038/nrn2762.

[10] Susan J. Sara. Sleep to Remember. *Journal of Neuroscience*, 37(3):457–463, January 2017. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0297-16.2017.

[11] Karola Kaefer, Michele Nardin, Karel Blahna, and Jozsef Csicsvari. Replay of Behavioral Sequences in the Medial Prefrontal Cortex during Rule Switching. *Neuron*, 106(1):154–165.e6, April 2020. ISSN 0896-6273. doi: 10.1016/j.neuron.2020.01.015.

[12] David R. Euston, Aaron J. Gruber, and Bruce L. McNaughton. The Role of Medial Prefrontal Cortex in Memory and Decision Making. *Neuron*, 76(6):1057–1070, December 2012. ISSN 0896-6273. doi: 10.1016/j.neuron.2012.12.002.

[13] N. S. Narayanan, N. K. Horst, and Mark Laubach. Reversible inactivations of rat medial prefrontal cortex impair the ability to wait for a stimulus. *Neuroscience*, 139 (3):865–876, January 2006. ISSN 0306-4522, 1873-7544. doi: 10.1016/j.neuroscience.2005.11.072.

[14] Sophie Tronel and Susan J. Sara. Blockade of NMDA Receptors in Prelimbic Cortex Induces an Enduring Amnesia for Odor–Reward Associative Learning. *Journal of Neuroscience*, 23(13):5472–5476, July 2003. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.23-13-05472.2003.

[15] James L. McClelland, Bruce L. McNaughton, and Randall C. O'Reilly. Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102(3):419–457, 1995. ISSN 1939-1471. doi: 10.1037/0033-295X.102.3.419.

[16] William Beecher Scoville and Brenda Milner. LOSS OF RECENT MEMORY AFTER BILATERAL HIPPOCAMPAL LESIONS. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20(1):11–21, February 1957. ISSN 0022-3050.

[17] L.R. Squire and C.B. Cave. The hippocampus, memory, and space. *Hippocampus*, 1 (1050-9631 (Print)):269–271, July 1991.

[18] G. Buzsáki. Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, 31(3):551–570, January 1989. ISSN 0306-4522. doi: 10.1016/0306-4522(89)90423-5.

[19] M.P. Witter. Organization of the entorhinal-hippocampal system: A review of current anatomical data. *Hippocampus*, 3 Spec No(1050-9631 (Print)):33–44, 1993.

[20] H. Eichenbaum, C. Stewart, and R.G. Morris. Hippocampal representation in place learning. *J Neurosci*, 10(0270-6474 (Print)):3531–3542, November 1990.

[21] H. Eichenbaum. Hippocampus: Cognitive processes and neural representations that underlie declarative memory. *Neuron*, 44(0896-6273):109–120, September 2004.

[22] Matthew L. Shapiro and Howard Eichenbaum. Hippocampus as a memory map: Synaptic plasticity and memory encoding by hippocampal neurons. *Hippocampus*, 9 (4):365–384, 1999. ISSN 1098-1063. doi: 10.1002/(SICI)1098-1063(1999)9:4<365::AID-HIPO4>3.0.CO;2-T.

[23] Thomas J. McHugh, Matthew W. Jones, Jennifer J. Quinn, Nina Balthasar, Roberto Coppari, Joel K. Elmquist, Bradford B. Lowell, Michael S. Fanselow, Matthew A. Wilson, and Susumu Tonegawa. Dentate Gyrus NMDA Receptors Mediate Rapid Pattern Separation in the Hippocampal Network. *Science*, 317(5834):94–99, July 2007. doi: 10.1126/science.1140263.

[24] Michael A. Yassa and Craig E. L. Stark. Pattern separation in the hippocampus. *Trends in neurosciences*, 34(10):515–525, October 2011. ISSN 0166-2236. doi: 10.1016/j.tins.2011.06.006.

[25] M.P. Witter, F.G. Wouterlood, P.A. Naber, and T. van Haeften. Anatomical organization of the parahippocampal-hippocampal network. *Ann.N.Y.Acad.Sci.*, 911(0077-8923 (Print)):1–24, June 2000.

[26] D. Marr, David Willshaw, and Bruce McNaughton. Simple Memory: A Theory for Archicortex. In Lucia Vaina, editor, *From the Retina to the Neocortex: Selected Papers of David Marr*, pages 59–128. Birkhäuser, Boston, MA, 1991. ISBN 978-1-4684-6775-8. doi: 10.1007/978-1-4684-6775-8_5.

[27] Lee A. Cenquizca and Larry W. Swanson. Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Research Reviews*, 56(1):1–26, November 2007. ISSN 0165-0173. doi: 10.1016/j.brainresrev.2007.05.002.

[28] Thérèse M. Jay and Menno P. Witter. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, 313(4):574–586, 1991. ISSN 1096-9861. doi: 10.1002/cne.903130404.

[29] Candela Sánchez-Bellot, Rawan AlSubaie, Karyna Mishchanchuk, Ryan W. S. Wee, and Andrew F. MacAskill. Two opposing hippocampus to prefrontal cortex pathways for the control of approach and avoidance behaviour. *Nature Communications*, 13 (1):339, January 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-27977-7.

[30] Ivan Soltesz and Attila Losonczy. CA1 pyramidal cell diversity enabling parallel information processing in the hippocampus. *Nature Neuroscience*, 21(4):484–493, April 2018. ISSN 1546-1726. doi: 10.1038/s41593-018-0118-0.

[31] Julia Rosauro Clarke, Martín Cammarota, Agnès Gruart, Iván Izquierdo, and José María Delgado-García. Plastic modifications induced by object recognition memory processing. *Proceedings of the National Academy of Sciences*, 107(6): 2652–2657, February 2010. doi: 10.1073/pnas.0915059107.

[32] Agnès Gruart, María Dolores Muñoz, and José M. Delgado-García. Involvement of the CA3–CA1 Synapse in the Acquisition of Associative Learning in Behaving Mice. *Journal of Neuroscience*, 26(4):1077–1087, January 2006. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2834-05.2006.

[33] S. J. Martin, P. D. Grimwood, and R. G. M. Morris. Synaptic Plasticity and Memory: An Evaluation of the Hypothesis. *Annual Review of Neuroscience*, 23 (Volume 23, 2000):649–711, March 2000. ISSN 0147-006X, 1545-4126. doi: 10.1146/annurev.neuro.23.1.649.

[34] Donna M. Martin, Jennifer M. Skidmore, Sharon E. Fox, Philip J. Gage, and Sally A. Camper. Pitx2 Distinguishes Subtypes of Terminally Differentiated Neurons in the Developing Mouse Neuroepithelium. *Developmental Biology*, 252(1):84–99, December 2002.

[35] D.O. Hebb. *The Organization of Behavior*. Wiley, New York, 1949.

[36] T. V. P. Bliss and G. L. Collingridge. A synaptic model of memory: Long-term potentiation in the hippocampus. *Nature*, 361(6407):31–39, January 1993. ISSN 1476-4687. doi: 10.1038/361031a0.

[37] T. V. P. Bliss and T. Lømo. Long-lasting potentiation of synaptic transmission in the dentate area of the anaesthetized rabbit following stimulation of the perforant path. *The Journal of Physiology*, 232(2):331–356, 1973. ISSN 1469-7793. doi: 10.1113/jphysiol.1973.sp010273.

[38] Y. Y. Huang, P. V. Nguyen, T. Abel, and E. R. Kandel. Long-lasting forms of synaptic potentiation in the mammalian hippocampus. *Learning & Memory*, 3(2-3): 74–85, September 1996. ISSN 1072-0502, 1549-5485.

[39] Guilherme Neves, Sam F. Cooke, and Tim V. P. Bliss. Synaptic plasticity, memory and the hippocampus: A neural network approach to causality. *Nature Reviews Neuroscience*, 9(1):65–75, January 2008. ISSN 1471-0048. doi: 10.1038/nrn2303.

[40] Jonathan R. Whitlock, Arnold J. Heynen, Marshall G. Shuler, and Mark F. Bear. Learning Induces Long-Term Potentiation in the Hippocampus. *Science*, 313(5790): 1093–1097, August 2006. doi: 10.1126/science.1128134.

[41] R. G. M. Morris, E. Anderson, G. S. Lynch, and M. Baudry. Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5. *Nature*, 319(6056):774–776, February 1986. ISSN 0028-0836, 1476-4687. doi: 10.1038/319774a0.

[42] Kazu Nakazawa, Thomas J. McHugh, Matthew A. Wilson, and Susumu Tonegawa. NMDA receptors, place cells and hippocampal spatial memory. *Nature Reviews Neuroscience*, 5(5):361–372, May 2004. ISSN 1471-0048. doi: 10.1038/nrn1385.

[43] Pojeong Park, Arturas Volianskis, Thomas M. Sanderson, Zuner A. Bortolotto, David E. Jane, Min Zhuo, Bong-Kiun Kaang, and Graham L. Collingridge. NMDA receptor-dependent long-term potentiation comprises a family of temporally overlapping forms of synaptic plasticity that are induced by different patterns of stimulation. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633): 20130131, January 2014. doi: 10.1098/rstb.2013.0131.

[44] Joe Z Tsien, Patricio T Huerta, and Susumu Tonegawa. The Essential Role of Hippocampal CA1 NMDA Receptor–Dependent Synaptic Plasticity in Spatial Memory. *Cell*, 87(7):1327–1338, December 1996. ISSN 0092-8674. doi: 10.1016/S0092-8674(00)81827-9.

[45] Yasunori Hayashi, Song-Hai Shi, José A. Esteban, Antonella Piccini, Jean-Christophe Poncer, and Roberto Malinow. Driving AMPA Receptors into Synapses by LTP and CaMKII: Requirement for GluR1 and PDZ Domain Interaction. *Science*, 287 (5461):2262–2267, March 2000. doi: 10.1126/science.287.5461.2262.

[46] Eric R. Kandel. The Molecular Biology of Memory Storage: A Dialogue Between Genes and Synapses. *Science*, 294(5544):1030–1038, November 2001. doi: 10.1126/science.1067020.

[47] H. Wigstrom and B. Gustafsson. Postsynaptic control of hippocampal long-term potentiation. *J.Physiol (Paris)*, 81(0021-7948 (Print)):228–236, 1986.

[48] Miquel Bosch and Yasunori Hayashi. Structural plasticity of dendritic spines. *Current Opinion in Neurobiology*, 22(3):383–388, June 2012. ISSN 0959-4388. doi: 10.1016/j.conb.2011.09.002.

[49] Eric R. Kandel, Yadin Dudai, and Mark R. Mayford. The Molecular and Systems Biology of Memory. *Cell*, 157(1):163–186, March 2014. ISSN 0092-8674. doi: 10.1016/j.cell.2014.03.001.

[50] Adrian J. Duszkiewicz, Colin G. McNamara, Tomonori Takeuchi, and Lisa Genzel. Novelty and Dopaminergic Modulation of Memory Persistence: A Tale of Two Systems. *Trends in Neurosciences*, 42(2):102–114, February 2019. ISSN 0166-2236, 1878-108X. doi: 10.1016/j.tins.2018.10.002.

[51] Kimberly A. Kempadoo, Eugene V. Mosharov, Se Joon Choi, David Sulzer, and Eric R. Kandel. Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. *Proceedings of the National Academy of Sciences*, 113(51):14835–14840, December 2016. doi: 10.1073/pnas.1616515114.

[52] Jefferson Menezes, Niége Alves, Sidnei Borges, Rafael Roehrs, Jociane de Carvalho Myskiw, Cristiane Regina Guerino Furini, Ivan Izquierdo, and Pâmela B. Mello-Carpes. Facilitation of fear extinction by novelty depends on dopamine acting on D1-subtype dopamine receptors in hippocampus. *Proceedings of the National Academy of Sciences*, 112(13):E1652–E1658, March 2015. doi: 10.1073/pnas.1502295112.

[53] Tomonori Takeuchi, Adrian J. Duszkiewicz, Alex Sonneborn, Patrick A. Spooner, Miwako Yamasaki, Masahiko Watanabe, Caroline C. Smith, Guillén Fernández, Karl Deisseroth, Robert W. Greene, and Richard G. M. Morris. Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*, 537(7620):357–362, September 2016. ISSN 1476-4687. doi: 10.1038/nature19325.

[54] Sho Yagishita, Akiko Hayashi-Takagi, Graham C.R. Ellis-Davies, Hidetoshi Urakubo, Shin Ishii, and Haruo Kasai. A critical time window for dopamine actions on the structural plasticity of dendritic spines. *Science*, 345(6204):1616–1620, September 2014. doi: 10.1126/science.1255514.

[55] Sadegh Nabavi, Rocky Fox, Christophe D. Proulx, John Y. Lin, Roger Y. Tsien, and Roberto Malinow. Engineering a memory with LTD and LTP. *Nature*, 511 (7509):348–352, July 2014. ISSN 1476-4687. doi: 10.1038/nature13294.

[56] Alessio Attardo, James E. Fitzgerald, and Mark J. Schnitzer. Impermanence of dendritic spines in live adult CA1 hippocampus. *Nature*, 523(7562):592–596, July 2015. ISSN 1476-4687. doi: 10.1038/nature14467.

[57] Tracey J. Shors and Louis D. Matzel. Long-term potentiation: What's learning got to do with it? *Behavioral and Brain Sciences*, 20(4):597–614, December 1997. ISSN 1469-1825, 0140-525X. doi: 10.1017/S0140525X97001593.

[58] A. M. Taylor, T. Bus, R. Sprengel, P. H. Seeburg, J. N. P. Rawlins, and D. M. Bannerman. Hippocampal NMDA receptors are important for behavioural inhibition but not for encoding associative spatial memories. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1633):20130149, January 2014. doi: 10.1098/rstb.2013.0149.

[59] Colin Lever, Tom Wills, Francesca Cacucci, Neil Burgess, and John O'Keefe. Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*, 416(6876):90–94, March 2002. ISSN 1476-4687. doi: 10.1038/416090a.

[60] Matthew A. Wilson and Bruce L. McNaughton. Dynamics of the Hippocampal Ensemble Code for Space. *Science*, 261(5124):1055–1058, August 1993. doi: 10.1126/science.8351520.

[61] Kareem Abdou, Mohammad Shehata, Kiriko Choko, Hirofumi Nishizono, Mina Matsuo, Shin-ichi Muramatsu, and Kaoru Inokuchi. Synapse-specific representation of the identity of overlapping memory engrams. *Science*, 360(6394):1227–1231, June 2018. doi: 10.1126/science.aat3810.

[62] György Buzsáki. Neural Syntax: Cell Assemblies, Synapsembles, and Readers. *Neuron*, 68(3):362–385, November 2010. ISSN 0896-6273. doi: 10.1016/j.neuron.2010.09.023.

[63] John J. Hopfield and David W. Tank. Computing with Neural Circuits: A Model. *Science*, 233(4764):625–633, August 1986. doi: 10.1126/science.3755256.

[64] Sheena A. Josselyn, Stefan Köhler, and Paul W. Frankland. Finding the engram. *Nature Reviews Neuroscience*, 16(9):521–534, September 2015. ISSN 1471-0048. doi: 10.1038/nrn4000.

[65] Susumu Tonegawa, Michele Pignatelli, Dheeraj S Roy, and Tomás J Ryan. Memory engram storage and retrieval. *Current Opinion in Neurobiology*, 35:101–109, December 2015. ISSN 0959-4388. doi: 10.1016/j.conb.2015.07.009.

[66] Susumu Tonegawa, Xu Liu, Steve Ramirez, and Roger Redondo. Memory Engram Cells Have Come of Age. *Neuron*, 87(5):918–931, September 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.08.002.

[67] Aleena R. Garner, David C. Rowland, Sang Youl Hwang, Karsten Baumgaertel, Bryan L. Roth, Cliff Kentros, and Mark Mayford. Generation of a Synthetic Memory Trace. *Science*, 335(6075):1513–1516, March 2012. doi: 10.1126/science.1214985.

[68] Anthony Holtmaat and Pico Caroni. Functional and structural underpinnings of neuronal assembly formation in learning. *Nature Neuroscience*, 19(12):1553–1562, December 2016. ISSN 1546-1726. doi: 10.1038/nn.4418.

[69] Thomas Rogerson, Denise J. Cai, Adam Frank, Yoshitake Sano, Justin Shobe, Manuel F. Lopez-Aranda, and Alcino J. Silva. Synaptic tagging during memory allocation. *Nature Reviews Neuroscience*, 15(3):157–169, March 2014. ISSN 1471-0048. doi: 10.1038/nrn3667.

[70] Jose Viosca, Mikel Lopez de Armentia, Dragana Jancic, and Angel Barco. Enhanced CREB-dependent gene expression increases the excitability of neurons in the basal amygdala and primes the consolidation of contextual and cued fear memory. *Learning & Memory*, 16(3):193–197, March 2009. ISSN 1072-0502, 1549-5485. doi: 10.1101/lm.1254209.

[71] Gaël Daoudal and Dominique Debanne. Long-Term Plasticity of Intrinsic Excitability: Learning Rules and Mechanisms. *Learning & Memory*, 10(6):456–465, November 2003. ISSN 1072-0502, 1549-5485. doi: 10.1101/lm.64103.

[72] Flavio Donato, Santiago Belluco Rompani, and Pico Caroni. Parvalbumin-expressing basket-cell network plasticity induced by experience regulates adult learning. *Nature*, 504(7479):272–276, December 2013. ISSN 1476-4687. doi: 10.1038/nature12866.

[73] Johannes J. Letzkus, Steffen B. E. Wolff, Elisabeth M. M. Meyer, Philip Tovote, Julien Courtin, Cyril Herry, and Andreas Lüthi. A disinhibitory microcircuit for associative fear learning in the auditory cortex. *Nature*, 480(7377):331–335, December 2011. ISSN 1476-4687. doi: 10.1038/nature10674.

[74] Tom Hartley, Colin Lever, Neil Burgess, and John O'Keefe. Space in the brain: How the hippocampal formation supports spatial cognition. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 369(1635):20120510, February 2014. doi: 10.1098/rstb.2012.0510.

[75] J. O'Keefe. Place units in the hippocampus of the freely moving rat. *Exp.Neurol.*, 51(0014-4886):78–109, April 1976.

[76] L.L. Colgin, E.I. Moser, and M.B. Moser. Understanding memory through hippocampal remapping. *Trends Neurosci.*, 31(0166-2236 (Print)):469–477, September 2008.

[77] Dmitriy Aronov, Rhino Nevers, and David W. Tank. Mapping of a non-spatial dimension by the hippocampal–entorhinal circuit. *Nature*, 543(7647):719—722, 2017. ISSN 0028-0836. doi: 10.1038/nature21692.

[78] Howard Eichenbaum and Neal J. Cohen. Can We Reconcile the Declarative Memory and Spatial Navigation Views on Hippocampal Function? *Neuron*, 83(4):764—770, 2014. ISSN 08966273. doi: 10.1016/j.neuron.2014.07.032.

[79] John Huxter, Neil Burgess, and John O'Keefe. Independent rate and temporal coding in hippocampal pyramidal cells. *Nature*, 425(6960):828–832, October 2003. ISSN 1476-4687. doi: 10.1038/nature02058.

[80] R. U. Muller and J. L. Kubie. The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 7(7):1951–1968, July 1987. ISSN 0270-6474. doi: 10.1523/JNEUROSCI.07-07-01951.1987.

[81] Isabel A. Muzzio, Liat Levita, Jayant Kulkarni, Joseph Monaco, Clifford Kentros, Matthew Stead, Larry F. Abbott, and Eric R. Kandel. Attention Enhances the Retrieval and Stability of Visuospatial and Olfactory Representations in the Dorsal Hippocampus. *PLOS Biology*, 7(6):e1000140, June 2009. ISSN 1545-7885. doi: 10.1371/journal.pbio.1000140.

[82] Emma R. Wood, Paul A. Dudchenko, and Howard Eichenbaum. The global record of memory in hippocampal neuronal activity. *Nature*, 397(6720):613–616, February 1999. ISSN 1476-4687. doi: 10.1038/17605.

[83] Charlotte N Boccara, Michele Nardin, Federico Stella, Joseph O'Neill, and Jozsef Csicsvari. The entorhinal cognitive map is attracted to goals. *Science*, 363(6434): 1443–1447, 2019. ISSN 0036-8075. doi: 10.1126/science.aav4837.

[84] D. Dupret and D.N. Abrous. [A new chapter in the field of memory: hippocampal neo-neurogenesis]. *Biol.Aujourdhui.*, 204(2105-0678 (Print)):113–129, 2010.

[85] Alexandra Mansell Kaufman, Tristan Geiller, and Attila Losonczy. A Role for the Locus Coeruleus in Hippocampal CA1 Place Cell Reorganization during Spatial Reward Learning. *Neuron*, 105(6):1018–1026.e4, March 2020. ISSN 0896-6273. doi: 10.1016/j.neuron.2019.12.029.

[86] Mark H. Plitt and Lisa M. Giocomo. Experience-dependent contextual codes in the hippocampus. *Nature Neuroscience*, 24(5):705–714, May 2021. ISSN 1546-1726. doi: 10.1038/s41593-021-00816-6.

[87] David M. Smith and Sheri J. Y. Mizumori. Learning-Related Development of Context-Specific Neuronal Responses to Places and Events: The Hippocampal Role in Context Processing. *Journal of Neuroscience*, 26(12):3154–3163, March 2006. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.3234-05.2006.

[88] Guo Wanjia, Serra E. Favila, Ghootae Kim, Robert J. Molitor, and Brice A. Kuhl. Abrupt hippocampal remapping signals resolution of memory interference. *Nature Communications*, 12(1):4816, August 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-25126-0.

[89] George Dragoi and György Buzsáki. Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron*, 50(1):145–157, April 2006. ISSN 0896-6273. doi: 10.1016/j.neuron.2006.02.023.

[90] Doyun Lee, Bei-Jung Lin, and Albert K. Lee. Hippocampal Place Fields Emerge upon Single-Cell Manipulation of Excitability During Behavior. *Science*, 337(6096): 849–853, August 2012. doi: 10.1126/science.1221489.

[91] J J Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8): 2554–2558, April 1982. doi: 10.1073/pnas.79.8.2554.

[92] B. L. McNaughton and R. G. M. Morris. Hippocampal synaptic enhancement and information storage within a distributed memory system. *Trends in Neurosciences*, 10(10):408–415, January 1987. ISSN 0166-2236. doi: 10.1016/0166-2236(87)90011-7.

[93] A. Treves and E.T. Rolls. Computational constraints suggest the need for two distinct input systems to the hippocampal CA3 network. *Hippocampus*, 2(1050-9631 (Print)):189–199, April 1992.

[94] Tom J. Wills, Colin Lever, Francesca Cacucci, Neil Burgess, and John O'Keefe. Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science*, 308(5723):873–876, May 2005. doi: 10.1126/science.1108905.

[95] Xu Liu, Steve Ramirez, Petti T. Pang, Corey B. Puryear, Arvind Govindarajan, Karl Deisseroth, and Susumu Tonegawa. Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature*, 484(7394):381–385, April 2012. ISSN 1476-4687. doi: 10.1038/nature11028.

[96] L. T. Thompson and P. J. Best. Long-term stability of the place-field activity of single units recorded from the dorsal hippocampus of freely behaving rats. *Brain Research*, 509(2):299–308, February 1990. ISSN 0006-8993. doi: 10.1016/0006-8993(90)90555-P.

[97] Clifford G Kentros, Naveen T Agnihotri, Samantha Streater, Robert D Hawkins, and Eric R Kandel. Increased Attention to Spatial Context Increases Both Place Field Stability and Spatial Memory. *Neuron*, 42(2):283–295, April 2004. ISSN 0896-6273. doi: 10.1016/S0896-6273(04)00192-8.

[98] Melissa E. Wang, Ellen G. Wann, Robin K. Yuan, Manuel M. Ramos Álvarez, Squire M. Stead, and Isabel A. Muzzio. Long-Term Stabilization of

Place Cell Remapping Produced by a Fearful Experience. *Journal of Neuroscience*, 32(45):15802–15814, November 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0480-12.2012.

[99] Roland Zemla, Jason J. Moore, Maya D. Hopkins, and Jayeeta Basu. Task-selective place cells show behaviorally driven dynamics during learning and stability during memory recall. *Cell Reports*, 41(8):111700, November 2022. ISSN 2211-1247. doi: 10.1016/j.celrep.2022.111700.

[100] Yuichiro Hayashi. NMDA Receptor-Dependent Dynamics of Hippocampal Place Cell Ensembles. *Journal of Neuroscience*, 39(26):5173–5182, June 2019. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.0243-19.2019.

[101] Jae Sung Lee, John J. Briguglio, Jeremy D. Cohen, Sandro Romani, and Albert K. Lee. The Statistical Structure of the Hippocampal Code for Space as a Function of Time, Context, and Value. *Cell*, 183(3):620–635.e22, October 2020. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2020.09.024.

[102] Yaniv Ziv, Laurie D. Burns, Eric D. Cocker, Elizabeth O. Hamel, Kunal K. Ghosh, Lacey J. Kitch, Abbas El Gamal, and Mark J. Schnitzer. Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*, 16(3):264–266, March 2013. ISSN 1546-1726. doi: 10.1038/nn.3329.

[103] Nathan B. Danielson, Jeffrey D. Zaremba, Patrick Kaifosh, John Bowler, Max Ladow, and Attila Losonczy. Sublayer-Specific Coding Dynamics during Spatial Navigation and Learning in Hippocampal Area CA1. *Neuron*, 91(3):652–665, August 2016. ISSN 0896-6273. doi: 10.1016/j.neuron.2016.06.020.

[104] Thomas Hainmueller and Marlene Bartos. Parallel emergence of stable and dynamic memory engrams in the hippocampus. *Nature*, 558(7709):292–296, 2018. ISSN 0028-0836. doi: 10.1038/s41586-018-0191-2.

[105] Gregor-Alexander Pilz, Stefano Carta, Andreas Stäuble, Asli Ayaz, Sebastian Jessberger, and Fritjof Helmchen. Functional Imaging of Dentate Granule Cells in the Adult Mouse Hippocampus. *Journal of Neuroscience*, 36(28):7407–7414, July 2016. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.3065-15.2016.

[106] Hannah S. Wirtshafter and John F. Disterhoft. In Vivo Multi-Day Calcium Imaging of CA1 Hippocampus in Freely Moving Rats Reveals a High Preponderance of Place Cells with Consistent Place Fields. *Journal of Neuroscience*, 42(22):4538–4554, June 2022. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1750-21.2022.

[107] Michael E Rule, Timothy O'Leary, and Christopher D Harvey. Causes and consequences of representational drift. *Current Opinion in Neurobiology*, 58:141–147, October 2019. ISSN 09594388. doi: 10.1016/j.conb.2019.08.005.

[108] Yadin Dudai, Avi Karni, and Jan Born. The Consolidation and Transformation of Memory. *Neuron*, 88(1):20–32, October 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2015.09.004.

[109] Alessio Attardo, Ju Lu, Takashi Kawashima, Hiroyuki Okuno, James E. Fitzgerald, Haruhiko Bito, and Mark J. Schnitzer. Long-Term Consolidation of Ensemble

Neural Plasticity Patterns in Hippocampal Area CA1. *Cell Reports*, 25(3):640–650.e2, October 2018. ISSN 2211-1247. doi: 10.1016/j.celrep.2018.09.064.

[110] Emily A. Mankin, Fraser T. Sparks, Begum Slayyeh, Robert J. Sutherland, Stefan Leutgeb, and Jill K. Leutgeb. Neuronal code for extended time in the hippocampus. *Proceedings of the National Academy of Sciences*, 109(47):19462–19467, November 2012. doi: 10.1073/pnas.1214107109.

[111] Emily A. Mankin, Geoffrey W. Diehl, Fraser T. Sparks, Stefan Leutgeb, and Jill K. Leutgeb. Hippocampal CA2 Activity Patterns Change over Time to a Larger Extent than between Spatial Contexts. *Neuron*, 85(1):190–201, January 2015. ISSN 0896-6273. doi: 10.1016/j.neuron.2014.12.001.

[112] Joseph R. Manns, Marc W. Howard, and Howard Eichenbaum. Gradual Changes in Hippocampal Activity Support Remembering the Order of Events. *Neuron*, 56 (3):530–540, November 2007. ISSN 0896-6273. doi: 10.1016/j.neuron.2007.08.017.

[113] Alon Rubin, Nitzan Geva, Liron Sheintuch, and Yaniv Ziv. Hippocampal ensemble dynamics timestamp events in long-term memory. *eLife*, 4:e12247, December 2015. ISSN 2050-084X. doi: 10.7554/eLife.12247.

[114] Nitzan Geva, Daniel Deitch, Alon Rubin, and Yaniv Ziv. Time and experience differentially affect distinct aspects of hippocampal representational drift. *Neuron*, 111(15):2357–2366.e5, August 2023. ISSN 0896-6273. doi: 10.1016/j.neuron.2023.05.005.

[115] Dorgham Khatib, Aviv Ratzon, Mariell Sellevoll, Omri Barak, Genela Morris, and Dori Derdikman. Active experience, not time, determines within-day representational drift in dorsal CA1. *Neuron*, 111(15):2348–2356.e4, August 2023. ISSN 0896-6273. doi: 10.1016/j.neuron.2023.05.014.

[116] Garrett J Blair, Changliang Guo, Shiyun Wang, Michael S Fanselow, Peyman Golshani, Daniel Aharoni, and Hugh T Blair. Hippocampal place cell remapping occurs with memory storage of aversive experiences. *eLife*, 12:e80661, July 2023. ISSN 2050-084X. doi: 10.7554/eLife.80661.

[117] Hoonwon Lee and Bong-Kiun Kaang. How engram mediates learning, extinction, and relapse. *Current Opinion in Neurobiology*, 81:102723, August 2023. ISSN 0959-4388. doi: 10.1016/j.conb.2023.102723.

[118] Nathaniel R. Kinsky, William Mau, David W. Sullivan, Samuel J. Levy, Evan A. Ruesch, and Michael E. Hasselmo. Trajectory-modulated hippocampal neurons persist throughout memory-guided navigation. *Nature Communications*, 11(1):2443, May 2020. ISSN 2041-1723. doi: 10.1038/s41467-020-16226-4.

[119] Alexandra T. Keinath, Coralie-Anne Mosser, and Mark P. Brandon. The representation of context in mouse hippocampus is preserved despite neural drift. *Nature Communications*, 13(1):2415, May 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-30198-7.

[120] Shaul Druckmann and Dmitri B. Chklovskii. Neuronal Circuits Underlying Persistent Representations Despite Time Varying Activity. *Current Biology*, 22(22):2095–2103, November 2012. ISSN 0960-9822. doi: 10.1016/j.cub.2012.08.058.

[121] Uri Rokni, Andrew G. Richardson, Emilio Bizzi, and H. Sebastian Seung. Motor Learning with Unstable Neural Representations. *Neuron*, 54(4):653–666, May 2007. ISSN 0896-6273. doi: 10.1016/j.neuron.2007.04.030.

[122] Wickliffe C. Abraham and Anthony Robins. Memory retention – the synaptic stability versus plasticity dilemma. *Trends in Neurosciences*, 28(2):73–78, February 2005. ISSN 0166-2236. doi: 10.1016/j.tins.2004.12.003.

[123] Michael E Rule, Adrianna R Loback, Dhruva Raman, Laura N Driscoll, Christopher D Harvey, and Timothy O'Leary. Stable task information from an unstable neural population. *eLife*, 9:e51121, 2020. doi: 10.7554/elife.51121.

[124] Daniel Deitch, Alon Rubin, and Yaniv Ziv. Representational drift in the mouse visual cortex. *Current Biology*, 31(19):4327–4339.e6, October 2021. ISSN 0960-9822. doi: 10.1016/j.cub.2021.07.062.

[125] Laura N. Driscoll, Noah L. Pettit, Matthias Minderer, Selmaan N. Chettih, and Christopher D. Harvey. Dynamic Reorganization of Neuronal Activity Patterns in Parietal Cortex. *Cell*, 170(5):986–999.e16, August 2017. ISSN 0092-8674, 1097-4172. doi: 10.1016/j.cell.2017.07.021.

[126] Tyler D. Marks and Michael J. Goard. Stimulus-dependent representational drift in primary visual cortex. *Nature Communications*, 12(1):5169, August 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-25436-3.

[127] Carl E. Schoonover, Sarah N. Ohashi, Richard Axel, and Andrew J. P. Fink. Representational drift in primary olfactory cortex. *Nature*, 594(7864):541–546, June 2021. ISSN 1476-4687. doi: 10.1038/s41586-021-03628-7.

[128] K. Richard Ridderinkhof, Markus Ullsperger, Eveline A. Crone, and Sander Nieuwenhuis. The Role of the Medial Frontal Cortex in Cognitive Control. *Science*, 306 (5695):443–447, October 2004. doi: 10.1126/science.1100301.

[129] Matthew F. S. Rushworth, MaryAnn P. Noonan, Erie D. Boorman, Mark E. Walton, and Timothy E. Behrens. Frontal Cortex and Reward-Guided Learning and Decision-Making. *Neuron*, 70(6):1054–1069, June 2011. ISSN 0896-6273. doi: 10.1016/j.neuron.2011.05.014.

[130] Stan B. Floresco, Annie E. Block, and Maric T. L. Tse. Inactivation of the medial prefrontal cortex of the rat impairs strategy set-shifting, but not reversal learning, using a novel, automated procedure. *Behavioural Brain Research*, 190(1):85–96, June 2008. ISSN 0166-4328. doi: 10.1016/j.bbr.2008.02.008.

[131] Kevin A. Corcoran and Gregory J. Quirk. Activity in Prelimbic Cortex Is Necessary for the Expression of Learned, But Not Innate, Fears. *Journal of Neuroscience*, 27 (4):840–844, January 2007. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.5327-06.2007.

[132] Paul W. Frankland, Bruno Bontempi, Lynn E. Talton, Leszek Kaczmarek, and Alcino J. Silva. The Involvement of the Anterior Cingulate Cortex in Remote Contextual Fear Memory. *Science*, 304(5672):881–883, May 2004. doi: 10.1126/science.1094804.

[133] Irit Akirav and Mouna Maroun. Ventromedial Prefrontal Cortex Is Obligatory for Consolidation and Reconsolidation of Object Recognition Memory. *Cerebral Cortex*, 16(12):1759–1765, December 2006. ISSN 1047-3211. doi: 10.1093/cercor/bhj114.

[134] Anna Carballo-Márquez, Anna Vale-Martínez, Gemma Guillazo-Blanch, Meritxell Torras-Garcia, Núria Boix-Trelis, and Margarita Martí-Nicolovius. Differential effects of muscarinic receptor blockade in prelimbic cortex on acquisition and memory formation of an odor-reward task. *Learning & Memory*, 14(9):616–624, September 2007. ISSN 1072-0502, 1549-5485. doi: 10.1101/lm.597507.

[135] Yoshinori Izaki, Koji Hori, and Masahiko Nomura. Disturbance of rat lever-press learning by hippocampo-prefrontal disconnection. *Brain Research*, 860(1):199–202, March 2000. ISSN 0006-8993. doi: 10.1016/S0006-8993(00)02039-4.

[136] Wanda C. Leon, Martin A. Bruno, Simon Allard, Karim Nader, and A. Claudio Cuello. Engagement of the PFC in consolidation and recall of recent spatial memory. *Learning & Memory*, 17(6):297–305, June 2010. ISSN 1072-0502, 1549-5485. doi: 10.1101/lm.1804410.

[137] Ming-Gao Zhao, Hiroki Toyoda, Yong-Seok Lee, Long-Jun Wu, Shanelle W. Ko, Xue-Han Zhang, Yongheng Jia, Fanny Shum, Hui Xu, Bao-Ming Li, Bong-Kiun Kaang, and Min Zhuo. Roles of NMDA NR2B Subtype Receptor in Prefrontal Long-Term Potentiation and Contextual Fear Memory. *Neuron*, 47(6):859–872, September 2005. ISSN 0896-6273. doi: 10.1016/j.neuron.2005.08.014.

[138] Howard Eichenbaum. Prefrontal–hippocampal interactions in episodic memory. *Nature Reviews Neuroscience*, 18(9):547–558, September 2017. ISSN 1471-0048. doi: 10.1038/nrn.2017.74.

[139] Gareth R. I. Barker, Flora Bird, Victoria Alexander, and E. Clea Warburton. Recognition memory for objects, place, and temporal order: A disconnection analysis of the role of the medial prefrontal cortex and perirhinal cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(11):2948–2957, March 2007. ISSN 1529-2401. doi: 10.1523/JNEUROSCI.5289-06.2007.

[140] Owen Y. Chao, Joseph P. Huston, Jay-Shake Li, An-Li Wang, and Maria A. de Souza Silva. The medial prefrontal cortex-lateral entorhinal cortex circuit is essential for episodic-like memory and associative object-recognition. *Hippocampus*, 26(5):633–645, May 2016. ISSN 1098-1063. doi: 10.1002/hipo.22547.

[141] Darren K. Hannesson, John G. Howland, and Anthony G. Phillips. Interaction between Perirhinal and Medial Prefrontal Cortex Is Required for Temporal Order But Not Recognition Memory for Objects in Rats. *Journal of Neuroscience*, 24(19):4596–4604, May 2004. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.5517-03.2004.

[142] Ryan Place, Anja Farovik, Marco Brockmann, and Howard Eichenbaum. Bidirectional prefrontal-hippocampal interactions support context-guided memory. *Nature Neuroscience*, 19(8):992—994, 2016. ISSN 1097-6256. doi: 10.1038/nn.4327.

[143] H. Wang, Y. Hu, and J.Z. Tsien. Molecular and systems mechanisms of memory consolidation and storage. *Prog.Neurobiol.*, 79(0301-0082 (Print)):123–135, June 2006.

[144] J. G. Jenkins and K. M. Dallenbach. Obliviscence During Sleep and Waking. *The American Journal of Psychology*, 35:605–612, 1924. ISSN 1939-8298. doi: 10.2307/1414040.

[145] William Fishbein and Baruch M. Gutwein. Paradoxical sleep and memory storage processes. *Behavioral Biology*, 19(4):425–464, April 1977. ISSN 0091-6773. doi: 10.1016/S0091-6773(77)91903-4.

[146] P. Leconte and V. Bloch. [Deficiency in retention of conditioning after deprivation of paradoxical sleep in rats]. *Comptes Rendus Hebdomadaires Des Seances De l'Academie Des Sciences. Serie D: Sciences Naturelles*, 271(2):226–229, July 1970.

[147] Carlyle Smith. Sleep states and learning: A review of the animal literature. *Neuroscience & Biobehavioral Reviews*, 9(2):157–168, June 1985. ISSN 0149-7634. doi: 10.1016/0149-7634(85)90042-9.

[148] Bernard Hars, Elizabeth Hennevin, and Patrick Pasques. Improvement of learning by cueing during postlearning paradoxical sleep. *Behavioural Brain Research*, 18(3): 241–250, December 1985. ISSN 0166-4328. doi: 10.1016/0166-4328(85)90032-4.

[149] Pascal Ravassard, Al Mahdy Hamieh, Mickaël Antoine Joseph, Nicolas Fraize, Paul-Antoine Libourel, Léa Lebarillier, Sébastien Arthaud, Claire Meissirel, Monique Touret, Gaël Malleret, and Paul-Antoine Salin. REM Sleep-Dependent Bidirectional Regulation of Hippocampal-Based Emotional Memory and LTP. *Cerebral Cortex*, 26(4):1488–1500, April 2016. ISSN 1047-3211. doi: 10.1093/cercor/bhu310.

[150] Richard Boyce, Stephen D. Glasgow, Sylvain Williams, and Antoine Adamantidis. Causal evidence for the role of REM sleep theta rhythm in contextual memory consolidation. *Science*, 352(6287):812–816, May 2016. doi: 10.1126/science.aad5252.

[151] Isabel C. Hutchison and Shailendra Rathore. The role of REM sleep theta activity in emotional memory. *Frontiers in Psychology*, 6, October 2015. ISSN 1664-1078. doi: 10.3389/fpsyg.2015.01439.

[152] Larry R. Squire, editor. *Fundamental Neuroscience*. Elsevier/Academic Press, Amsterdam, fourth edition edition, 2013. ISBN 978-0-12-385871-9.

[153] L. Marshall, H. Helgadottir, M. Molle, and J. Born. Boosting slow oscillations during sleep potentiates memory. *Nature*, 444(1476-4687 (Electronic)):610–613, November 2006.

[154] Björn Rasch and Jan Born. About Sleep's Role in Memory. *Physiological Reviews*, 93(2):681–766, 2013. ISSN 0031-9333. doi: 10.1152/physrev.00032.2012.

[155] GyÖrgy BuzsÁk. Memory consolidation during sleep: A neurophysiological perspective. *Journal of Sleep Research*, 7(S1):17–23, 1998. ISSN 1365-2869. doi: 10.1046/j.1365-2869.7.s1.3.x.

[156] H.S. Kudrimoti, C.A. Barnes, and B.L. McNaughton. Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *J.Neurosci.*, 19(1529-2401):4090–4101, May 1999.

[157] W.E. Skaggs and B.L. McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271(0036-8075): 1870–1873, March 1996.

[158] Matthew A. Wilson and Bruce L. McNaughton. Reactivation of Hippocampal Ensemble Memories During Sleep. *Science*, 265(5172):676–679, July 1994. doi: 10.1126/science.8036517.

[159] Oxana Eschenko, Wiâm Ramadan, Matthias Mölle, Jan Born, and Susan J. Sara. Sustained increase in hippocampal sharp-wave ripple activity during slow-wave sleep after learning. *Learning & Memory*, 15(4):222–228, April 2008. ISSN 1072-0502, 1549-5485. doi: 10.1101/lm.726008.

[160] W. Ramadan, O. Eschenko, and S.J. Sara. Hippocampal Sharp Wave/Ripples during Sleep for Consolidation of Associative Memory. *Plos One*, 4(8), August 2009.

[161] Jens G. Klinzing, Niels Niethard, and Jan Born. Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, 22(10):1598–1610, October 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0467-3.

[162] C. Pavlides and J. Winson. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J.Neurosci.*, 9 (0270-6474 (Print)):2907–2918, August 1989.

[163] Valérie Ego-Stengel and Matthew A. Wilson. Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, 20 (1):1–10, 2010. ISSN 1098-1063. doi: 10.1002/hipo.20707.

[164] Gabrielle Girardeau, Karim Benchenane, Sidney I. Wiener, György Buzsáki, and Michaël B. Zugaro. Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, 12(10):1222–1223, October 2009. ISSN 1546-1726. doi: 10.1038/nn.2384.

[165] Antonio Fernández-Ruiz, Azahara Oliva, Eliezyer Fermino de Oliveira, Florbela Rocha-Almeida, David Tingley, and György Buzsáki. Long-duration hippocampal sharp wave ripples improve memory. *Science*, 364(6445):1082–1086, June 2019. doi: 10.1126/science.aax0758.

[166] Igor Gridchyn, Philipp Schoenenberger, Joseph O'Neill, and Jozsef Csicsvari. Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. *Neuron*, 106(2):291–300.e6, April 2020. ISSN 1097-4199. doi: 10.1016/j.neuron.2020.01.021.

[167] Yadin Dudai and Richard G.M. Morris. To consolidate or not to consolidate: What are the questions? In Johan J. Bolhuis, editor, *Brain, Perception, Memory: Advances in Cognitive Neuroscience*, page 0. Oxford University Press, November 2000. ISBN 978-0-19-852482-3. doi: 10.1093/acprof:oso/9780198524823.003.0011.

[168] Gustavo Deco and Viktor K. Jirsa. Ongoing Cortical Activity at Rest: Criticality, Multistability, and Ghost Attractors. *Journal of Neuroscience*, 32(10):3366–3375, March 2012. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.2523-11.2012.

[169] Karola Kaefer, Federico Stella, Bruce L. McNaughton, and Francesco P. Battaglia. Replay, the default mode network and the cascaded memory systems model. *Nature Reviews Neuroscience*, 23(10):628–640, October 2022. ISSN 1471-0048. doi: 10.1038/s41583-022-00620-6.

[170] N. K. Logothetis, O. Eschenko, Y. Murayama, M. Augath, T. Steudel, H. C. Evrard, M. Besserve, and A. Oeltermann. Hippocampal–cortical interaction during periods of subcortical silence. *Nature*, 491(7425):547–553, November 2012. ISSN 1476-4687. doi: 10.1038/nature11618.

[171] Gareth R. I. Barker and E. Clea Warburton. NMDA Receptor Plasticity in the Perirhinal and Prefrontal Cortices Is Crucial for the Acquisition of Long-Term Object-in-Place Associative Memory. *Journal of Neuroscience*, 28(11):2837–2844, March 2008. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.4447-07.2008.

[172] Rony Paz, Elizabeth P. Bauer, and Denis Paré. Learning-Related Facilitation of Rhinal Interactions by Medial Prefrontal Inputs. *Journal of Neuroscience*, 27(24):6542–6551, June 2007. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1077-07.2007.

[173] L. Nadel and M. Moscovitch. Memory consolidation, retrograde amnesia and the hippocampal complex. *Curr.Opin.Neurobiol*, 7(0959-4388):217–227, April 1997.

[174] Marlieke T. R. van Kesteren, Guillén Fernández, David G. Norris, and Erno J. Hermans. Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences*, 107(16):7550–7555, April 2010. doi: 10.1073/pnas.0914892107.

[175] David R. Euston, Masami Tatsuno, and Bruce L. McNaughton. Fast-Forward Playback of Recent Memory Sequences in Prefrontal Cortex During Sleep. *Science*, 318(5853):1147–1150, November 2007. doi: 10.1126/science.1148979.

[176] Adrien Peyrache, Mehdi Khamassi, Karim Benchenane, Sidney I. Wiener, and Francesco P. Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience*, 12(7):919–926, July 2009. ISSN 1546-1726. doi: 10.1038/nn.2337.

[177] Randolph F. Helfrich, Bryce A. Mander, William J. Jagust, Robert T. Knight, and Matthew P. Walker. Old Brains Come Uncoupled in Sleep: Slow Wave-Spindle Synchrony, Brain Atrophy, and Forgetting. *Neuron*, 97(1):221–230.e4, January 2018. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.11.020.

[178] Y. Isomura, A. Sirota, S. Ozen, S. Montgomery, K. Mizuseki, D.A. Henze, and G. Buzsaki. Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. *Neuron*, 52(0896-6273 (Print)):871–882, December 2006.

[179] Charles-Francois V Latchoumane, Hong-Viet V Ngo, Jan Born, and Hee-Sup S Shin. Thalamic Spindles Promote Memory Formation during Sleep through Triple Phase-Locking of Cortical, Thalamic, and Hippocampal Rhythms. *Neuron*, 95(2):424–435.e6, 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.06.025.

[180] M. Molle, O. Yeshenko, L. Marshall, S.J. Sara, and J. Born. Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. *J.Neurophysiol.*, 96 (0022-3077 (Print)):62–70, July 2006.

[181] J. O'Neill, B. Pleydell-Bouverie, D. Dupret, and J. Csicsvari. Play it again: Reactivation of waking experience and memory. *Trends Neurosci.*, 33(1878-108X (Electronic)):220–229, May 2010.

[182] Athanassios G Siapas and Matthew A Wilson. Coordinated Interactions between Hippocampal Ripples and Cortical Spindles during Slow-Wave Sleep. *Neuron*, 21(5): 1123–1128, November 1998. ISSN 0896-6273. doi: 10.1016/S0896-6273(00)80629-7.

[183] Anton Sirota, Jozsef Csicsvari, Derek Buhl, and György Buzsáki. Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences*, 100(4):2065–2069, February 2003. doi: 10.1073/pnas.0437938100.

[184] Bernhard P. Staresina, Til Ole Bergmann, Mathilde Bonnefond, Roemer van der Meij, Ole Jensen, Lorena Deuker, Christian E. Elger, Nikolai Axmacher, and Juergen Fell. Hierarchical nesting of slow oscillations, spindles and ripples in the human hippocampus during sleep. *Nature Neuroscience*, 18(11):1679–1686, November 2015. ISSN 1546-1726. doi: 10.1038/nn.4119.

[185] F.P. Battaglia, G.R. Sutherland, and B.L. McNaughton. Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions. *Learn.Mem.*, 11(1072-0502 (Print)):697–704, November 2004.

[186] Carlos Puentes-Mestril and Sara J. Aton. Linking Network Activity to Synaptic Plasticity during Sleep: Hypotheses and Recent Data. *Frontiers in Neural Circuits*, 11, September 2017. ISSN 1662-5110. doi: 10.3389/fncir.2017.00061.

[187] Graham H. Diering, Raja S. Nirujogi, Richard H. Roth, Paul F. Worley, Akhilesh Pandey, and Richard L. Huganir. Homer1a drives homeostatic scaling-down of excitatory synapses during sleep. *Science (New York, N.Y.)*, 355(6324):511–515, February 2017. ISSN 1095-9203. doi: 10.1126/science.aai8355.

[188] Andres D. Grosmark, Kenji Mizuseki, Eva Pastalkova, Kamran Diba, and György Buzsáki. REM Sleep Reorganizes Hippocampal Excitability. *Neuron*, 75(6):1001–1007, 2012. ISSN 0896-6273. doi: 10.1016/j.neuron.2012.08.015.

[189] Reto Huber, Hanna Mäki, Mario Rosanova, Silvia Casarotto, Paola Canali, Adenauer G. Casali, Giulio Tononi, and Marcello Massimini. Human cortical excitability increases with time awake. *Cerebral Cortex (New York, N.Y.: 1991)*, 23(2):332–338, February 2013. ISSN 1460-2199. doi: 10.1093/cercor/bhs014.

[190] Marion Kuhn, Elias Wolf, Jonathan G. Maier, Florian Mainberger, Bernd Feige, Hanna Schmid, Jan Bürklin, Sarah Maywald, Volker Mall, Nikolai H. Jung, Janine Reis, Kai Spiegelhalder, Stefan Klöppel, Annette Sterr, Anne Eckert, Dieter Riemann, Claus Normann, and Christoph Nissen. Sleep recalibrates homeostatic and associative synaptic plasticity in the human cortex. *Nature Communications*, 7(1): 12455, August 2016. ISSN 2041-1723. doi: 10.1038/ncomms12455.

[191] Hiroyuki Miyawaki and Kamran Diba. Regulation of Hippocampal Firing by Network Oscillations during Sleep. *Current biology: CB*, 26(7):893–902, April 2016. ISSN 1879-0445. doi: 10.1016/j.cub.2016.02.024.

[192] Giulio Tononi and Chiara Cirelli. Sleep and the Price of Plasticity: From Synaptic and Cellular Homeostasis to Memory Consolidation and Integration. *Neuron*, 81(1): 12–34, January 2014. ISSN 0896-6273. doi: 10.1016/j.neuron.2013.12.025.

[193] Vladyslav V. Vyazovskiy, Chiara Cirelli, Martha Pfister-Genskow, Ugo Faraguna, and Giulio Tononi. Molecular and electrophysiological evidence for net synaptic potentiation in wake and depression in sleep. *Nature Neuroscience*, 11(2):200–208, February 2008. ISSN 1546-1726. doi: 10.1038/nn2035.

[194] Luisa de Vivo, Michele Bellesi, William Marshall, Eric A. Bushong, Mark H. Ellisman, Giulio Tononi, and Chiara Cirelli. Ultrastructural evidence for synaptic scaling across the wake/sleep cycle. *Science (New York, N.Y.)*, 355(6324):507–510, February 2017. ISSN 1095-9203. doi: 10.1126/science.aah5982.

[195] Marcos G. Frank. Sleep and plasticity in the visual cortex: More than meets the eye. *Current Opinion in Neurobiology*, 44:8–12, June 2017. ISSN 1873-6882. doi: 10.1016/j.conb.2017.01.001.

[196] Wei Li, Lei Ma, Guang Yang, and Wen-Biao Gan. REM sleep selectively prunes and maintains new synapses in development and learning. *Nature Neuroscience*, 20 (3):427–437, 2017. ISSN 1097-6256. doi: 10.1038/nn.4479.

[197] Guang Yang, Cora Sau Wan Lai, Joseph Cichon, Lei Ma, Wei Li, and Wen-Biao Gan. Sleep promotes branch-specific formation of dendritic spines after learning. *Science (New York, N.Y.)*, 344(6188):1173–1178, June 2014. ISSN 1095-9203. doi: 10.1126/science.1249098.

[198] M. V. Ambrosini, A. G. Sadile, U. A. Gironi Carnevale, M. Mattiaccio, and A. Giuditta. The sequential hypothesis on sleep function. I. Evidence that the structure of sleep depends on the nature of the previous waking experience. *Physiology & Behavior*, 43(3):325–337, 1988. ISSN 0031-9384. doi: 10.1016/0031-9384(88)90196-5.

[199] Elizabeth Bostock, Robert U. Muller, and John L. Kubie. Experience-dependent modifications of hippocampal place cell firing. *Hippocampus*, 1(2):193–205, 1991. ISSN 1098-1063. doi: 10.1002/hipo.450010207.

[200] Joseph O'Neill, Timothy J. Senior, Kevin Allen, John R. Huxter, and Jozsef Csicsvari. Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, 11(2):209–215, February 2008. ISSN 1546-1726. doi: 10.1038/nn2037.

[201] M.A. Lopez-Fernandez, M.F. Montaron, E. Varea, G. Rougon, C. Venero, D.N. Abrous, and C. Sandi. Upregulation of Polysialylated Neural Adhesion Molecule in the dorsal hippocampus following contextual fear conditioning is involved in long-term memory formation. *J Neurosci.*, 27(17):4552–4561, 2007.

[202] Bapun Giri, Hiroyuki Miyawaki, Kenji Mizuseki, Sen Cheng, and Kamran Diba. Hippocampal Reactivation Extends for Several Hours Following Novel Experience.

The Journal of Neuroscience, 39(5):866–875, January 2019. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.1950-18.2018.

[203] Anuck Sawangjit, Carlos N. Oyanedel, Niels Niethard, Carolina Salazar, Jan Born, and Marion Inostroza. The hippocampus is crucial for forming non-hippocampal long-term memory during sleep. *Nature*, 564(7734):109–113, December 2018. ISSN 1476-4687. doi: 10.1038/s41586-018-0716-8.

[204] Nghia D. Nguyen, Andrew Lutas, Jesseba Fernando, Josselyn Vergara, Justin McMahon, Jordane Dimidschstein, and Mark L. Andermann. Cortical reactivations predict future sensory responses. *bioRxiv*, page 2022.11.14.516421, January 2022. doi: 10.1101/2022.11.14.516421.

[205] George Dragoi and Susumu Tonegawa. Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, 469(7330):397–401, January 2011. ISSN 1476-4687. doi: 10.1038/nature09633.

[206] Kamran Diba, Kourosh Maboudi, Bapun Giri, Hiroyuki Miyawaki, and Caleb Kemere. Retuning of hippocampal representations during sleep. https://www.researchsquare.com, November 2022.

[207] Andres D Grosmark and György Buzsáki. Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science (New York, N.Y.)*, 351(6280):1440–3, 2016. ISSN 0036-8075. doi: 10.1126/science.aad1935.

[208] William A Liberti, Jeffrey E Markowitz, L Nathan Perkins, Derek C Liberti, Daniel P Leman, Grigori Guitchounts, Tarciso Velho, Darrell N Kotton, Carlos Lois, and Timothy J Gardner. Unstable neurons underlie a stable learned behavior. *Nature Neuroscience*, 19(12):1665–1671, 2016. ISSN 1097-6256. doi: 10.1038/nn.4405.

[209] Blake A Richards and Paul W Frankland. The Persistence and Transience of Memory. *Neuron*, 94(6):1071–1084, 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.04.037.

[210] Francesca Schönsberg, Yasser Roudi, and Alessandro Treves. Efficiency of Local Learning Rules in Threshold-Linear Associative Networks. *Physical Review Letters*, 126(1):018301, January 2021. doi: 10.1103/PhysRevLett.126.018301.

[211] Niels Niethard and Jan Born. Back to baseline: Sleep recalibrates synapses. *Nature Neuroscience*, 22(2):149–151, 2019. ISSN 1097-6256. doi: 10.1038/s41593-018-0327-6.

[212] Hiroyuki Miyawaki, Brendon O. Watson, and Kamran Diba. Neuronal firing rates diverge during REM and homogenize during non-REM. *Scientific Reports*, 9(1):689, 2019. ISSN 2045-2322. doi: 10.1038/s41598-018-36710-8.

[213] Jesse J. Langille. Remembering to Forget: A Dual Role for Sleep Oscillations in Memory Consolidation and Forgetting. *Frontiers in Cellular Neuroscience*, 13:71, 2019. ISSN 1662-5102. doi: 10.3389/fncel.2019.00071.

[214] Masako Tamaki, Zhiyan Wang, Tyler Barnes-Diana, DeeAnn Guo, Aaron V. Berard, Edward Walsh, Takeo Watanabe, and Yuka Sasaki. Complementary contributions of non-REM and REM sleep to visual learning. *Nature Neuroscience*, 23(9):1150–1156, 2020. ISSN 1097-6256. doi: 10.1038/s41593-020-0666-y.

[215] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520, 2019. ISSN 1097-6256. doi: 10.1038/s41593-019-0460-x.

[216] Federico Stella, Peter Baracskay, Joseph O'Neill, and Jozsef Csicsvari. Hippocampal Reactivation of Random Trajectories Resembling Brownian Diffusion. *Neuron*, 102 (2):450–461.e7, April 2019. ISSN 0896-6273. doi: 10.1016/j.neuron.2019.01.052.

[217] R.P. Kesner, G. Farnsworth, and H. Kametani. Role of parietal cortex and hippocampus in representing spatial information. *Cereb.Cortex*, 1(1047-3211 (Print)): 367–373, September 1991.

[218] J. Csicsvari, H. Hirase, A. Czurkó, A. Mamiya, and G. Buzsáki. Fast network oscillations in the hippocampal CA1 region of the behaving rat. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 19(16):RC20, August 1999. ISSN 1529-2401.

[219] Joseph O'Neill, Timothy Senior, and Jozsef Csicsvari. Place-Selective Firing of CA1 Pyramidal Cells during Sharp Wave/Ripple Network Patterns in Exploratory Behavior. *Neuron*, 49(1):143–155, January 2006. ISSN 0896-6273. doi: 10.1016/j.neuron.2005.10.037.

[220] K.D. Harris, H. Hirase, X. Leinekugel, D.A. Henze, and G. Buzsaki. Temporal interaction between single spikes and complex spike bursts in hippocampal pyramidal cells. *Neuron*, 32(0896-6273):141–149, October 2001.

[221] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.

[222] J. Csicsvari, H. Hirase, A. Czurko, and G. Buzsaki. Reliability and state dependence of pyramidal cell-interneuron synapses in the hippocampus: An ensemble approach in the behaving rat. *Neuron*, 21(0896-6273):179–189, July 1998. doi: 10.1016/s0896-6273(00)80525-5.

[223] P.P. Mitra and B. Pesaran. Analysis of dynamic brain imaging data. *Biophys.J.*, 76 (0006-3495):691–708, February 1999.

[224] D.J. Thomson. Spectrum estimation and harmonic analysis. *Proc.IEEE*, 70:1055–1096, 1982.

[225] Kourosh Maboudi, Etienne Ackermann, Laurel Watkins de Jong, Brad E Pfeiffer, David Foster, Kamran Diba, and Caleb Kemere. Uncovering temporal structure in hippocampal output patterns. *eLife*, 7:e34467, June 2018. ISSN 2050-084X. doi: 10.7554/eLife.34467.

[226] E.J. Markus, C.A. Barnes, B.L. McNaughton, V.L. Gladden, and W.E. Skaggs. Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*, 4(1050-9631 (Print)):410–421, August 1994.

[227] William H. Alexander and Joshua W. Brown. Medial prefrontal cortex as an action-outcome predictor. *Nature Neuroscience*, 14(10):1338–1344, October 2011. ISSN 1546-1726. doi: 10.1038/nn.2921.

[228] Jane X. Wang, Zeb Kurth-Nelson, Dharshan Kumaran, Dhruva Tirumala, Hubert Soyer, Joel Z. Leibo, Demis Hassabis, and Matthew Botvinick. Prefrontal cortex as a meta-reinforcement learning system. *Nature Neuroscience*, 21(6):860–868, June 2018. ISSN 1546-1726. doi: 10.1038/s41593-018-0147-8.

[229] B.A. Kirchhoff, A.D. Wagner, A. Maril, and C.E. Stern. Prefrontal-temporal circuitry for episodic encoding and subsequent memory. *J.Neurosci.*, 20(0270-6474 (Print)):6173–6180, August 2000.

[230] Takashi Kitamura, Sachie K. Ogawa, Dheeraj S. Roy, Teruhiro Okuyama, Mark D. Morrissey, Lillian M. Smith, Roger L. Redondo, and Susumu Tonegawa. Engrams and circuits crucial for systems consolidation of a memory. *Science*, 356(6333):73–78, April 2017. doi: 10.1126/science.aam6808.

[231] Michael E. Ragozzino, Christine Wilcox, Mario Raso, and Raymond P. Kesner. Involvement of rodent prefrontal cortex subregions in strategy switching. *Behavioral Neuroscience*, 113(1):32–41, 1999. ISSN 1939-0084. doi: 10.1037/0735-7044.113.1.32.

[232] Nicolas Maingret, Gabrielle Girardeau, Ralitsa Todorova, Marie Goutierre, and Michaël Zugaro. Hippocampo-cortical coupling mediates memory consolidation during sleep. *Nature Neuroscience*, 19(7):959–964, July 2016. ISSN 1546-1726. doi: 10.1038/nn.4304.

[233] Gideon Rothschild, Elad Eban, and Loren M. Frank. A cortical–hippocampal–cortical loop of information processing during memory consolidation. *Nature Neuroscience*, 20(2):251–259, February 2017. ISSN 1546-1726. doi: 10.1038/nn.4457.

[234] Karim Benchenane, Adrien Peyrache, Mehdi Khamassi, Patrick L. Tierney, Yves Gioanni, Francesco P. Battaglia, and Sidney I. Wiener. Coherent Theta Oscillations and Reorganization of Spike Timing in the Hippocampal- Prefrontal Network upon Learning. *Neuron*, 66(6):921–936, June 2010. ISSN 0896-6273. doi: 10.1016/j.neuron.2010.05.013.

[235] Nghia D. Nguyen, Andrew Lutas, Oren Amsalem, Jesseba Fernando, Andy Young-Eon Ahn, Richard Hakim, Josselyn Vergara, Justin McMahon, Jordane Dimidschstein, Bernardo L. Sabatini, and Mark L. Andermann. Cortical reactivations predict future sensory responses. *Nature*, 625(7993):110–118, January 2024. ISSN 1476-4687. doi: 10.1038/s41586-023-06810-1.

[236] Daniel Levenstein, György Buzsáki, and John Rinzel. NREM sleep in the rodent neocortex and hippocampus reflects excitable dynamics. *Nature Communications*, 10(1):2478, 2019. ISSN 2041-1723. doi: 10.1038/s41467-019-10327-5.

[237] HaoRan Chang, Ingrid M. Esteves, Adam R. Neumann, Majid H. Mohajerani, and Bruce L. McNaughton. Cortical reactivation of spatial and non-spatial features coordinates with hippocampus to form a memory dialogue. *Nature Communications*, 14(1):7748, November 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-43254-7.

[238] Akihiro Goto and Yasunori Hayashi. Offline neuronal activity and synaptic plasticity during sleep and memory consolidation. *Neuroscience Research*, 189:29–36, April 2023. ISSN 0168-0102. doi: 10.1016/j.neures.2022.12.021.

[239] Julie Seibt and Marcos G. Frank. Primed to Sleep: The Dynamics of Synaptic Plasticity Across Brain States. *Frontiers in Systems Neuroscience*, 13, February 2019. ISSN 1662-5137. doi: 10.3389/fnsys.2019.00002.

[240] Kevin G. Guise and Matthew L. Shapiro. Medial Prefrontal Cortex Reduces Memory Interference by Modifying Hippocampal Encoding. *Neuron*, 94(1):183–192.e8, April 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.03.011.

[241] Pedro A. Feliciano-Ramos, Maria Galazo, Hector Penagos, and Matthew Wilson. Hippocampal memory reactivation during sleep is correlated with specific cortical states of the retrosplenial and prefrontal cortices. *Learning & Memory (Cold Spring Harbor, N.Y.)*, 30(9):221–236, September 2023. ISSN 1549-5485. doi: 10.1101/lm.053834.123.

[242] Pascal Fries. A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10):474–480, October 2005. ISSN 1364-6613. doi: 10.1016/j.tics.2005.08.011.

[243] João D. Semedo, Amin Zandvakili, Christian K. Machens, Byron M. Yu, and Adam Kohn. Cortical Areas Interact through a Communication Subspace. *Neuron*, 102 (1):249–259.e4, April 2019. ISSN 0896-6273. doi: 10.1016/j.neuron.2019.01.026.

[244] Robert S. Blumenfeld and Charan Ranganath. Prefrontal cortex and long-term memory encoding: An integrative review of findings from neuropsychology and neuroimaging. *The Neuroscientist: A Review Journal Bringing Neurobiology, Neurology and Psychiatry*, 13(3):280–291, June 2007. ISSN 1073-8584. doi: 10.1177/1073858407299290.

[245] Thomas J. Delaney and Cian O'Donnell. Fast-local and slow-global neural ensembles in the mouse brain. *Network Neuroscience*, 7(2):731–742, June 2023. ISSN 2472-1751. doi: 10.1162/netn_a_00309.

[246] Xin Liu, Chi Ren, Yichen Lu, Yixiu Liu, Jeong-Hoon Kim, Stefan Leutgeb, Takaki Komiyama, and Duygu Kuzum. Multimodal neural recordings with Neuro-FITM uncover diverse patterns of cortical–hippocampal interactions. *Nature Neuroscience*, 24(6):886–896, June 2021. ISSN 1546-1726. doi: 10.1038/s41593-021-00841-5.

[247] Guosong Hong and Charles M. Lieber. Novel electrode technologies for neural recordings. *Nature Reviews Neuroscience*, 20(6):330–345, June 2019. ISSN 1471-0048. doi: 10.1038/s41583-019-0140-6.

[248] Douglas A. Ruff and Marlene R. Cohen. Simultaneous multi-area recordings suggest that attention improves performance by reshaping stimulus representations. *Nature Neuroscience*, 22(10):1669–1676, October 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0477-1.

[249] Evren Gokcen, Anna I. Jasper, João D. Semedo, Amin Zandvakili, Adam Kohn, Christian K. Machens, and Byron M. Yu. Disentangling the flow of signals between populations of neurons. *Nature Computational Science*, 2(8):512–525, August 2022. ISSN 2662-8457. doi: 10.1038/s43588-022-00282-5.