# The structure of hippocampal activity during REM sleep

Thesis by

Andreas Hoenselaar

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy


California Institute of Technology

Pasadena, California

2015

(Defended August 27, 2014)

Meinen Eltern gewidmet

# Acknowledgments

I would like to thank my advisor, Thanos Siapas, for many years of support and guidance. I also wish to thank my thesis committee members, Michael Dickinson, Michael Elowitz, Pietro Perona and Michael Roukes, for valuable input and help in bringing this research project to conclusion. I am grateful for the wonderful environment created by all members of the Siapas lab. Without my friends at Caltech and outside of Caltech, the past six years would have been severly lacking in variety, distraction and fun. Nicholas Hoh, I wish to thank you for being an infinte source of love, inspiration and support.

Finally, I wish to thank my parents — for everything.

# Abstract

The hippocampus is a brain structure critical for the formation of long-term episodic memories. The current predominant theory is that memories are gradually established across neocortical networks under the influence of hippocampal activity. This process of memory consolidation is conjectured to occur during sleep, which is characterized by two different modes of activation: slow-wave sleep (SWS) and rapid eye movement (REM) sleep. The functional roles of these two different sleep states remain unknown. Paradoxically, REM sleep exhibits the main features of awake activity, and is the stage of sleep when most dreams occur. Despite decades of study, the organization and function of REM sleep activity remains poorly understood. The goal of this thesis is to achieve a deeper quantitative understanding of the patterns of firing in area CA1 of the hippocampus during REM sleep using chronic multi-tetrode recordings from freely behaving and naturally sleeping rats. Our analysis shows that CA1 neurons significantly elevate their firing rate for periods that are short in relation to the duration of the REM sleep episode. Furthermore, for the majority of neurons, there is exactly one such burst per REM episode. This leads to lower overall firing rates and sparser population activity in CA1 compared to SWS. The time of onset of these bursts defines a natural order of firing across the population of recorded neurons within each REM episode. We demonstrate that this order does not repeat across REM episodes. Our results suggest that CA1 neurons are activated in random sequences across REM episodes, resulting in sparse patterns with only a small fraction of neurons active at any given time.

# Contents

# List of Figures

# Chapter 1

# Introduction

Many lines of evidence have shown that the hippocampus is critical for the formation of long-term episodic memories [51, 11]. The currently predominant theory is that episodic memories are gradually established in distributed neocortical networks under the influence of hippocampal activity. This process is conjectured to occur in two stages: encoding of memories in hippocampal circuits during awake behavior, and consolidation of memories in neocortical circuits during sleep [31, 8, 16].

Consistent with this hypothesis, the hippocampus exhibits two drastically different modes of activation during wakefulness and sleep. In particular, locomotion and attentive wakefulness are characterized by the presence of theta oscillations in the electrocorticogram and local field potential (LFP). The theta rhythm in the 4 Hz to 10 Hz range strongly modulates spiking across the hippocampal network [9]. By contrast, during quiet wakefulness and slow-wave sleep (SWS), theta oscillations disappear, and local field potentials exhibit large irregular activity (LIA) interspersed with synchronous population bursts (see Figure 1.2). During sleep, SWS regularly alternates with rapid eye movement (REM) sleep, which exhibits the main electrophysiological signatures of wakefulness: theta oscillations, and firing of hippocampal neurons that is modulated by the phase of the theta rhythm.

Since its discovery in 1953 [2], REM sleep has remained the most enigmatic state of the brain. Consistent with its awake-like activity patterns, REM represents the state of sleep when the majority of dreams occur. The functional role of REM sleep remains unknown, as previous studies have resulted in conflicting and controversial evidence [47, 42]. All mammals exhibit REM sleep, and this brain state dominates early development, arguing for its critical importance in brain development and function. After REM sleep deprivation, an increase of REM sleep occurs (REM rebound), further highlighting the importance of this

state, as the brain actively tries to compensate for lost REM sleep. Surprisingly, monoamine oxidase inhibitors (MAOIs), which have been widely used as antidepressants, appear to suppress REM sleep, arguing against a vital role for this brain state [54]. Furthermore, some REM sleep deprivation studies found evidence for memory impairment, while others have not [47, 27]. REM sleep deprivation studies have been hard to interpret, as they are confounded by non-specific effects, such as stress and fatigue [22].

A complementary approach is to examine the structure of brain activity in REM sleep of naturally sleeping animals [30, 7]. The goal of this thesis is to achieve a deeper quantitative understanding of the patterns of firing in area CA1 of the hippocampus during REM sleep using chronic multi-tetrode recordings from freely behaving and naturally sleeping rats. In the rest of the introduction, we describe in more detail the anatomy, physiology, and previous work relating to this question to set the context for describing our results on the structure of CA1 activity in REM sleep and its implications for our understanding of this stage of sleep.

## 1.1 Anatomy and major pathways

What follows is an extremely abbreviated description of the hippocampal network. For further details on this vast subject, we refer to work by Andersen et al. [1]. The hippocampus can be thought of as a layer of archicortex rolled into a cylinder and then bent into a C shape, which can most readily be appreciated in coronal sections of the rodent brain. The two tips of the C shape are referred to as the septal pole (rostrodorsal) and the temporal pole (caudoventral). Given these two reference points, which define the long axis of the hippocampus, the main body of the hippocampus curves outward laterally. When cut orthogonally to the long axis, thus revealing the transverse axis of the hippocampus, three cytoarchitectonically different regions can be distinguished: CA1, CA2 and CA3. Figure 1.1 also shows the dentate gyrus and the subiculum, two structures that share the three-layer architecture of archicortex and are part of the hippocampal region but not the hippocampus proper. Furthermore, the entorhinal cortex is delineated as part of the hippocampal formation despite its six-layer isocortical architecture due to its highly developed connectivity with all areas of the hippocampal region. Areas CA1, CA2 and CA3 can be readily identified by the very dense packing of pyramidal cell bodies (cf. Figure 1.1). Of course,

other cell types, including a variety of interneurons, are also present in the cornu ammonis subregions.

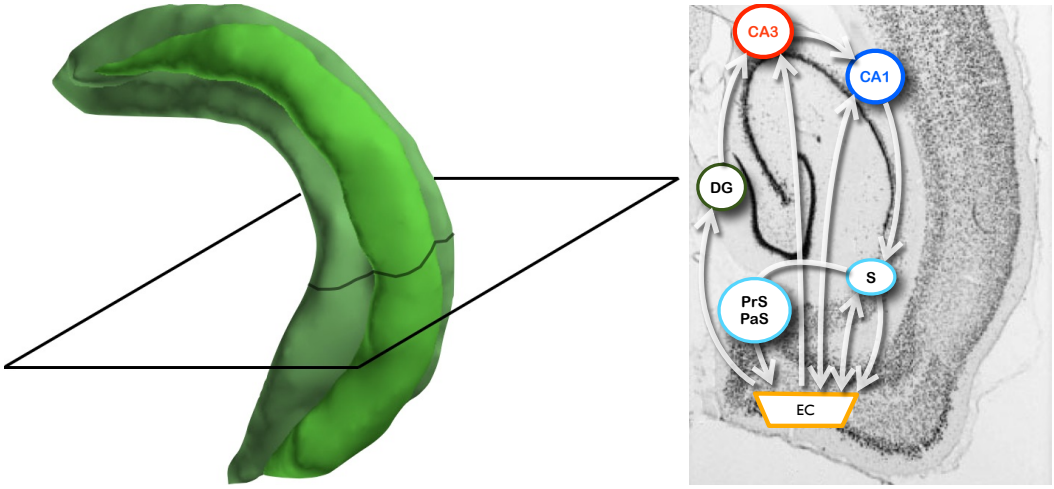


**Figure 1.1:** Left panel: Shape of the rodent hippocampus in an anterolateral view. Cut plane of a horizontal section is indicated. Right panel: Horizontal histological section illustrating the dense cell layer of the cornu ammonis. A circuit diagram of major hippocampal areas is superimposed. DG: dentate gyrus, EC: entorhinal cortex, S: subiculum, PrS: Presubiculum, PaS: Parasubiculum

Most sensory information enters the hippocampal formation via entorhinal cortex. Superficial layers II and III of entorhinal cortex project to dentate gyrus, CA1 and the subiculum via the perforant pathway, named aptly because its fibers have to travel through subicular cell layers to reach their target. Granule cells in the dentate synapse on pyramidal cells in CA3 via the mossy fibres. Area CA3 is characterized by a rich recurrent network architecture that is hypothesized to form the substrate for an associative memory network [4, 36]. Axons of CA3 pyramidal cells form the Schaffer Collaterals, which synapse on CA1 principal cells.

Area CA1 and subiculum are considered the two major output structures of the hippocampal network. CA1 projects directly to cortical areas but also contributes massively to input into the subiculum. The subiculum sends most of its output to deep layers V and VI of entorhinal cortex with additional projections to perirhinal cortex, retrosplenial cortex and medial prefrontal cortex and subcortical targets (mammillary nucleus, nucleus accumbens and the thalamus). For this thesis, we are going to focus on recordings from area CA1. It is the most well-studied area of the hippocampal formation, and of particular interest due to its rich output projections. Area CA1, located at the interface between the

hippocampus and cortical networks, is a logical target to study memory consolidation, a process that is believed to involve the flow of information from the hippocampus to cortical areas. In addition, dorsal CA1 is easily approachable for multi-tetrode drives, and the dense cell layer creates an environment amenable to stable high-yield recordings.

## 1.2 Neural activity patterns in the hippocampus

### 1.2.1 Neural activity during awake behavior

Extracellular recordings from the hippocampus, in particular area CA1, exhibit several very distinct features. The local field potential (LFP) — low-frequency components of the signal below 200 Hz that are associated with the summed synaptic potentials of the neural population in the proximity of the recording electrode — are heavily dominated by the theta rhythm. This oscillation in the frequency range from 4 to 10 Hz is a defining signature of the hippocampal LFP during exploratory behavior [20, 53]. Hippocampal neurons are known to be modulated by the theta rhythm, and often fire preferentially at certain theta phases [10]. Since theta oscillations can also be observed in cortical and subcortical structures [25, 14], they have been ascribed a role in the synchronization of information flow within and outside of the hippocampus [13]. One particular instantiation of this concept is the chunking of spatial information into individual theta cycles [21].

In the early 1970s, John O'Keefe discovered that many principal cells in area CA1 emit spikes only when the animal's location is within a particular region of its environment [37]. Outside of this region, the cells are almost completely silent. This close tie between location and firing rate led to the term *place cell* for pyramidal cells that exclusively fire in a particular *place field.*

A large population of place cells with place fields that evenly span an environment constitute a neural representation of space. This notion of the hippocampus as a substrate for self-location was first proposed by John O'Keefe and Lynn Nadel [38].

It is estimated that in any environment, about 20% of pyramidal cells in CA1 participate in the spatial representation. Different environments are represented by different place cell subsets. One remarkable feature is the relative stability of this spatial map. When reintroducing an animal into the same environment repeatedly, there is significant overlap between the populations that are active during each exposure [55]. Furthermore, place

fields maintain stable positions and — to some extent — peak firing rates across multiple exposures.

A close tie exists between the theta rhythm and place cell activity during place field traversals. As mentioned above, most hippocampal neurons have a preference to emit spikes in a certain theta phase range. The preferred phase of firing in place cells changes as the animal progresses through the field and shifts to earlier and earlier theta phase values. This phenomenon of *phase precession* was initially reported by John O'Keefe and Michael Recce [39]. The transition through the range of phase values follows a stereotypical pattern. As the animal enters the place field, spikes are emitted late in the theta cycle, i.e. close to the peak at phase values between $\frac{3}{2}\pi$ and $2\pi$. The preferred phase then shifts towards the trough at $\pi$ and reaches values around 0 by the time the animal exits from the place field.

While the role of phase precession has not yet been understood, several appealing hypotheses illustrate functional benefits of phase precession both in the context of spatial and memory-related roles of the hippocampus. An animal traversing a path will cause a population of place cells along this path to activate sequentially. The order of this sequence is determined by the spatial relationship of place fields along the path. The temporal characteristics of this sequence are determined by running speed and place field size. Given the typical running speed of rodents, the activity of place cells with overlapping place fields may easily be separated by hundreds of milliseconds or seconds.

Intriguingly, phase precession provides a mechanism to map the sequence of cell activation from relatively slow behavioral timescales onto a compressed representation in each individual theta cycle. This is due to the fact that the relative spiking order of place cells with overlapping place fields will be maintained even though both cells change their preferred theta phase from cycle to cycle. If place field A precedes field B on the path, then more progress has already been made through field A than through field B. Consequently, cell A's theta phase preference has precessed further than that of cell B, causing it to emit spikes earlier than cell B on each individual theta cycle. Such a temporally compressed representation of space provides more information about the animal's position and path to a downstream target of place cell populations [48]. It has been shown that incorporating phase information into a Bayesian decoder reduces the reconstruction error of the animal's position [24]. Whether phase information can be used by a downstream area depends on whether the LFP theta component, or a signal that is highly coherent with it, is accessible by

the target. Medial prefrontal cortex and the ventral striatum are known to exhibit a theta rhythm that is highly coherent with hippocampal theta, at least temporarily, and dependent on task demands [3, 5]. However, there is no conclusive evidence that the information encoded by theta phase is actually used in any way.

Furthermore, in the compressed representation of space, cells with overlapping place fields repeatedly emit spikes in a fixed relative order. If the place cells under consideration have place fields of similar size and approximately equal phase precession slopes, the phase offset translates to a constant temporal delay on the order of no more than several tens of milliseconds between spikes from different neurons. If the cross-correlogram for a pair of cells with overlapping place fields is calculated, a local maximum that is shifted relative to zero lag, as well as asymmetry around zero lag, can be observed. Spike-timing-dependent plasticity is effective at this timescale [6] and can — after repeated traversals — strengthen the neural trace associated with the path taken by the animal, providing a substrate for the rapid formation of spatial memories or episodic memories tied to spatial cues.

Finally, phase precession is hypothesized to facilitate decision-making by providing a look-ahead mechanism for downstream targets. On a single theta cycle, spikes represent the most recent past (very early in the cycle), the animal's current position, and the animal's future position (very late in the cycle). A downstream neuron that receives input from a subset of neurons that participate in a theta cycle thus gains access to the animal's putative future location and its neural representation.

### 1.2.2 Neural activity during sleep

According to current thinking, new memories are formed and temporarily stored in the hippocampus, which is well-suited for this purpose due to its highly plastic synapses. Yet, it is known that the hippocampus is not the locus of long-term storage for memories. Lesion studies, such as in the famous patient H.M., have demonstrated that extensive mediotemporal lobe damage can cause anterograde amnesia, the inability to form new memories [46, 33]. In H.M., existing memories were affected proportionally to their recency [49]. Very old memories were unaffected by the lesioning event, whereas more recent memories were affected more severely (temporally graded retrograde amnesia). These observations provide support for the idea that, over time, memories are transferred from the hippocampal formation to other brain structures for long-term storage through the process of memory consolidation.

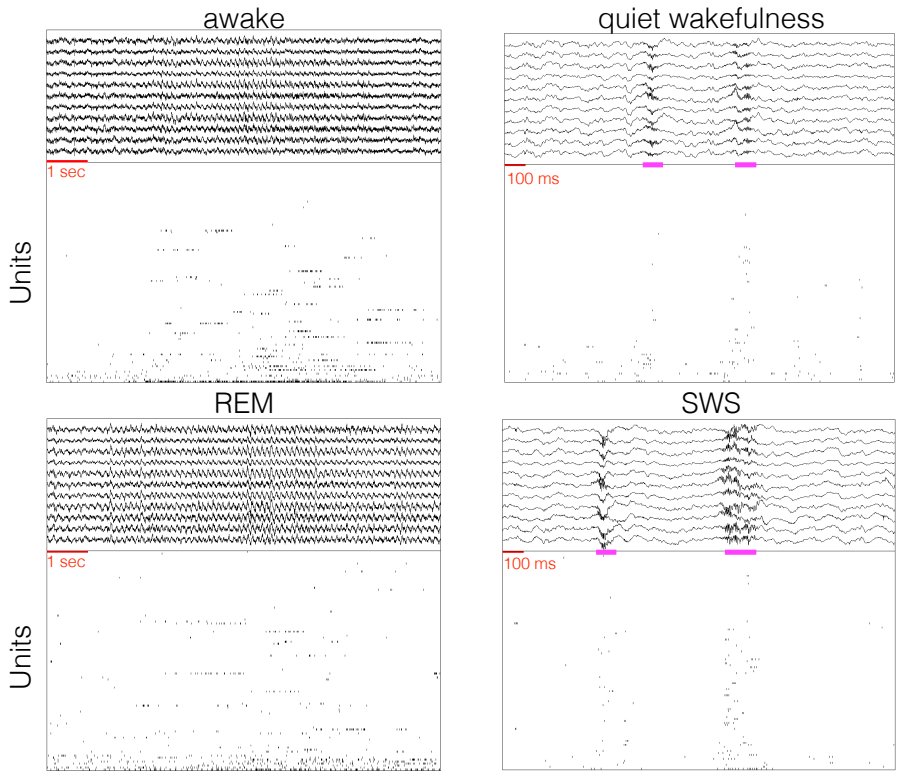


**Figure 1.2:** Hippocampal activity patterns in different behavioral states. Each panel contains several LFP traces in the top, and raster plots with one tick mark per spike below. LFPs in REM and awake are dominated by strong theta oscillations. Spiking is sparse, but place field traversals are visible as brief bursts. Firing in REM is similarly sparse with bursts that are less concentrated. LFPs in quiet wakefulness and SWS don't show strong theta modulation, but exhibit sharp wave-ripple complexes (highlighted). Spiking is highly synchronous and driven by ripples.

Cortical areas are believed to be the locus of long-term memory storage.

Neural activity observed during sleep differs from patterns seen in the awake animal. In rodents, two major modes of sleep are distinguished: slow-wave sleep and rapid eye movement (REM) sleep. This deviates from the classification of sleep in humans, where three or four separate stages of non-REM sleep are commonly delineated in addition to REM sleep [43].

Neural activity during REM sleep is very similar to the awake state. The local field potential is dominated by strong theta oscillations; firing is sparse with little synchronous activity. The hallmarks of slow-wave sleep are slow high-amplitude oscillations that can be recorded from an EEG or the local field potential. Sharp-wave ripples, an aperiodic large-amplitude deflection that is followed by a high-frequency oscillation in the 100 Hz to

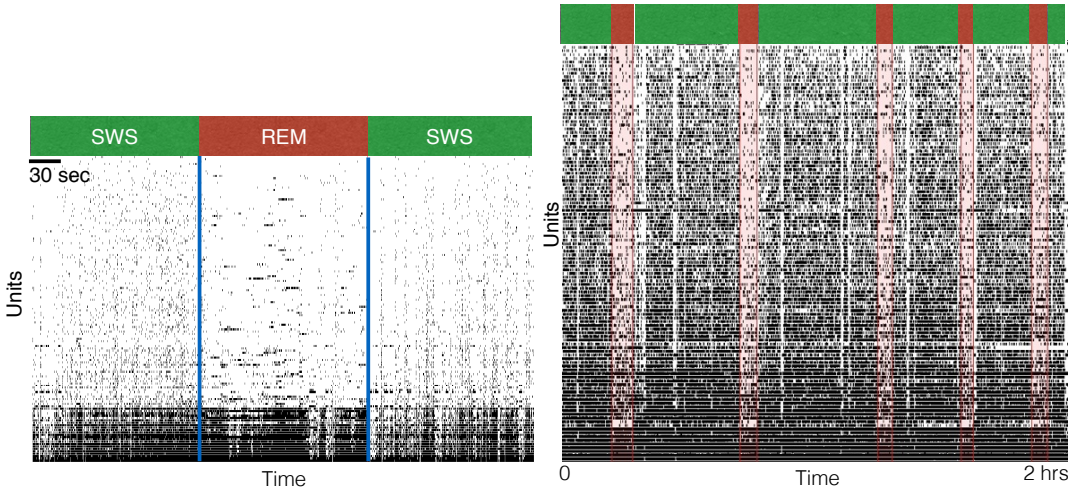


**Figure 1.3:** Activity in CA1 is strongly modulated by behavioral state. The left panels shows the differences on a short time scale. SWS raster plots exhibit vertical stripes due to ripple-associated population bursts. Spiking in REM is very sparse. Concentrated bursts of activity can be observed. The right panel illustrates the alternation between SWS and REM epochs over the course of two hours. Pyramidal units are shown at the top of raster plots, interneurons with very high firing rates at the bottom.

250 Hz range, drive synchronous bursts of activity in CA1 (see Figure 1.3).

Ripples occur during slow-wave sleep and quiet wakefulness, and are tightly coupled to the experience-dependent replay of awake activity. Entire spike trains recorded during awake behavior are replayed in a temporally compressed manner [28, 29, 35] during ripples. The reactivation occurs at a much faster speed than during the awake experience (by a factor of 10 to 20).

## 1.3 Electrophysiological recordings in the hippocampus

Different paradigms are applied to extracellular recordings in order to gather different kinds of information. Acute recordings allow for the sampling of several cells in quick succession, but severely limit the effective duration of data acquisition from an individual neuron. They are most commonly performed under anesthesia and head-fixation, a factor that restricts the usefulness in the study of natural behaviors. Chronic recordings, on the other hand, where the recording electrodes are permanently implanted in the brain, facilitate the monitoring of neural activity for weeks and months during natural awake behavior and sleep of the animal.

Given the supposed role of the hippocampus for memory formation, it is of paramount

importance to acquire data from neural populations and observe changes in their activity in response to new stimuli, physical environments and tasks. To achieve this goal, techniques to reliable identify and isolate individual neurons from extracellular recordings are of paramount importance. A vast trove of work has addressed the issues of spike detection, spike sorting and introduced measures to assess the quality of the output from these processing steps. Comparatively fewer studies describe methods to link identified single units across gaps in the recordings.

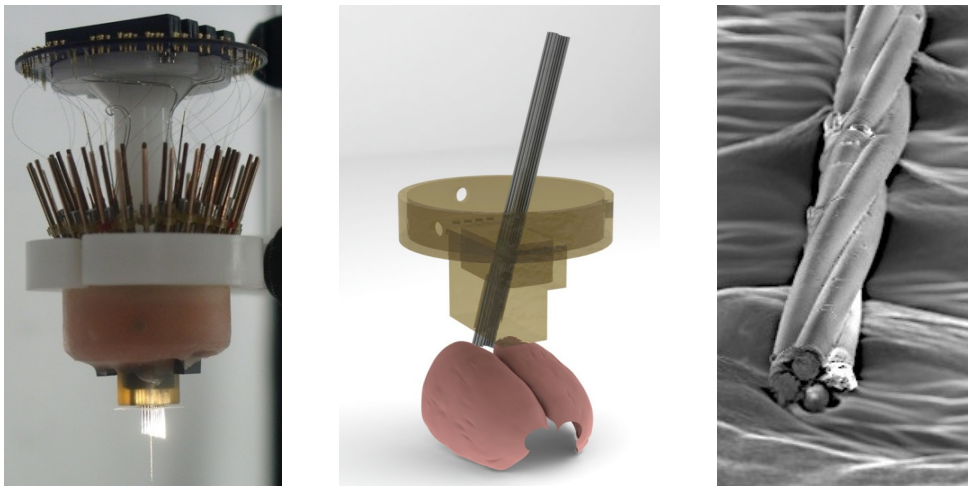

**Figure 1.4:** Left panel: Tetrode drive for chronic implantation. Electronic interface at the top, tetrodes visible as thin wires, posts with adjustable shuttles below. Plastic tip with guide tubes at the bottom. A single tetrode is protruding from a guide tube. Center panel: CAD model of a custom-designed tip. Right panel: SEM image of a tetrode. The four exposed and electroplated tips are visible at the bottom of the image.

Modern preparations for chronic data acquisition from tetrodes in rodents are based on drive bodies that house between sixteen and forty recording electrodes. Each electrode is attached to a shuttle assembly that can be raised and lowered via a ball screw in small increments — 80 µm per screw turn is a common thread pitch for this application. Guide tubes at the bottom of the drive body steer each electrode to its target coordinates in the brain. The stiffness of correctly manufactured tetrodes is sufficient to travel along a straight path after exiting from the guide tube. The drives we use for long-term recordings and custom-made tips to target specific areas are shown in Figure 1.4.

# Chapter 2

# Results

## 2.1 Neural activity of single CA1 units in REM sleep

Population activity in pyramidal CA1 neurons during rapid eye movement (REM) sleep differs significantly from that in slow-wave sleep (SWS) and awake periods (cf. Figure 1.3). Whereas slow-wave sleep is marked by ripple-associated population bursts, we rarely observed synchronized activity in multiple neurons during REM. Most units fired very sparsely and exhibited a single, relatively short mode in which the firing rate was significantly elevated. Outside of this mode, the cell was often silent, or spikes trickled at a very low rate.

### 2.1.1 Firing rates

Firing rates of principal cells in CA1 during REM sleep were low (median 0.096 Hz, interquartile range 0.129 Hz) and, for most cells, lower than in SWS (median 0.207 Hz, interquartile range 0.135 Hz). Firing rate distributions exhibited heavy right tails in both SWS and REM sleep (Figure 2.1(a)). In 60.7 % of cells, the distribution of firing rates in REM episodes was significantly different from the distribution in SWS episodes, and had a lower median (Wilcoxon rank-sum test, $p < 0.05$). In 4.6 % of cells, the distribution of firing rates in REM episodes was significantly different from the distribution in SWS episodes, but had a higher median (Wilcoxon rank-sum test, $p < 0.05$). The extent of rate modulation between SWS and REM varied heavily from cell to cell. The firing rate ratio $\frac{f(\mathrm{SWS})}{f(\mathrm{REM})}$ had a median of 2.2 but reached 7.8 at the 90th percentile, while outliers up to 25 were observed (Figure 2.1(b)), indicating strong attenuation in a subset of cells during REM sleep.

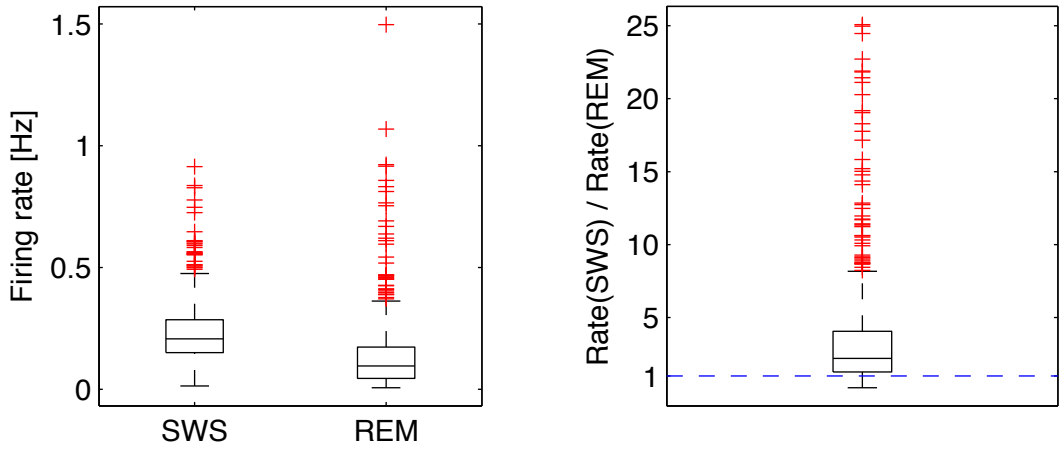


(a) Box plots of mean firing rates in REM sleep and SWS

(b) Mean firing rate ratio $\frac{f(\mathrm{SWS})}{f(\mathrm{REM})}$

**Figure 2.1:** Firing rates in REM sleep are very low, and significantly lower than in SWS.

### 2.1.2 Temporal activity profile

**Qualitative observations** Three features stood out in raster plots of REM sleep spike trains: complete absence of spiking, "trickling" of spikes at a very low rate, and brief bursts of elevated firing rates with durations of 5 s to 10 s. In the majority of REM episodes, only a single occurrence of high firing rate bursts could be detected per unit. Figure 2.2 illustrates these features. Long periods of silence are evident in raster plots, and the firing rate histograms are dominated by the 0 Hz-bin, particularly in Figure 2.2(c). Trickling is most pronounced in 2.2(a) outside of the activity mode, and in 2.2(b). It is represented in firing rate histograms in bins below 3 Hz to 5 Hz. A single mode of high activity is evident in Figures 2.2(a) and 2.2(c) in the form of 15 s-intervals with firing rates of up to 15 Hz. Activity modes can also be distinguished in firing rate histograms as additional density modes in the range above 10 Hz.

**Parametric model of temporal profile** In order to gain a better understanding of firing patterns observed in REM sleep, a terse, quantitative description of the major features was derived from spike trains. The accurate representation of millisecond-scale dynamics was neglected in favor of descriptors for the evolution of firing rates on longer intervals of several seconds. Based on the qualitative features described above, a mixture model was chosen to

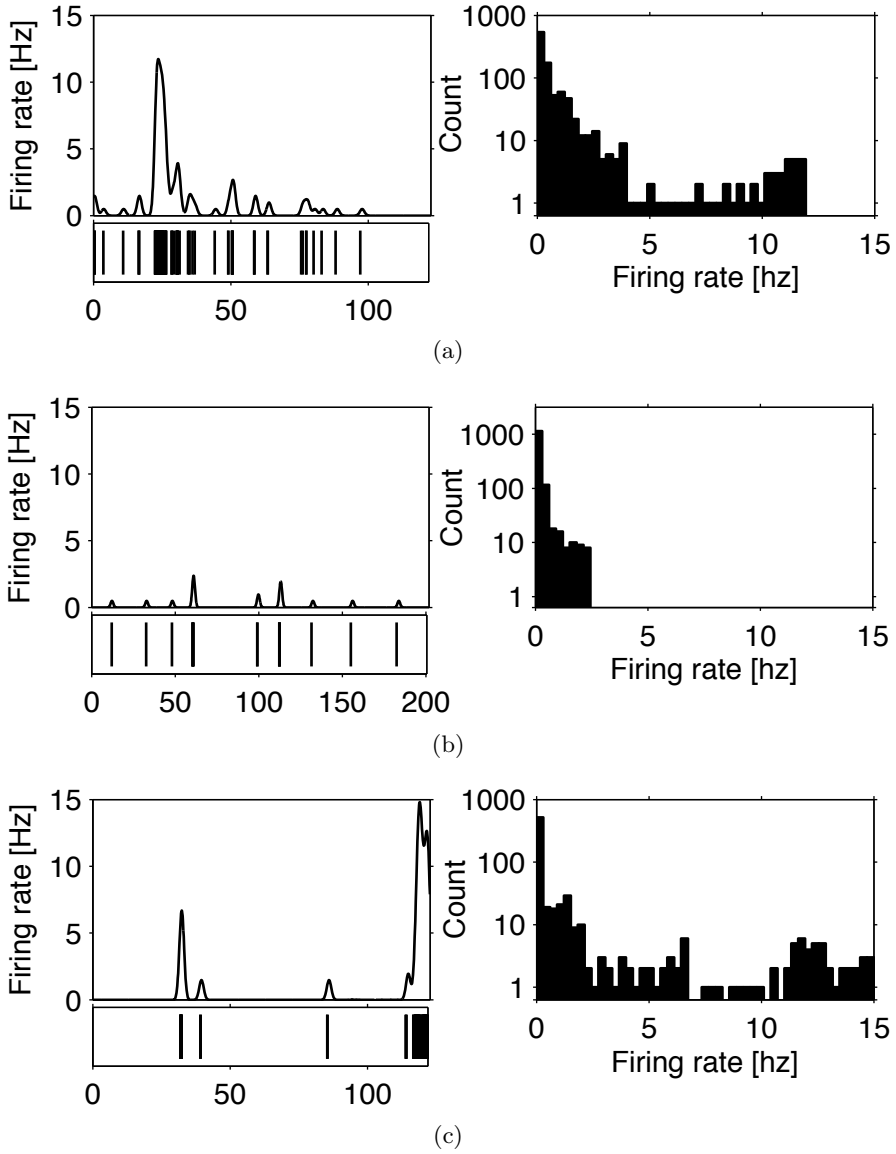


**Figure 2.2:** Major features of spike trains in REM (periods of total silence, trickling of spikes at a low rate, and brief bursts with high firing rates) are illustrated in three examples. Left panels: raster plots with tick marks for each spike and smoothed firing rates (Gaussian window, $\sigma = 5\,\mathrm{s}$). Right panels: Firing rate histograms of smoothed traces in the left panels.

model the temporal firing rate profile.

The mixture model is supported on $[t_{\text{start}}, t_{\text{end}}]$, the interval of the REM episode under consideration. The probability density function of spike emission on a unit interval is modeled as the weighted sum of uniform distributions. The first component of the mixture model always spans the entire REM episode but has, like all other components, adjustable weight. The model is parameterized by the triplet $(w_k, t_{k,1}, t_{k,2})$ for each component $k$: its weight $w_k$ and its support $[t_{k,1}, t_{k,2}]$, a subinterval of the model support $[t_{\text{start}}, t_{\text{end}}]$. Mixture models were fitted using an iterative optimization routine that adds components until the model distribution function and the distribution function of the counting process derived from spike times pass the Kolmogorov-Smirnov test at a significance level of 0.05. A more detailed description of the fitting process is given in Section 4.3.4.

Several examples of REM spike trains and their respective fitted mixture models are shown in Figure 2.3. The probability density of the model aims to match the smoothed firing rate function (Gaussian window, $\sigma = 5\,\text{s}$). Trickling is captured by the uniform background component that is present in all fitted models (see Figures 2.3(a) and 2.3(b)). Modes with high firing rates introduce an additional component with high weight into the model (Figures 2.3(a) and 2.3(c)). Some REM episodes have more than one mode, as shown in Figure 2.3(d), and require more than two components to accurately reflect spike train dynamics.

**Number and duration of mixture components** Uniform mixture models were fit to spike trains of pyramidal CA1 units observed over the course of 154 REM episodes (5h 54m total duration) in three animals. In each animal, single units were isolated and held for several days. For each cell that was recorded from and that fulfilled isolation criteria defined in Section 4.3.2, every REM episode of at least 60 s duration was counted as an observation.

A histogram of the number of mixture components $K$ is shown in Figure 2.4. As 14 % of REM spike trains do not contain any spikes, these instances were counted in a separate bin as having zero components. Two-component models were fitted to 41.8 % of spike trains, whereas a sufficient fit could be accomplished with a single component in 27.0 % of cases. Only 17.2 % of mixture models required three or more components. The histogram in Figure 2.4(a) includes spike trains with fewer than 5 spikes, i.e., an average firing rate below 0.07 Hz. We also investigated how this distribution changes when only episodes with

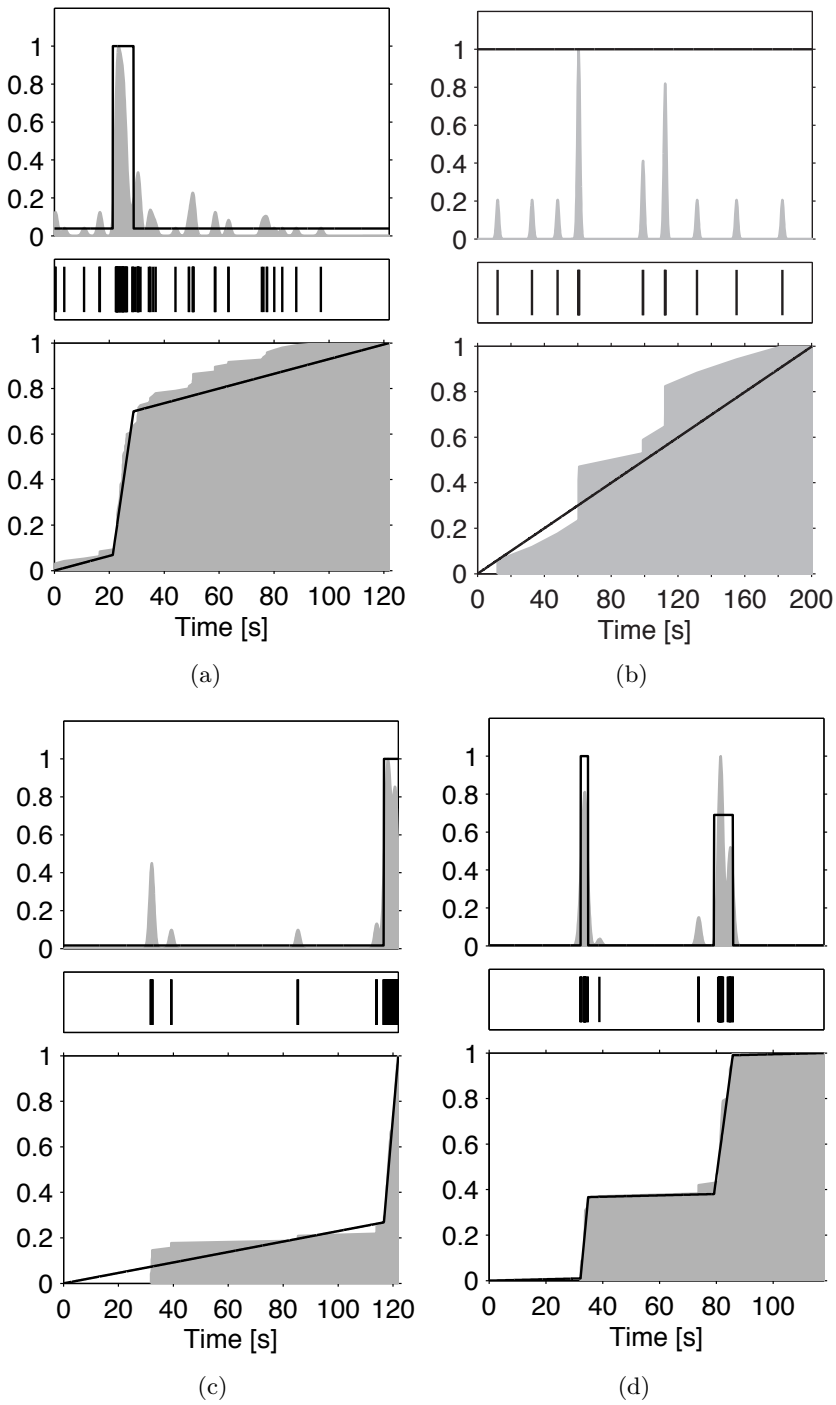


**Figure 2.3:** Uniform mixture models accurately capture spike train dynamics at coarse timescales. Top panels: smoothed firing rates, normalized to their peak value, shaded in gray. Spike emission probability densities of fitted uniform mixture models as black lines (scaled to match firing rates). Center panels: spike rasters with tick marks for each spike. Bottom panels: empirical cumulative densities of the spike train shaded in gray, cumulative densities of the mixture models overlaid as black lines.

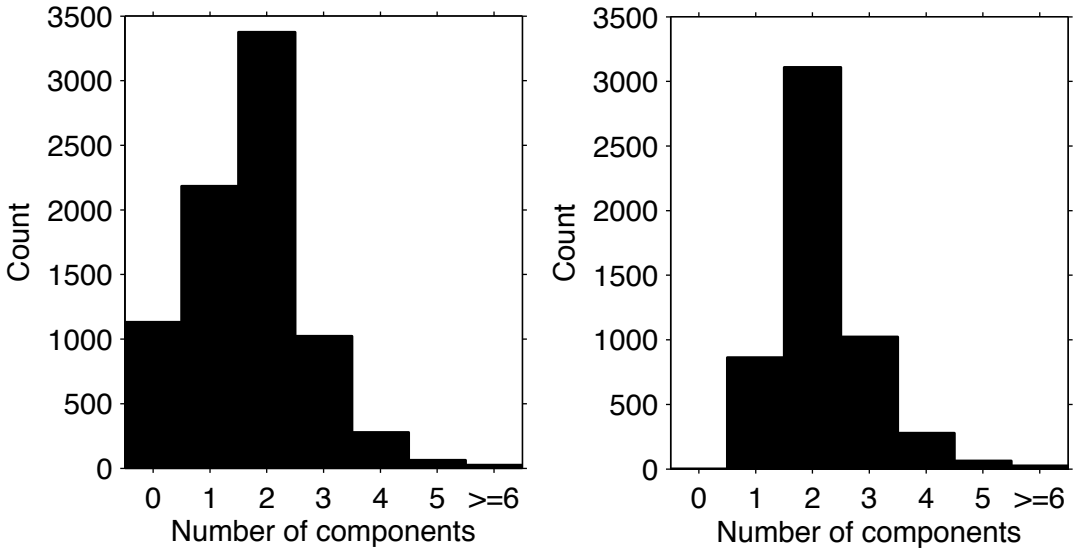


(a) Counting all episodes (b) Counting episodes with at least 5 spikes

**Figure 2.4:** Most spike trains are well-fit by mixture models with two components, in particular when episodes with very few spikes are excluded. Counts are relative to 8079 (units × REM episodes)

a higher level of activity are analyzed. Model fits restricted to spike trains with at least 5 spikes emphasize that the majority of spike trains (58.0 %) are well-fit by a two-component model, as illustrated in Figure 2.4(b). These numbers are consistent with our previous observations identifying low-frequency background firing and a single high-frequency mode of activity in most spike trains. Next, we are addressing the question whether the activity mode is indeed dominant, i.e., whether it is a component of short duration and high weight that significantly exceeds the weight of other mixture components.

**Presence of a dominant component** We call a mixture component dominant if it has short duration relative to the REM episode but high weight. Using the normalized component duration $d_k = \frac{t_{k,2}-t_{k,1}}{t_{\text{end}}-t_{\text{start}}}$, we introduce the *simple score*

$$s_k = \sqrt{w_k} - \sqrt{d_k}. \tag{2.1}$$

Components with high weight but short duration receive a high simple score, wider components a slightly lower score. If the model has a single background component, it receives the score 0. The simple score, a quantity that characterizes individual mixture components,

is used to derive a contrast score $c$ for the entire model:

$$c = \max_k s_k - \max_{j,j\neq k} s_j \tag{2.2}$$

In models with exactly two components, the distribution of the contrast score is heavily shifted towards high contrast scores (median: 1.27, interquartile range: 0.59). If models with two or more components are considered, a new mode emerges in the range $c \leq 0.3$, and indicates that activity modes have similar simple scores if there is more than one mode.

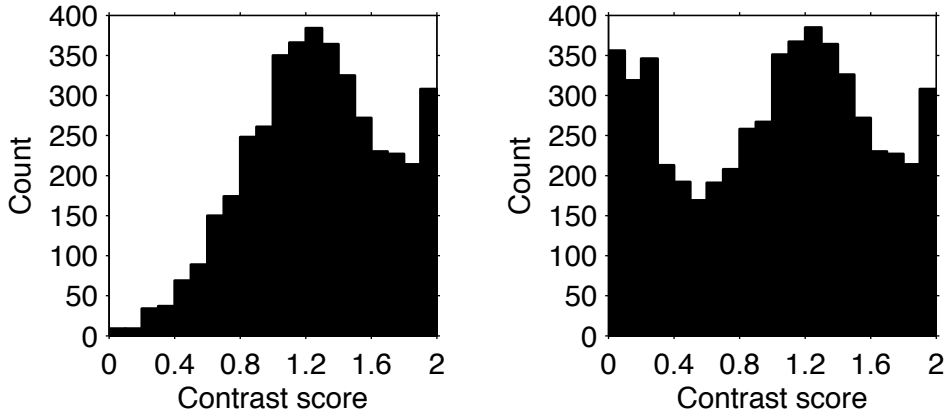


(a) Counting models with 2 components

(b) Counting episodes with at least 2 components

**Figure 2.5:** Activity modes have high contrast scores compared to the background.

The duration of the dominant mode, i.e., the model component with the highest simple score, is shown in Figure 2.6. Most dominant modes are indeed short relative to the episode duration (median duration: 4.2 s, median relative duration: 3.9 % of the episode length). In 78.6 % of models with at least two components, the dominant mode was concentrated on 15 % or less of the model support.

In summary, we have established that most pyramidal cells emit spikes at a significantly elevated rate for periods that are short in relation to the duration of the REM sleep episode. Furthermore, in the majority of episodes, there is exactly one such burst. These observations led us to hypothesize that the dominant activity mode may have functional significance at the population level.

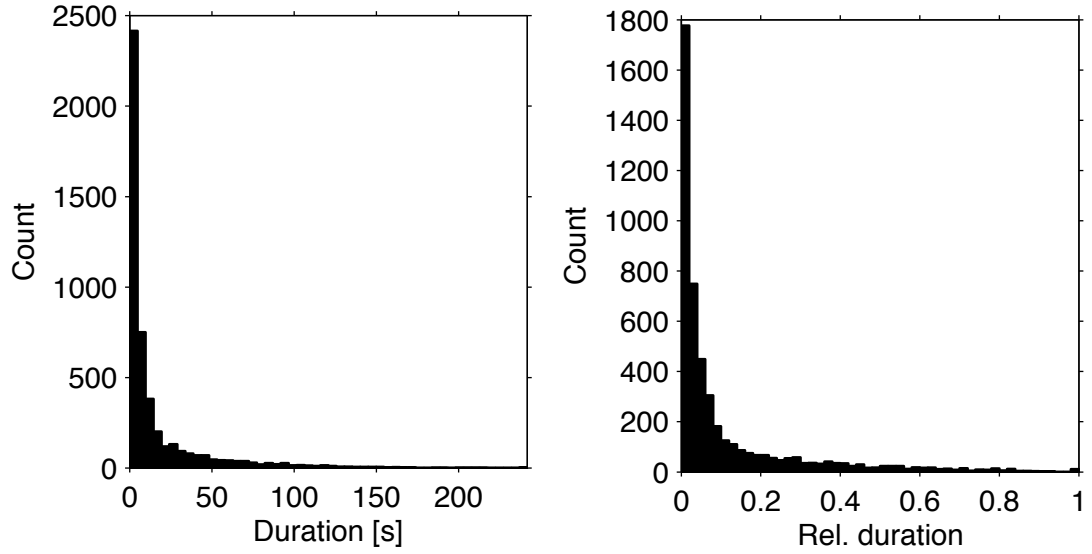


(a) Measured in seconds

(b) Measured relative to the REM episode duration

**Figure 2.6:** The majority of dominant mixture model components are of less than 5 s duration. The distribution of durations is approximately exponential. The dominant component is the one with the highest simple score; only models with at least two components are considered.

## 2.2 CA1 population activity in REM sleep

Activity in the hippocampus undergoes spatial modulation that can be observed at the level of single cells. Only analysis at the population level, however, reveals how a representation of an entire environment is created. The temporally patterned activity of a place cell population constitutes a trace of the animal's path through that environment. Concise timing relationships during ripple-associated replay events are believed to make these traces more persistent through neural plasticity mechanisms [28]. The results in Section 2.1 demonstrate that spiking during REM is extremely sparse and often concentrated into a single, brief bout. The temporal relationships between these modes at the population level determine whether neural plasticity mechanisms can become active in REM sleep. To address this question, we examined the firing order for features that are repeated in multiple REM episodes.

### 2.2.1 Ordering REM episodes by their dominant mode

Uniform mixture models were fit to the same database of REM episodes used in Section 2.1.2. For each spike train, the time at the center of the dominant mode — the

mixture component that maximizes the simple score $s_k$ — was calculated. This time was used to impart a sequential order on the spike trains belonging to a single REM episode. Spike trains with zero spikes were excluded from this procedure. Models with a single uniform component were moved to the end of the ordering. Raster plots of several reordered REM episodes, as well as false-color images generated from their mixture model probability densities, are shown in Figure 2.7. It is evident that the progression of modes with high spike density from cell to cell forms a diagonal band of dense spiking that spans the entire REM episode. There is a strong contrast between the spiking band and surrounding areas, where spiking is rare.

The temporal distribution of dominant modes over the REM episode appears roughly uniform, but also exhibits periods of agglomeration. In 21.4 % of episodes, the uniformity hypothesis had to be rejected (chi-square goodness of fit test, $p < 0.05$). Uniformity in a pooled dataset, created by aggregating the normalized incidence times of all 4764 dominant modes, also had to be rejected (chi-square goodness of fit test, $p = 1.92 \times 10^{-6}$ and one-sample Kolmogorov-Smirnov test, $p = 6.76 \times 10^{-6}$). Periods of stronger input drive or release of inhibition may be responsible for the concentrated appearance of activity modes in multiple cells. This point also relates to the underlying question of whether the observed spike train properties result independently from the properties of individual cells, or whether there is a mechanism at the population level that controls the appearance and timing of activity modes.

### 2.2.2 Repeated patterns in the temporal ordering

We investigated whether the sequential order of dominant modes is ever repeated in multiple REM episodes. Evidence for such repetitions would be consistent with the replay of previous experience and memory consolidation concepts. The sequential order defined by the occurrence of dominant modes was extracted from one *source* REM episode, and then applied to all other *candidate* episodes within the same dataset. An example outcome of this procedure is demonstrated in Figure 2.8. The distinct diagonal structure vanishes when the permutation of the source episode is applied to candidate episodes. The uniqueness of the order is also shown in Figure 2.9 for eleven REM episodes. The only consistent structure is visible along the diagonal, where spike trains are permuted by the order extracted from the same episode. Off the diagonal, where the order derived from one episode is applied to

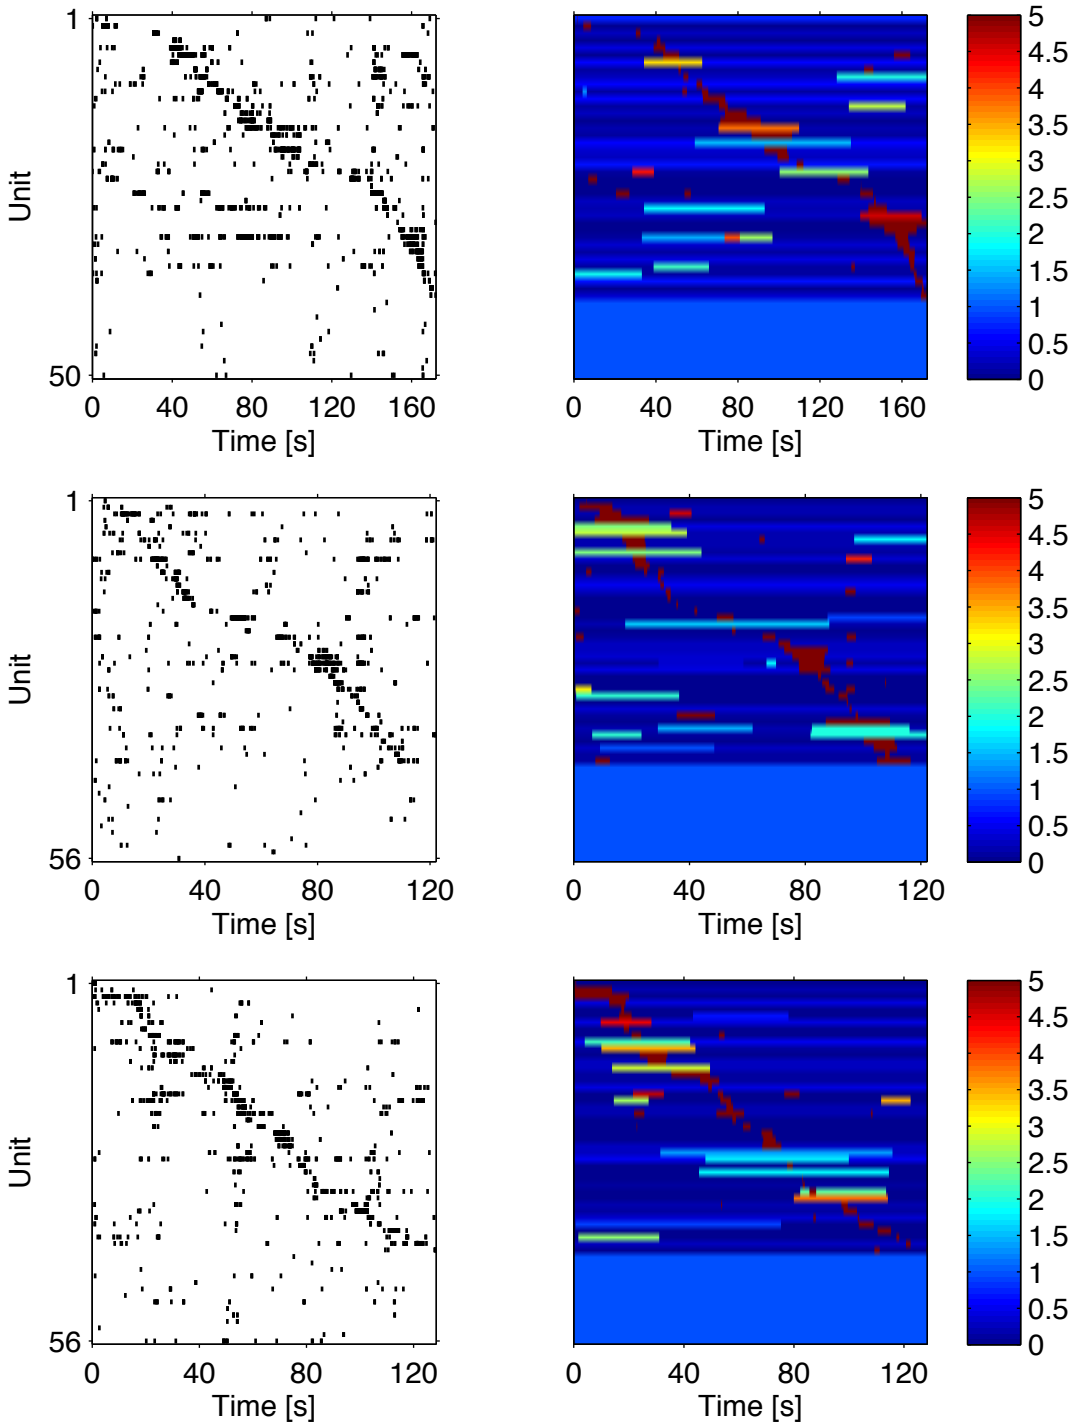


**Figure 2.7:** The timing of dominant modes defines a natural order for REM episodes. Spike trains for three representative episodes after applying the order are shown in the left panels. Probability densities of mixture models normalized relative to the density of a uniform distribution in the right panels. Models with a single component were moved to the end of each sequence.

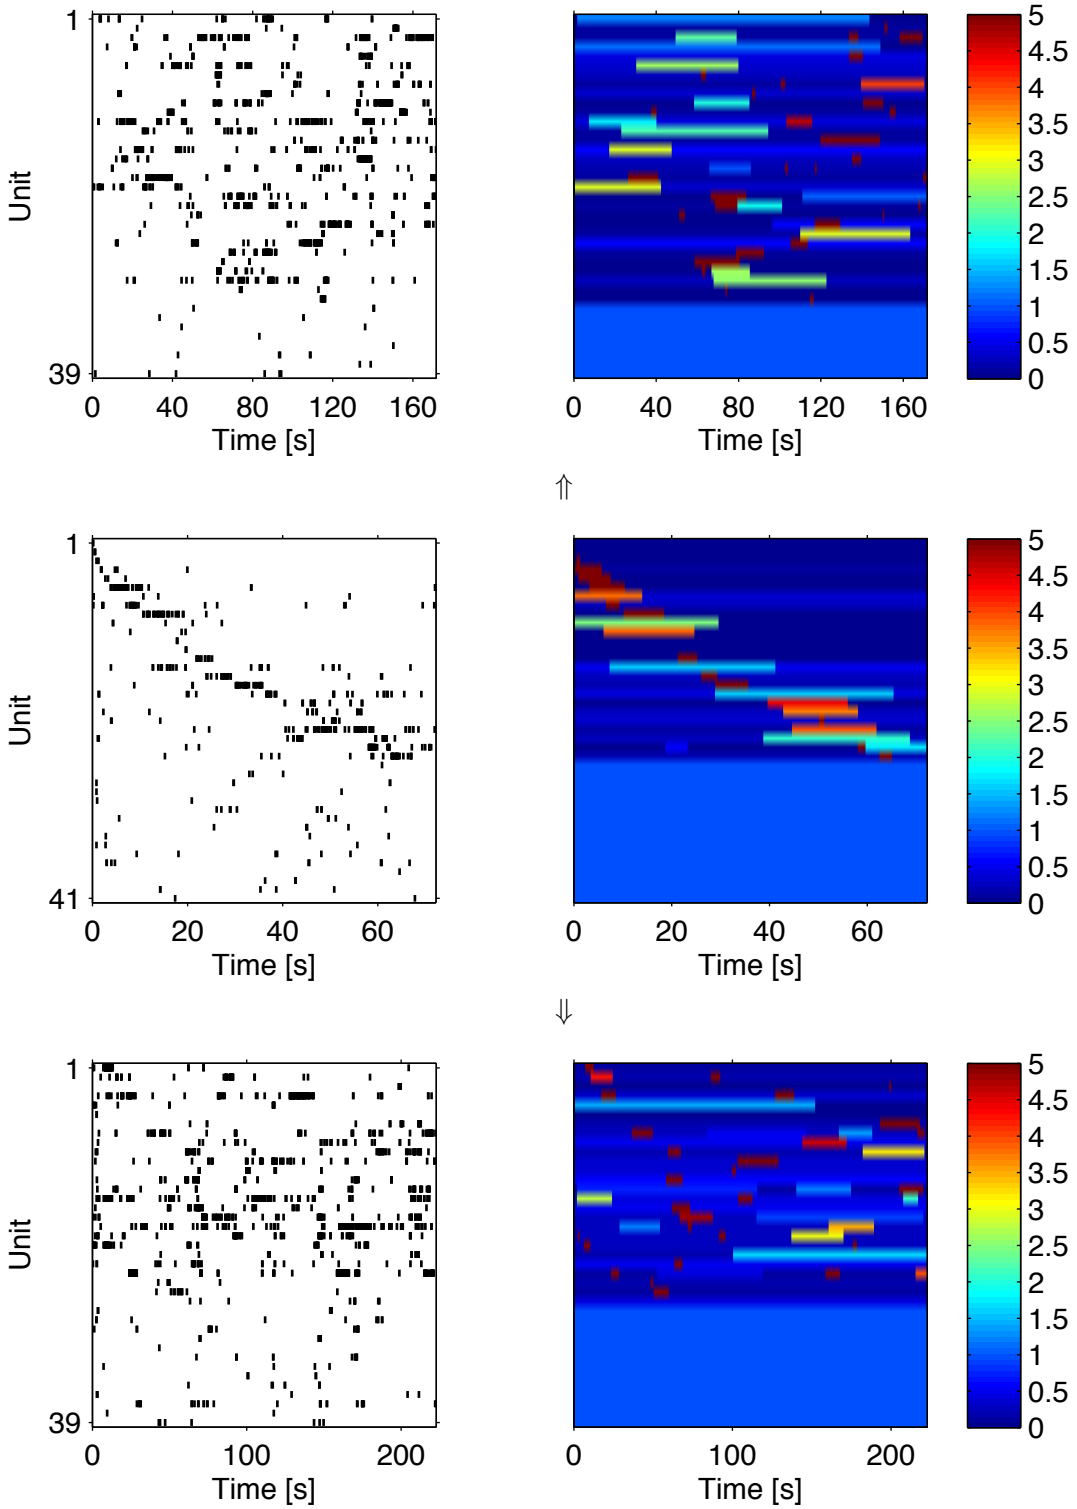


**Figure 2.8:** The order defined by the dominant mode is not repeated. When the sequential order inferred from one REM episode (middle panel) is applied to the preceding (top panel) and subsequent (bottom panel) episodes, the diagonal structure vanishes almost entirely.

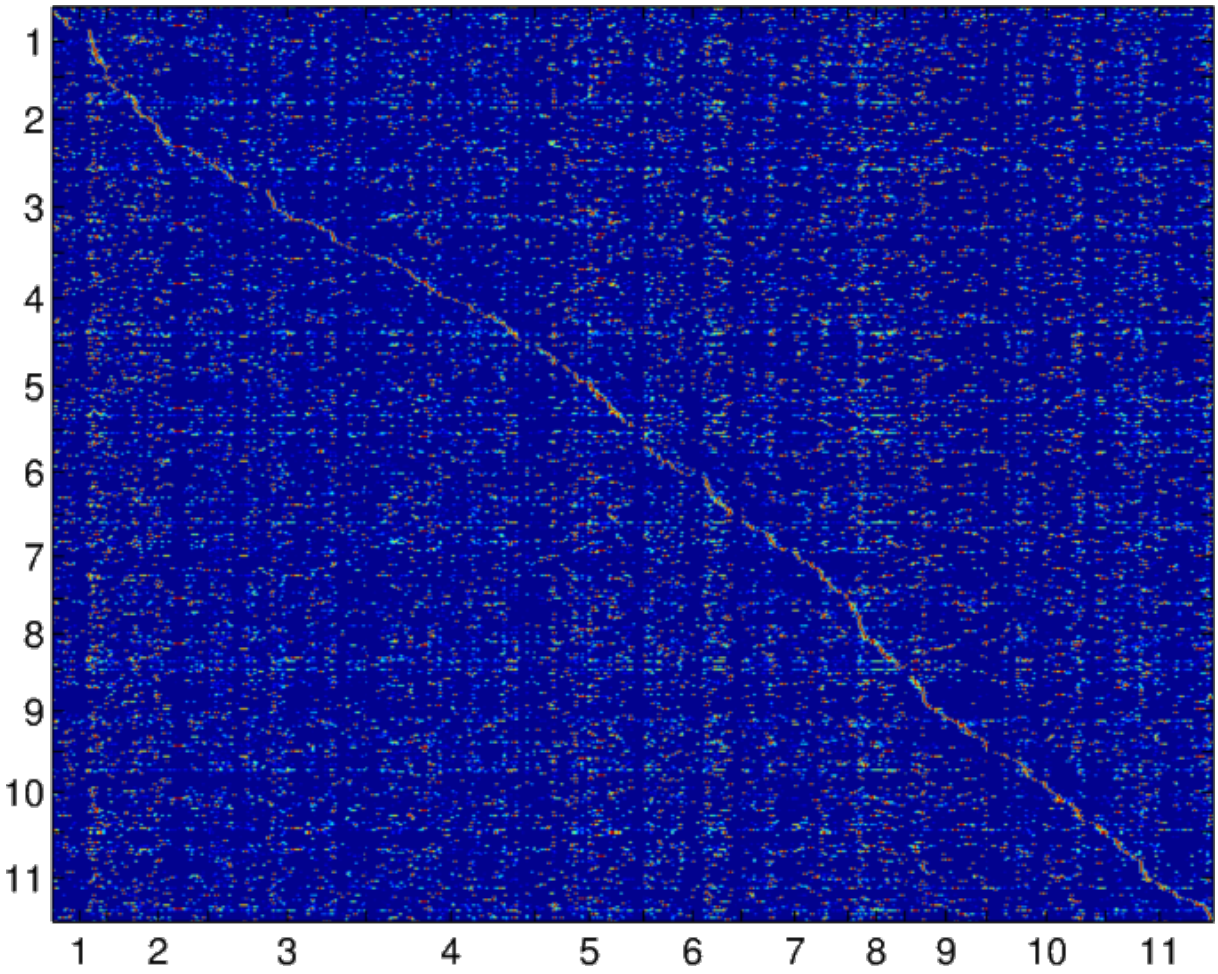


**Figure 2.9:** The structure of dominant modes is unique to each episode. Applying the order inferred from one episode to any other episode does not recover the diagonal structure. Shown is an $11 \times 11$ matrix where the submatrix at location $(i, j)$ contains the smoothed firing rates of episode $j$ permuted by the order of dominant modes in episode $i$.

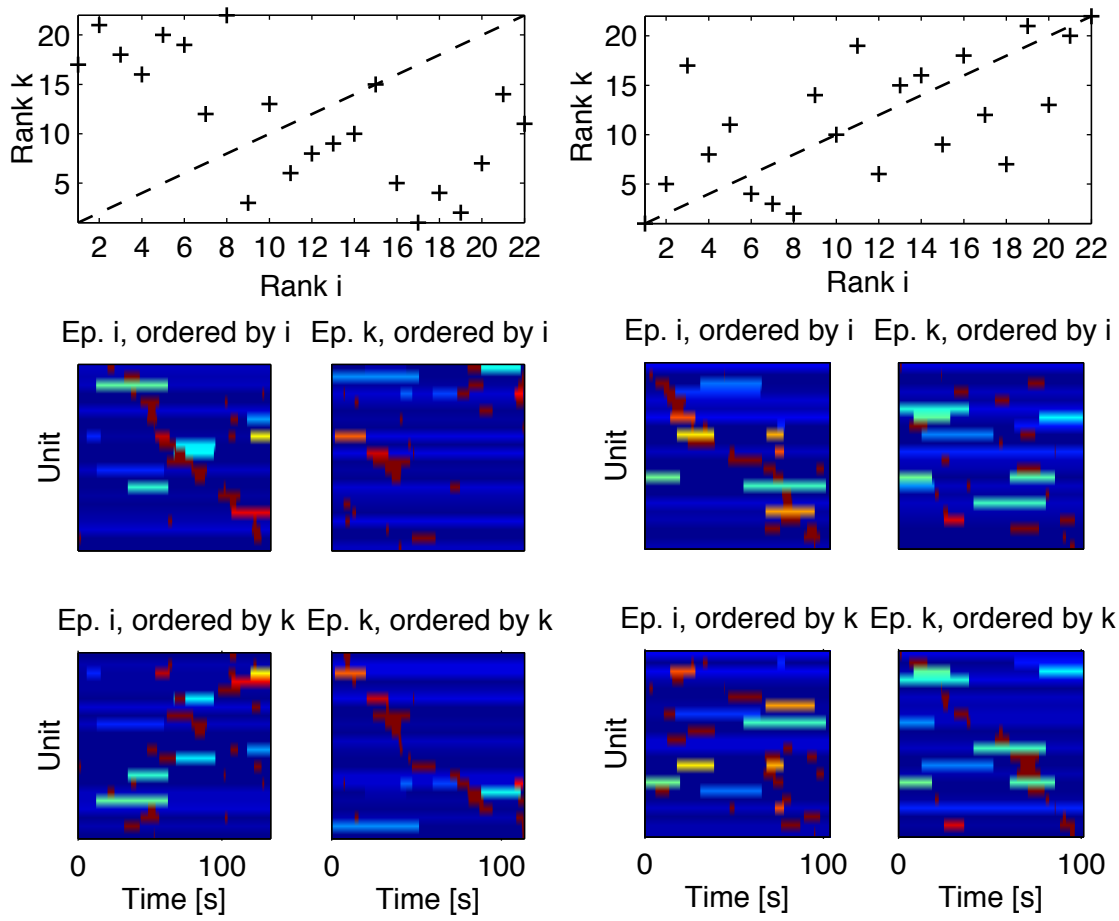


**Figure 2.10:** Two pairs of REM episodes where the dominant-mode-induced orders are significantly rank-correlated ($\tau = -0.43$ for the left pair, $\tau = 0.45$ for the right pair). Top panels: Scatter plot of spike train ranks. Bottom panels: Mixture model densities when applying the native and non-native order.

a different episode, no streaks of sequential activation can be observed.

While the order was not found to be highly conserved in different episodes, this observation does not exclude weaker effects, e.g., the partial ordering of a small subpopulation that persists for several REM episodes. We approached this point with two different methods. First, we focused on the question whether significant rank correlations exist between pairs of episodes. However, the null hypothesis of no correlation could only be rejected for 24 out of 2023 pairs (1.2 % using Kendall's tau rank correlations, $p < 0.01$). For significant pairs, the rank correlation strength was in the range $0.32 \leq |\tau| \leq 0.59$. Two examples of episode pairs with strong rank correlations are shown in Figure 2.10. We performed a second analysis to test whether the dominant modes of pairs of neurons appear in a particular order with higher frequency than expected. To maximize the power of this test, we relied on

our methods for linking units across multiple recordings (cf. Section 4.3.3) that allow us to incorporate a large number of REM episodes. If the order of unit activation were random, the relative occurrence of activity modes in two units would follow the distribution of a fair coin toss when one observation is made per episode. Hence, a sign test was used to detect pairs that heavily deviate from the expected outcome. A total of 5900 pairs of well-isolated units were observed for varying numbers of REM episodes (25th percentile: 4, median: 7 and 75th percentile: 16 after excluding all observations without spikes or without a dominant mode). The null hypothesis of a fair coin toss was rejected for 2.7 % of pairs (158 out of 5900) at $p < 0.05$. After accounting for the fact that the sign test has insufficient power to reject the null hypothesis for fewer than six observations and excluding those pairs, the null hypothesis was rejected for 4.5 % of neuron pairs (158 out of 3505).

The data presented provide little evidence to the hypothesis that structure in the order of unit activity is preserved across multiple episodes of REM sleep. Our analysis has found significant rank correlations between episodes to be rare, refuting the idea that unit order is ever repeated at the population level or even partially conserved. Examination of the firing order of pairs of neurons has also failed to uncover motifs that repeat at higher frequency than expected under the assumption of random permutations. Without concordant pairs of neurons, there is also no evidence for repeated patterns that involve larger populations of neurons. Showing conclusively that the order is indeed random in every single episode is unattainable given the vast space of permutations and the number of REM episodes that can be monitored with current technology. Yet, we have found no evidence that particular sequences of activity — consistent with the replay of awake experiences — are favored and reactivated repeatedly in REM sleep.

### 2.2.3 Validation of dominant-mode-based order

In order to validate that the results presented here are not a spurious result of using the dominant mode in fitted uniform mixture models as the sorting criterion, we considered two alternative methods. All methods are aimed at using the incidence times of high firing rate bursts to define a sequential order for spike trains. They differ in the specific method used to map the spike train to a rank:

- **Peak rate** Spike trains were converted into firing rate traces by binning at 100 ms

resolution and smoothing (Gaussian window, $\sigma = 5\,\mathrm{s}$). The rank of a spike train was determined by the incidence time of the peak firing rate.

- **Multidimensional scaling (MDS)** Spike trains were binned at 1 s resolution, and cross-correlations between all pairs of spike trains were calculated at lags up to $\pm(t_{\mathrm{end}} - t_{\mathrm{start}})$, corresponding to the duration of the REM episode. The lag values at which the pairwise correlations reached their highest magnitude were assembled into a distance matrix. All spike trains were mapped into a one-dimensional space via multidimensional scaling of the distance matrix. Sorting the coordinate vector yielded ranks for all spike trains.

Spike trains in 154 REM episodes were ordered according to three methods (mixture model, peak rate and multidimensional scaling). Similarity of the permutations was quantified using Kendall's tau rank correlations, and is shown in Table 2.1. The calculated rank correlations between the different sorting criteria were very high (mean between 0.83 and 0.87), and show that extracting the dominant mode from fitted mixture models results in permutations that are very similar to ones derived directly from smoothed spike trains.

| | mixture model | peak rate | MDS |
|---|---|---|---|
| mixture model | – | | |
| peak rate | $0.86 \pm 0.08$ | – | |
| MDS | $0.83 \pm 0.11$ | $0.87 \pm 0.10$ | – |

**Table 2.1:** Kendall's tau rank correlations for three different sorting methods

### 2.2.4 Replay in REM sleep

The data that have been presented so far do not support a role for REM sleep in memory consolidation via repeated reactivation of awake activity patterns. Bursts similar to those seen during place field traversals are too infrequent, and the relative temporal order of these bursts was shown to be indistinguishable from random permutations.

However, prior work by Louie and Wilson has argued in favor of coarse timescale replay of multi-unit activity patterns recorded during awake behavior in REM sleep [30]. To support this claim, multi-unit spike trains from REM sessions of more than 60 s duration were converted into a template matrix and compared to spiking of the same units while the animal was performing a spatial task in a circular environment.

**Methodology** REM spike trains from $n$ single units were binned at 1 s resolution and smoothed with a Gaussian window ($\sigma = 1.5$ s) to create a template matrix of dimension $n \times d_{\text{REM}}$, where $d_{\text{REM}}$ is the duration of the REM episode in seconds. Spike trains from the same $n$ cells recorded during awake behavior were binned and smoothed in an identical manner to form a larger $n \times d_{\text{RUN}}$ matrix. The template REM matrix was slid across the RUN matrix in 1 s-increments, and at each step a correlation coefficient between the REM template and the overlapping part of the RUN matrix was calculated. In order to quantify the significance of the correlation coefficient at each shift, a set of 50 shuffled versions of the REM template matrix was generated using four different shuffle methods (see Section 4.3.6). The correlation coefficients between the RUN segment at a particular time point and all shuffled versions of the REM template defined a distribution of expected correlation coefficients for each of the four shuffle methods. The means and standard deviations of these distributions allowed for the conversion of correlation coefficients to z-scores. Significance was evaluated using the lowest of the four z-scores. The analysis was performed for a range of temporal scaling factors (SF) from 0.3 to 3 to detect replay at faster or slower rates. Scaling factors greater than 1 correspond to slowed down replay during REM, i.e., the REM template matrix was matched to a RUN segment of duration $\frac{d_{\text{REM}}}{\text{SF}}$.

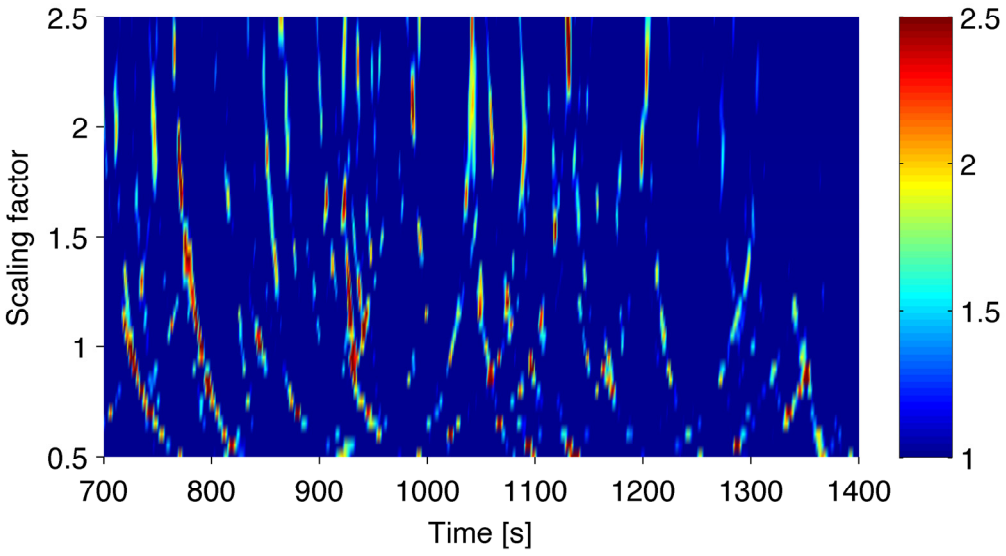


**Figure 2.11:** Map of z-scored correlation coefficients between a REM template and a longer RUN session. The time axis indicates the time point in the RUN recording where the REM template was matched. The scaling factor represents the amount of temporal compression applied to the REM template before matching.

**Replication of results** Following the methods described by Wilson and Louie, we performed the same analysis on our recordings. There were differences with regard to the behavioral paradigm as our animals performed an alternation task on a T-maze, or ran laps on a linear track (see Section 4.2). To obtain better statistics for shuffled spike trains, we evaluated the correlation coefficient for 10 000 instead of 50 replicates. Qualitatively, our findings were consistent with the reported data. We visualized the matrix of z-scored correlation coefficients for all possible shifts and scaling factors in Figure 2.11. In this representation, we identified islands of high correlation coefficients — putative replay episodes — by thresholding at a minimum z-score of 1. We observed a significant number of such events and the temporal distribution and shape of these islands was similar to that reported by Louie and Wilson. However, we were unable to reproduce the periodicity of putative replay episodes that is synchronized to trials of the behavioral paradigm. This discrepancy is likely due to the less periodic nature of place cell sequences in T-maze traversals compared to a circular environment in which the animal performs clockwise laps.

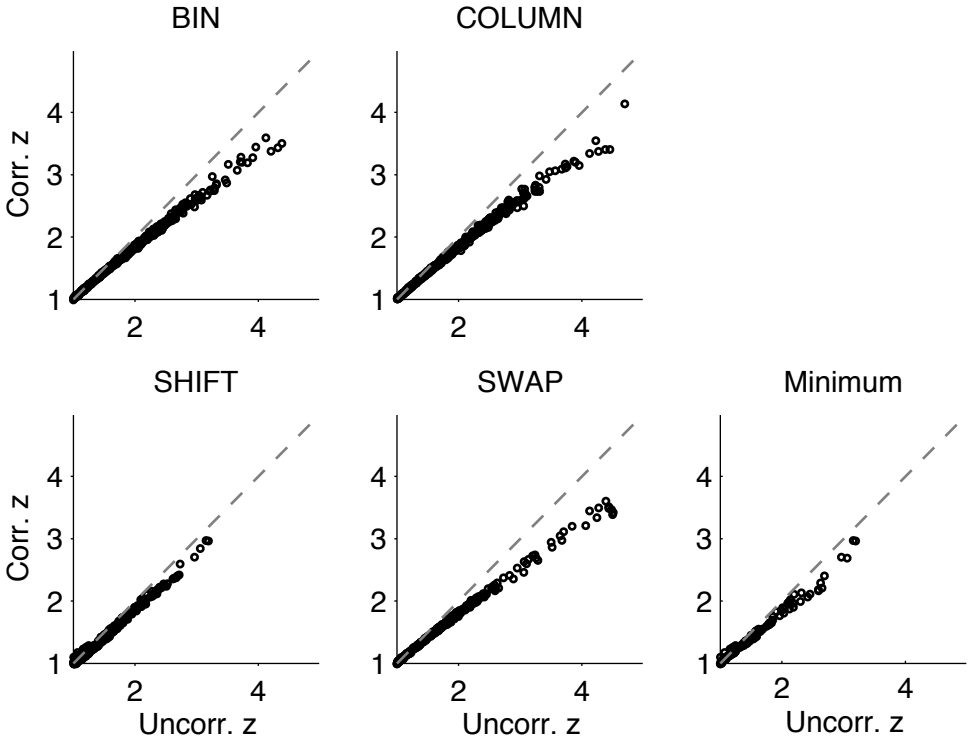


**Figure 2.12:** Correction for z-scores derived from shuffle distributions obtained by applying four different methods (BIN, COLUMN, SHIFT, SWAP). The bottom right panel shows the effect of correction on the minimum of the four scores, the quantity that is used to identify significant matches.

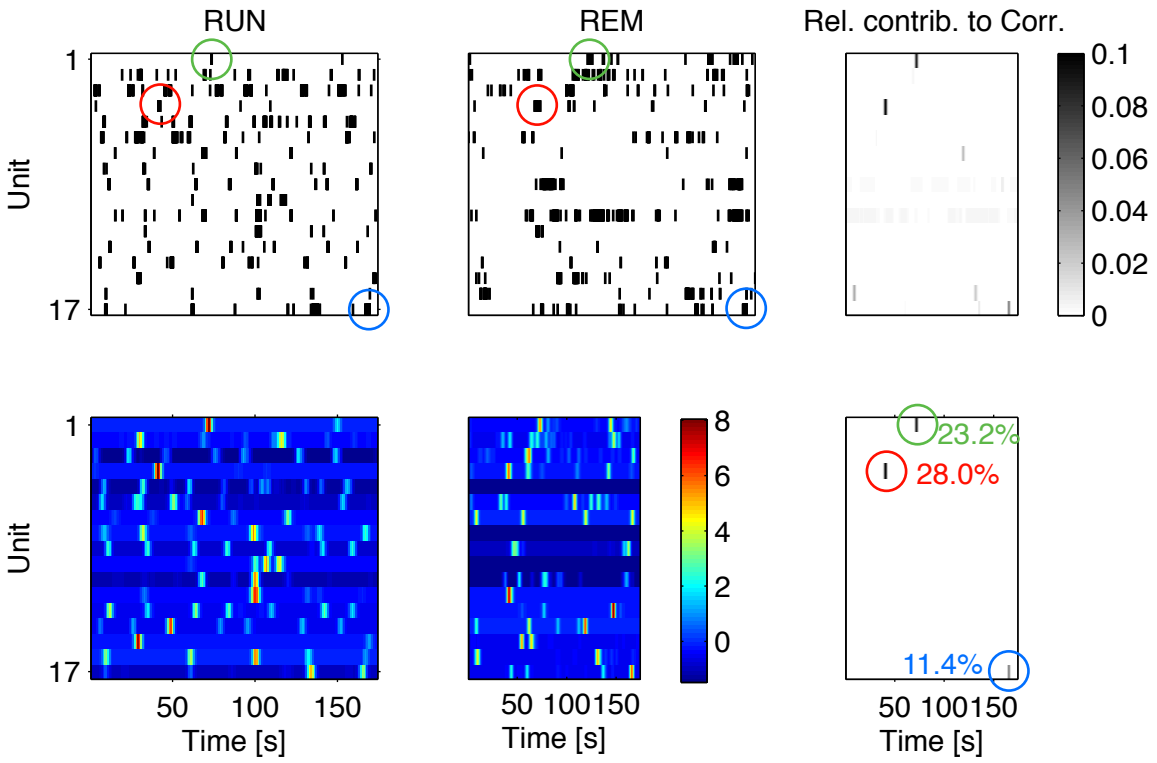


**Figure 2.13:** Multi-unit correlation scores can be heavily influenced by a small number of matching bins. In this highly significant match (z-score 3.3), three groups of overlapping bins (13 bins total, 0.4 % of bins) contribute more than 60 % of the correlation coefficient. Left column: Matched run segment. Center column: REM template. Top row: Raster plots with one tick per spike. Bottom row: Binned, smoothed and normalized spike counts used for the calculation of the correlation coefficient. Right column: Relative contribution to the correlation score. The three highest contributions are highlighted in red, green and blue.

**Criticism** The distributions used as the test statistic for the calculation of z-scores were found to have significant positive skewness. We inspected the shuffle distributions for 1528 data points (tuples of shift and scaling factor) and 10 000 replicates per data point. All 1528 distributions had positive skewness for each of the four shuffle methods. Strong positive skewness of 0.3 or larger was present in almost all distributions (BIN: 99.5 %, COLUMN: 99.8 %, SHIFT: 100.0 %, SWAP: 85.1 %). The strong positive skewness in the shuffle distributions causes the standard deviation to be underestimated and exaggerates z-scores. With 10 000 shuffles at each data point, we applied the probability integral transform to convert the skewed distribution to a normal distribution and reestimated z-scores. The difference between corrected and uncorrected scores is shown in Figure 2.12.

Matches with high scores ($z \geq 2.33$, the critical value for a one-tailed z-test at the 0.01 significance level) were found to be dominated by a very small number of bins. An example

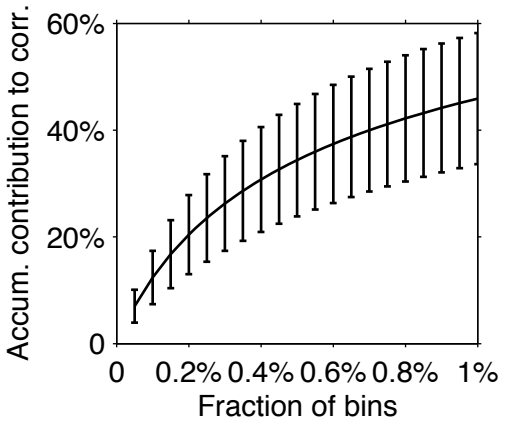


**Figure 2.14:** A small number of bins dominate the multi-unit correlation coefficient. The cumulative relative contribution is plotted against the fraction of bins. Error bars indicate $\pm 1$ standard deviation. Statistics over 5196 matches that exceed a z-score of 2.33.

of a highly significant match ($z = 3.3$, $p < 4.9 \times 10^{-4}$) is shown in Figure 2.13. The critical observation is the absence of similar firing patterns in most of the cells for the entire three-minute episode. Yet, thirteen bins in three groups that contain spikes in both the REM template and the RUN segment are sufficient to bias the correlation score and mark the match as a highly significant replay episode. The cumulative contributed correlation score as a function of the number of bins was estimated from a set of over 5196 significant matches (see Figure 2.14). On the average, the one percent of bins with the largest contributions account for 45.9 % of the correlation coefficient.

**Conclusion** Despite differences in the behavioral paradigms used for awake recordings, we succeeded in replicating the fundamental results reported by Louie and Wilson. There are indeed periods where the multi-neuron correlation coefficient, as defined by the authors, reaches values that are high enough to be considered significant when transformed into a z-score. However, the reference distributions that were created by correlating shuffled REM spike trains with RUN data are severely skewed, resulting in exaggerated z-scores. If the comparison at each time point and temporal scaling factor were phrased in terms of a hypothesis test, the number of matches that exceed the critical value should decrease after compensating for the skewed distributions.

We are also concerned that the multi-unit correlation coefficient is easily biased by a small number of bins with similar spike counts in RUN and REM. The spike patterns of most cells in putative REM replay events that receive a significant correlation score bear

little resemblance to the identified segment in the RUN session. Features that do match are usually confined to short periods and few cells. While there is no reason to discount the validity or significance of these matches, the qualitative features are not in agreement with the concept of replay. A correlation measure that penalizes mismatches between RUN and REM more severely may be better suited to the detection of replay events.

# Chapter 3

# Discussion

Rapid eye movement sleep is often referred to as "paradoxical sleep" due to the inconsistency of an outwardly deep sleep state (muscle atonia, reduced control of breathing and heart rate [17, 50]) and neural activity that strongly resembles awake patterns [41]. The function of REM is not well-understood to date, even though it has received a lot of attention since the discovery that most dreams occur during REM. A large number of animal studies point to a role of REM sleep for memory. In the two-stage model of memory formation, information is temporarily stored in the highly plastic circuits of the hippocampus and eventually transferred to cortical areas through the process of memory consolidation [52]. With external inputs shut down, sleep may be ideally suited to reconfigure circuits and rebalance synaptic weights. Experimental evidence for this hypothesis largely stems from two strategies: sleep deprivation experiments, and a comparison of neural activity in awake behavior and sleep.

Studies characterizing the effects of REM sleep deprivation have brought forth conflicting evidence [47, 27], and are hard to interpret due to the confounding effects of stress [22]. The surprising observation that total suppression of REM sleep caused by certain psychopharmaceuticals is not associated with persistent negative effects argues against an essential physiological role for REM sleep [54].

A complementary approach is to examine the structure of brain activity in naturally sleeping animals during episodes of REM. One study showed an experience-dependent reversal of the firing phase of hippocampal neurons with respect to the theta rhythm [40]. Since spikes occurring at different theta phases may change the direction of synaptic plasticity [23], this finding suggests an experience-specific reshaping of CA1 synapses. However, subsequent studies failed to replicate this finding [34]. Another study demonstrated

experience-specific replay of CA1 patterns during REM sleep [30]. As replay of spike patterns first experienced during awake behavior is an essential component of the consolidation process [31, 19], this finding would lend strong support to the involvement of REM in memory consolidation. However, the results of this study have not been replicated in full or built upon in recent work.

In this study, we quantified basic properties of the spiking patterns recorded from populations of pyramidal CA1 neurons over the course of many hours of REM sleep. Spiking in REM was found to be very sparse and at lower mean rates compared to slow-wave sleep. High firing rates are confined to brief bursts ($\leq 5\,\mathrm{s}$), of which only one is observed in the majority of episodes. The timing of this burst imposes a natural sequential order on the population of active cells. We found no evidence that this ordering is ever repeated for the entire observed population or a subset of cells.

This finding is at odds with the reported replay of CA1 activity in REM sleep [30]. Following the methods reported by Louie and Wilson, we were able to replicate the basic results and plots (cf. Section 2.2.4). However, we believe that the conclusions drawn from this study should be reevaluated. A z-scored correlation coefficient was used to quantify the significance of replay events. Severe skewness in the underlying distributions causes an inflation of reported z-scores. An inspection of the highest-scoring replay events revealed that the correlation coefficient is heavily dominated by a small number of matching time bins. In conclusion, the findings reported by Louie and Wilson are not necessarily inconsistent with our data and interpretation thereof. The similarity of behavioral epochs and REM episodes discovered by the correlative analysis method may not warrant the term "replay" though.

We propose the following parsimonious explanation for CA1 activity patterns observed in REM sleep: At a random point in time, pyramidal cells enter an ON state, emit a single burst of activity, and remain silent for the remainder of the epoch. A population of such cells gives rise to a random sequence of activity in each REM episode. The cellular and network underpinnings of this mechanism are currently not understood.

The sequential activation of cells provides for a limited amount of temporal overlap, and minimizes the number of spikes emitted in the window of plasticity ($\leq 50\,\mathrm{ms}$) for pairs of neurons. Downstream targets of CA1, for example the subiculum and the deep layers of entorhinal cortex, are unlikely to receive strong coincident inputs. Analysis of

cross-correlations between pairs of cells in awake behavior and REM sleep could provide additional evidence against a role for REM sleep in memory consolidation through replay.

*In vivo* whole-cell recordings, of pyramidal cells during natural sleep could open an experimental avenue to a deeper understanding of the system. In patch-clamp recordings it would be possible to distinguish the release of inhibition from excitatory currents as the driving factor behind the mode of activity. *In vivo* whole-cell recordings are difficult to perform, especially under conditions of natural sleep, but have been attained [44].

# Chapter 4

# Materials and Methods

## 4.1 Surgery and electrophysiological recordings

Electrophysiological recordings were performed with tetrode drives that house 24 individually adjustable tetrodes. Male Long-Evans rats (Charles River laboratories, age range 3 to 5 months, weight range 300 g to 450 g at the time of surgery) were chosen for drive implantation under stereotactical control. In two animals, all 24 tetrodes targeted the pyramidal cell layer of dorsal CA1 in the right hippocampus. In one animal, 12 tetrodes targeted area CA1 in the right and left hemispheres each, but only data recorded from the right hemisphere was used for analysis.

After surgery, tetrodes were slowly lowered towards the target area over the course of several days, relying on depth information from the microdrives and established electrophysiological signatures for guidance. In order to optimize unit yield in the thin and extraordinarily dense pyramidal cell layer of CA1, micro-adjustments in steps of 20 µm to 40 µm were carefully executed, and followed by rest periods to allow for tissue relaxation. Once unit yields of five to ten well-isolated pyramidal cells were achieved on most tetrodes, further adjustments were minimized or even completely halted for several weeks. This strategy was aimed at maximizing recording stability, i.e., the ability to track well-isolated units across many recording days.

Voltage signals were fed into unity gain buffers on the headstage (Neuralynx HS-27) and transported via flexible fine-wire tether cables to the recording hardware. In one recording setup, the signals were passed through analog filters with a passband from 1 Hz to 6000 Hz, amplified by a factor of 2000 and digitized by National Instruments PXI-4472 cards (24-bit resolution, −10 V to 10 V input range) at a sampling rate of 25 kHz. In a second hardware

configuration, signals from the unity gain buffer were digitized by National Instruments PXI-4498 cards[1] (24-bit resolution, −316.23 mV to 316.23 mV effective input range, AC-coupled at 0.1 Hz) at 25 kHz sampling rate. In both cases, electrode signals were referenced against a skull screw implanted above the contralateral cerebellum.

NIH guidelines for the use of laboratory animals were strictly followed. All procedures were reviewed and approved by the Caltech Institutional Animal Care and Use Committee.

## 4.2 Behavioral experiments

Experiments were performed in three different rooms that are laid out in a similar manner. The animal was housed in a sleep box, and kept there before and after behavioral experiments, as well as during adjustment or other periods of sleep. The walls of the sleep box were opaque, preventing the animal from seeing other parts of the room, including the tracks or open fields used during experiments.

Animals were kept in a state of water-deprivation, and rewarded during the experiments with small quantities of water. The total water consumption was carefully monitored and logged to keep the animal motivated and guarantee a healthy level of fluid intake. Track runs were performed on an elevated linear track with boxes at each end that house a port for water delivery. In the "linear track" paradigm, the animal was rewarded for each complete lap from the left end box to the right end box and vice versa. The second paradigm used for track runs was an alternation task on a T-shaped elevated track. Reward ports were located at the bottom of the stem, and at the end of the left and right arm respectively. Only the traversal order [left arm → base → right arm] triggered a reward at the right end box. Only the the traversal order [right arm → base → left arm] rewarded the animal at the left end box. At the base of the stem, a large reward quantity could be earned for the sequences [base → left arm → base] and [base → right arm → base]; a small reward quantity was delivered for any other traversal sequence that eventually led the animal into the stem end box. At any reward port, water was delivered only once after entering the box. In order to potentially earn a second reward, the animal had to leave the box, enter a different segment of the T-maze and return to the end box. Whether a reward was actually delivered was governed by the schedule described above. Reward delivery and logging of

[1] PXI-4498 cards are based on an oversampling delta-sigma ADC architecture with digital anti-aliasing filters that are automatically adjusted according to the selected sampling rate.

the animal's activity were performed by a custom state machine implemented in National Instruments LabView that received input from laser-based interrupters integrated into the tracks. Additional light-based interrupters were built into the water ports to detect licking activity.

In order to monitor the animal position, three color LEDs were affixed to the headstage in an L-shaped pattern. Digital overhead cameras (PointGrey Flea and Flea 2 models as well as Basler scout models) with a resolution of $1024 \times 768$ acquired frames at a rate of 30 fps and emitted one TTL pulse per frame to allow for the synchronization with the neural data stream. National Instruments Vision software was used to correct for geometric distortion of the varifocal wide-angle lenses and map from pixel space to physical coordinates. Given the mounting distance, camera resolution and distortion correction, LEDs were tracked with a best-case resolution of 2 mm to 3.5 mm.

## 4.3 Processing of neural data

Neural data was acquired with ADC hardware described in Section 4.1 using custom National Instruments LabView software, and streamed to disk for later offline analysis. Data processing and analyses were performed in MATLAB (MathWorks, Natick, MA).

### 4.3.1 Spike detection and sorting

Broadband voltage signals from tetrodes were upsampled by a factor of two to 50 kHz and filtered with a linear-phase FIR bandpass from 600 Hz to 6000 Hz, the frequency band that contains the most energy of typical action potential waveforms but excludes low-frequency field oscillations and high-frequency noise. In order to emphasize the rapid deflection that occurs at the beginning of action potentials, the discretized version of the nonlinear Teager-Kaiser energy operator (NEO) [26] was applied to the voltage traces:

$$\psi\left(x(t)\right) = x^2(t) - x(t-1)x(t+1) \tag{4.1}$$

The transformed signals were smoothed with a Gaussian kernel to reduce the incidence rate of multiple detections of the same spike due to noise. After taking the sample-wise maximum across the four channels, a threshold was applied, and all segments that exceeded

the threshold were considered as putative spikes. Within each super-threshold segment, the sample index of the maximum value was designated as the time of occurrence of the spike, and used to align the waveform. The waveform was taken from the signal that had only gone through upsampling and bandpass filtering but not the application of NEO and subsequent steps. For each spike, samples that correspond to the time window $[-0.48\,\text{ms}, 1.00\,\text{ms}]$ relative to the spike time were used to represent the waveform. After downsampling the signal back to the original rate of 25 kHz, each putative spike was stored in a matrix $w_{c,t}$ of dimension $4 \times 38$.

Spike sorting — the process of assigning spikes to their putative units of origin — is usually performed in a feature space of lower dimensionality than the one defined by the full waveform on all channels (152 with the extraction method described above). Principal component analysis was used to project the 38-sample waveform on each channel into a three-dimensional feature space. Feature vectors from all four channels were concatenated to form the final 12-dimensional space in which spike sorting was performed.

In the hierarchy of common machine learning paradigms, sorting of extracellular spikes is posed as an unsupervised clustering problem. We chose to model the probability density of spikes in the feature space using a mixture model. In this approach, the density is modeled as the weighted linear combination of $K$ components that are distributed according to a common and well-understood distribution, e.g., the normal distribution:

$$p(x) = \sum_{k=1}^{K} \alpha_k p_k(x) \text{ with } \alpha_k \in (0, 1] \tag{4.2}$$

The assumption of Gaussianity is often not well-justified in neural data [18]. In part, this is due to heavier tails observed in the empirical distribution of spikes emitted by a single neuron. To address the insufficient fit of the normal distribution, we chose the multivariate t-distribution to model component densities, as it has heavier tails than the normal distribution for finite degrees of freedom. The density function $p_k(x)$ for a single model component $k$ in the mixture density 4.2 is then given by

$$p_k(x) = \mathcal{T}_{\mu_k, C_k, \nu}(x) \tag{4.3}$$

where $\mathcal{T}_{\mu_k, C_k, \nu}(x)$ is the density function of the multivariate t-distribution, parameterized

by its mean $\mu_k$, the scale matrix $C_k$ and the degrees of freedom $\nu$.

When mixture models are applied to the problem of spike sorting, the intuition is that, after the model has been fitted to the data, there is going to be a one-to-one correspondence between well-isolated single units and mixture components. A certain number of components of the model are expected to capture spikes from several units (multi-unit activity, MUA); this is typically the case for low-amplitude units.

The model, as described, is stationary in time and cannot explicitly model waveform amplitude variations that occur during long recordings. Calabrese and Paninski address this problem by allowing mixture component means to vary gently over time [12]. The time series of component means $\mu_k(t)$ is inferred using a Kalman filter approach. As rapid variations of spike amplitudes are not expected in correctly performed recordings, a Gaussian prior is used to penalize cluster drift. This measure also prevents the model from overfitting. With these modifications in place, it is more appropriate to refer to the model as a mixture of Kalman filters (MoK) than a mixture of t-distributions.

Closed-form equations for the maximum-likelihood parameters of mixture models given a dataset are generally not available. Instead, the Expectation Maximization (EM) algorithm [15] is used to iteratively optimize the parameters towards a local maximum of the model likelihood. A comprehensive overview of finite mixture models and the use of the EM algorithm is given by McLachlan and Peel [32]. The equations used for the E-step and M-step were taken from Calabrese and Paninski [12], with modifications for computational efficiency. In particular, allowing cluster means to shift for every single spike leads to excessive runtimes on long recordings with up to three million spikes. Instead, we chose to keep component means constant across one-minute segments of the recording. The fitting algorithm was implemented in MATLAB, and is based on an open source implementation[2] by Alexander Ecker (Max Planck Institute for Biological Cybernetics, Tübingen, Germany, Baylor College of Medicine, Houston, TX) and James Cotton (Baylor College of Medicine, Houston, TX).

### 4.3.2 Evaluation of clustering quality

After fitting an MoK model to the data, with inspection and modification by the user as necessary, spikes were assigned to clusters using soft and hard assignment rules based on

[2] https://github.com/aecker/moksm

the posterior probability. The likelihood of a spike $x$ at time $t$ to have been generated by component $k$ of mixture model $\mathcal{M}_{K,\mu,C,\nu}$ is given by

$$\alpha_k p_k(x,t) = \alpha \mathcal{T}_{\mu_k(t),C_k,\nu}(x) \tag{4.4}$$

A probabilistic clustering can be derived by calculating the posterior probability $\tau_k(x)$ of belonging to a certain component $k$:

$$\tau_k(x,t) = \frac{\alpha_k \mathcal{T}_{\mu_k(t),C_k,\nu}(x)}{\sum_{j=1}^{K} \alpha_j \mathcal{T}_{\mu_j(t),C_j,\nu}(x)} \tag{4.5}$$

Assigning each spike to the component with the highest posterior probability defines a hard assignment of spikes to clusters, i.e., putative single units in the context of spike sorting:

$$z(x,t) = \arg\max_k \tau_k(x,t) \tag{4.6}$$

By comparing soft assignments with hard assignments, estimates of the false-positive and false-negative rates of the clustering for a dataset $X = \{(x_1,t_1),\ldots,(x_n,t_n)\}$ can be derived. First, we partition the dataset into $k$ clusters based on the hard assignment, such that cluster $C_k$ contains all data points $(x,t)$ where $z(x,t) = k$, and $n_k = |C_k|$ is the number of spikes that have been assigned to cluster $k$. Then, we define

$$\mathrm{FP}_k = \frac{1}{n_k} \sum_{j \neq k} \sum_{(x,t) \in C_k} \tau_j(x,t) \tag{4.7}$$

$$\mathrm{FN}_k = \frac{1}{n_k} \sum_{j \neq k} \sum_{(x,t) \in C_j} \tau_k(x,t) \tag{4.8}$$

In practice, applying the criterion $(\mathrm{FP}_k + \mathrm{FN}_k) \leq 0.1$ usually selects for well-isolated units. Because $\mathrm{FP}_k$ and $\mathrm{FN}_k$ are continuous quantities, stricter selection criteria can be applied for analyses that are particularly sensitive to isolation quality.

We also relied on other measures of contamination that have been adopted in the community. We chose the "isolation distance" and "L-ratio," introduced by Schmitzer-Torbert et al. [45], and routinely calculated them for all spike-sorted units. Furthermore, clean clusters exhibit a pronounced dip in the auto-correlogram at short lags up to 2 ms, as the refractory period of neurons prohibits the observation of shorter interspike intervals (ISIs). Hence, we

quantified $k_{1.5}$ and $k_{2.5}$, the fraction of interspike intervals below 1.5 ms and 2.5 ms for all single units used for analysis purposes. Both ISI contamination measures had extremely low values (95th percentile of $k_{1.5}$ at 0.16 %, 95th percentile of $k_{2.5}$ at 2.46 %).

### 4.3.3 Linking of spike-sorted units across recording gaps

After spike sorting via the MoK model and assessment of cluster contamination, well-isolated single units were linked to their matches in preceding and subsequent recordings whenever possible. For the purpose of this section, we are going to refer to a *source* dataset and a *target* dataset under the assumption that the target dataset was recorded after the source dataset.

PCA projection vectors vary from dataset to dataset, causing the source and target mixture models to live in different feature spaces. In order to successfully compare the clusterings defined by the two mixture models, the source model $\mathcal{M}_{\text{source}}$ has to be projected into the target feature space, and vice versa. First, an approximate rotation matrix was established by projecting all source waveforms onto the target PCA basis vectors and calculating the Moore Penrose pseudoinverse from the feature representation of the same set of waveforms in both spaces. This rotation matrix was applied to the mean vectors and covariance matrices of the source mixture model. Additional EM steps were performed to correct for the approximate nature of the rotation matrix, and yielded the final projected model $\mathcal{M}'_{\text{source}}$. To clarify, this model was fit to the set of source waveforms, but exists in the same feature space as the target model. The analogous procedure was used to map the target model into the source feature space and create $\mathcal{M}'_{\text{target}}$.

Finally, the projected models were used to generate spike assignments for the "foreign" dataset, i.e., a $\mathcal{M}'_{\text{source}}$-derived clustering of the target dataset and a $\mathcal{M}'_{\text{target}}$-derived clustering of the source dataset. Confusion matrices were populated by comparing the spike assignments of $\mathcal{M}'_{\text{source}}$ and $\mathcal{M}_{\text{target}}$, as well as $\mathcal{M}'_{\text{target}}$ and $\mathcal{M}_{\text{source}}$. Putative matches were established between the highest-scoring cluster pairs. The experimenter inspected the matches in a feature-rich graphical user interface and performed comparisons of the feature representation, waveform shapes and autocorrelograms.

### 4.3.4 Analysis of REM spike trains

Mixture models with uniform components were fitted to spike trains on the time interval $[t_{\text{start}}, t_{\text{end}}]$, which also represents the support of the $K$-component mixture model with density

$$p(t) = \sum_{k=1}^{K} \alpha_k p_k(t) \text{ with } \alpha_k \in (0, 1], \tag{4.9}$$

and $p_k(t) \sim \mathcal{U}(a_k, b_k)$ is the uniform distribution density on the interval $[a_k, b_k]$, a subinterval of $[t_{\text{start}}, t_{\text{end}}]$. The first component is constrained to fully span the REM episode, i.e., $a_1 = t_{\text{start}}$ and $b_1 = t_{\text{end}}$.

Once models were fitted using an optimization procedure for each $K$ from 1 to 10. The chosen objective function was the difference between the empirical cumulative density function (cdf) of the spike train (estimated via the Kaplan-Meier method) and the cdf of the mixture model, sampled in steps of 100 ms on $[t_{\text{start}}, t_{\text{end}}]$. For a given $K$, the uniform mixture model is fully parameterized by $\alpha_k$ for $k = 1, \ldots, K$ and $(a_k, b_k)$ for $k = 2, \ldots, K$. An initial, valid set of parameters was randomly chosen and optimized in the least-squares sense using MATLAB's `lsqcurvefit` function. To escape from local minima, the optimization was randomly restarted ten times. The model with the fewest components that passed the Kolmogorov-Smirnov test at a significance level of 0.05 was chosen as the final fit.

### 4.3.5 Semi-automatic sleep stage detection

Behavioral state was assessed semi-automatically using three features. A hippocampal tetrode with high spectral power in the theta band (4 Hz to 10 Hz) and good theta modulation during lap running was selected. Spectral power in the theta band (4 Hz to 10 Hz) and delta band (1 Hz to 4 Hz) were extracted and temporally smoothed (Gaussian window, $\sigma = 0.5$ s for theta and $\sigma = 1$ s for delta). The energy in the EMG signal recorded from neck muscles was highpass-filtered (cutoff frequency 100 Hz) and smoothed with a Gaussian window ($\sigma = 250$ ms).

In a custom, MATLAB-based user interface, the experimenter used visualizations of the three extracted features and tagged episodes as "SWS," "REM" or "awake." The output of this supervised step was used to train a multinomial, logistic regression on theta power, delta power and the logarithm of neck EMG power (processed as described above). The

trained model was then applied to other datasets as a classifier. Its output was sampled at approximately 4 Hz and smoothed to reduce noise-induced rapid osciallations between two states. Segments shorter than 5 s were removed, and segments of the same type separated by less than 10 s were merged. Holes that remained in the segmentation after this step were filled using nearest-neighbor interpolation.

Due to the stringent quality criteria for the segmentation of REM sleep segments in this thesis, the segmentation of all REM segments was manually verified, and splits or merges applied where necessary.

### 4.3.6 Detection of replay episodes

For a detailed description of the methods, we refer to Louie and Wilson [30].

REM spike trains from single units were binned at 1 s resolution and smoothed with a Gaussian window ($\sigma = 1.5$ s) to create a template matrix $T$ of dimension $N \times D$, where $D$ is the duration of the REM episode in seconds and $N$ the numer of cells. The entry $T_{nd}$ contains the spike count in the $d$th bin of cell $n$. The template correlation coefficient $C$ was used to quantify the similarity between the template $T$ and a RUN matrix $R$:

$$C = \frac{\sum_{d=1}^{D}\sum_{n=1}^{N}\left(\frac{T_{nd}}{\mathrm{RMS}(T_n)} - \overline{T}\right)\left(\frac{R_{nd}}{\mathrm{RMS}(R_n)} - \overline{R}\right)}{N \cdot D \cdot \sigma_T \cdot \sigma_R} \tag{4.10}$$

with

$$\overline{T} = \frac{1}{N \cdot D}\sum_{d=1}^{D}\sum_{n=1}^{N} T_{nd}$$

$$\mathrm{RMS}(T_n) = \sqrt{\frac{1}{D}\sum_{d=1}^{D} T_{nd}^2}$$

$$\sigma_T = \sqrt{\frac{1}{N \cdot D}\sum_{d=1}^{D}\sum_{n=1}^{N}\left(\frac{T_{nd}}{\mathrm{RMS}(T_n)} - \overline{T}\right)^2}$$

and analgous definitions for the RUN matrix $R$.

The significance of the correlation coefficient was evaluated with respect to shuffled versions of the REM template matrix. Four different shuffle methods were employed:

- **BIN** Each row of the matrix is shuffled by applying a random permutation. Independent permutations are applied to each cell.

- **COLUMN** Columns of the matrix are randomly permuted. Temporal alignment across cells is maintained.
- **SWAP** Rows of the matrix are randomly exchanged. Cell identities are reassigned, but the temporal structure of the spike trains remains untouched.
- **SHIFT** A circular shift is applied to each row of the matrix. The shift distance is chosen randomly for each cell.

# Bibliography

[1] P. Andersen, R. Morris, D. Amaral, T. Bliss, and J. O'Keefe. *The Hippocampus Book*. Oxford Neuroscience Series. Oxford University Press, USA, 2006. ISBN 9780199723164.

[2] E. Aserinsky and N. Kleitman. Regularly occurring periods of eye motility, and concomitant phenomena, during sleep. *Science*, 118(3062):273–4, Sep 1953.

[3] K. Benchenane, A. Peyrache, M. Khamassi, P. L. Tierney, Y. Gioanni, F. P. Battaglia, and S. I. Wiener. Coherent theta oscillations and reorganization of spike timing in the hippocampal–prefrontal network upon learning. *Neuron*, 66(6):921–36, Jun 2010.

[4] M. R. Bennett, W. G. Gibson, and J. Robinson. Dynamics of the CA3 pyramidal neuron autoassociative memory network in the hippocampus. *Philos Trans R Soc Lond, B, Biol Sci*, 343(1304):167–87, Jan 1994.

[5] J. D. Berke, M. Okatan, J. Skurski, and H. B. Eichenbaum. Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron*, 43(6):883–96, Sep 2004.

[6] G. Q. Bi and M. M. Poo. Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neurosci*, 18(24):10464–72, Dec 1998.

[7] M. P. Brandon, A. R. Bogaard, C. M. Andrews, and M. E. Hasselmo. Head direction cells in the postsubiculum do not show replay of prior waking sequences during sleep. *Hippocampus*, 22(3):604–18, Mar 2012.

[8] G. Buzsáki. Two-stage model of memory trace formation: a role for "noisy" brain states. *Neuroscience*, 31(3):551–70, Jan 1989.

[9] G. Buzsáki. Theta oscillations in the hippocampus. *Neuron*, 33(3):325–40, Jan 2002.

[10] G. Buzsáki and E. Eidelberg. Phase relations of hippocampal projection cells and interneurons to theta activity in the anesthetized rat. *Brain Res*, 266(2):334–9, May 1983.

[11] G. Buzsáki and E. I. Moser. Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, 16(2):130–8, Feb 2013.

[12] A. Calabrese and L. Paninski. Kalman filter mixture model for spike sorting of non-stationary data. *J Neurosci Methods*, 196(1):159–69, Mar 2011.

[13] L. L. Colgin. Mechanisms and functions of theta rhythms. *Annual review of neuroscience*, 36:295–312, Jul 2013.

[14] L. V. Colom, B. R. Christie, and B. H. Bland. Cingulate cell discharge patterns related to hippocampal EEG and their modulation by muscarinic and nicotinic agents. *Brain Res*, 460(2):329–38, Sep 1988.

[15] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

[16] S. Diekelmann and J. Born. The memory function of sleep. *Nat Rev Neurosci*, 11(2): 114–26, Feb 2010.

[17] N. J. Douglas, D. P. White, C. K. Pickett, J. V. Weil, and C. W. Zwillich. Respiration during sleep in normal man. *Thorax*, 37(11):840–4, Nov 1982.

[18] M. S. Fee, P. P. Mitra, and D. Kleinfeld. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non-Gaussian variability. *J Neurosci Methods*, 69(2):175–88, Nov 1996.

[19] P. W. Frankland and B. Bontempi. The organization of recent and remote memories. *Nat Rev Neurosci*, 6(2):119–30, Feb 2005.

[20] J. D. Green and A. A. Arduini. Hippocampal electrical activity in arousal. *Journal of Neurophysiology*, 17(6):533–57, Nov 1954.

[21] A. S. Gupta, M. A. A. van der Meer, D. S. Touretzky, and A. D. Redish. Segmentation of spatial experience by hippocampal θ sequences. *Nat Neurosci*, 15(7):1032–9, Jul 2012.

[22] J. A. Horne and M. J. McGrath. The consolidation hypothesis for REM sleep function: stress and other confounding factors–a review. *Biol Psychol*, 18(3):165–84, May 1984.

[23] P. T. Huerta and J. E. Lisman. Bidirectional synaptic plasticity induced by a single burst during cholinergic theta oscillation in CA1 in vitro. *Neuron*, 15(5):1053–63, Nov 1995.

[24] O. Jensen and J. E. Lisman. Position reconstruction from an ensemble of hippocampal place cells: contribution of theta phase coding. *J Neurophysiol*, 83(5):2602–9, May 2000.

[25] R. Jung and A. Kornmüller. Eine Methodik der Ableitung lokalisierter Potentialschwankungen aus subcorticalen Hirngebieten. *Archiv für Psychiatrie und Nervenkrankheiten*, 109(1):1–30, 1938. ISSN 0003-9373.

[26] J. Kaiser. On a simple algorithm to calculate the 'energy' of a signal. In *IEEE International Conference Acoustic Speech Signal Processing*, 1990.

[27] A. Karni, D. Tanne, B. S. Rubenstein, J. J. Askenasy, and D. Sagi. Dependence on REM sleep of overnight improvement of a perceptual skill. *Science*, 265(5172):679–82, Jul 1994.

[28] A. K. Lee and M. A. Wilson. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6):1183–94, Dec 2002.

[29] A. K. Lee and M. A. Wilson. A combinatorial method for analyzing sequential firing patterns involving an arbitrary number of neurons based on relative time order. *J Neurophysiol*, 92(4):2555–73, Oct 2004.

[30] K. Louie and M. A. Wilson. Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. *Neuron*, 29(1):145–56, Jan 2001.

[31] J. L. McClelland, B. L. McNaughton, and R. 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. *Psychol Rev*, 102(3):419–57, Jul 1995.

[32] G. McLachlan and D. Peel. *Finite mixture models*. John Wiley & Sons, 2004.

[33] B. Milner, S. Corkin, and H.-L. Teuber. Further analysis of the hippocampal amnesic syndrome: 14-year follow-up study of H.M. *Neuropsychologia*, 6(3):215–234, Sep 1968.

[34] K. Mizuseki, K. Diba, E. Pastalkova, and G. Buzsáki. Hippocampal CA1 pyramidal cells form functionally distinct sublayers. *Nature Neuroscience*, 14(9):1174–81, Sep 2011.

[35] Z. Nádasdy, H. Hirase, A. Czurkó, J. Csicsvari, and G. Buzsáki. Replay and time compression of recurring spike sequences in the hippocampus. *The Journal of Neuroscience*, 19(21):9497–507, Nov 1999.

[36] T. Nakashiba, D. L. Buhl, T. J. McHugh, and S. Tonegawa. Hippocampal CA3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron*, 62 (6):781–7, Jun 2009.

[37] J. O'Keefe and J. Dostrovsky. The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res*, 34(1):171–5, Nov 1971.

[38] J. O'Keefe and L. Nadel. *The hippocampus as a cognitive map*. Oxford University Press, 1978. ISBN 9780198572060.

[39] J. O'Keefe and M. L. Recce. Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3):317–30, Jul 1993.

[40] G. R. Poe, D. A. Nitz, B. L. McNaughton, and C. A. Barnes. Experience-dependent phase-reversal of hippocampal neuron firing during REM sleep. *Brain Res*, 855(1): 176–80, Feb 2000.

[41] B. Rasch and J. Born. About sleep's role in memory. *Physiol Rev*, 93(2):681–766, Apr 2013.

[42] B. Rasch, J. Pommer, S. Diekelmann, and J. Born. Pharmacological REM sleep suppression paradoxically improves rather than impairs skill memory. *Nature Neuroscience*, 12(4):396–7, Apr 2009.

[43] A. Rechtschaffen, A. Kales, University of California, Los Angeles. Brain Information Service, and NINDB Neurological Information Network. *A Manual of Standardized Terminology, Techniques and Scoring System for Sleep Stages of Human Subjects*. Publication. Brain Information Service/Brain Research Institute, University of California, 1968.

[44] M. Rudolph, M. Pospischil, I. Timofeev, and A. Destexhe. Inhibition determines membrane potential dynamics and controls action potential generation in awake and sleeping cat cortex. *Journal of Neuroscience*, 27(20):5280–90, May 2007.

[45] N. Schmitzer-Torbert, J. Jackson, D. Henze, K. Harris, and A. D. Redish. Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience*, 131(1): 1–11, Jan 2005.

[46] W. B. Scoville and B. Milner. Loss of recent memory after bilateral hippocampal lesions. *J Neurol Neurosurg Psychiatr*, 20(1):11–21, Feb 1957.

[47] J. M. Siegel. The REM sleep-memory consolidation hypothesis. *Science*, 294(5544): 1058–63, Nov 2001.

[48] W. E. Skaggs, B. L. McNaughton, M. A. Wilson, and C. A. Barnes. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2):149–72, Jan 1996.

[49] E. Smith and S. Kosslyn. *Cognitive Psychology: Mind and Brain*. Pearson/Prentice Hall, 2007. ISBN 9780131825086.

[50] F. Snyder, J. A. Hobson, D. F. Morrison, and F. Goldfrank. Changes in respiration, heart rate, and systolic blood pressure in human sleep. *J Appl Physiol*, 19:417–22, May 1964.

[51] L. Squire. Memory and the hippocampus: A synthesis from findings with rats, monkeys, and humans. *Psychological Review*, Jan 1992.

[52] L. R. Squire and P. Alvarez. Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current opinion in neurobiology*, 5(2):169–77, Apr 1995.

[53] C. H. Vanderwolf. Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalogr Clin Neurophysiol*, 26(4):407–18, Apr 1969.

[54] R. P. Vertes. Memory consolidation in sleep; dream or reality. *Neuron*, 44(1):135–48, Sep 2004.

[55] Y. Ziv, L. D. Burns, E. D. Cocker, E. O. Hamel, K. K. Ghosh, L. J. Kitch, A. E. Gamal, and M. J. Schnitzer. Long-term dynamics of CA1 hippocampal place codes. *Nat Neurosci*, 16(3):264–6, Mar 2013.