Accessing the Diversity of Hippocampal Ripples

Manfredi Castelli

Trinity College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Hilary Term 2025

Acknowledgements

Personal and Institutional

First, I would like to thank my supervisor, Dr. David Dupret, for making me feel welcome and valued — both personally and scientifically. I am especially grateful for the insights he shared throughout my DPhil, and for teaching me how to communicate scientific ideas clearly, effectively, and appealingly. I am also deeply thankful to Dr. Vítor Lopes dos Santos, who taught me more about data analysis (and jazz harmony!) than I ever expected. Always generous with his time, he was quick to make a bet, with a smile, to show me why my results were wrong. Many thanks to all members of the Dupret Lab for their feedback on my work, for creating such a collaborative and stimulating research environment, and for valuing my input on their own projects. I am also grateful to Prof. Renaud Lambiotte for his guidance and helpful discussions on the topological analyses presented in this thesis. Special thanks to Ben Micklem, who helped me develop the summary schematic at the end of this thesis (Figure 7.1), and to the wider community at the MRC Brain Network Dynamics Unit for their support throughout my time here.

This work was supported by a Medical Research Council (MRC) UK studentship (MC_ST_BNDU_2019) and the MRC Transition from DPhil to First Position Award. I also thank Trinity College for generous grants that funded conference travel and the final stages of thesis submission.

On a personal note, I want to thank some of the closest friends I've made in Oxford — my 'Trinity family': Ivan, Jack, Ishbel, Tobias, Swathi, and Joffrey. Thank you for turning even the worst days into something better, always filled with laughter. I'm also deeply grateful to my local 'Italian family' — Melissa, Enrico, Giuseppe, and Luca — for making me feel at home, even while being far from it.

Finally, none of this would have been possible without the constant support of my family. I am especially grateful to my grandparents, and in particular to nonno Melchi, who taught me to be curious about the world. I am profoundly thankful to my sister Emily and to my parents, who have supported me throughout my life and given me every opportunity to pursue my dreams.

'Ciascuno di noi si crede uno ma non è vero: è tanti.'

'Each of us believes he is one, but that's not true: he is many.'

— Luigi Pirandello

Specific contributions

Of the 18 mice included in this thesis, I recorded data from one animal (implanted with a silicon probe), while the remaining recordings were conducted by Vítor Lopes-dos-Santos. All surgeries and implants were performed by certified PIL holders in the lab.

The scripts for preprocessing raw multichannel data, generating shuffling controls, and estimating recording depth (see Sections 2.1.5, 2.1.6, 2.2, 2.3, 2.4.5, and 6.2.1) were developed by Vítor Lopes-dos-Santos, who also performed spike sorting and unit isolation for the tetrode dataset. Giuseppe Pietro Gava developed the scripts for constructing neuronal coactivity graphs (see Section 2.4.3) and for generating the mean difference estimation plots (see Section 2.5).

Parts of the main conceptual and analytical framework described throughout this thesis are included in a preprint manuscript recently published (Castelli et al., 2025). The analytical approach applied to dentate spikes in the second half of Chapter 4 was developed as part of a published article (McHugh et al., 2024) and is available online as an open-source tool (Castelli, 2025). All material presented in this thesis related to these studies originates from the data and analytical tools described above. The other co-authors of the aforementioned published work have given me permission to include the material in this thesis.

The conceptualisation of the studies presented in this thesis, all written content, and the analyses are solely my own work.

Abstract

The brain is capable of retaining prior knowledge while continuously incorporating new information. How neuronal populations support this dual function — updating internal representations without disrupting existing ones — remains a fundamental biological question. Hippocampal ripples have been identified as key events supporting offline processing and memory consolidation. These short-lived, highly synchronised patterns reactivate waking experiences during sleep and rest. Yet, it remains unclear whether, at the level of individual ripple events, the associated laminar currents, neuronal firing patterns, and reactivation content are consistent across events or flexibly tuned on a ripple-by-ripple basis.

In this thesis, I addressed this question by profiling individual ripple events using laminar current source density recordings in the mouse hippocampus during sleep and rest. I identified two ripple profiles — $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples — characterised by current sinks in *stratum radiatum* and *stratum lacunosum-moleculare*, respectively. The distinct laminar signatures of these ripple types suggest they may arise from separate hippocampal circuit mechanisms. Each profile was associated with characteristic spectral features and LFP waveforms. To further support a circuit-level distinction, I showed that the two ripple types differentially recruited CA1 and CA3 neurons, giving rise to distinct motifs of millisecond-timescale coactivity and different topological organisations of population activity. I then investigated how current diversity mapped onto functional diversity in reactivation content. $\mathrm{Rad}^{\mathrm{sink}}$ ripples integrated recent motifs of waking coactivity, combining superficial and deep CA1 principal cells into denser, higher-dimensional patterns that exhibited stable reactivation throughout sleep. In contrast, $\mathrm{LM}^{\mathrm{sink}}$ ripples contained core motifs of prior coactivity, engaging primarily deep cells into sparser, lower-dimensional patterns that gradually drifted over time, updating pre-existing content with recent experience.

Together, these findings suggest that ripple-by-ripple diversity enables parallel reactivation channels that support both the integration of recent experience and the gradual updating of prior representations.

Contents

Appendices

List of Figures

List of Tables

List of Abbreviations

REM	rapid eye movement
nREM	non-rapid eye movement
CA	Cornu Ammonis
CA1/2/3	Dorsal regions of the CA1, CA2, and CA3 of the hippocampus
MEC	Medial Entorhinal Cortex
Or	*stratum oriens* in CA1
Pyr	*stratum pyramidale* in CA1
Rad	*stratum radiatum* in CA1
LM	*stratum lacunosum moleculare* in CA1
DG	*dentate gyrus*
GCL	granule cell layer in the *dentate gyrus*
LFP	Local Field Potential
SW(Rs)	Sharp Wave (Ripples)
CSD	Current Source Density
Rad$^{\text{sink}}$Ripple	Ripple with a current sink in *stratum radiatum*
LM$^{\text{sink}}$Ripple	Ripple with a current sink in *stratum lacunosum moleculare*
PC(A)	Principal Component (Analysis)
GLM	Generalised Linear Model
MI	Mutual Information
SI	Structure Index
SD	Standard Deviation
IQR	Interquartile Range
CI	Confidence Interval
vs.	versus
#	Number or Count

The sharp-wave ripple complex has become the first definite bio-marker for cognitive operations.

— György Buzsáki (Buzsáki, 2015)

1

Introduction

Contents

The brain has a remarkable capacity to retain and use past experience while continuously integrating new information. It is widely accepted that these processes benefit from offline periods such as sleep and quiet rest. One historical anecdote illustrating this idea comes from chemist Dmitri Mendeleev, who described the moment of insight that led to the periodic table: "*I saw in a dream a table where all the elements fell into place as required. Awakening, I immediately wrote it down on a piece of paper.*"

The cognitive benefits of offline states have been closely linked to the hippocampus, a region that plays a central role in memory (Buzsáki & Moser, 2013; Eichenbaum et al., 1999; O'Keefe & Nadel, 1978; Scoville & Milner, 2000). In particular, the hippocampus plays a critical role not only in spatial navigation and episodic memory encoding, but also in linking events across time and supporting inferential reasoning — enabling information processing that extends beyond the domain of physical space (Barron et al., 2020; Buzsáki, 1989; Buzsáki & Moser, 2013; Dusek & Eichenbaum, 1997; Eichenbaum et al., 1999; O'Keefe & Nadel, 1978; Tavares et al., 2015).

These diverse memory functions are supported by a rich repertoire of neural oscillations that coordinate hippocampal network activity (Buzsáki, 1986, 1989; Colgin, 2016; Fernandez-Ruiz et al., 2023). Neural oscillations are temporal fluctuations in the local field potential (LFP) that help organise the timing of neuronal firing, enabling efficient information processing through dynamic interactions within and across brain regions (Buzsáki, 2006; Buzsáki & Draguhn, 2004). In the hippocampus, these oscillations are spatially structured across distinct sublayers and have been extensively studied as markers of circuit-level computation (Buzsáki et al., 1983; Colgin, 2016; Lopes-dos-Santos et al., 2023). Distinct hippocampal frequency bands are associated with specific memory operations, including encoding, retrieval, inference, and consolidation (Barron et al., 2020; Colgin, 2016; C. Liu et al., 2023; Lopes-dos-Santos et al., 2018).

Among these rhythms, fast oscillations occurring during offline states — known as ripples — have emerged as one of the most synchronised population events in the brain (Buzsáki et al., 1983; Buzsáki et al., 1992; O'Keefe & Nadel, 1978; Ylinen et al., 1995). Unlike other characteristic frequency bands, such as theta oscillations (8–12 Hz), ripples have been observed in the hippocampus of every mammalian species studied to date, including rodents, cats, bats, monkeys, and humans (Bragin et al., 1999; Buzsáki et al., 2013). For this reason, Buzsáki, 2015 described ripples as the first definitive biomarker of cognitive operations. Despite their highly synchronous nature, ripples are under tight inhibitory control at millisecond timescales (Stark et al., 2014), and small alterations in this control can transform physiological ripples into pathological "fast ripples" — a biomarker of epileptic activity (Bragin et al., 1999; Valero et al., 2017).

Hippocampal ripples are transient network events defined by synchronous firing of CA1 neurons and by high-frequency oscillations (100–250 Hz) in the LFPs recorded from the *stratum pyramidale* of CA1 (Buzsáki et al., 1992; Ylinen et al., 1995). These are typically accompanied by a sharp negative deflection in the LFPs of *stratum radiatum*, known as a sharp wave, together forming the sharp-wave ripple (SWR) complex (see Section 1.2) (Buzsáki et al., 1983; Buzsáki, 1986; Buzsáki et al., 1992; Sullivan et al., 2011; Ylinen et al., 1995).

The oscillatory profile of SWRs stands in contrast to the dynamics observed during active wakefulness. During exploratory behaviours such as locomotion, sniffing, and rearing, hippocampal activity is dominated by theta oscillations (8–12 Hz), especially in CA1 (Figure 1.1(A)) (Buzsáki et al., 1983; Colgin, 2016). By contrast, during immobility, grooming, quiet wakefulness, and non-REM (nREM) sleep, theta is replaced by irregular, slow-wave oscillations (i.e. slow-wave sleep) interspersed with SWRs (Figure 1.1(B)) (Buzsáki et al., 1983; Buzsáki, 1986; Ylinen et al., 1995). This theta–SWR dichotomy reflects two distinct brain states: a volitional-exploratory mode characterised by sparse neuronal recruitment, and a consummatory-offline mode marked by highly synchronised neural activity (Figure 1.1). Together, these states underpin both the encoding of experience and its offline consolidation (Buzsáki, 1989) (see also Section 1.1). Notably, theta oscillations are also present during REM sleep (Figure 1.1(C)), a state in which SWRs are almost entirely absent (Buzsáki, 2015; Colgin, 2016). This suggests that theta and SWRs function as mutually exclusive network modes within the hippocampus (Buzsáki, 2015; Buzsáki, 1989).

However, even if theta and SWRs alternate in time, they are functionally interdependent. Each plays a critical role in shaping the other, and together they support the coordination of memory-related processing across brain states (Buzsáki, 2015; Buzsáki, 1989). Importantly, the coactivity patterns expressed during SWRs are not independent from those structured during theta oscillations in wakefulness. In fact, one of the defining features of SWRs is their ability to re-express offline the neuronal assemblies that were previously encoded during exploratory theta

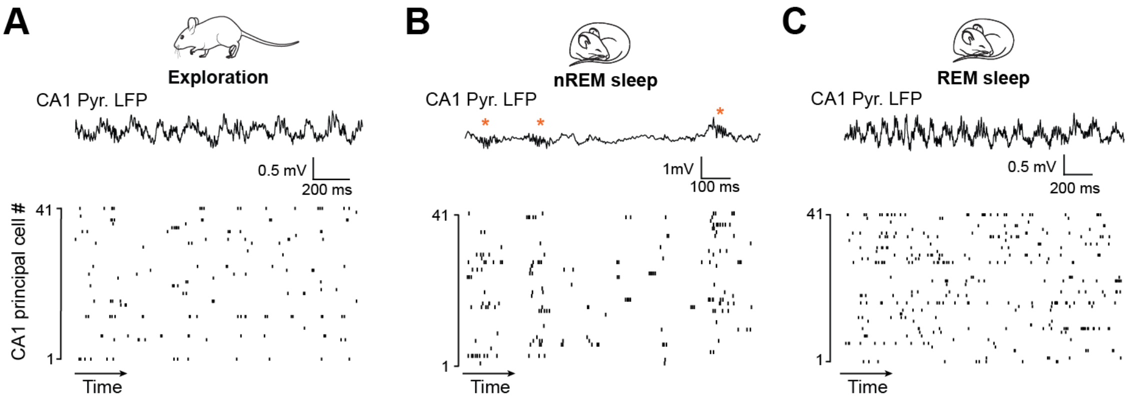

Figure 1.1: State-dependent hippocampal activity across waking, nREM, and REM sleep.
(A–C) Raw LFP signal from the CA1 pyramidal layer (top) and raster plot of CA1 principal cells (bottom), with each row representing a different cell. The panels show different brain states: wakefulness during exploration (A), nREM sleep (B), and REM sleep (C). Note the different time scales across panels. Theta cycles are visible during both exploration and REM sleep, but are absent during nREM sleep; during nREM, ripples (highlighted here by red asterisks) are present and synchronise the firing of CA1 principal cells.
The data shown were obtained from one sleep session in a mouse implanted with tetrodes (see Chapter 2).

cycles (Wilson & McNaughton, 1994). This forms the basis of the long-standing hypothesis that SWRs support systems-level memory consolidation — a topic I expand on in the following section (Buzsáki, 1989; Klinzing et al., 2019).

While SWRs have often been treated as a homogeneous phenomenon, recent work suggests they may exhibit structured variability — raising the possibility that different ripple types contribute differently to memory processes and other functions. In the remainder of this Chapter, I review current knowledge about hippocampal ripples, their functions, their circuit-level generation, and the emerging evidence for ripple diversity, setting the stage for the central question of this thesis: how diverse are hippocampal ripples, and what is the functional significance of this diversity?

1.1 Ripples and their function in memory and beyond

One of the earliest mentions of hippocampal ripples appears in the work of O'Keefe and Nadel, 1978, who noted their presence but admitted that their function was still unknown. At the time, the authors even speculated — somewhat ironically in light of subsequent findings — that "ripples may be summated abortive spikes either in the proximal dendrites or in the soma of the pyramidal cells". In the decades that followed, however, extensive research has shown that ripples are not abortive spikes, but instead constitute a key mechanism supporting memory consolidation (Buzsáki, 2015; Buzsáki, 1989; Joo & Frank, 2018; Wilson & McNaughton, 1994; Ylinen et al., 1995).

Sleep has long been shown to play a critical role in memory consolidation (Buzsáki, 2015), and over the past few decades, research has converged on what is often referred to as the two-stage model of memory formation (Buzsáki, 1989). According to this framework, during theta-dominated states such as waking exploration, cell assemblies form as sensory information is encoded and processed. In a second stage, during offline states such as rest or sleep, the absence of ongoing sensory input allows these assemblies to be reactivated and stabilised without interference from new sensory stimuli, thereby facilitating memory consolidation (Buzsáki, 1989; Klinzing et al., 2019). Buzsáki, 1989 proposed that distinct network dynamics within the hippocampal circuit underlie each stage of the model. During theta oscillations, granule cells in the dentate gyrus (DG) are highly active, relaying sensory information from the medial entorhinal cortex (MEC) to CA3. This activity results in weak synaptic potentiation of CA3 neurons, effectively tagging them as candidates for later consolidation. During the second stage — typically occurring offline — CA3 neurons fire synchronously, driven by transient disinhibition from subcortical structures. This synchronous burst, known as a sharp wave, reactivates the previously potentiated CA3 neurons, drives strong activation in CA1, and triggers ripple oscillations (Stark et al., 2014; Sullivan et al., 2011). The high-frequency activity of ripples is associated with long-term potentiation (LTP) in CA1, which serves to stabilise memory traces into longer-lasting representations.

Building on this, Wilson and McNaughton, 1994 demonstrated that the CA1 coactivity structure observed during waking exploration is preserved during post-exploration sleep, but not during pre-exploration sleep. This temporal asymmetry — where post-sleep activity resembles waking more than pre-sleep — has been termed *reactivation*, referring to the reinstatement of waking coactivity patterns during subsequent offline periods. This finding laid the foundation for decades of research implicating SWRs as critical events for the offline reorganisation, stabilisation, and processing of hippocampal activity (Buzsáki, 2015; Dupret et al., 2010; Girardeau & Lopes-dos-Santos, 2021; Joo & Frank, 2018; Klinzing et al., 2019; O'Neill et al., 2008).

In addition to reactivating coactivity, another key feature of SWRs is that the recruited neurons often fire in a sequential manner. These sequences form temporally compressed representations of firing patterns that occurred during waking exploration, supporting the idea that ripples promote memory *replay* (Diba & Buzsáki, 2007; Foster & Wilson, 2006; Nádasdy et al., 1999). A recent study has proposed that these two ripple features — reactivation and replay — may serve distinct mnemonic functions: with replay supporting predictive coding, and reactivation supporting associative learning (C. Liu et al., 2023).

Importantly, recent work has shown that reactivation and replay are not uniformly distributed across all CA1 principal cells. Superficial and deep CA1 pyramidal cells, defined by their radial position (close to *stratum radiatum* and *stratum oriens*, respectively), exhibit distinct behaviours across states. During exploration, deep CA1 cells tend to fire more and show rigid activity patterns, by being similarly tuned across different environments (Berndt et al., 2023; Gava et al., 2021, 2024; Grosmark & Buzsáki, 2016; Harvey et al., 2023; Lopes-dos-Santos et al., 2023; Valero et al., 2015). In contrast, superficial cells fire less but exhibit flexible activity, allowing them to distinguish between different contexts (Berndt et al., 2023; Gava et al., 2021, 2024; Grosmark & Buzsáki, 2016; Harvey et al., 2023; Lopes-dos-Santos et al., 2023; Valero et al., 2015). Interestingly, during ripple events, this pattern of activity shifts: superficial cells become more active than their deep counterparts. Building on this, Harvey et al. (2023) showed that both reactivation and replay are predominantly driven by superficial CA1 cells, reinforcing their relevance for memory plasticity. Furthermore, Harvey et al. demonstrated that deep and superficial CA1 cells engage distinct hippocampo-cortical assemblies during reactivation, suggesting that they may contribute to different aspects of memory consolidation through separate brain-wide circuits.

Beyond these correlational findings, multiple studies have demonstrated a causal link between ripples and memory. Disrupting ripples impairs subsequent memory performance (Chang et al., 2025; Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Roux et al., 2017; van de Ven et al., 2016), whereas prolonging ripple duration has been shown to enhance memory retention (Fernández-Ruiz et al., 2019).

Additionally, ripples are not limited to mnemonic processing. They have been implicated in a variety of other functions, including planning, inference, and even metabolic regulation (Barron et al., 2020; Diba & Buzsáki, 2007; Foster & Wilson, 2006; Tingley et al., 2021). These diverse roles may be supported by the capacity of ripples to coordinate brain-wide activity, acting as a temporal window for hippocampal interaction with distributed cortical and subcortical circuits (Nitzan et al., 2022; Peyrache et al., 2011; Ramirez-Villegas et al., 2015).

To summarise, SWRs play a central role in memory formation by reactivating coactivity patterns and replaying experience-related sequences during offline periods. These dynamics appear to be preferentially driven by superficial CA1 cells, which show greater plasticity than their deep counterparts. Disrupting ripple events impairs memory, while enhancing them improves retention. Moreover, ripples support functions beyond memory, highlighting their broader role in orchestrating large-scale network dynamics.

But how does the hippocampus generate ripple events in the first place? To explore this question, it is important to first characterise the features of ripples and their underlying circuit dynamics. In the next two sections, I first describe the canonical SWR complex, focusing on its electrophysiological signatures, and then review models of ripple generation, from the classic CA3-driven view to broader circuit-level mechanisms involving intra- and extra-hippocampal inputs.

1.2 The canonical sharp-wave ripple (SWR) complex

The term "ripples" is sometimes used interchangeably with "sharp-wave ripples" (SWRs), typically referring to the canonical ripple profile observed as part of the SWR complex (Buzsáki, 2015). However, defining events as ripples is not equivalent to defining them as SWRs. The two terms refer to distinct electrophysiological phenomena, which originate from different layers of CA1 and reflect separate underlying mechanisms (Buzsáki et al., 1983; Buzsáki, 1986; Buzsáki et al., 1992; Stark et al., 2014).

Sharp waves (SWs) are characterised by large negative polarity deflections in CA1 *stratum radiatum.* These events typically manifest in the lower frequency range (5–40 Hz) and are short-lived, with an average duration of 40–110 ms (Buzsáki et al., 1983; Buzsáki, 1986). During periods of immobility, rest, or sleep, SWs are often accompanied by high-frequency oscillations, referred to as ripples, within the CA1 *stratum pyramidale*, together forming the SWR complex (Figure 1.2(A)) (Buzsáki, 2015; Buzsáki et al., 1983; Buzsáki, 1986; Buzsáki et al., 1992; Ylinen et al., 1995).

Importantly, sharp waves and ripples also differ in their underlying circuit mechanisms (Sullivan et al., 2011; Ylinen et al., 1995). Sharp waves, which emerge in *stratum radiatum*, reflect excitatory input from CA3 to CA1 via the Schaffer collaterals (Figure 1.2(A); Buzsáki et al., 1983). In contrast, findings suggest that ripple oscillations are not simply transmitted to CA1 from upstream regions like CA3, but are instead generated locally within CA1 itself through interactions between pyramidal cells and interneurons (Stark et al., 2014; Sullivan et al., 2011; Ylinen et al., 1995).

Despite their distinct nature, SWs and ripples interact. First, they tend to co-occur. When ripple events are averaged based on their peak timestamps in *stratum pyramidale*, a concurrent SW component typically emerges in *radiatum* (Figure 1.2), defining one of the canonical features of the SWR profile (Buzsáki et al., 1992; Ylinen et al., 1995). Second, the amplitude of the SW in *radiatum* positively correlates with the frequency of the ripple oscillation in the pyramidal layer (Sullivan et al., 2011). Nonetheless, it is crucial to recognise that the two events are dissociable. Interestingly, in a minority of cases, ripples occur without detectable SWs, and vice versa (Navas-Olive et al., 2020; Ylinen et al., 1995).

As I will further discuss in Chapter 7, this distinction matters for interpreting previous research on ripples. Many studies that analyse 'ripples' in fact focus exclusively on SWRs, requiring the co-occurrence of both ripple and SW components for event detection (e.g., Chang et al., 2025; Harvey et al., 2023; Sebastian et al., 2023). While this approach captures the majority of ripple events, it effectively biases the analysis toward a specific subset — namely, those with a SW in *stratum radiatum*, and thus strong CA3 input. Thus, when making claims about ripple mechanisms or functions, it is important to consider whether such studies are truly generalisable to all ripple events.

The canonical SWR profile has also been studied using current source density (CSD) analysis (Buzsáki et al., 1983; Mitzdorf, 1985; Sullivan et al., 2011; Ylinen et al., 1995). CSD methods estimate the spatial distribution of transmembrane currents, distinguishing between sinks — interpreted as active glutamatergic input currents — and sources, which reflect passive return currents. Applied to laminar CA1 recordings, CSD offers qualitative insights into the origin of synaptic inputs. In fact, CA1 receives anatomically segregated inputs (Valero & de la Prida, 2018): CA2 preferentially projects to *stratum oriens*; CA3 to *stratum radiatum*; and medial entorhinal cortex layer III (MEC3) to *stratum lacunosum-moleculare* (Figure 1.2)).

The average CSD profile during ripples reveals a large current sink in *radiatum* (corresponding to the SW) and a concurrent source in *lacunosum-moleculare*, reinforcing the interpretation that canonical SWRs primarily reflect CA3 input to CA1 (Figure 1.2(B)) (Sullivan et al., 2011; Ylinen et al., 1995). In addition, weaker sinks can sometimes be observed in *stratum oriens* (Ylinen et al., 1995), suggesting a possible role for CA2 inputs in ripple generation (Oliva et al., 2016; Sebastian et al., 2023).

In summary, the canonical SWR is defined as a fast oscillatory event in CA1 *stratum pyramidale*, accompanied by a sharp-wave and a current sink in *stratum radiatum* (Figure 1.2), highlighting strong CA3 input (Buzsáki et al., 1983). However, this raises important questions: what happens during ripple events that are not accompanied by sharp waves? Do they still depend on CA3, or might other input pathways become dominant? Is CA3 the sole structure influencing ripple occurrence in CA1?

In the next section, I review evidence that ripple generation and features are influenced by additional intra- and extra-hippocampal circuits, and explore how these regions modulate ripple occurrence and timing beyond the classical CA3-driven model.

1.3 Ripple generation: from canonical models to circuit diversity

In the section above, I presented the consensus view of canonical SWRs, in which CA3 is considered the key region in ripple generation. Another key question is whether regions other than CA3 can also influence the initiation or characteristics of ripple events. In this context, two broad hypotheses regarding ripple generation may be formulated:

1. **Homogeneous ripples**: all ripples are primarily driven by CA3, and variability in inputs from other regions is either negligible, effectively cancelled

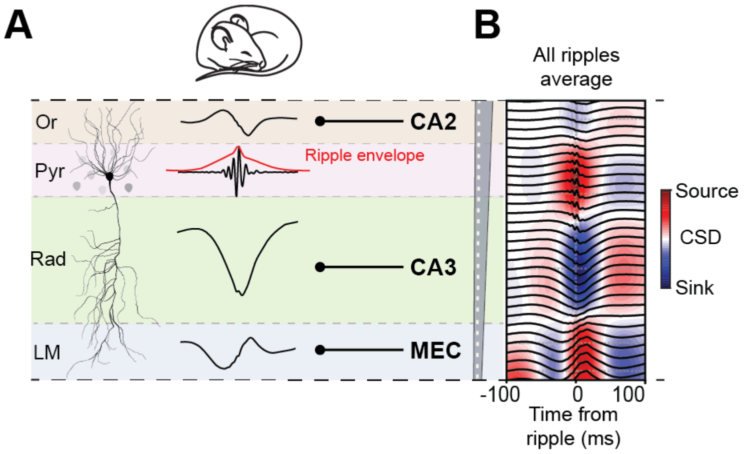

Figure 1.2: The sharp-wave ripple complex and associated current sinks within layers of CA1.
(A) Schematic showing the average LFP waveforms (black traces) during ripples across different layers of CA1: *oriens* (Or), *pyramidale* (Pyr), *radiatum* (Rad), and *lacunosum-moleculare* (LM), alongside the main excitatory inputs to each layer (CA2, CA3 and medial entorhinal cortex (MEC)) (Buzsáki, 2015; Valero & de la Prida, 2018). The ripple envelope (i.e., the instantaneous ripple-band power) is shown in the pyramidal layer (red trace). Note the distinct LFP waveforms across layers, particularly the well-known sharp-wave in *radiatum*, where CA3 inputs arrive.
(B) Mean current source density (CSD; colour-coded map) and corresponding LFP waveform (black traces) during ripples recorded in *stratum pyramidale* using a silicon probe spanning the CA1 somato-dendritic axis. Note the prominent current sink and sharp-wave in *stratum radiatum*.
The data shown were obtained from one sleep session in a mouse implanted with silicon probes (see Chapter 2).

out, or contributes only as noise.

2. **Diverse ripples**: different ripples originate from distinct upstream input profiles — some may be driven by CA3, while others may be influenced by intra- or extra-hippocampal regions — suggesting functional diversity among ripple events.

In this section, I first review studies that support the canonical CA3-driven model and the homogeneous ripple hypothesis. I then present recent evidence supporting the idea that regions beyond CA3 can also influence ripple occurrence and shape their features, consistent with the diverse ripple hypothesis.

Ripple generation in the canonical CA3→CA1 model

Buzsáki, 2015 and Stark et al., 2014 describe ripple generation as a consequence of strong, synchronous excitatory input — detected as a sharp wave (SW) in *stratum radiatum* — to populations of CA1 pyramidal cells. This excitatory drive results in a burst of pyramidal spiking, which in turn recruits reciprocally connected interneurons, most likely parvalbumin-expressing (PV) basket cells. The resulting feedback loop of excitation and reciprocal inhibition paces the activity of pyramidal cells at high frequencies, forming a self-sustaining oscillatory pattern that is read out in the LFP as a ripple (Stark et al., 2014). Thus, while the ripple itself may emerge from intrinsic circuit properties within the CA1 pyramidal layer, a key prerequisite — and perhaps the more interesting aspect to investigate — is the presence of a large tonic excitatory input to the apical dendrites of CA1 pyramidal cells. Supporting this, CA1 isolated *in vitro* is unable to generate ripple-associated sharp waves (Colgin et al., 2004).

On average, ripples are accompanied by sharp waves and by current sinks in *stratum radiatum*, the layer targeted by CA3 projections via the Schaffer collaterals (Figure 1.2) (Buzsáki, 2015; Valero & de la Prida, 2018). These electrophysiological features, together with the anatomical positioning of CA3 inputs, have led to the hypothesis that CA3 plays a pivotal role in initiating ripples by delivering large excitatory currents to the apical dendrites of CA1 pyramidal cells (Buzsáki et al., 1983; Buzsáki, 1986; Sullivan et al., 2011). In line with this, studies have demonstrated that CA3 pyramidal cells fire in bursts during SWRs — coupled to the sharp wave — and that this precedes the burst of ripple-associated spiking in CA1 (Buzsáki, 1986; Oliva et al., 2016; Sullivan et al., 2011). Interestingly, CA3 pyramidal cells also burst at high rates during SWs, but these fast oscillations are not transmitted cycle-by-cycle to CA1, further highlighting that ripple generation is intrinsic to CA1 (Sullivan et al., 2011; Ylinen et al., 1995).

These observations establish the SW–associated burst in *radiatum* as a likely trigger of ripple-generating dynamics in CA1. But can this burst be attributed to feedforward excitation from upstream regions — for instance, the entorhinal cortex? Several studies suggest this is not the case. SWs persist after removal of the neocortex or lesioning of the entorhinal cortex, indicating that they are primarily intra-hippocampal in origin (Buzsáki et al., 1983; Ylinen et al., 1995). Moreover, SWs are observed *in vitro* in isolated hippocampal slices (Buzsáki, 2015; Colgin et al., 2004). This aligns with the idea that CA3, with its highly recurrent connectivity, can autonomously generate synchronous population bursts (Buzsáki, 2015; Buzsáki, 1989).

But what regulates the conditions under which CA3 generates such synchronous bursts? One key factor is neuromodulation. During active wakefulness and exploration, the cholinergic tone from medial septal projections is high, and this suppresses SWR occurrence (Buzsáki, 2015; Vandecasteele et al., 2014). Conversely, when cholinergic input decreases — during consummatory behaviours or nREM sleep — the recurrent excitatory dynamics within CA3 are disinhibited, creating favourable conditions for SW generation (Viney et al., 2013; Zhang et al., 2021).

How is this disinhibition mediated? Local interneurons in both CA1 and CA3 have been suggested to play a key role in shaping these windows of heightened excitability. While the firing dynamics of different interneuron types during ripples are diverse — such as PV-expressing basket and bistratified cells — PV-expressing axo-axonic cells in both CA1 and CA3 consistently reduce their firing during SWRs (Klausberger & Somogyi, 2008; Viney et al., 2013). Notably, CA3 axo-axonic cells are innervated by GABAergic neurons from the medial septum, which has led to the hypothesis that medial septal GABAergic cells inhibit axo-axonic interneurons in CA3, resulting in disinhibition of CA3 pyramidal cells and increasing the likelihood of synchronous bursts that manifest as SWs (Viney et al., 2013). From this point of view, Buzsáki, 2015 described SWs as events that are “released” when subcortical neuromodulatory systems withdraw their inhibitory influence on the hippocampus, rather than being actively induced.

Once the SW from CA3 reaches CA1, it triggers the ripple oscillation by transiently synchronising CA1 pyramidal neurons, and is followed by a post-ripple hyperpolarisation that scales with the SW amplitude, which effectively imposes a refractory period that limits the immediate recurrence of ripple events (English et al., 2014).

In short, this view suggests that the high recurrent connectivity of CA3, gated by subcortical neuromodulation and local disinhibition, enables the synchronous firing (i.e., SWs) required to initiate ripples via feedforward excitation of CA1.

Beyond CA3: Other contributors to ripple occurrence and structure

If the hypothesis discussed above — where ripple generation depends entirely on CA3 — were strictly correct, then removing CA3 output should completely abolish ripples. However, Nakashiba et al., 2009 challenged this view using CA3-TeTX inducible transgenic mice, in which CA3 output can be selectively silenced. Remarkably, ripple events still occurred, suggesting that ripple generation does not depend exclusively on CA3. Instead, these “surviving” ripples may be triggered by other inputs to the apical dendrites of CA1 pyramidal cells, pointing to alternative pathways involved in ripple initiation.

So what are these other inputs that might contribute to ripple generation or shape ripple features? Over the last two decades, multiple studies have challenged the idea that CA3 is the sole driver of ripples. Other regions — including CA2, the medial entorhinal cortex (MEC), and even cortical oscillations — appear to influence the timing and characteristics of CA1 ripples.

CA2 is a major input to CA1, projecting predominantly to *stratum oriens* (Valero & de la Prida, 2018), and also receives inputs from CA3. Oliva et al., 2016 showed that deep CA2 pyramidal cells ramp up their activity before CA1 SWRs, and that superficial CA2 cells exhibit ripple-like synchronous firing even before CA1 and CA3 synchronous ripple-firing. This suggests that CA2 may help trigger CA1 ripples, especially during awake SWRs, whereas CA3 appears to play a more prominent role during sleep (Oliva et al., 2016). Similarly, Imbrosci et al., 2021 found that about 15% of ripples — mainly in the ventral hippocampus — appear to originate in the subiculum and then propagate backward to CA1 and CA3, and forward to MEC.

Another key intra-hippocampal region is the dentate gyrus (DG), which provides major input to CA3 (Valero & de la Prida, 2018). Interestingly, during rest and sleep, the DG also gives rise to large, synchronous events of granule cell activity known as dentate spikes (DSs), which have recently been shown to play a role in memory consolidation, similar to ripples (Farrell et al., 2024; McHugh et al., 2024). In the context of ripple generation, McHugh et al., 2024 reported that a subset of DSs coincide with CA1 ripples, potentially suggesting that the DG can bias ripple timing — either indirectly via CA3, or through a shared upstream input to both DG and CA1, likely originating in the MEC.

Previous reports have indeed shown that MEC can influence ripple features and timing. MEC projects directly to CA1 *stratum lacunosum-moleculare* from layer III (MEC3), and indirectly to *stratum radiatum* via the trisynaptic pathway (MEC2→DG→CA3→CA1) (Valero & de la Prida, 2018). Several studies have shown that MEC inputs influence ripple expression. Yamamoto and Tonegawa, 2017 found that blocking MEC3 inputs to CA1 reduced the number of ripple bursts and impaired spatial coverage of replay — specifically during awake ripples, not during slow-wave sleep. They also disrupted CA3→CA1 projections and showed that ripple occurrence and replay were affected in both behavioural states, reaffirming CA3's central role during SWRs. In line with this, C. Liu et al., 2023 showed that disrupting MEC inputs impaired the temporal organisation of CA1 ripple sequences — critical for predictive coding — while preserving the overall coactivity patterns and associated memory content. Together, these results suggest that MEC inputs influence not only the likelihood of ripple occurrence, but also the fine-grained structure of ripple content — particularly during wakefulness.

Taken together, the studies reviewed so far reveal that, depending on behavioural state, the occurrence and content of ripples are shaped by inputs from regions beyond CA3, including MEC, CA2, DG, and subiculum.

In addition to influencing ripple initiation and content, recent work suggests that the timing of ripple events — that is, when they occur — is also modulated by subcortical and cortical inputs. For instance, sleep spindles — 10–15 Hz oscillations generated in the thalamus — can influence ripple timing (Girardeau & Lopes-dos-Santos, 2021; Peyrache et al., 2011; Sirota et al., 2003). While ripples and spindles occur at different frequency bands, ripple events — and the firing of CA1 neurons — are phase-locked to the troughs of spindles (Sirota et al., 2003), suggesting a dynamic interplay between hippocampal and thalamocortical activity during sleep, potentially mediated by the DG (Sullivan et al., 2011).

Similarly, cortical UP and DOWN states during nREM sleep also influence ripple timing. UP states are characterised by high cortical activity, whereas DOWN states are marked by reduced cortical firing and the presence of delta waves (1–4 Hz cortical oscillations) (Peyrache et al., 2011). CA1 SWRs have been shown to occur more frequently during UP states and at the DOWN→UP transition, suggesting that cortical dynamics can gate the timing of ripples (Ji & Wilson, 2007; Peyrache et al., 2011; Sirota et al., 2003). This cortical control of ripple timing may be important for hippocampal–cortical communication: ripples occur during windows when cortical activity is reduced (DOWN states), enabling hippocampal output to efficiently reach the cortex (Peyrache et al., 2011).

In summary, while CA3 plays a major role in ripple generation, the occurrence, timing, and features of CA1 ripples are shaped by additional inputs from CA2, subiculum, DG, and MEC, as well as thalamic spindles and cortical slow oscillations. Together, these findings support the diverse ripple hypothesis, which proposes that ripples are not homogeneous events, but are instead modulated by distinct upstream inputs that influence their characteristics. For instance, one ripple might be primarily driven by CA3, another might originate in CA2, and yet another might be shaped by MEC input.

In the next section, I review recent studies that support this idea — that ripples are diverse — and highlight current knowledge gaps in the field. This sets the stage for the work presented in this thesis.

1.4 Ripple diversity: evidence, limitations, and open questions

In the sections above, I reviewed studies that challenge the canonical view of ripples as homogeneously originating from CA3, instead suggesting that ripple activity is diverse, with multiple regions influencing — or even generating — individual ripple events. This view aligns with a broader theme in hippocampal physiology: the hippocampus often leverages internal diversity, even within seemingly homogeneous states, to support distinct computational goals. For example, Lopes-dos-Santos et al., 2018 showed that during exploratory behaviour, individual CA1 theta cycles are not all the same. On a cycle-by-cycle basis, the hippocampus engages different gamma sub-bands, likely reflecting distinct circuit mechanisms to either encode new information or retrieve existing memories. If theta is dynamically modulated on a cycle basis, it seems natural and coherent to consider that ripples too are modulated on a ripple-by-ripple basis — supporting the idea of ripple diversity.

Most previous research on ripples, however, is based on averaging their LFP and spectral properties, potentially masking important ripple-by-ripple differences. This methodology may have unintentionally biased the consensus view, painting ripples as a uniform class of events — canonically generated in CA3, reactivating recently formed memories, and exhibiting stereotyped spectral profiles. In recent years, a few studies have directly challenged this assumption and provided empirical evidence that ripples are indeed diverse: in their LFP waveform, in how they engage single neurons, and in their broader brain-wide correlates.

For instance, Sebastian et al., 2023 applied tools from topological data analysis to show that ripple LFP waveforms in mice do not fall into discrete categories, but instead map onto a continuous low-dimensional space (i.e., a manifold). This manifold encodes not only spectral properties such as ripple frequency and amplitude, but also upstream input origin. To test this, the authors optogenetically triggered ripples by activating the axon terminals of either CA3 or CA2 pyramidal cells projecting to CA1, and found that these two conditions produced distinguishable ripple waveforms within the manifold — suggesting that the diversity in waveform shape carries information about which inputs initiated the ripple.

Similarly, Ramirez-Villegas et al., 2015 showed in rhesus macaques that ripple waveforms cluster into different subtypes with distinct spectral features and durations. They further combined LFP recordings with functional magnetic resonance imaging (fMRI) and found that these ripple subtypes mapped differently onto global brain-wide activity — some showing stronger neocortical engagement, while others aligning more with subcortical activity. This work not only supports ripple diversity in non-rodent models but also links it to differences in global output and potential function.

At the cellular level, Valero et al., 2017 combined juxtacellular and intracellular recordings in rats to show that ripple-related firing is not indiscriminate: specific subsets of CA1 pyramidal cells are selectively recruited in different ripple events. Interestingly, this selectivity was absent in chronic epilepsy models, suggesting that ripple-specific recruitment is tightly controlled and functionally meaningful.

What about population dynamics? Do different ripples engage different CA1 assemblies? X. Liu et al., 2022 addressed this by combining LFPs and calcium imaging in mice. They found that some ripples were spatiotemporally local within CA1, while others propagated more globally across the region, and these differences were associated with the recruitment of distinct cell assemblies. This reinforces the idea that ripples with distinct spatial signatures are linked to distinct mnemonic processing.

Building on this, Chang et al., 2025 showed that ripple diversity is not only structural but also functional. By recording from CA1 while tracking pupil diameter as a proxy for global neuromodulatory state, they found that ripples occurring during small-pupil (low arousal) periods were essential for consolidating recently acquired memories, whereas those during large-pupil (high arousal) states were more relevant for familiar memories. Using closed-loop optogenetics, they showed that selectively disrupting ripples during small- or large-pupil periods impaired the corresponding memory type — recent or familiar, respectively — suggesting that ripple diversity reflects a broader organisation of nREM sleep, which the authors refer to as “sleep microstructure”.

Together, these studies point to a clear conclusion: ripples are not uniform. Their diversity spans spectral features, cell recruitment, spatial extent, upstream inputs, and mnemonic function.

However, despite the progress made, a number of key questions remain unanswered. For instance, Sebastian et al. (2023) showed that distinct ripple waveforms are associated with distinct CSD profiles and upstream inputs (e.g., CA2 vs. CA3 stimulation), but did not directly test whether CA3-triggered ripples are actually associated with stronger sinks in *stratum radiatum*, as would be expected from canonical SWR features (Sullivan et al., 2011; Ylinen et al., 1995). That is, although they linked input pathways to ripple waveform and spectral features, they did not relate these triggered ripples back to established CSD signatures of SWRs, such as *radiatum* sinks for CA3-driven events. Similarly, the role of *stratum oriens* sinks in CA2-triggered ripples remains unclear.

Another gap is the over-focus on CA1 pyramidal cells. While studies like Valero et al., 2017 and X. Liu et al., 2022 have shown selective recruitment within CA1, the canonical view argues that ripples originate in CA3 (Buzsáki, 1986). Thus, a full picture of ripple diversity should also show distinct recruitment patterns in CA3, not just in CA1.

A more fundamental limitation is the definition of what counts as a ripple. Most of the studies cited above have focused on SWRs, not ripples per se. That is, ripple detection required the co-occurrence of a CA1 ripple and a sharp wave in *stratum radiatum*, effectively excluding events where the ripple lacked a clear SW. Valero et al., 2017, Sebastian et al., 2023, and Chang et al., 2025 all employed such criteria. In other cases, frequency cutoffs were used to exclude slower ripples — e.g., X. Liu et al., 2022 discarded events with frequency below 120 Hz. However, ripples with frequencies below this threshold still match the consensus definition of hippocampal ripples (A. A. Liu et al., 2022). If the goal is to understand ripple diversity, should these slower or SW-less ripples be disregarded? Arguably not. A proper investigation into ripple diversity should avoid overly restrictive assumptions and thresholds at the preprocessing stage.

In the next section, I describe how I addressed these limitations and constructed the analytical framework used in this thesis to study ripple diversity in an unbiased and comprehensive way.

1.5 A data-driven approach to investigate ripple diversity

The dominant framework for ripple generation emerging from previous work holds that bursts of synchronised activity in CA3 give rise to sharp waves and current sinks in *stratum radiatum*, which in turn trigger ripple oscillations in CA1. However, as discussed in the previous section, growing evidence suggests that ripple occurrence and characteristics are modulated by various intra- and extra-hippocampal inputs, strongly indicating that ripples are not homogeneous but instead diverse. This diversity appears to map onto distinct spectral features, CA1 cell recruitment patterns, and memory content.

Nevertheless, most previous studies have focused on a restricted subset of events — namely canonical SWRs — leaving open key questions about ripple diversity. In particular, the field lacks a detailed understanding of how ripple-by-ripple variability relates to upstream input, both in terms of laminar current profiles within CA1 and the recruitment of specific neurons, especially in CA3.

In this thesis, I present a new analytical framework aimed at capturing and characterising ripple-by-ripple diversity, spanning from laminar current flow in CA1 to network topology in CA1 and CA3, and finally to the memory content expressed by different ripple events. Crucially, this framework is based on minimal preprocessing assumptions and follows a purely data-driven approach.

To do so, I set out to answer the following questions:

1. **What input currents to CA1 shape ripple-by-ripple variability?**
   To answer this, I combined silicon probe recordings I performed in sleeping mice with additional recordings previously acquired in the lab. Using this combined dataset, I applied unsupervised methods to extract and classify laminar current source density (CSD) profiles on a ripple-by-ripple basis. In Chapter 3, I describe how these CSD profiles relate to distinct spectral features and ripple waveforms.

2. **How do hippocampal neurons respond to different ripple profiles?**
   To explore the circuit dynamics underlying ripple diversity, I first established a mapping between the current profiles observed in the silicon probe dataset and ripple events recorded with tetrodes, using a machine learning approach (Chapter 4). I then analysed an existing dataset of multichannel tetrode recordings from CA1 and CA3, previously acquired in the lab, to investigate how ripples with different current profiles engage single neurons, cell pairs, and population-level activity across both regions (Chapter 5).

3. **What content do different ripple profiles convey?**
   In Chapter 6, I assess how coactivity patterns expressed during previous wakeful exploration are reactivated during distinct ripple events. I examine whether different ripple profiles engage different subpopulations of CA1

pyramidal cells, and whether they correspond to different types of memory content.

Chapter 2 provides an overview of the experimental, analytical, and statistical methodologies used throughout the thesis. Finally, Chapter 7 summarises the key findings, situates them within the broader context of existing literature, and outlines future directions for the study of ripple diversity.

2

Experimental Procedures

Contents

2.1 Experimental methodology

2.1.1 Animals

These experiments used adult (4-6 months old) C57BL/6J wild-type mice (Charles River Laboratories, UK). Animals were housed with their littermates up until the start of the experiment. All mice were held in IVCs, with wooden chew sticks and nestlets in a dedicated housing facility with a 12/12 h light/dark cycle (lights on at 07:00), 19-23°C ambient temperature, and 40-70 % humidity. They had free access to water and food *ad libitum* throughout the experiment. Experimental procedures were performed on mice in accordance with the Animals (Scientific Procedures) Act, 1986 (United Kingdom), with final ethical review by the Animals in Science Regulation Unit of the UK Home Office.

2.1.2 Surgical procedure

All surgical procedures were performed under deep anaesthesia using isoflurane (0.5-2%) and oxygen (2 l/min), with analgesia provided before (0.1 mg/kg vetergesic) and after (5 mg/kg metacam) surgery.

For silicon probe recordings, mice were implanted with a single-shank silicon probe (Table 2.1) under stereotaxic control in reference to bregma, using central coordinates -2.0 mm anteroposterior from bregma, +1.7 mm lateral from bregma, and an initial depth of 1.5 mm ventral from the brain surface to span the somato-dendritic axis of CA1 principal cells and reach the DG. Following the implantation, the exposed parts of the silicon probe were covered with Vaseline® Healing Jelly, after which its plastic drive was secured to the skull using dental cement and stainless-steel anchor screws inserted into the skull. Two of the anchor screws, both above the cerebellum, were attached to a 50 μm tungsten wire (California Fine Wire) and served as ground. For the recordings, the silicon probe was positioned along the radial axis of CA1 pyramidal cells, using the rotations applied to its holding screw.

For tetrode recordings, mice were similarly implanted with a single microdrive containing 14 independently movable tetrodes, each positioned to target the *stratum pyramidale* of either CA1 or CA3 in the dorsal hippocampus. Tetrodes were constructed by twisting together four insulated tungsten wires (12 μm diameter, California Fine Wire) which were briefly heated to bind them together into a single bundle. Each tetrode was loaded in one cannula attached to a 6 mm long M1.0 screw to enable its independent manipulation of depth. The drive was implanted under stereotaxic control in reference to bregma using the following coordinates. For CA1 pyramidal cell layer tetrodes, the span was between AP[1] -2.0 to -2.4 mm and ML[2] 1.6 to 2.3 mm. For CA3 pyramidal cell layer tetrodes, the span was between AP -1.8 to -2.2 mm and ML 2.0 to 2.7 mm. The initial depth of the tetrodes during the implantation surgery was 1.0 mm ventral from the brain surface. The distance between neighbouring tetrodes was 350 μm. Following the implantation, the exposed parts of the tetrodes were covered with paraffin wax, after which the drive was secured to the skull using dental cement and stainless-steel anchor screws inserted into the skull. Two of the anchor screws, both above the cerebellum, were attached to a 50 μm tungsten wire (California Fine Wire) and served as ground. For the recordings, each tetrode was lowered along the vertical axis to reach either the CA1 or CA3 pyramidal layers, using the rotations applied to its tetrode cannula-holding screw and the electrophysiological profile of the local field potentials in the hippocampal ripple frequency band, with the final depth position subsequently confirmed by histology of anatomical tracks.

n	Implant
1	NeuroNexus A1x32-5mm-25-177-H32_21mm
1	NeuroNexus A1x64-edge-6mm-20-177-H64LP_30mm
3	Cambridge NeuroTech ASSY-236 H3 Chronic 64-Molex

Table 2.1: Silicon probe implants details: The mice used were of strain C57Bl6/J with Research Resource Identifier (RRID) IMSR_JAX:000664. The table summarises the number of mice (n) implanted with silicon probes and the specific probe models used.

---

[1] Anteroposterior.

[2] Mediolateral.

2.1.3 Recording procedure

Following implantation surgery, mice were allowed to recover for at least seven days. They were then familiarised with the recording procedure, being handled daily in a dedicated towel, connected to the recording cable, and exposed to the sleep box for at least 30 minutes per day for a minimum of four days. During this period, both silicon probes and tetrodes were gradually lowered toward the target cell layers. For silicon probe implants, probes were positioned to target the CA1 region, and once the appropriate depth was reached, they were left in place for the remaining days of the experiment. For tetrode implants, tetrodes were adjusted every day to target the CA1 and CA3 pyramidal cell layers. Electrophysiological profiles, including local field potential characteristics such as sharp-wave ripples and gamma oscillations, were used to guide daily placement. On each recording day, tetrodes were lowered into the target layers in the morning to capture ensemble spiking activity and left in place for approximately 1.5-2 hours before recordings began that day. At the end of the recording day, tetrodes were raised by approximately 150 μm to prevent mechanical damage to the hippocampal layers. The following morning, tetrodes were re-adjusted to locate new cells, minimising the likelihood of recording from the same neurons across days.

Each recording day began with a baseline sleep/rest session (pre-exploration sleep/rest), followed by an exploration session to finish with a post-exploration sleep/rest session (post-exploration sleep/rest). The environments used in these recordings were either open-field arenas for exploration sessions (e.g., 41 cm diameter cylinder, 41 × 41 cm square box; all with 30 cm high walls) or the sleep box (12 × 12 × 28 cm; containing sawdust bedding and nesting material). After placing the mouse in the sleep box, experimenters monitored the animal's movements and real-time raw electrophysiological signals to confirm that the mouse had started to sleep/rest. In the absence of electromyographic or any other signals to categorise sleep stages, I refer to these offline periods of extended immobility as sleep/rest. Exploration sessions lasted $\sim$30 minutes; sleep/rest sessions lasted $\sim$60-90 minutes.

All experiments were conducted under dim light conditions (∼20 lux) with low-level background white noise (∼50 dB). In total, the silicon probe dataset included 38 sleep sessions [mean duration (IQR): 82.4 (59.5–105.9) minutes per session] from 5 mice. The tetrode dataset included 244 sleep sessions [mean duration (IQR): 68.8 (48.0–90.2) minutes per session] from 13 mice, 11 of which had dual-site recordings from CA1 and CA3, and 2 with CA1 recordings only.

2.1.4 Acquisition of multichannel data and tracking of animal position

Extracellular signals were recorded using an integrated circuit mounted on the animal's head (model RHD2164, Intan Technologies; http://intantech.com/products_RHD2000.html), which provided a frequency response from 0.09 Hz to 7.60 kHz during the amplification stage. The amplified and filtered signals were digitised at a sampling rate of 20 kHz. These digitised signals were stored alongside additional data streams, including digital pulses indicating the animal's position (via transistor-transistor logic) and signals from a three-axis accelerometer integrated into the head-mounted device, which measured head movements and provided an additional measure of the animal's movement. Positional data were obtained using an overhead colour camera (https://github.com/kevin-allen/positrack/wiki), which tracked the movement of LED clusters in three distinct colours affixed to the electrode assembly. These positional signals were captured at a rate of 39 frames per second.

2.1.5 Spike detection and unit isolation

Spike sorting and unit isolation used an automated clustering approach, leveraging Kilosort (https://github.com/cortex-lab/KiloSort) within the SpikeForest framework (https://github.com/flatironinstitute/spikeforest), as outlined in Pachitariu et al., 2016 and Magland et al., 2020. For data acquired using tetrodes, KiloSort's algorithm was adapted to limit templates to channels associated with a specific tetrode bundle and to exclude all other recording channels. Data from all sessions recorded within a single day were concatenated and processed collectively, enabling continuous cell tracking across the day. The clusters generated were manually confirmed by examining cross-channel spike waveforms, auto-correlation histograms, and cross-correlation histograms. Units selected for analysis consistently exhibited stable spike waveforms, a distinct refractory period in their auto-correlation histograms, and no refractory periods in cross-correlation histograms with other units throughout the day.

As highlighted throughout the following Chapters, all neuronal spikes-related analyses in this thesis are based on data acquired using tetrode recordings, rather than silicon probes, reflecting the lab's methodology.

2.1.6 Principal cell versus interneuron classification

Hippocampal principal cells and interneurons were distinguished using features of their spike waveforms, as described previously (Lopes-dos-Santos et al., 2023). Briefly, waveform consistency for each unit was evaluated using the waveform with the maximum amplitude across tetrode channels for each cluster. To quantify the prominence of a unit's mean waveform amplitude relative to its spike-to-spike variability, I computed a waveform score:

$$wv_{score} = \sqrt{\sum_{i=1}^{n} \frac{(w_i/\sigma_i)^2}{n}} \tag{2.1}$$

, where w_i is the value of the mean waveform at sample i, σ_i is the standard deviation at sample i across all spikes, and n is the number of waveform samples. This metric reflects the relative magnitude of the mean waveform amplitude compared to variability across spikes. Units with a waveform score above 0.75 and fewer than 2% refractory period violations (intervals < 2 ms in the inter-spike interval distribution) were included for further analysis. Putative interneurons and principal cells were then classified based on the trough-to-peak latency of their waveforms. In a prior dataset of 4,000 neurons, the trough-to-peak latency exhibited a bimodal distribution. A one-dimensional, two-component Gaussian mixture model was fitted to this distribution, and the intersection of the two components was used as the classification threshold: units with latencies above the threshold were classified as putative principal cells, and those below as putative interneurons. I applied the same inclusion criteria to the principal cells and interneurons in the tetrode dataset of this study. In total, this study includes 2,196 CA1 principal cells and 1,325 CA3 principal cells, with 408 CA1 interneurons and 333 CA3 interneurons, recorded across 83 days in 13 mice [mean number of cells per recording day (IQR): CA1 principal cells, 26.8 (15.3 - 38.0); CA3 principal cells, 19.8 (14.0 - 25.0); CA1 interneurons, 5.0 (3.3 - 6.0); CA3 interneurons, 5.0 (3.0 - 6.0)]. For analyses involving principal cell correlations or distances between population vectors across ripple types, I addressed the issue of highly sparse spike trains and imbalanced ripple counts across groups by using sleep sessions with at least 250 ripples from each group ($\mathrm{Rad}^{\mathrm{sink}}$ or $\mathrm{LM}^{\mathrm{sink}}$) and five principal cells with an average firing rate of at least 0.25 Hz over the entire recording day. This criterion was met by 208 sleep/rest sessions (across 13 mice) for CA1 principal cells and 171 sessions for CA3 principal cells (across 11 mice), resulting in 1,580 CA1 and 866 CA3 principal cells.

2.2 Local field potential analysis

2.2.1 Preprocessing of Local Field Potential (LFP) signals

LFP signals underwent initial filtering using an 8th-order Chebyshev type I anti-aliasing filter, applied to the wide-band signals sampled at 20 kHz. These filtered signals were then down-sampled to a rate of 1,250 Hz, employing the decimate function within `scipy`'s `signal` submodule (version 1.11.2).

2.2.2 Ripple event detection

LFPs were first referenced to a channel without CA1 ripples. This differential signal underwent dual stage-filtering: through a ripple-specific bandpass filter (80-250 Hz, 4th-order Butterworth Filter), and then through a high-frequency bandpass filter (200-500 Hz, 4th-order Butterworth Filter). Instantaneous signal characteristics, including envelopes and phases, were derived using the Hilbert transform. Ripple events were identified by detecting envelope peaks within the ripple band that exceeded a threshold of five times the median value. In instances of multiple peaks within a 20-ms window, only the peak with the highest amplitude was considered. For each event, the onset and offset points were identified - relying on the envelope's decrease to below half of the established threshold. Analysis extended to quantifying the ripple cycle count within each event by examining phase shifts, with cycle calculation based on the unwrapped phase difference between event onset and offset as previously described (Lopes-dos-Santos et al., 2023). The mean frequency of each event was calculated by dividing the total cycle count by the event's duration. Finally, I validated candidate ripple events using four criteria: (1) Ripple band power in the detection channel, calculated as the squared mean amplitude, needed to be double that of the reference channel, ensuring detected events were prominent; (2) The mean frequency of detected events must exceed 80 Hz; (3) Each event should contain at least four complete ripple cycles; (4) Ripple band power should be at least double that of the control high-frequency band. Note that this detection method, unlike those used in previous studies (see Section 1.4), does not require candidate ripples to be accompanied by a sharp wave in *stratum radiatum* — i.e., it allows me to detect ripples without restricting the analysis to canonical SWRs. In total, $942,331$ ripples were detected across all datasets. All detected ripples were included in the analyses unless stated otherwise, as appropriate [mean number of detected ripples per sleep/rest session (IQR): 3,341.6 (2,130.3–4,413.5)].

For specific analyses requiring the prevention of contamination by overlapping events — which could otherwise bias average signals — only isolated ripples were used. Isolated ripples were defined as ripple events with no other detected ripple occurring within ± 250 ms. This restriction applies to all ripple-triggered average signals, including the triggered-average CSD (Figures 3.1 and 3.4) and peri-ripple neuronal responses (Figures 5.3 and 5.4). Unless stated otherwise, all ripple-triggered average signals were referenced to the ripple envelope peak.

2.2.3 Determination of CA1 pyramidal layer reference channel

To identify the optimal reference channel for ripple events and theta oscillations within the CA1 pyramidal layer, I computed a ripple band score for each channel. This score was calculated by dividing the power in the ripple band ($80 - 250$ Hz) by the power in an adjacent frequency range ($70 - 300$ Hz), using a Welch spectrum (4-second Hann windows overlapping by 50%). The channel with the highest score was set as a reference channel within the CA1 pyramidal layer.

2.2.4 Extraction of theta oscillations from LFPs

To isolate theta oscillations from the LFP data in exploration sessions, I employed the masked Empirical Mode Decomposition method (Deering & Kaiser, 2005; Huang et al., 1998) as implemented in Quinn, Lopes-dos-Santos, Dupret, et al., 2021. With this, I adopted the mask sift procedure using specific mask frequencies set at 350, 200, 70, 40, 30, and 7 Hz, following the parameters optimised in Quinn, Lopes-dos-Santos, Huang, et al., 2021 grounded in Fosso and Molinas, 2018. For each mask, the amplitude was set to three times the standard deviation of the input signal. This procedure decomposes each LFP signal into oscillatory components termed intrinsic mode functions (IMFs) from faster to lower frequency components. Upon completion of this procedure with the aforementioned parameters, six IMFs and a residue were computed, with IMF-6 effectively isolating theta oscillations.

To delineate individual theta cycles, I began by pinpointing peaks and troughs (i.e., the local maxima and minima, respectively) of the obtained theta IMF. The residue of the LFP not captured by the six IMFs was defined as the lower frequency component of the signal and its envelope was used as the amplitude threshold for retaining peaks and troughs for the next step. We then defined each peak-trough-peak sequence as a candidate theta cycle. We took as valid cycles sequences having their peak-trough and trough-peak intervals falling within the 31 to 100 ms range (corresponding to the half period of cycles with frequencies ranging from

∼16 to 4 Hz); and peak-to-peak distance was between 71 ms (equivalent to ∼14 Hz) and 200 ms (equivalent to 5 Hz).

For each validated cycle, I found six control points: the zero-crossing prior to the first peak, the peak itself, the subsequent zero-crossing post the first peak, the trough, and the zero-crossing following the trough. Then, I computed the instantaneous theta phase for each timestamp through a linear interpolation of the control points (Belluscio et al., 2012; Lopes-dos-Santos et al., 2018).

2.2.5 Wavelet spectrograms

To visualise the spectral components of LFPs (i.e., frequency and corresponding power), I used spectrograms. The spectrograms in Figure 3.5 were generated using the complex Morlet Wavelet Transform. For this analysis, 50 logarithmically spaced frequencies were selected, spanning from 2 Hz to 300 Hz unless otherwise specified. Each wavelet kernel was L1-normalised, meaning the sum of the absolute values of the elements in the kernel was set to 1. This normalisation ensured that the wavelet preserved the relative amplitudes of individual frequency components without introducing artificial amplification or attenuation.

2.3 Current Source Density (CSD) analysis

Current source density (CSD) analysis provides insight into the spatial distribution of extracellular currents across different layers of CA1 by computing relative differences in LFP signals (Lopes-dos-Santos et al., 2023; Mitzdorf, 1985; Sullivan et al., 2011). This method was applied to ripple-triggered LFP recordings obtained with linear silicon probes to study current flow within CA1 (e.g., see Figures 3.1 and 3.2).

CSD analysis distinguishes two primary components: sinks (blue), which indicate active depolarizing currents generated by glutamatergic input pathways, and sources (red), which reflect either passive return currents or active hyperpolarization mediated by GABAergic inhibition. These components allow for the characterisation of layer-specific activity in the hippocampus.

2.3.1 Computing CSD

To examine the spatial distribution of currents during ripples (see Chapter 3), LFP signals were first aligned to the ripple-band envelope peak. I then calculated the CSD at each channel n and time point using the formula:

$$\mathrm{CSD}_n = -(\mathrm{LFP}_{n-1} - 2 \cdot \mathrm{LFP}_n + \mathrm{LFP}_{n+1}) \tag{2.2}$$

where $n-1$ and $n+1$ represent the channels directly above and below channel n, respectively. This formula computes the second spatial derivative of LFPs across channels, effectively highlighting regions of active current flow. This approach provided a CSD profile for each ripple event, revealing the spatial organisation of currents within CA1.

2.3.2 Defining CA1 laminar structure

For accurate layer-specific analysis, I assigned CSD channels to the *oriens*, *radiatum*, and *lacunosum-moleculare* layers based on ripple and sharp-wave laminar profiles, combined with electrode spacing, following the methodology presented in Lopes-dos-Santos et al., 2023. To ensure uniform spatial resolution across silicon probes with different channel spacings (Table 2.1), I applied Gaussian kernel smoothing with a standard deviation (SD) of $50\,\mu$m.

2.4 Analytical details of data analyses

2.4.1 Principal component analysis (PCA)

In several analyses throughout this thesis, I employed PCA for dimensionality reduction. Briefly, given a matrix X with dimensions $(n \times p)$, where n represents the number of observations (e.g., number of ripples) and p the number of dimensions (e.g., number of neurons), PCA is performed by applying singular value decomposition (SVD) to:

$$X = U\Sigma V^\top \tag{2.3}$$

Here, U (with dimensions $(n \times n)$) is the matrix of left singular vectors, representing the directions of variability in the data for each observation. Σ (with dimensions$(n \times p)$) is a diagonal matrix with the singular values σ_i, which are related to the variance captured by each principal component. $V^\top$ (with dimensions $(p \times p)$) is the matrix of right singular vectors, where each column corresponds to a principal component (PC). These right singular vectors project the original matrix X onto a new set of axes, i.e., the principal components, which capture most of the variance in the data.

To quantify the contribution of each PC to the overall variance, I computed the variance explained by each PC as:

$$\text{Variance Explained} = \frac{\sigma_i^2}{\sum_{j=1}^{p} \sigma_j^2} \tag{2.4}$$

where σ_i^2 is the variance explained by the ith PC. The PC that explains the highest variance is denoted as PC_1.

2.4.2 Peri-event time histograms (PETHs)

In Figures 5.3 and 5.4, to visualise and measure how hippocampal cells responded during ripples, I constructed peri-event time histograms (PETHs) over 400-ms windows, covering 200 ms on either side of the envelope peak of isolated ripples (see Section 2.2.2), with a bin width of 0.8 ms. For each cell group (e.g., principal cells or interneurons in CA1 or CA3), I first computed the raw firing rate responses during Rad^{sink} and LM^{sink} ripples. These responses were then smoothed using a Gaussian kernel (SD = 5 ms) to determine the peak firing rate, defined as the maximum firing rate within a 50-ms window centred on the ripple peak (Figures 5.3 and 5.4). To visualise responses across all recorded cells, I z-scored the responses relative to their mean and standard deviation during Rad^{sink} ripples and applied an additional smoothing step using a Gaussian kernel (SD = 5 ms).

2.4.3 Coactivity analysis of CA1 and CA3 principal cells

To analyse the overall structure of coactive principal cells during ripples, I computed coactivity graphs as previously described in Gava et al., 2024. Below, I describe how the procedure for constructing neuronal graphs and calculating coactivity metrics from that study was adapted for the analyses in this thesis.

Construction of neuronal coactivity graphs

To analyse the coactivity of CA1 (or CA3) principal cells during ripples, I constructed corresponding neuronal graphs. For each sleep session, I first computed the spike count of each principal cell within a 50-ms window centred around the time of each ripple envelope peak. I then compiled these ripple-triggered activity profiles into a matrix with dimensions defined by the number of cells and ripple events, ($N_{\text{cells}} \times N_{\text{ripples}}$). I constructed separate matrices for events classified as Rad^{sink} and LM^{sink} ripples.

Using these matrices, I constructed coactivity graphs by comparing the ripple-event rates between all pairs of principal cells. To control for the shared influence of general network activity between pairs of neurons (i, j), I fitted a generalised linear model (GLM) to obtain the regression coefficient β_{ij} :

$$x_j \sim \beta_{ij} x_i + \alpha_{ij} P \tag{2.5}$$

where x_j and x_i are the z-scored ripple-nested spike trains of individual neurons j (the target) and i (the predictor), and P represents the summed activity of the other $N - 2$ neurons:

$$P = \sum_{n=0}^{N-\{i,j\}} x_n \tag{2.6}$$

In Equation 2.5, α_{ij} weights the influence of the population contribution on the activity of the target neuron j. Thus, the recorded neurons and their coactivity defined the nodes and edges of the coactivity graphs for each sleep session. I characterised each neuronal graph by its adjacency matrix, A, which is a $N \times N$ square matrix encapsulating the ripple-associated pairwise coactivity interactions across the network. This results in a signed and weighted graph with no self-connections:

$$A = \begin{pmatrix} \beta_{0,0} & \cdots & \beta_{0,N} \\ \vdots & \ddots & \vdots \\ \beta_{N,0} & \cdots & \beta_{N,N} \end{pmatrix} \tag{2.7}$$

with $\beta_{i,i} = 0 \; \forall i \, \text{in} \, N$. A schematic of this method, along with an example of a ripple adjacency matrix, is shown in Figure 5.11.

Single-neuron coactivity strength

From the coactivity graphs, I computed the mean coactivity of each neuron with the rest of the population of cells recorded that day. In Figure 5.12, I define the single-neuron coactivity strength as the average pairwise activity correlation of a given node with the other nodes in the weighted graph ($n = 1,580$ CA1 and 866 CA3 principal cells).

As a reference, the strength in a weighted graph can be compared to the degree in a binary graph, which accounts for the number of neighbours of the node. Here, the strength S_i of a node i is the average across all the weights β_{ij} of the edges connected to that node:

$$S_i = \frac{\sum_{j=0}^{N} \beta_{ij}}{N} \tag{2.8}$$

where N is the number of neurons j that node i projects to.

2.4.4 Assessing model predictive performance

In this thesis, I use a combination of machine learning models to test relationships between variables for both regression (e.g., predicting the CSD value in the *stratum radiatum* from the pyramidal LFP waveform, see Section 4.2.1) and classification (e.g., discriminating Rad$^{\text{sink}}$ vs. LM$^{\text{sink}}$ ripples based on ripple identity, see Section 5.2.4). This section outlines the primary methods used to evaluate model performance.

In general, a model tests the statistical relationship or 'correlation' between an independent variable X and a dependent variable y, where y depends on the values of X. The domain of y determines the appropriate algorithm: if $y \in \mathbb{R}$ (i.e., y is a continuous variable), regression models are used to predict its values based on X. Conversely, if y is a discrete variable (e.g., categorical labels such as ripple type), classification models might be more appropriate.

Prior to model training (unless stated otherwise), the dataset was split into 80% training and 20% testing subsets ($X_{\text{train}}, X_{\text{test}}, y_{\text{train}}, y_{\text{test}}$). Models were trained using the training set and evaluated on the test set. This process was repeated over 20 cross-validation iterations, and the average performance across these iterations was used as the model's accuracy.

Regression

Throughout this thesis, all regression analyses employed linear regression from the `linear_model` module of `scikit-learn`. To assess model performance on the testing set (y_{test}), in each cross-validation iteration, the trained model generated predictions on X_{test}, producing y_{pred}. Model performance was then evaluated using Pearson correlation (computed with the `scipy.stats` module) between the true and predicted values (y_{test} and y_{pred}). A correlation value of 1 indicated a perfect prediction, while a value of -1 signified a perfect prediction but in the opposite direction. A correlation value of 0 implied no correlation, meaning the model could not make reliable predictions.

These regression models were used in Sections 4.2.1 and 6.2.4.

Classification

To classify y from X, I exclusively used linear models, such as Linear Discriminant Analysis (LDA) and logistic regression, implemented in the `discriminant_analysis` and `linear_model` modules of `scikit-learn`.

In each cross-validation iteration (unless stated otherwise), I balanced the number of observations across classes by downsampling to match the class with the fewest samples. This ensured that the model was not biased toward the more frequent class. After training the models on the training set, performance was assessed using accuracy and mutual information (MI) between y_{test} and y_{pred} (computed using the `scikit-learn`'s `metrics` module).

The MI score quantifies the dependence between two categorical variables, similar to accuracy, but normalised relative to chance level. In a binary classification task (e.g., distinguishing between two ripple types), a chance-level prediction (50% accuracy) corresponds to an MI score of 0 bits, while a perfect prediction (100% accuracy) corresponds to an MI score of 1 bit, as illustrated in Figure 2.1.

These classification models were applied in Section 4.2.3 (accuracy-based evaluation) and Section 5.2.4 (MI-based evaluation).

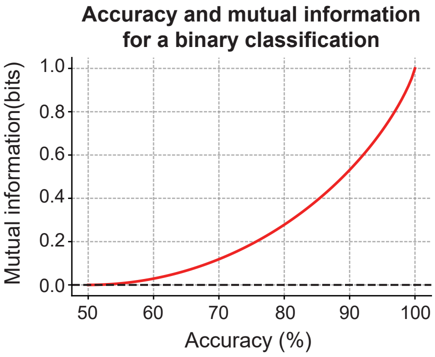

Figure 2.1: Relationship between mutual information and classification accuracy.
Plot illustrating the relationship (red line) between mutual information (MI) and accuracy in a binary classification task. The *dashed line* represents the chance level. Notably, 0 bits of mutual information corresponds to 50% accuracy (chance level).

2.4.5 Spikes shuffling control

In the analyses outlined in Sections 5.2.4, 5.2.8, and 6.2.3, I used a shuffling procedure to estimate chance levels in classifiers distinguishing between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples based on neural activity. To control for both individual neuron firing rates and overall population activity during ripples, I used a shuffling method that preserved these features while disrupting neuronal coactivity patterns.

To do so, the original spike matrix, where each row represents a neuron and each column represents a ripple event, was shuffled to maintain the sum of spikes for each neuron (i.e., individual neuron firing rates) and the total spike count across all neurons for each time bin (i.e., overall population activity). This procedure began by calculating the total spike count for each neuron across all ripples and generating a list of indices corresponding to the times (i.e. the ripples' indices) of these spikes. The list of spike indices was then randomly shuffled to ensure a uniform redistribution of spikes while maintaining the original firing rate constraints. The shuffled spike indices were iteratively reassigned to a new matrix, ensuring that the total spike counts per neuron (row) and per ripple (column) remained consistent with the original matrix. When applied to binary spike matrices (i.e., active/inactive neurons; see Section 5.2.8), this procedure automatically controlled for the number of ripples in which a cell was active and the overall population sparsity in each ripple.

This approach effectively controlled for individual neuron firing rates and overall population activity while disrupting coactivity patterns within each time bin. This shuffling control was employed in Figures 5.9, 5.17 and 6.2(C).

2.5 Statistical analysis

Data analyses were conducted using Python versions 3.6 and 3.10, incorporating the following packages: `scikit-learn` (Pedregosa et al., 2011), `NetworkX` (Hagberg et al., 2008), `NumPy` (Harris et al., 2020), `SciPy` (Virtanen et al., 2020), `Statsmodels` (Seabold & Perktold, 2010), `Matplotlib` (Hunter, 2007), `Pandas` (McKinney, 2010), and `Seaborn` (Waskom, 2021). Statistical comparisons included t-tests against a mean (for comparing a distribution to a fixed value), permutation tests (see subsections below) for comparing two conditions, and one-way ANOVA for multiple conditions, followed by a Tukey post-hoc test.

All confidence intervals (95% CI) in this thesis were calculated via bootstrapping with 100,000 resamples unless stated otherwise. For each interval, data were resampled randomly with replacement, and the 2.5th and 97.5th percentiles of the bootstrapped distribution defined the lower and upper bounds of the CI.

Bootstraps and permutation tests

Bootstrap tests were employed to compare two conditions, x and y, accommodating both paired and unpaired comparisons. The bootstrapped mean difference (Δ_{boot}) was calculated for each iteration as either an absolute value or a percentage relative to one of the variables. This process involved resampling the data 100,000 times (unless stated otherwise) with replacement. For paired comparisons, indices were resampled to preserve the pairing structure, while for unpaired comparisons, each condition was resampled independently.

The bootstrapped mean difference for a single iteration i was computed as:

$$\Delta_{\text{boot}_i} = \overline{x_{\text{boot}_i}} - \overline{y_{\text{boot}_i}}$$

where x_{boot_i} and y_{boot_i} are the i^{th}bootstrap resamples of x and y, respectively, and $\overline{x_{\text{boot}_i}}$ and $\overline{y_{\text{boot}_i}}$ represent their respective means.

P-values were computed numerically under the null hypothesis of zero difference. For one-sided tests, the p-value was calculated as the proportion of all bootstrap samples where the effect size (Δ_{boot}) was either greater than or less than zero, depending on the test direction. For two-sided tests, the p-value was determined by multiplying the smaller proportion of bootstraps below or above zero by two:

$$p_{\text{two-sided}} = 2 \cdot \min\left(\frac{\#(\Delta_{\text{boot}} < 0)}{N}, \frac{\#(\Delta_{\text{boot}} > 0)}{N}\right)$$

where N is the total number of bootstrap iterations. Throughout the thesis, unless stated otherwise, p-values were computed using two-sided tests.

Mean difference estimation plots

Mean difference estimation plots, also known as 'Gardner-Altman' plots, were used for visualisation to illustrate both the distributions of individual groups and the effect sizes by comparing mean or median differences across groups (e.g., Figures 3.5(C,D) and 5.9). These plots provided a clear visual representation of the data, complementing the numerical statistical analyses. The resampling approach used to generate the mean differences in these plots was similar to the bootstrap procedure described in the 'Bootstraps and Permutation Tests' Section. Specifically, data were resampled 5,000 times (unless stated otherwise) with replacement to calculate the bootstrapped mean differences and their corresponding confidence intervals. Each plot consisted of two panels. The top panel displayed raw data distributions with group means ± the standard error of the mean (SEM), while the bottom panel showed the mean or median differences relative to a reference group. Black dots represented the bootstrapped mean or median, black ticks indicated the 95% confidence intervals, and bootstrapped error distribution curves illustrated the variability. Importantly, p-values for statistical significance were not derived from the mean difference plots. Instead, to keep consistency across analyses, p-values were computed using the bootstrap method described in the 'Bootstraps and Permutation Tests' Section.

The real problem with the averaging procedure, of course, is that the state of the brain is constantly changing. [...]. State variability is, to a large extent, internally coordinated. This "correlated brain noise," as it is often referred to, might be critically important because it is a potential source of mental operations.

— György Buzsáki (Buzsáki, 2006)

3

Profiling individual hippocampal ripples using their laminar currents

Contents

3.1 Introduction

As discussed in the first chapter of this thesis, hippocampal ripples are often studied as a uniform class of events, primarily through averaging approaches. However, their underlying network dynamics may exhibit meaningful variability (Chang et al., 2025; Ramirez-Villegas et al., 2015; Sebastian et al., 2023; Valero et al., 2017), similar to what has been observed in awake theta cycles (Lopes-dos-Santos et al., 2018). For instance, ripples have traditionally been associated with current sinks in *stratum radiatum* (Buzsáki et al., 1983; Buzsáki, 1986; Sullivan et al., 2011; Ylinen et al., 1995). Here, I hypothesised that the laminar organisation of hippocampal inputs, which terminate in distinct CA1 layers, may result in ripple-associated current flow exhibiting heterogeneous circuit mechanisms reflected in distinct laminar current profiles.

The objective of this chapter is to characterise the laminar current signatures of individual ripples and determine whether distinct subtypes emerge based on their CSD features. To address this, I analysed ripple-aligned CSD signals obtained from silicon probe recordings spanning the CA1 radial axis. By examining the variability in ripple-nested CSD signals, I identified the dominant features shaping these ripple profiles and assessed whether ripples naturally segregate into distinct subtypes. The results indicate that, while ripple-associated currents form a continuum, two predominant patterns emerge: $\mathrm{Rad}^{\mathrm{sink}}$ ripples, which exhibit a strong sink in *stratum radiatum*, and $\mathrm{LM}^{\mathrm{sink}}$ ripples, characterised by a prominent sink in *stratum lacunosum-moleculare*. These profiles consistently appeared across animals and recording sessions.

Finally, I examine how these ripple subtypes relate to distinct spectral properties, including differences in ripple waveform, frequency, and power.

3.2 Methods

3.2.1 Extracting single-ripple CSD signatures

To extract the CSD profile shaping each individual ripple, I computed the mean CSD signal (see Section 2.3) within a 50-ms window around the ripple envelope peak for each channel. This yielded a curve describing the average CSD signals for each ripple event, which I define as the CSD signature of that event. The purpose of this analysis was to extract the CSD profile associated with each ripple event, enabling the characterisation of ripple-specific laminar features.

3.2.2 Principal component analysis of single-ripple CSD signatures

To analyse the variance across CSD signature profiles and identify dominant features associated with ripple events, I applied principal component analysis (PCA) to the CSD signatures (see Section 2.4.1). The principal component (PC) that explains the highest variance represents the most dominant or recurring CSD profile, as it captures the majority of variability across ripple events.

Since the sign of PCs is assigned arbitrarily during the decomposition shown in Equation 2.3, I adjusted the sign of the first PC to ensure consistency in its interpretation across different recording sessions. Specifically, the sign was adjusted so that the first PC consistently exhibited negative weights within *stratum radiatum*.

3.2.3 Identification of Rad^{sink} and LM^{sink} profiles from CSD signatures

To investigate the variations in CSD profiles captured by PC1 across different ripple events, I categorised the distribution of PC1 strengths into three groups, focusing on the extremes of the distribution to capture the most pronounced CSD signatures. Since PC1 was adjusted to exhibit negative weights in *stratum radiatum* — indicating that a higher PC1 value corresponds to a stronger current sink in this layer — ripples with PC1 strengths exceeding the 70th percentile were classified as Rad^{sink} events. Conversely, ripples with PC1 strengths below the 30th percentile were designated as LM^{sink} events. The remaining ripples were placed in a 'baseline' category (e.g., used in Figures 3.5 and 4.4).

3.3 Diversity of ripples laminar current profiles

3.3.1 Characterising the diversity of ripple CSD profiles

To investigate the diversity of hippocampal ripples, I began by analysing simultaneously recorded LFPs across all layers along the somato-dendritic axis of CA1, from *stratum oriens* to the hippocampal fissure, using silicon probes (Table 2.1). This recording setup enabled both ripple detection in the LFPs of *stratum pyramidale* and Current Source Density (CSD) analysis of laminar currents during sleep/rest (Figures 3.1 and 3.2; 50.9 total hours of sleep/rest from 5 mice; 38 sleep/rest sessions; mean duration (95% CI): 58.3 (46.9 - 69.7) minutes per sleep/rest session).

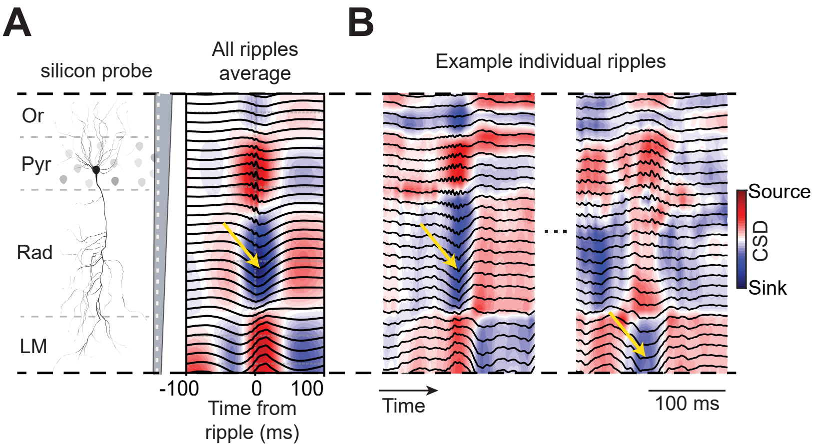

Figure 3.1: CSD profiles of individual ripples vs. their grand average pattern. (A) Mean current source density (CSD; colour-coded map) and local field potential (LFP) waveform (black traces) for ripples recorded in *stratum pyramidale* using a silicon probe spanning the parallel neural layers of CA1 (*oriens* [Or], *pyramidale* [Pyr], *radiatum* [Rad], and *lacunosum-moleculare* [LM]). A yellow arrow highlights the location of a current sink in *stratum radiatum.*
(B) Example CSD and LFP waveforms from two individual ripples recorded during the same sleep session, each exhibiting a distinct current sink profile: the left example features a sink in *stratum radiatum*, while the right one shows a sink in *stratum lacunosum-moleculare.*

Consistent with previous findings (Buzsáki, 2015; Lopes-dos-Santos et al., 2023; Sullivan et al., 2011; Ylinen et al., 1995), the average ripple profile displayed a current sink in *stratum radiatum* and a current source in *stratum lacunosum-moleculare* (Figure 3.1(A)). However, individual ripples exhibited considerable variability in their CSD profiles (Figures 3.1(B) and 3.2(A,B)). Most ripples featured a current sink in *stratum radiatum* (Figures 3.1(B) and 3.2(A)), aligning with the grand average CSD profile (Figure 3.1(A)). The strength of this *stratum radiatum* sink positively correlated with the strength of a concurrent sink in *stratum oriens* (Figures 3.2(C,D)). In addition to these canonical events, some ripples instead exhibited a strong current sink in *stratum lacunosum-moleculare* and a concurrent source in *stratum radiatum* (Figures 3.1(B) and 3.2(B)), revealing deviations from the canonical average profile.

This variability raised a central question: does the heterogeneity in ripple-associated currents carry functional significance? If so, how can I investigate this variability on a single-ripple basis, and how does this reflect the overall distribution of ripple events?

To address this, I extracted the 'CSD signature' of each ripple, defined as the laminar CSD signal within a 50 ms window centred on the ripple power peak (see Section 3.2.1). This approach allowed each ripple to be represented as a curve characterising its laminar current profile along the CA1 somato-dendritic axis (Figure 3.3(A) and 3.2(A,B)). To understand how these CSD signatures were distributed across ripples, I applied Principal Component Analysis (PCA) (see Section 3.2.2). The first principal component (PC1) explained 80% of the variance across all ripples and was highly consistent across sleep sessions (Figure 3.3(B,C)). Importantly, PC1 closely resembled the grand average CSD, with large negative weights (i.e., a current sink) in *stratum radiatum* and positive weights (i.e., a current source) in *stratum lacunosum-moleculare*, confirming that ripples, on average, are characterised by a current sink in radiatum and a source in lacunosum-moleculare (Figure 3.1(A) and 3.3(C)).

To investigate how individual ripples deviated from this canonical profile, I projected each ripple's CSD signature onto PC1, obtaining a distribution of scores reflecting how similar each ripple was to the grand average profile (Figure 3.3(D–E)). Ripples with high positive scores exhibited strong *radiatum* sinks, while ripples with negative scores featured a *radiatum* source and a sink in *lacunosum-moleculare* (Figure 3.3(C–E)). Consistent with previous studies (Chang et al., 2025; Sebastian et al., 2023), these PC1 projections delineated a continuum rather than discrete clusters (Figure 3.3(B,C)). However, within this continuum, two recurring profiles consistently emerged at opposite extremes across mice. One profile displayed the canonical strong current sink in *stratum radiatum*, which I refer to as $\mathrm{Rad}^{\mathrm{sink}}$ ripples

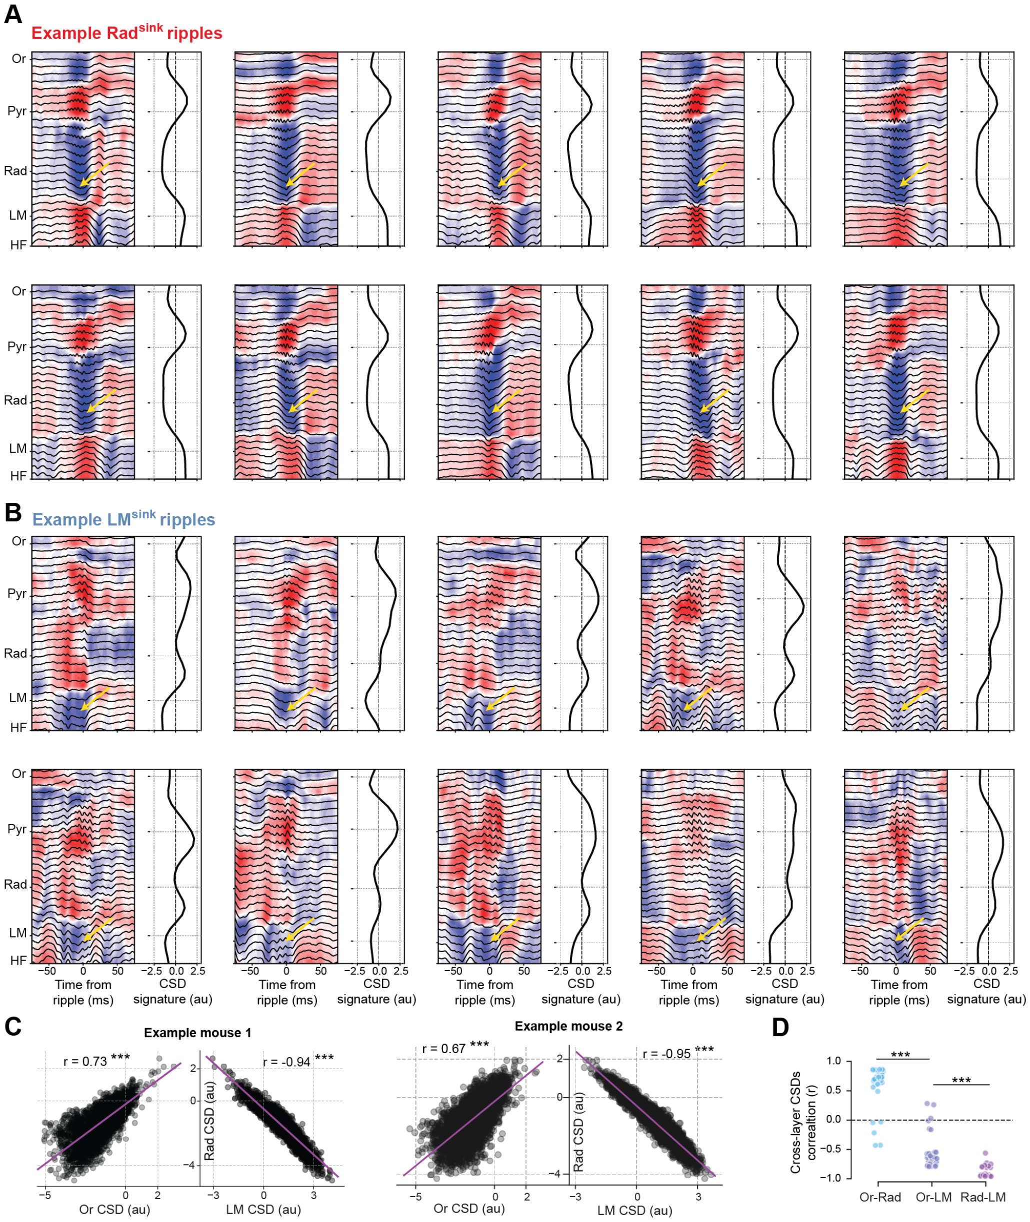

Figure 3.2: Comparison of CSD signatures for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. (A–B) Raw CSD signals (*left*) alongside corresponding LFP waveforms (*left*, black traces), and the computed CSD signature (*right*, black curve) for ten Rad$^{\text{sink}}$ ripples (A) and ten LM$^{\text{sink}}$ ripples (B). Yellow arrows indicate the location of current sinks in either *stratum radiatum* or *stratum lacunosum-moleculare*, depending on the ripple type.

(Figure 3.3(D-F)). In contrast, the other profile exhibited a stronger current sink in *stratum lacunosum-moleculare*, which I refer to as LM$^{\text{sink}}$ ripples (Figure 3.3(D-F)).

These findings demonstrate that hippocampal ripples vary in their laminar currents and that event-by-event analysis reveals a continuum of variability, including ripples with prominent *lacunosum-moleculare* sinks, which are not evident in the

(C) Relationship between ripple-associated CSD values in *stratum oriens*, *stratum radiatum*, and *stratum lacunosum-moleculare*. Each dot represents a single ripple, with its coordinates indicating the mean CSD values across the three layers during an example sleep session. Purple lines denote the best linear fit, with Pearson correlation coefficients quantifying the relationship between ripple-CSDs in each pair of layers.
(D) Distribution of mean Pearson correlation coefficients for ripple-CSDs across different CA1 layers. Each dot represents a sleep session ($p = 5 \times 10^{-48}$; one-way ANOVA; all pairwise comparisons $p < 10^{-6}$, Tukey post-hoc test; $n = 54$ sleep/rest sessions from 5 mice).
$^{***}p < 0.001$.

grand average profile (Figure 3.1).

3.3.2 Spectral and waveform differences between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

In the previous section, I demonstrated that, on a ripple-by-ripple basis, laminar currents vary significantly but with two dominant, recurring patterns: Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. This raises an important question: Do these groups correspond to two intrinsically distinct ripple profiles, or do they arise from averaging artifacts or biological noise? In this section, I address this question by showing that these two profiles emerge consistently, regardless of averaging references, and that they are associated with distinct ripple features such as *stratum pyramidale* LFP waveforms, ripple frequency, and power. These findings demonstrate that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples are indeed distinct events that likely originate differently within the hippocampus.

First, I tested whether these profiles consistently emerged regardless of how CSDs were referenced (e.g., ripple peak, onset, offset). I recomputed the triggered averages of Rad$^{\text{sink}}$, baseline, and LM$^{\text{sink}}$ ripples using different reference points, including the ripple envelope onset, offset, and the strongest current sinks in *stratum radiatum* and *stratum lacunosum-moleculare* (Figure 3.4). Regardless of the reference method used, the distinct Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ profiles consistently emerged: Rad$^{\text{sink}}$ ripples were associated with the strongest sinks in *radiatum*, while LM$^{\text{sink}}$ ripples exhibited the strongest sinks in *lacunosum-moleculare*. As expected,

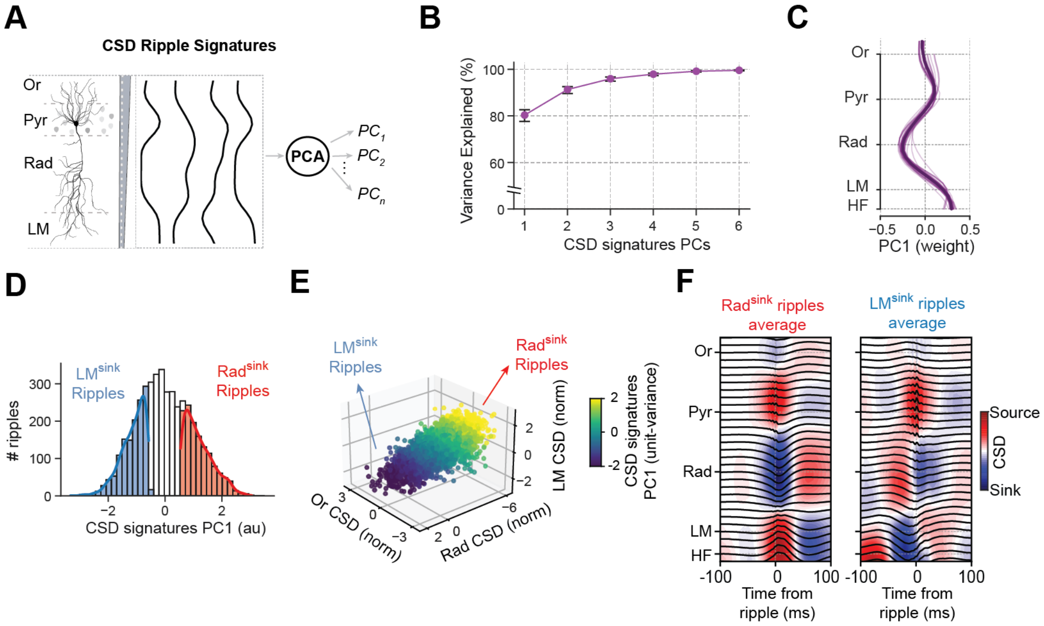

Figure 3.3: Isolating Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples using PCA of laminar CSD profiles.
(A) PCA was applied to the CSD signatures spanning the CA1 radial axis to reduce dimensionality and identify the CSD-signature components that account for the most variance across individual ripple events.
(B) Mean variance explained by each principal component (PC) of the CSD signatures. The first principal component (PC1) captures the majority of the variance [mean variance explained by PC1 (95% CI): 80.3 (77.6 - 82.6)%]. *Black error bars* indicate the 95% confidence interval.
(C) Laminar distribution of PC1 weights, including the *hippocampal fissure* (HF), showing prominent contributions from sinks in *stratum radiatum*, consistent with the average CSD profile reported in Figure 3.1. *Thinner lines* represent individual sleep/rest sessions, while the *thicker line* denotes the mean PC1 across all sessions.
(D) Distribution of PC1 strength across ripples from an example sleep session, highlighting notable variability across events. Ripples at the lower and upper extremes of the distribution (bottom and top 30%) correspond to strong current sinks in *stratum lacunosum-moleculare* (LM$^{\text{sink}}$ ripples, blue) and *stratum radiatum* (Rad$^{\text{sink}}$ ripples, red), respectively. Ripples falling within the central portion (white bins) of the distribution were categorised as 'baseline'.

baseline ripples, positioned between these two extremes, displayed a bias toward *radiatum* sinks. This outcome was expected, as the grand-average CSD profile reflects *radiatum* sinks, meaning the middle portion of the distribution (i.e. baseline ripples) is inherently skewed toward *radiatum* sinks (Figures 3.1 and 3.3).

Having established that these profiles were not artifacts of averaging, I next investigated whether they were associated with distinct ripple characteristics. In

(E) 3D scatter plots of ripples from a single sleep/rest session, where each point represents an individual ripple with coordinates corresponding to its instantaneous CSD values in *stratum oriens*, *stratum radiatum*, and *stratum lacunosum-moleculare*. CSD values are normalised to unit variance. Points are colour-coded by PC1 strength, which, as expected, explains the majority of the variance in this lower-dimensional space characterised by a continuum. Ripples with stronger PC1 values align with the two profiles identified in (D), corresponding to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(F) Average CSD and LFP waveforms for ripples with strong current sinks in either *stratum radiatum* (Rad$^{\text{sink}}$ ripples) or *stratum lacunosum-moleculare* (LM$^{\text{sink}}$ ripples) from one sleep session.
For (B),(C): $n = 54$ sleep/rest sessions from 5 mice.

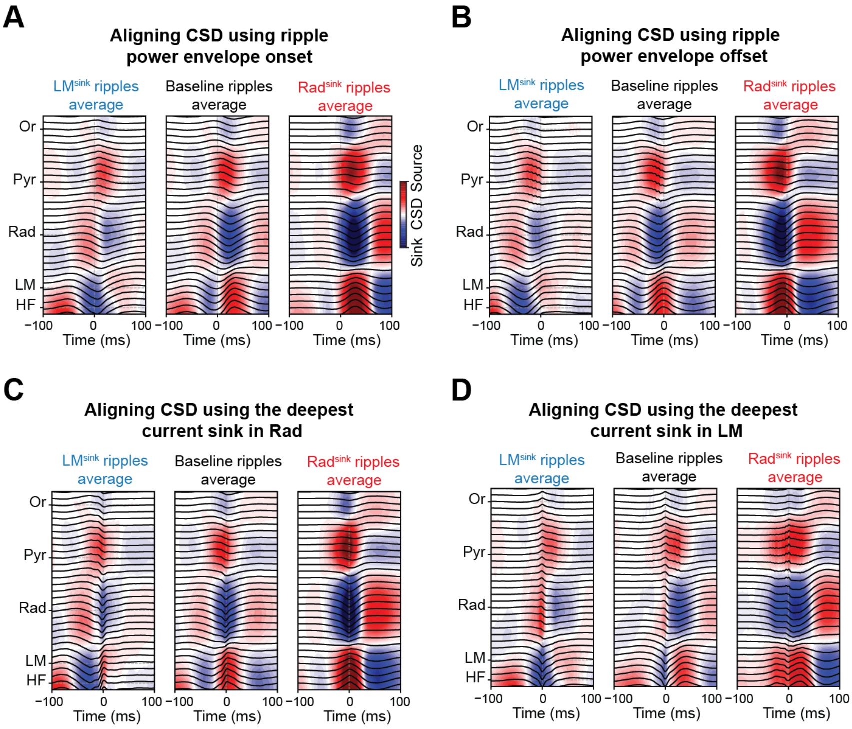

Figure 3.4: Extraction of Rad$^{\text{sink}}$and LM$^{\text{sink}}$ripple profiles using different triggering methods.
(A-D) Averaging CSD signals using different timestamps within the ripple event revealed that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ profiles consistently emerged, independent of how the LFP and CSD signals were referenced. Displayed are the triggered average CSD and LFP signals across CA1 layers during LM$^{\text{sink}}$, baseline, and Rad$^{\text{sink}}$ ripples from one example sleep session. Notably, the Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripple profiles remained consistent whether LFP and CSD signals were referenced to the ripple envelope onset (A), offset (B), or to ripple current sinks located in *stratum radiatum* (C) or *stratum lacunosum-moleculare* (D). To prevent contamination of the averages by overlapping events, only isolated ripples (i.e., those with no other ripple occurring within ± 250 ms) were included in this analysis.

terms of LFP waveforms, Rad$^{\text{sink}}$ ripples exhibited the well-described negative *radiatum* sharp-wave, while LM$^{\text{sink}}$ ripples displayed a pyramidal layer positive deflection accompanied by a *lacunosum-moleculare* negative deflection (Figure 3.5(A)). LM$^{\text{sink}}$ events also exhibited lower ripple frequency [Figure 3.5(B,C); mean frequency (95% CI): Rad$^{\text{sink}}$, 147.3 (145.7 - 149.1) Hz; LM$^{\text{sink}}$, 124.5 (122.9 - 126.1) Hz; $p < 10^{-5}$; paired bootstrap test] and lower ripple power (Figure 3.5(B,D); $p < 10^{-5}$; paired bootstrap test).

Importantly, these results are not trivial, as the CSD signatures and PCA classification were entirely agnostic to ripple waveforms, frequency, and power; they only contained information about the CSD signals. The fact that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples are associated with distinct ripple characteristics confirms that they do not arise from biological noise or averaging artifacts but instead represent intrinsically distinct event types within the hippocampus.

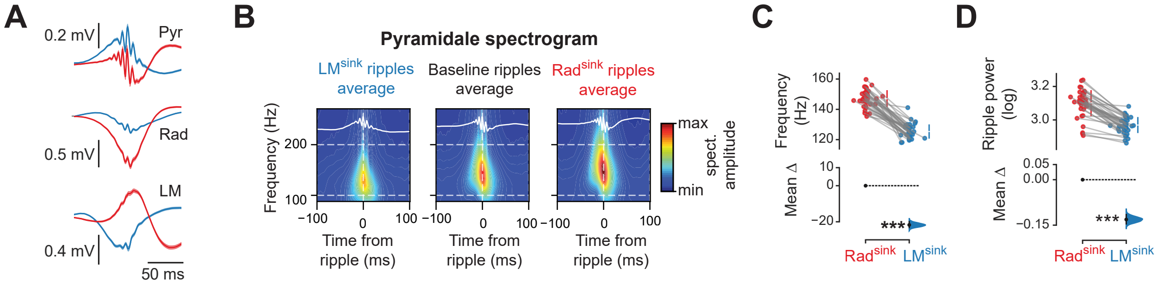

Figure 3.5: Spectral properties of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A) Average waveforms for Rad$^{\text{sink}}$ (red) and LM$^{\text{sink}}$ (blue) ripples recorded from LFPs in *stratum pyramidale* (Pyr), *stratum radiatum* (Rad), and *stratum lacunosum-moleculare* (LM). To highlight differences in ripple frequency, LFP waveforms were aligned to the peak amplitude of each ripple event. Shaded regions represent the 95% confidence interval.
(B) Spectrograms (spect.) of CA1 pyramidal LFPs from a representative mouse across different ripple types. From left to right, panels show the average spectrogram (with the LFP waveform overlaid in white) for LM$^{\text{sink}}$, Baseline, and Rad$^{\text{sink}}$ ripples within a ±100-ms window centred on the ripple-envelope peak. Note that both ripple power and mean frequency increase progressively from LM$^{\text{sink}}$ to Rad$^{\text{sink}}$ ripples.
(C) Estimation plot illustrating the effect size of the difference in mean ripple frequency between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. The top panel shows raw data distributions (each dot represents a sleep session), while the bottom panel depicts the mean difference (*black dot*, mean; *black ticks*, 95% confidence intervals; *coloured area*, bootstrapped error distribution).
(D) Same as (C), but for ripple power, which is significantly higher in Rad$^{\text{sink}}$ ripples.
For (A),(C) and (D): $n = 54$ sleep sessions from 5 mice. $^{***}p < 0.001$.

3.4 Discussion

In this chapter, I have demonstrated that ripples exhibit distinct laminar current profiles, indicating that the currents underlying ripple events are not homogeneous but instead display significant heterogeneity. Within a continuum of CSD profiles, two distinct patterns consistently emerged: Rad^{sink} and LM^{sink} ripples.

Previous studies have shown that pyramidal LFP waveforms during ripples form a continuum (Chang et al., 2025; Sebastian et al., 2023). In this thesis, I took a different approach: rather than assuming that waveform differences primarily drive ripple heterogeneity, I reasoned that these differences emerge from the underlying current sources and sinks (i.e., laminar currents), which in turn shape ripple waveforms. My findings are consistent with Sebastian et al., 2023 in showing that waveform variability correlates with distinct CSD profiles and with changes in ripple power and frequency. However, in contrast to previous work that starts from waveform features, I first identified the underlying laminar currents and then examined their impact on ripple expression. This current-first approach — focusing on features independent of ripple waveform — offers a new perspective on the organisation of ripple variability.

A key result of this chapter is that PCA consistently revealed Rad^{sink} and LM^{sink} profiles in the data. As described in Section 2.4.1, PCA extracts axes of maximal variance in the dataset, ensuring that the most prominent features in ripple-associated CSD signals are identified. Here, the first principal component (PC1) alone explained 80% of the variability across all CSD signatures, along which Rad^{sink} and LM^{sink} ripples naturally emerged. This suggests that despite the overall variability in the dataset, these two profiles are consistently present on average, as reflected in the proportion of variance explained. PCA provides a robust statistical framework for describing the data, as it orders components by their contribution to variance. In contrast, other algorithms, such as independent component analysis (ICA), while effective for isolating independent signal sources, do not prioritise variance explanation, making them less suited for the data-driven discovery of dominant ripple profiles.

To classify ripples into these subtypes, I applied a percentile-based approach, defining Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples as the upper and lower 30$^{\text{th}}$percentiles of PC1 scores, respectively. This was necessary because the distribution of CSD profiles forms a continuum rather than discrete clusters, requiring a data-driven threshold to distinguish the most extreme cases. While the exact cut-off is arbitrary, it enables targeted analysis of the most distinct ripple subtypes. Importantly, this classification was cross-validated by examining spectral features, confirming that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples differ systematically in frequency and power. Although alternative methods, such as clustering approaches, could be explored in future studies, any classification method would ultimately require a similar decision about where to define boundaries within a continuous distribution.

Rad$^{\text{sink}}$ ripples are characterised by large current sinks in *stratum radiatum* and prominent sharp-wave deflections in this layer. This profile aligns with canonical ripple patterns described in previous studies (Buzsáki et al., 1983; Buzsáki, 1986; Sullivan et al., 2011; Ylinen et al., 1995) (see Section 1.2). Notably, *radiatum* is the sublayer where CA3 projects to CA1 (Valero & de la Prida, 2018), suggesting that Rad$^{\text{sink}}$ currents and sharp-wave components may primarily originate from CA3 inputs (Buzsáki et al., 1983; Sullivan et al., 2011). This supports the widely accepted hypothesis that CA3 plays a crucial role in ripple generation (Buzsáki et al., 1983; Buzsáki, 1986, 1989), a topic I will explore further in Chapter 5. However, across sleep sessions and animals, I consistently observed a second class of ripples characterised by current sinks in *stratum lacunosum-moleculare* rather than *stratum radiatum*. These ripples were also associated with weak or absent sharp-wave deflections in *stratum radiatum* (Figure 3.5(A)). Given that *radiatum* is the primary target of CA3 inputs, this suggests that CA3 may be less involved in LM$^{\text{sink}}$ ripples compared to Rad$^{\text{sink}}$ ripples, a hypothesis I will investigate in Chapter 5. Instead, LM$^{\text{sink}}$ ripples were characterised by a large negative deflection in *stratum lacunosum-moleculare*, the layer targeted by MEC3 projections (Valero

& de la Prida, 2018), resembling the sharp-wave typically observed in *stratum radiatum* during Rad$^{\text{sink}}$ ripples.

This raises an interesting question about the definition of a sharp-wave. Previous studies have defined sharp-waves as transient negative deflections in the *radiatum* LFP during ripples, attributed to CA3 inputs (Buzsáki et al., 1983; Buzsáki, 1986). However, a debate remains as to whether a sharp-wave should be defined purely based on LFP waveform morphology or whether it specifically requires association with the CA3→*radiatum* pathway. In other words, is a 'sharp-wave' any brief, sharp negative LFP deflection, or must it originate from CA3? If the definition is purely morphological, then the deflection observed in *lacunosum-moleculare* during LM$^{\text{sink}}$ ripples could also be considered a 'sharp-wave', potentially reflecting MEC3 inputs. This would imply that both Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples are associated with sharp-waves, but occurring in different layers: *radiatum* for Rad$^{\text{sink}}$ and *lacunosum-moleculare* for LM$^{\text{sink}}$. The interpretation of these findings may be debated, particularly regarding the definition of sharp-waves in relation to ripple heterogeneity. Regardless, the key finding of this chapter is that two distinct ripple profiles consistently emerge during sleep, differing in their CSD patterns, LFP waveforms, ripple frequency, and power.

By establishing a reliable classification of ripple subtypes, the next step is to investigate how hippocampal neurons respond to these different ripples. Addressing this requires overcoming a technical challenge since most of my lab's spike data comes from tetrode recordings, which lack CSD information. The question then becomes: can Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples be distinguished without relying on CSD signals but instead using the distinct features identified in this chapter? In the next chapter, I explore analytical frameworks that enable the identification of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples in the tetrode dataset, making it possible to analyse neuronal activity in relation to these ripple types later in Chapter 5.

All models are wrong, but some are useful.

— George E. P. Box (Box, 1979)

4

Bridging silicon probe and tetrode data with machine learning models

Contents

4.1 Introduction

The objective of this short chapter is to establish an analytical framework for identifying and analysing ripples (or other events) that are distinguishable only through their laminar current profiles, particularly in tetrode recordings where electrodes are not arranged along the somato-dendritic axis. This effectively aims to bridge the analytical gap between silicon probe and tetrode methodologies.

As explained in Chapter 3, Rad^{sink} and LM^{sink} ripples were classified based on their distinct laminar current profiles (Figure 3.3). However, most spike sorting and unit isolation in my lab (see Section 2.1.5) is performed using tetrode implants, which do not capture laminar profiles and thus cannot directly estimate CSDs. This means that I have no spike data available from the silicon probe datasets where these ripple classifications were originally established. Furthermore, the lab has an extensive dataset of tetrode recordings from CA1 and CA3 during sleep and rest, but these cannot be analysed in relation to Rad^{sink} and LM^{sink} ripples due to the lack of corresponding CSD information. Ideally, these two ripple types could be analysed in any dataset, independent of the specific recording methodology used. However, this presents a key methodological challenge: how can Rad^{sink} and LM^{sink} ripples be identified in the absence of direct CSD measurements?

Since Rad^{sink} and LM^{sink} ripples exhibit distinct LFP waveforms in *stratum pyramidale* (see Section 3.3.2), I investigate whether these waveform features alone provide a reliable basis for classification. Specifically, I leverage ground truths established with silicon probe data to develop an alternative classification approach applicable to tetrode recordings. This approach is based on the relationship between CSD profiles and LFP waveforms, allowing me to test whether LFP waveform differences are sufficient to distinguish ripple types without direct access to CSDs.

Building on the insights from Chapter 3, I first explore the extent to which ripple-by-ripple variations in *stratum pyramidale* LFP waveforms reflect differences in their underlying CSD signals. I then present a classifier that can distinguish Rad^{sink} and LM^{sink} ripples using only the *pyramidale* LFP waveforms. Finally, I extend this analytical framework to other hippocampal event types that also rely on CSDs, such as dentate spikes of type I and II, testing the generalisability of this approach.

4.2 Methods

4.2.1 Explained variance of ripple LFP waveforms by laminar CSD

To quantify how much of the ripple-by-ripple variability in LFP waveforms recorded from the CA1 *stratum pyramidale* could be explained by underlying CSD dynamics across CA1 layers — from *stratum oriens* to *stratum lacunosum-moleculare* — I trained linear regression models. Specifically, for each individual ripple, the model was trained to predict the strength and sign of the CSD in each CA1 layer based on the corresponding *stratum pyramidale* LFP waveform. Each ripple-nested *stratum pyramidale* LFP trace was low-pass filtered through a Butterworth filter (4th order, with a cut-off frequency of 30 Hz)to isolate the low-frequency component of the LFP waveform, which I refer to as the 'sharp-wave' component. Importantly, I hypothesised that this 'sharp-wave' component contained all the information to distinguish the underlying current profiles, based on previous findings (Figure 3.5). Then, these filtered LFP waveforms were standardised through z-score transformation to ensure uniformity in variance and mean across the signals of all ripples.

I then applied PCA to reduce the dimensionality of the 200 ms LFP traces down to five principal components (PCs), which explained more than 80% of the variance of all ripple-LFP waveforms (see Figure 4.1(E)). Similarly, the CSD signals for each CA1 layer (*strata oriens, pyramidale, radiatum, and lacunosum-moleculare*) were normalised by dividing each CSD signal by its standard deviation calculated over all ripples. This normalisation allowed maintaining the polarity information of the CSD signals while ensuring comparability across different magnitudes.

Subsequently, I employed linear regression models to predict the normalised CSD in each CA1 layer from the dimensionality-reduced LFP signals during ripple events. To validate the robustness and generalisability of our models, I performed cross-validation 20 times (80% training, 20% testing) on the LFP and CSD data. For each iteration, the model was fitted to the training set and then evaluated on the testing set, thereby obtaining a coefficient of determination R^2 for each run. The coefficient of determination, R^2, was calculated as:

$$R^2_{\text{layer}} = 1 - \frac{\sum_{i=1}^{n} \left(y_{i,\text{layer}} - \hat{y}_{i,\text{layer}}\right)^2}{\sum_{i=1}^{n} \left(y_{i,\text{layer}} - \bar{y}_{\text{layer}}\right)^2} \tag{4.1}$$

where $y_{i,\text{layer}}$ and $\hat{y}_{i,\text{layer}}$ respectively represent the actual and predicted CSD value for the i-th ripple in the specified layer; $\bar{y}_{\text{layer}}$ is the mean of the actual CSD values in the specified layer, and n is the total number of ripples in each recording session.

The variance explained by each model was then determined by averaging the R^2 values across all 20 cross-validation iterations, providing a measure of how well the CSD can be predicted from the LFP signals across the different CA1 layers. The chance level of explained variance of the LFP waveforms was determined by shuffling the true CSD values 500 times and then computing the variance explained by the shuffled data.

4.2.2 Structure index of ripple LFP waveforms

To further assess the relationship between ripple-triggered LFP signals and laminar CSD signals in CA1, I employed the structure index (SI) measure (Sebastian et al., 2024). The SI quantifies whether two ripples with similar LFP waveforms are also associated with similar CSD values, effectively testing whether a latent variable (i.e. the CSD in a CA1 layer) explains the observed similarities in high-dimensional data (i.e. the LFP waveform).

Unlike the coefficient of determination (R^2), which measures the proportion of variance in CSD that can be explained by a linear mapping from LFP waveforms, SI evaluates the consistency of ripple events across different CSD values. That is, while R^2 assesses how well variations in LFPs can be used to predict variations in CSD, SI provides a complementary perspective by determining whether ripples with similar LFP waveforms correspond to similar underlying CSD profiles.

For each sleep/rest session, I first applied the same pre-processing steps to the LFP and CSD signals as described earlier (see Section 4.2.1). I then computed the SI for each layer by comparing the ripple-triggered LFP waveforms recorded in the pyramidal cell layer with the CSD magnitudes recorded in each CA1 layer during ripples. An SI of 0 indicates that the latent variable (CSD) does not explain the similarity in LFP waveforms, while an SI of 1 indicates a perfect correlation, meaning the CSD fully accounts for the observed similarities in waveforms.

To assess the significance of these results, I generated chance levels by shuffling the CSD values 500 times.

4.2.3 Single-ripple CSD profile prediction from pyramidal LFP waveform

I applied Linear Discriminant Analysis (LDA) to predict both the polarity (sign) and the magnitude of the *lacunosum-moleculare* (LM) current for each ripple event based on the corresponding *stratum pyramidale* LFP time course. Ripple LFPs were extracted and pre-processed as described in the Section 4.2.1 ($n =$ total $117,658$ isolated ripples from 5 mice). To classify each ripple event into the Rad^{sink} or LM^{sink} category, the corresponding LM CSD was computed within a 50-ms window centred on the ripple peak and averaged over three adjacent probe channels: one in *lacunosum-moleculare* and the channels immediately above and below it. To ensure comparability across sleep sessions while retaining information about current polarity, the ripple CSD was normalised by the standard deviation of all ripple events (i.e. unit-variance CSD). Ripple events with a normalised negative CSD in the lowest 30thpercentile were classified as LM^{sink} ripples, while those in the top 30thpercentile were classified as Rad^{sink} ripples. Ripples with CSD values within the 30-70thpercentile range were classified as 'baseline' ripples, characterised by relatively small current magnitudes. A schematic of this pipeline is shown in Figure 4.1.

To assess the robustness of the LDA model, I employed leave-one-out cross-validation, excluding each mouse from the training set in turn and evaluating accuracy on the left-out mouse. To address class imbalance, I downsampled each class to match the size of the smallest class, repeating this balancing procedure 1,000 times. For each permutation, a null model was also trained with shuffled training labels, disrupting the relationship between *stratum pyramidale* LFP waveforms and LM currents to estimate chance-level accuracy. In each iteration, the accuracy of both the true and null models was computed on the testing data for the left-out mouse. This process yielded a chance-normalised accuracy for each left-out mouse, calculated by comparing the mean accuracy of the true models to that of the null models (Figure 4.2(D)).

To test whether ripple frequency added significant information to the LDA classification beyond the 'sharp-wave' component alone, I included the ripple frequency of each event as an additional feature alongside the first five principal components (PCs) of the LFP waveform to train the LDA (Figure 4.3).

When this model was applied to the tetrode dataset (see Section 4.3.3 and Chapters 5,6), the pyramidal layer tetrode used for ripple classification was selected as the one whose average LFP waveform was most similar to the average waveform from the LDA's training set, i.e., the average across the $117,658$ ripples from the silicon probe dataset.

4.2.4 Dentate spikes detection and classification into type I and II from CSD

DSs were detected from the LFP signals recorded from silicon probe channels in the dentate gyrus (DG) granule cell layer (GCL), following the procedure described in McHugh et al., 2024. To detect DSs, I first subtracted from the GCL LFP a reference LFP recorded from a channel in *stratum oriens*. The GCL LFP was then band-pass filtered (1-200 Hz, 4$^{\text{th}}$-order Butterworth filter). Candidate DSs were identified as LFP peaks that exceeded a threshold of six times the median absolute

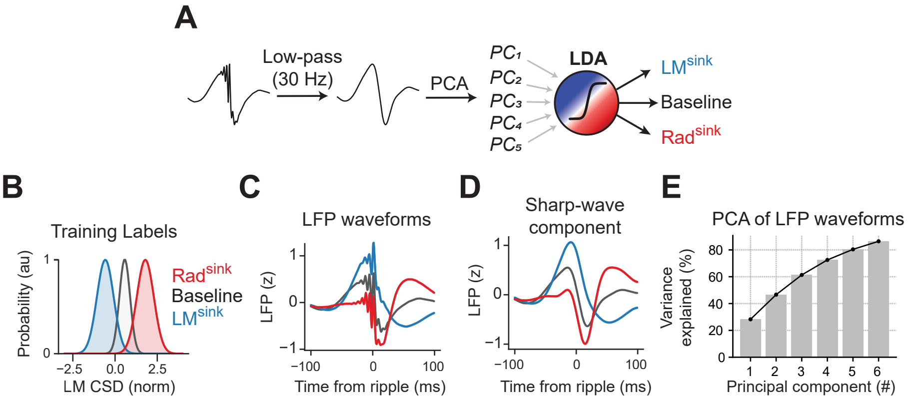

Figure 4.1: Pipeline for training an LDA model to classify ripples with distinct CSD profiles from pyramidal LFP waveforms.
(A-E) Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripple profiles were predicted based on ripple waveforms recorded in *stratum pyramidale* LFPs. (A) Schematic of the LDA model used to classify Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples from pre-processed LFP waveforms. Ripple events were assigned to three categories (Rad$^{\text{sink}}$, Baseline, and LM$^{\text{sink}}$) based on digitised current values in *lacunosum-moleculare* (B), the layer with the highest predictive power (see Figure 4.2(A,B)). (C) The red, gray, and blue distributions represent Rad$^{\text{sink}}$, Baseline, and LM$^{\text{sink}}$ ripples, respectively. For each ripple, pyramidal LFP waveforms were z-scored (C) and low-pass filtered to extract the 'sharp-wave' component (A,C). (D) The triggered average LFP waveforms and corresponding 'sharp-wave' components for the three ripple classes, as defined by LM CSD values, are shown. (E) The filtered waveforms were then decomposed using PCA, with the top five PCs collectively explaining more than 80% of the variance across all ripples. (A) These five PCs were then used to train the LDA classifier to predict the ripple class (Rad$^{\text{sink}}$, Baseline, or LM$^{\text{sink}}$).

deviation of the signal [total number of detected DSs: $15,085$ events; mean number of DSs per recording day (IQR): 1,676.1 (1,213.0 - 2,415.0); $n = 3$ mice]. The DS event time was defined as the time of the highest peak within each candidate event.

To classify DSs into type I and II, I used their CSD profiles as previously reported (Lopes-dos-Santos et al., 2023). Briefly, I computed the CSD during detected DS events at their peak, across channels spanning from the pyramidal layer to the deepest DG channel. After applying PCA to these CSD signals, the projections of each DS event onto the first two principal components (PC1 and PC2) were used as input to a two-component Gaussian Mixture Model. This consistently identified two distinct CSD profiles in the DG. In agreement with previous studies (Bragin et al., 1995; Dvorak et al., 2021; Farrell et al., 2024), events featuring a strong current sink in the outer molecular layer of the DG were classified as DS1, while events with a stronger sink closer to the granule cell layer were classified as DS2 (see Figure 4.5).

4.2.5 Dentate spikes classifier

To determine whether GCL LFP waveform differences could be used to classify DS1 versus DS2, I applied a similar framework to the one described in Section 4.2.3. Briefly, I extracted LFP traces from the DG GCL in a time window spanning from -100 ms to +100 ms relative to the DS peak. Each LFP waveform was then pre-processed by applying a low-pass filter (30 Hz, 4^{th}-order Butterworth filter) and z-scoring. To reduce dimensionality, I applied PCA, selecting the top five PCs, which together explained 80% of the total variance. These five PCs, along with the DS type labels derived from the CSD classification, were used to train an LDA model. The model was cross-validated 1,000 times using an 80% training and 20% testing split. In each iteration, I controlled for class imbalance by down-sampling the class with fewer examples and computed the accuracy of the model on each test set. A schematic of this analytical framework is shown in Figure 4.6(A).

4.3 Discriminating Rad^{sink} and LM^{sink} from the LFP waveforms

4.3.1 Variance of ripple LFP waveforms captured by LM CSD

To investigate whether I could discriminate Rad^{sink} and LM^{sink} ripples based on the *stratum pyramidale* LFP waveforms, I first needed to establish whether the underlying CSD signals in the CA1 layers could predict distinct *pyramidale* LFP waveforms above chance. In particular, I aimed to determine which CA1 layer had the highest predictive power in shaping LFP waveforms based on its underlying CSD magnitude and polarity. In other words, in which layer of CA1 does a change in CSD correspond to significant changes in the *stratum pyramidale* LFP waveform? This is crucial, as labelling ripples as Rad^{sink}, baseline, or LM^{sink} requires ensuring that the selected variable contains enough information to differentiate CSD profiles.

To address this, I quantified how much of the variance in LFP waveforms could be explained by the CSD in each layer. A high explained variance value would indicate that fluctuations in CSD magnitude and polarity in a given layer are associated with significantly distinct LFP waveforms. As mentioned in Section 1.3, ripples are thought to originate locally within the CA1 pyramidal layer (Stark et al., 2014), whereas the lower-frequency components of the waveform — referred to here as the 'sharp-wave' component — may reflect more informative features related to distinct synaptic inputs and current sinks. Based on this and previous observations in Figure 3.5, I hypothesised that the 'sharp-wave' component would contain most of the information necessary to distinguish between different CSD profiles. Therefore, prior to computing the explained variance and in all subsequent analyses, ripple LFP waveforms were pre-processed using a low-pass filter to isolate this component (30 Hz; Butterworth filter; see Section 4.2.1).

When I measured how much of the ripple waveform variance was explained by the CSD in the Or., Pyr., Rad., and LM layers, I found that while all layers significantly predicted LFP waveforms, the LM CSD was the best predictor (Figure 4.2(A)). This result indicates that for an equivalent CSD change across layers, the associated LFP waveforms were most dissimilar when the change occurred in the LM layer.

To cross-validate this result, I applied an additional method based on measuring the structure index (SI) of the LFP waveforms relative to the underlying CSDs (Sebastian et al., 2024). The SI quantifies how different CSD values in a given layer correspond to distinct *stratum pyramidale* LFP waveforms, providing an alternative approach to the CSD variance explained. This analysis confirmed that the LM CSD remained the best predictor of LFP waveform differences (Figure 4.2(B)). These findings suggest that LM CSD could provide the ground truth used to train models to classify Rad^{sink}, baseline, and LM^{sink} ripples.

4.3.2 Distinguishing Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ from the Pyr. LFP waveform

Above, I demonstrated that the LM CSD, more than other layers, shapes the variability of *stratum pyramidale* LFP waveforms on a ripple-by-ripple basis. However, this still leaves open the central question of this chapter: do these *pyramidale* LFP waveforms contain enough information to reliably distinguish between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples?

To test this, I examined whether the 'sharp-wave' component of the LFP waveform could distinguish the distinct current profiles described in the previous chapter: Rad$^{\text{sink}}$, baseline, and LM$^{\text{sink}}$ ripples. The decision to isolate this low-frequency component was motivated by the analyses presented in the previous section. To address this, I trained a linear discriminant analysis (LDA) classifier on low-pass filtered *stratum pyramidale* LFP waveforms to predict the sign and digitised strength of the LM CSD, which, as demonstrated above, defines the three ripple profiles (Rad$^{\text{sink}}$, LM$^{\text{sink}}$, and baseline) (Figures 4.1 and 4.2(C); see also Section 4.2.3).

I employed a leave-one-out approach, where the model was trained on data from all but one mouse, which was then used for testing. The LDA significantly predicted ripple profiles above chance level, suggesting that the *stratum pyramidale* LFP waveform contains substantial information about single-ripple CSD profiles [Figure 4.2(D); mean accuracy gain over shuffle controls (95% CI): 55 $(38.5-66.2)\%$; $p = 0.0014$, one-tailed t-test; model cross-validated over 5 mice].

To further validate the model's predictions, I measured the CSD sign and strength in *radiatum* and *lacunosum-moleculare* during each leave-one-out iteration of training. This was done to ensure that the classification aligned with the expected physiological profiles: Rad$^{\text{sink}}$ ripples should exhibit a stronger current sink in the *radiatum*, while LM$^{\text{sink}}$ ripples should exhibit a stronger current sink in the *lacunosum-moleculare*. Comparing these CSD strengths across classified ripple groups confirmed that the model effectively discriminated the three profiles: Rad$^{\text{sink}}$ ripples were characterised by the strongest current sink in the *radiatum* compared to the other groups, while LM$^{\text{sink}}$ ripples were associated with strong current sinks in the *lacunosum-moleculare* (Figure 4.2(E).

In Section 3.3.2, I also showed that, in addition to differences in LFP waveform, Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples exhibit distinct ripple frequencies (Figure 3.5). However, since the LDA classifier was trained on the 'sharp-wave' components of the LFP waveforms (Figure 4.1), it had no access to information about ripple frequency. Nevertheless, the LDA could significantly distinguish ripple types based solely on these LFP waveforms (Figure 4.2).

Would adding frequency information improve the overall classification performance? To address this question, I included ripple frequency as an additional feature in the LDA input. While this led to a slight improvement in model performance [mean accuracy gain over shuffle controls (95% CI): 75.8% (55.4 – 89.5)%; model cross-validated over 5 mice], the improvement was not significantly different from the performance achieved using only the 'sharp-wave' components (Figure 4.3).

These results demonstrate that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples can be significantly discriminated using the information contained in the *stratum pyramidale* ripple LFP waveform, without relying on silicon probe implants. This approach can therefore be applied to the tetrode datasets to infer how hippocampal cells respond to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples (see Section 4.3.3 and Chapter 5).

4.3.3 Application and validation of the classifier on tetrode recordings

Having established that LFP waveforms in *stratum pyramidale* reliably predict ripple CSD profiles, I next applied the trained classifier to tetrode-recorded ripples. I tested whether the LDA model, trained on silicon probe data, could successfully classify ripples recorded with tetrodes based on their LFP waveform characteristics.

Applying this model to the tetrode dataset identified $256,273$ Rad$^{\text{sink}}$ and $211,505$ LM$^{\text{sink}}$ ripples from 13 mice [mean number of ripples per sleep/rest session (IQR): Rad$^{\text{sink}}$, 1,062 (677.8 - 1,417.0); LM$^{\text{sink}}$, 866.8 (490.8 - 1,146.3)]. Notably,

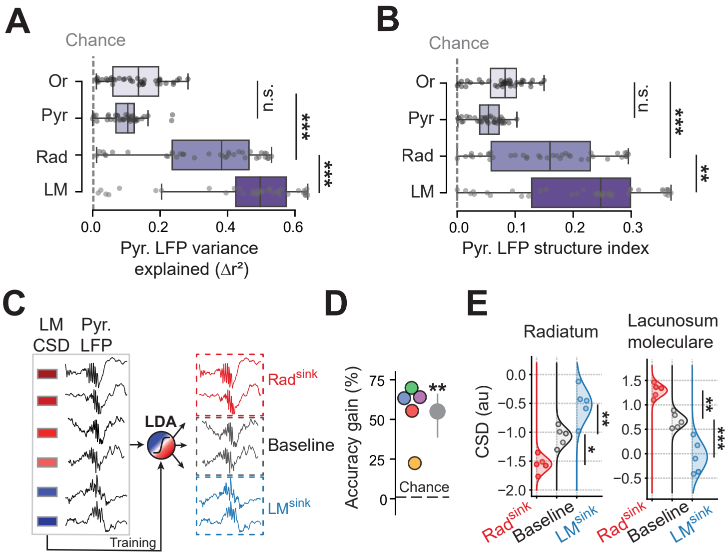

Figure 4.2: Distinguishing Rad$^{\text{sink}}$and LM$^{\text{sink}}$ripples using their pyramidal layer LFP waveforms.
(A) Boxplot showing the variance in CA1 *stratum pyramidale* ripple LFP waveforms explained by CSD signals from different layers. Values were normalised relative to chance by subtracting the variance explained by surrogate models (see Section 4.2.1). Each dot represents a sleep session. The CSD in *lacunosum-moleculare* accounts for the largest amount of LFP waveforms variance (p $= 1.2 \times 10^{-24}$, one-way ANOVA; Rad vs. LM: p $= 5.7 \times 10^{-4}$, Tukey post-hoc). The vertical gray dashed line marks the chance level.
(B) Same as (A), but displaying the structure index (SI) (Sebastian et al., 2024) instead of the explained variance. The SI is normalised by subtracting the average SI obtained from surrogate models (see Section 4.2.2). The LM CSD contributed significantly more to the LFP structure than other layers (p $= 2.7 \times 10^{-17}$, one-way ANOVA; SI in LM CSD vs. Rad CSD: p $= 2.4 \times 10^{-3}$, Tukey post-hoc). Each dot represents a sleep session. The vertical gray dashed line indicates chance level.
(C) Schematic representation of the LDA model predicting CSD signals in LM based on ripple LFP waveforms recorded in *stratum pyramidale*. Each ripple LFP trace, along with its corresponding CSD magnitude and sign, was used as input for the LDA classifier to differentiate between Rad$^{\text{sink}}$, baseline, and LM$^{\text{sink}}$ ripples.
(D) LDA classification accuracy relative to chance. Each dot represents one mouse excluded from the training set (leave-one-out approach). The dashed line denotes chance level, while the gray dot and vertical ticks indicate the mean and 95% CI, respectively.
(E) *Radiatum* and *lacunosum-moleculare* CSD values for LDA-classified Rad$^{\text{sink}}$, baseline, and LM$^{\text{sink}}$ ripples. For each leave-one-out model, the trained classifier was applied to the excluded mouse, and the mean CSD for each predicted ripple category was computed. Each dot represents a left-out mouse. LDA models effectively distinguished ripples characterised by a strong *radiatum* current sink, those with a baseline CSD profile, and those featuring a prominent *lacunosum-moleculare* sink (*radiatum* CSD: p $= 2.6 \times 10^{-5}$; *lacunosum-moleculare* CSD: $p = 3.1 \times 10^{-6}$, one-way ANOVA with Tukey post-hoc).
For (A),(B): $n = 54$ sleep/rest sessions from 5 mice; for (D),(E): $n = 5$ mice. $^{\text{n.s.}}p > 0.05$, $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

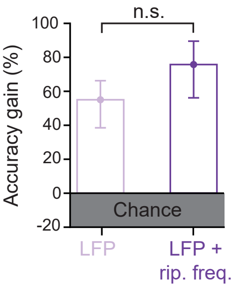

Figure 4.3: Information in 'sharp-wave' components of LFP waveforms versus ripple frequency for ripple type classification.
Mean gain in accuracy from chance level of LDA models distinguishing ripple types. LDAs were trained either using only the 'sharp-wave' components (i.e. the low-frequency components) of the LFP waveforms from the pyramidal layer (left bar) or using the same LFP features plus the ripple frequency of each event (right bar). Adding ripple frequency slightly improved model performance, but this increase was not significantly greater than using the 'sharp-wave' components alone ($p = 0.084$, bootstrap test). Error bars represent the 95% confidence interval; *shaded area* indicates chance level.
$n = 5$ mice. $^{\text{n.s.}}p > 0.05$,

ripples classified as Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ in the tetrode dataset closely matched those recorded with silicon probes in both LFP waveform characteristics and ripple frequency, further validating the classification approach (Figure 4.4).

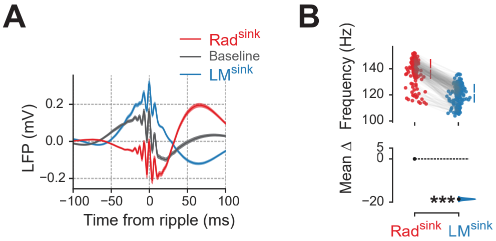

Figure 4.4: Application and cross-validation of the LDA model on the tetrode dataset.
(A) Using the ripple-type classifier validated on the silicon probe dataset, I classified ripples in the tetrode-recorded dataset. The retrieved LFP waveforms, referenced to the peak ripple amplitude, for Rad$^{\text{sink}}$, baseline, and LM$^{\text{sink}}$ ripples closely matched those obtained from the silicon probe dataset (Figure 3.5(A)).
(B) Estimation plot illustrating the effect size for the difference in ripple frequency between events classified as Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ in the tetrode dataset. The observed frequency differences align with those previously identified in the silicon probe dataset, confirming the LDA's reliability [see Figures 3.5 (B,C); mean frequency (95% CI): Rad$^{\text{sink}}$, 139.1 (138.9 - 140.9) Hz; LM$^{\text{sink}}$, 122.1 (120.1 - 122.1) Hz; $p < 10^{-5}$; paired bootstrap test].
$n = 244$ sleep/rest sessions from 13 mice. $^{***}p < 0.001$.

4.4 Distinguishing dentate spikes (DSs) of type I and II from LFP waveforms

In this part of the chapter, I demonstrate that the framework previously described (see Section 4.3) can also be applied to distinguish other events and not just ripples, such as DSs of type I and II. DSs of type I and II are large-amplitude peaks in the LFP of the granule cell layer (GCL) of the dentate gyrus (DG) and, similarly to ripples, have been implicated in memory consolidation during offline brain states (McHugh et al., 2024). These two DS types exhibit distinct behavioural correlates, with type II triggering brain-wide responses (Farrell et al., 2024). Therefore, distinguishing DSs is essential to prevent misinterpreting their distinct functional roles.

Previous studies have shown that DSs can be classified based on their CSD profiles into type I or type II (Bragin et al., 1995; Farrell et al., 2024; Lopes-dos-Santos et al., 2023), and that these classifications correspond to distinct LFP waveforms in the DG granule cell layer (McHugh et al., 2024). Following the same approach as in the previous section, I investigated whether a silicon probe implant is necessary for distinguishing DSs of type I and II or if the DG LFP waveform alone contains sufficient information for classification. To address this, I trained a model to classify DSs using single-channel LFP recordings from the GCL. The model has been published online as an open-source tool (Castelli, 2025) and a previous version of this model was used in McHugh et al., 2024.

4.4.1 Distinguishing DSs I and II from GCL LFP waveforms

First, I detected DSs from the LFPs of the DG GCL and then classified them into type I or type II based on their CSD profiles (Figure 4.5(A)). Consistent with previous findings (Bragin et al., 1995; Dvorak et al., 2021; Farrell et al., 2024), DS1 events were characterised by a strong current sink in the outer molecular layer of the DG, whereas DSs with a stronger sink closer to the granule cell layer were classified as DS2 (total number of classified events: DS1, $7,135$; DS2, $7,950$)). Beyond their distinct CSD profiles, DSs of type I and type II also exhibit characteristic LFP waveform differences: DS1 events are more symmetrical around their peak, whereas DS2 events display a pronounced negative deflection following the peak, resulting in a lower post-event LFP value compared to the pre-event baseline (Figure 4.5(B)).

To determine whether these waveform differences could be used for classification, I applied a similar framework to the one described in the previous section. Briefly, I trained a linear discriminant classifier to distinguish DS1 from DS2 based on the GCL LFP waveform (Figure 4.6(A)). The classifier significantly distinguished the two DS types above chance level (chance-level accuracy = 0.5) [Figure 4.6(B); mean accuracy across cross-validation iterations (IQR): 0.819 (0.816 − 0.822); $p = 6.2 \times 10^{-37}$; one-tailed t-test]. These results suggest that the GCL LFP waveforms contain enough information to discriminate DS1 and DS2 without requiring CSD signals.

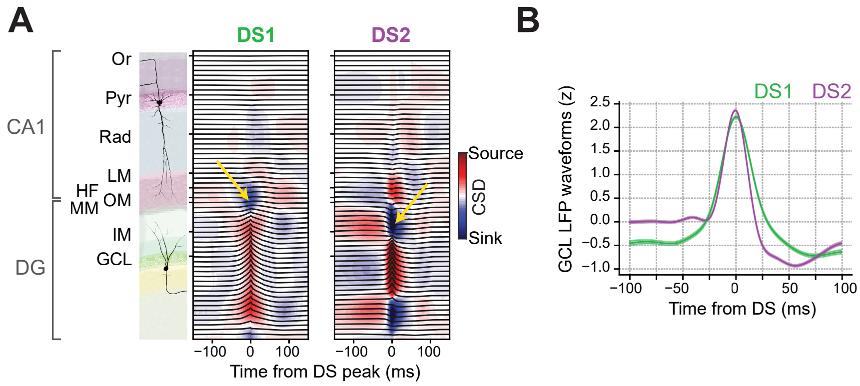

Figure 4.5: DSs of type I and II LFP waveforms.
(A) Average current source density (CSD) and local field potential (LFP) waveforms from one recording day for DS1 (left) and DS2 (right), recorded in the granule cell layer (GCL) of the dentate gyrus in the dorsal hippocampus using a silicon probe spanning CA1 and DG. Yellow arrows indicate the distinct locations of current sinks. Hippocampal layers: *oriens* (Or), *pyramidale* (Pyr), *radiatum* (Rad), *lacunosum-moleculare* (LM), hippocampal fissure (HF), outer (OM), middle (MM), and inner (IM) and *granulare* (GCL).
(B) Average GCL waveforms for DS1 and DS2, extracted from the corresponding LFP traces in (A). DS1 is shown in green, and DS2 in purple.

4.5 Discussion

In this chapter, I demonstrated that events distinguishable based on their CSD profiles, such as Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples or DS1 and DS2, and thus classified using silicon probes, can also be identified using alternative features, such as

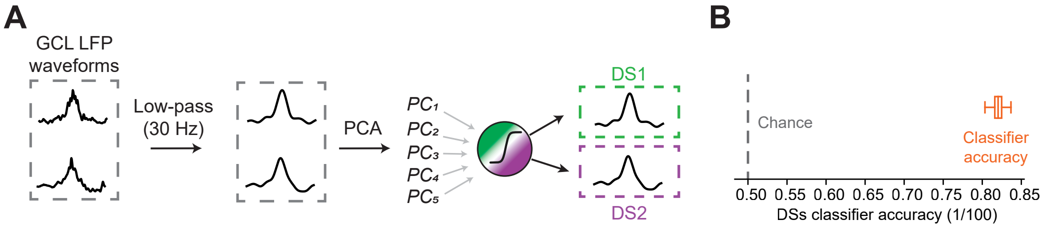

Figure 4.6: DS1 vs DS2 type classifier.
(A) Schematic of the LDA model used to classify DS1 and DS2 based on the GCL LFP waveform. The GCL waveforms were pre-processed by z-scoring and low-pass filtering. These filtered traces were then decomposed using PCA, retaining the five PCs that explained more than 80% of the variance. The resulting PC scores were used as input features for the LDA model to predict the DS type.
(B) Boxplot of cross-validation accuracies ($n = 20$ iterations) of the trained model shown in (A). The dashed line indicates the chance level (chance accuracy = 50%).

the LFP waveform from a single channel. These analytical frameworks, which leverage machine learning models, reduce dependence on specific recording implants, effectively bridging the gap between silicon probe and tetrode implants. Moreover, this approach extends beyond differentiating ripple types and can be applied to other classification tasks, such as distinguishing DS types.

CSD signals play a critical role in defining specific electrophysiological features. However, estimating CSDs requires probe implants with vertically arranged recording sites, as described in Equation 2.2. When channels are not vertically distributed - such as in tetrodes or probes inserted at an angle - laminar information is lost, making direct CSD computation unfeasible. In some experimental contexts, silicon probes may not be the optimal choice. For instance, if a study aims to examine neuronal responses in the MEC during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, the anatomical positioning of CA1 and MEC prevents simultaneous recordings from both regions using a single vertically implanted probe. In contrast, tetrodes offer more flexible placement: one tetrode can be positioned in the CA1 pyramidal layer for ripple classification based on LFP waveforms, while the remaining recording sites can be placed bilaterally in the MEC, maximising the number of recorded cells in the region of interest.

Additionally, due to the anatomical structure of the hippocampus, particularly CA1, vertically arranged channels may not capture a high number of cells per channel. In contrast, tetrode implants, with their flexibility to span a broader medial-lateral extent of the CA1 pyramidal layer and CA3, provide higher spatial coverage, enabling the isolation of more cells per recording site compared to silicon probes. By applying machine learning models, distinct events can be efficiently classified using a single channel instead of dedicating approximately 30 recording sites from a silicon probe (Table 2.1) solely for event discrimination. This analytical framework optimises channel allocation, allowing for the recording of more neurons or the targeting of additional brain regions without being constrained by implant design or 'sacrificing' multiple recording channels for classification purposes.

A key advantage of this approach is its applicability to pre-recorded datasets, independent of surgical constraints. Once trained on ground-truth CSD data from silicon probes, these models can be applied to existing datasets, facilitating a more efficient and reproducible approach to data analysis. By leveraging pre-recorded datasets, this method significantly reduces the demands of data collection, eliminating the need for new surgeries, recordings, and preprocessing. Applying this framework to a tetrode dataset from 13 mice enabled the classification of over 400,000 ripples, effectively saving more than two years of experimental work that would have been required to obtain comparable data through new recordings. This highlights the potential for repurposing existing electrophysiological data to address novel research questions, optimising both time and resources.

In terms of specific considerations regarding the $\mathrm{Rad}^{\mathrm{sink}}$ vs. $\mathrm{LM}^{\mathrm{sink}}$ ripples model, I have demonstrated that an LDA classifier can significantly distinguish $\mathrm{Rad}^{\mathrm{sink}}$, baseline, and $\mathrm{LM}^{\mathrm{sink}}$ ripples using LFP waveforms from a single channel. These results are robust for several reasons. First, the model is agnostic to the known CSD differences, such as the stronger sink in *stratum lacunosum-moleculare* during $\mathrm{LM}^{\mathrm{sink}}$ ripples compared to baseline or $\mathrm{Rad}^{\mathrm{sink}}$ ripples.This is because, in a classification framework, classes (i.e. ripple types) are treated as independent labels, so the model has no prior knowledge of the underlying $\mathrm{LM}^{\mathrm{sink}}\rightarrow$ baseline $\rightarrow$ $\mathrm{Rad}^{\mathrm{sink}}$ gradient — unlike regression, which captures such structure. Yet, the fact that the LDA model retrieves the expected gradient of CSD strengths in *stratum* *radiatum* and *lacunosum-moleculare* - with baseline ripples always falling between Rad^{sink} and LM^{sink} - directly from the LFP waveforms, supports the findings from Chapter 3 that ripple features exist along a continuum. Future studies could also explore hyper-parameter searches to optimise parameters such as the duration of LFP waveforms, the number of PCs used as input to the classifier, or even the model architecture itself (e.g., testing more complex classifiers such as support vector machines (SVMs) or convolutional neural networks (CNNs) to assess whether the simpler LDA is sufficient, particularly in terms of generalisation).

Furthermore, I trained the classifier only on the 'sharp-wave' component of the LFP, obtained via low-pass filtering (see Section 4.2.3). Thus, the classifier had no information about the ripple frequency, which was filtered out. The fact that the model still retrieved the expected differences in ripple frequency, validated against silicon probe ground truths (see Section 4.3.3), suggests an intrinsic link between 'sharp-wave' components and ripple frequency. This aligns with previous reports showing a correlation between ripple LFP waveforms and frequency (Chang et al., 2025; Sebastian et al., 2023). However, my approach is more robust since, unlike prior studies that used ripple-band specific LFP waveforms (and therefore included explicit frequency information), my analysis demonstrates that 'sharp-wave' components alone contain sufficient information to distinguish ripple types and predict their frequency and amplitude. In fact, while there was a slight increase in classification accuracy when adding ripple frequency as an additional feature to the LDA inputs, this improvement was not significantly greater than using the 'sharp-wave' component alone (Figure 4.3). One might argue that even this slight increase could be informative to the model. However, keeping the model agnostic to ripple frequency has the advantage of allowing cross-validation of the classifier's output by examining the ripple frequencies of putative Rad^{sink} and LM^{sink} ripples (Figure 4.4), thus providing a way to ensure that the model is not biased toward false positives. This highlights a key advantage of my approach over previous studies, where conclusions about waveform–frequency relationships may be partly circular, as frequency information is embedded in the waveforms used for classification (Sebastian et al., 2023).

Another important consideration concerns the choice of 'sharp-wave' components used to train the LDA classifier. Lopes-dos-Santos et al., 2023 showed that ripple-triggered waveforms vary with recording depth within the pyramidal layer: channels closer to *stratum oriens* tend to show more upward deflections at the ripple peak (similar to $\mathrm{LM}^{\mathrm{sink}}$), while those closer to *stratum radiatum* exhibit reversed polarity, more characteristic of $\mathrm{Rad}^{\mathrm{sink}}$. One might argue that using 'sharp-wave' components from a single channel to classify ripple types could be biased by channel depth - for instance, a ripple might be classified as $\mathrm{LM}^{\mathrm{sink}}$ in one channel but as $\mathrm{Rad}^{\mathrm{sink}}$ in another. However, this is controlled for in the silicon probe dataset by consistently selecting the channel at the centre of the pyramidal layer, identified using electrophysiological landmarks (see Section 2.3).

Regarding tetrode datasets, I took multiple steps to mitigate this potential bias. First, I selected, for each tetrode, the channel with the waveform most similar to that observed at the pyramidal centre in silicon probes. Second, ripple detection was performed across all CA1 pyramidal channels simultaneously, ensuring that a ripple detected in one channel was visible across all others. Third, I observed both ripple types within the same recording site (see Chapters 3 and 4.2.3), supporting the notion that $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples are intrinsically different, rather than artifacts of channel selection. Such a bias would only become problematic if ripple detection were performed on a single channel, in which case the relative position (closer to *oriens* or *radiatum*) could bias the detection toward $\mathrm{LM}^{\mathrm{sink}}$ or $\mathrm{Rad}^{\mathrm{sink}}$ ripples, respectively. For this reason, throughout this thesis, I deliberately avoid comparing the proportion of $\mathrm{Rad}^{\mathrm{sink}}$ vs. $\mathrm{LM}^{\mathrm{sink}}$ ripples in relation to behaviour or other functional correlates, as this ratio could indeed be confounded by tetrode positioning.

Finally, as this model is used for all subsequent analyses, an important question is whether potential depth biases in tetrode placement could undermine the results later presented in Chapters 5 and 6. On the contrary, if anything, such biases would introduce noise and make the observed differences between ripple types harder to detect, rather than creating false positives. For instance, later in Chapter 5, I show that CA1 principal cells fire more strongly during $\mathrm{Rad}^{\mathrm{sink}}$ than $\mathrm{LM}^{\mathrm{sink}}$ ripples. If on one recording day all tetrodes were deeper (closer to *oriens*, potentially favouring $\mathrm{LM}^{\mathrm{sink}}$ classification) and on another day closer to *radiatum* (potentially favouring $\mathrm{Rad}^{\mathrm{sink}}$), such between-session variability would introduce noise, reducing the ripple-type effects within sessions. Thus, the fact that strong and consistent differences between ripple types emerged despite such potential variability strengthens the conclusion that $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples are intrinsically different and not artefacts of classification or electrode placement.

To further validate that classification accuracy was not biased by depth, I performed a complementary analysis, reported in the Appendix A, showing that LDA predictions remained consistent across different recording depths within the pyramidal layer.

In the next chapter, I apply this LDA-based ripple classifier to the tetrode dataset to examine how CA1 and CA3 cells respond to $\mathrm{Rad}^{\mathrm{sink}}$ versus $\mathrm{LM}^{\mathrm{sink}}$ ripples.

But it is not simply the number of neurons that matters. Instead, it is the connectivity and the connectivity-confined communication that largely determines the share single neurons have in brain computations. It is much like the smartness issue with us humans.

— György Buzsáki (Buzsáki, 2006)

5

Recruitment of CA1 and CA3 neurons during Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples

Contents

5.1 Introduction

As mentioned in Chapter 1, hippocampal ripples are among the most synchronous events in the brain, characterised by highly coordinated CA1 activity within short time windows (Buzsáki et al., 1983; Buzsáki, 1986; Buzsáki et al., 1992; Ylinen et al., 1995). These events facilitate the reactivation of coactivity patterns from waking experience, supporting offline processing and stabilisation during sleep (Girardeau et al., 2009; Nádasdy et al., 1999; Wilson & McNaughton, 1994). While ripple-related memory functions are well documented, the mechanisms underlying the generation of individual ripples remain not fully understood. In particular, as discussed in Chapter 1, there is substantial diversity in how inputs to CA1 shape ripple characteristics and timing. Yet how this input diversity relates to variability in the recruitment of hippocampal neurons at the single-ripple level remains largely unexplored.

The prevailing consensus, as discussed in Chapter 1, is that CA3 plays a key role in ripple generation, primarily through its highly recurrent network and strong feedforward excitation to CA1 — a process reflected by the sharp wave component in *stratum radiatum* (Buzsáki et al., 1983; Buzsáki, 1986). This framework implies that CA3 activity is essential for ripple initiation and that all ripples should be preceded by a sharp wave in *radiatum*. However, as discussed in Section 1.4, recent studies have challenged this view by showing that ripples are diverse and do not all homogeneously originate from CA3. In line with this, I showed in Chapter 3 that a subset of ripples — LM$^{\text{sink}}$ ripples — exhibit significantly weaker sharp waves and a strong current sink in *lacunosum-moleculare*. This observation suggests that these events may rely on an alternative, or at least less CA3-dependent, mechanism of generation.

Supporting the idea that ripple generation is influenced by multiple upstream regions, Sebastian et al., 2023 demonstrated that optogenetic stimulation of CA3 and CA2 differentially shapes the ripple LFP waveform in CA1 *stratum pyramidale*. Given that ripple waveforms exhibit distinct profiles (as shown in Chapters 3 and 4), it follows that differences in ripple-associated CSD patterns may reflect distinct upstream activity. If laminar current profiles indeed capture variation in CA3 engagement, then LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ ripples may differentially recruit the hippocampal circuit.

A key prediction arising from these observations is that CA3 neurons should be less active, or play a reduced role, during LM$^{\text{sink}}$ ripples compared to Rad$^{\text{sink}}$ ripples. If true, this would indicate that ripple events are not all generated by the same mechanisms but rather that distinct ripple subtypes engage hippocampal circuits differently. To investigate this, it is therefore necessary to examine not only CA1 (where ripples occur) but also CA3, the primary upstream region thought to trigger them.

Building on this reasoning, in this chapter, I investigate how hippocampal neurons in CA1 and CA3, recorded using tetrodes, are differentially recruited during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. To explore how these distinct CSD profiles shape neuronal activity, I examine differences in recruitment across multiple levels. I first quantify single-cell responses to each ripple type, then assess coactivity at the level of cell pairs, and finally evaluate population-level dynamics.

5.2 Methods

5.2.1 Preferred ripple phase and phase coherence

To measure how much CA1 neurons couple to the local ripple oscillations, I band-passed filtered the local field potentials (LFPs) between 90 and 300 Hz using a 4th-order Butterworth filter and extracted the instantaneous phase of the ripple signal throughout each ripple event, from onset to offset. For each CA1 principal cell and interneuron, I calculated the probability of spiking relative to the local ripple phase, using the phase signal recorded from the tetrode where that cell was recorded from. The phase range was divided into 24 equally spaced bins between 0 and 2π, and spike-phase probabilities were computed separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples (see Figure 5.6). From these distributions, I determined the preferred ripple phase as the angular component of the mean resultant vector, while the mean phase coherence was obtained from the magnitude of this vector, representing the strength of phase locking to the ripple oscillations (see Figure 5.7).

5.2.2 Response similarity between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

To assess the similarity in firing responses of individual principal cells between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, I computed the Pearson correlation between the firing rates of each CA1 and CA3 principal cell, or the likelihood of firing at least one spike during a ripple event, across the two ripple types (Figures 5.8(A,B)). A high correlation coefficient indicated a similar response to both ripple types, such that a cell exhibiting a high firing rate (or high spiking probability) during Rad$^{\text{sink}}$ ripples showed comparable activity during LM$^{\text{sink}}$ ripples. To compare response similarity between CA1 and CA3 (Figures 5.8(C,D)), I accounted for differences in cell numbers between the two regions by performing 1,000 permutations. In each permutation, I randomly sampled five principal cells from each region and computed the response similarity score. For each sleep session, the mean across these permutations provided an estimate of response similarity for both CA1 and CA3. This analysis included sleep sessions with at least five principal cells recorded in the respective CA region (n = 208 CA1 and 171 CA3 sleep sessions; Figures 5.8).

5.2.3 Interneurons to principal cells firing ratio during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

In Figure 5.5, I quantified the relative firing rate of interneurons to principal cells during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. For both CA1 and CA3, I computed the mean firing rate of individual interneurons during ripples ($rate_{interneurons}$) and normalised it by the mean firing rate of all principal cells ($rate_{principals}$) recorded on the same day:

$$Ratio = \log_{10}\left(\frac{rate_{interneurons}}{rate_{principals}}\right) \tag{5.1}$$

Positive values of this ratio indicated a relatively higher firing rate of interneurons compared to principal cells. This analysis was conducted separately for Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples, and the resulting ratios were statistically compared to assess differences in interneuron-principal cell engagement between the two ripple types.

5.2.4 Discriminating population activity during Rad$^{\mathrm{sink}}$ vs. LM$^{\mathrm{sink}}$ ripples

To assess whether the overall structure of principal cell population activity significantly differed between Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples, I trained a logistic regression model to classify ripple identity (i.e., whether an event was a Rad$^{\mathrm{sink}}$ or LM$^{\mathrm{sink}}$ ripple) based on the z-scored ripple-nested population vectors (PVs) containing the number of spikes discharged by CA1 or CA3 principal cells (Figure 5.2). To control for the number of predictors across CA1 and CA3 when comparing these two hippocampal regions, I trained these models using multiples of five principal cells (i.e., 5, 10, 15, $\cdots$, N, where N represents the maximum number of cells divisible by 5). For each step in the number of cells, I performed 200 permutations, randomly selecting cells for the model. Each model was cross-validated 20 times (80% training, 20% testing), with accuracy measured as the mutual information between the true ripple classes and model predictions (0 bit = chance; 1 bit = perfect prediction; see Section 2.4.4).

To account for class imbalance, I matched the number of events in each ripple class by resampling the larger class to match the size of the smaller class in each permutation. Additionally, for each step in the number of cells, I trained a separate model using surrogate PVs, in which the coactivity of principal cells was shuffled while preserving both the individual cell firing rates and the population firing rate within each ripple (see method Section 2.4.5). For each sleep session, I computed the mean accuracy (mutual information) across the 200 permutations as a function of the number of cells used for training. This procedure was repeated for the shuffle control models.

These analyses were performed separately for CA1 and CA3. In Figures 5.9 and Table 5.1, I compared the performance of models trained with 15 principal cells, which included 129 sleep sessions for CA1 and 89 for CA3.

In Figure 5.1 and 5.10, I conducted an alternative analysis in which ripple identity was classified based on the full temporal pattern of spiking activity surrounding each ripple, rather than using ripple-nested PVs. In each cross-validation iteration, I extracted the spike train of each cell within a ±200-ms window around the ripple peak and smoothed the temporal pattern using a Gaussian kernel (SD = 3 ms). This produced a matrix for each cell with dimensions ($N_{\text{ripples}} \times N_{\text{bins}}$), where N_{ripples} represents the number of ripples (ensuring an equal number of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples to maintain class balance), and N_{bins} corresponds to the number of time bins in the ±200-ms window. To extract dominant temporal patterns, I applied PCA to this matrix, retaining the first two principal components (PCs) for each cell. This yielded a matrix of dimensions ($N_{\text{bins}} \times 2$), representing the strength of the first two PCs across all ripples. These matrices were then concatenated across all cells, forming a larger matrix of dimensions ($N_{\text{ripples}} \times (2 \cdot N_{\text{cells}})$), where N_{cells} is the total number of CA1 (or CA3) principal cells recorded that day. Each PC was z-scored, and the resulting matrix was used as input for the classifier described above. A schematic of this method is shown in Figure 5.1. This procedure was repeated across all 20 cross-validation iterations, ensuring that PCA was independently validated alongside model training.

5.2.5 Structural balance

The interaction of positive and negative edges within a network determines both stable and unstable relationship patterns (Gava et al., 2024). I computed the structural balance of hippocampal graphs in Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples to assess the stability and coherence of underlying neuronal coactivity patterns. In this

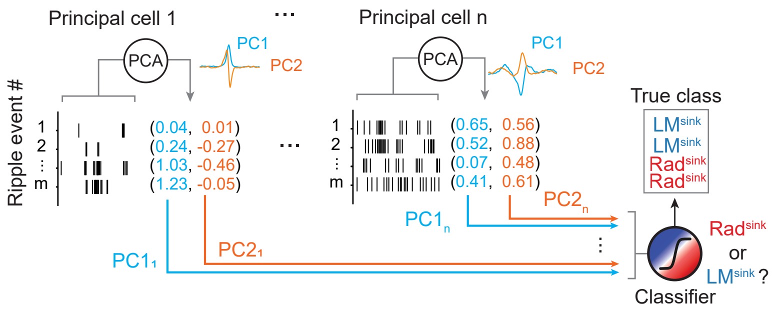

Figure 5.1: Classifying Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples based on spiking temporal dynamics.
Schematic representation of the method used to classify Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples based on the temporal features of spike trains surrounding ripple events (see Section 5.2.4). For each principal cell, the first two principal components (PCs) of its spike-time responses were computed independently. The PCs from all recorded principal cells were then concatenated and used as input features for a logistic regression classifier to predict whether a given ripple corresponded to a Rad$^{\text{sink}}$ or LM$^{\text{sink}}$ event. The two example cells shown here represent real data from a CA1 principal cell (left) and a CA3 principal cell (right).

network analysis, triads of neurons were classified as either balanced or unbalanced (Figure 5.13(B)). This classification follows principles from social networks, where a triad is considered balanced if it consists of either three positive edges (e.g., "the friend of my friend is my friend") or two negative edges (e.g., "the enemy of my friend is my enemy"). Any triad with an odd number of negative relationships is deemed unbalanced, as conflicting pathways between the same pair of nodes would lead to signal interference. I classified all triads in Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ coactivity graphs based on the signs of their three edges. The structural balance of a graph was defined as the proportion of its triads that were classified as balanced (Figures 5.13(B-D); n = 208 sleep sessions (graphs) for CA1 and 171 sleep sessions (graphs) for CA3).

5.2.6 Population-level sparsity

The sparsity (S) of a population firing vector (x) was calculated using the Gini index (Gini, 1921; Hurley & Rickard, 2009; McHugh et al., 2022) as:

$$S = \frac{\sum_{i=1}^{N}(2i - N - 1)x_i}{N\sum_{i=1}^{N} x_i} \tag{5.2}$$

where x represents the PV of spike counts for each principal cell, ordered in ascending rank, within a 50-ms window centred at the ripple peak (either Rad$^{\text{sink}}$ or LM$^{\text{sink}}$). N denotes the total number of simultaneously recorded principal cells, and i represents the rank of spike count in ascending order. PVs with more evenly distributed spike counts exhibit a lower Gini index (indicating lower sparsity), whereas those where spike counts are concentrated in a few neurons have a higher Gini index (indicating higher sparsity). In Figures 5.14(C) and 5.15(B), I report the sparsity of CA1 and CA3 principal cells respectively (n = 208 sleep/rest sessions for CA1 and 171 for CA3).

5.2.7 Cross-population vector inclusion of activity motifs

To quantify the degree of inclusion between the coactivity patterns in LM$^{\text{sink}}$ versus Rad$^{\text{sink}}$ ripples, I computed the overlap among their PVs. Each ripple-nested PV captured the instantaneous firing activity of CA1 or CA3 principal cells within a 50-ms window centred on the ripple peak. To construct a binary representation, I assigned a value of 1 to neurons that fired at least one spike within the window and 0 otherwise. This approach enabled the construction of an adjacency matrix, where each entry quantified the overlap between the activity patterns of different PVs (see Figure 5.16).

Given two population vectors, $PV_m = [p_{m1}, p_{m2}, \cdots, p_{mn}]$ and $PV_q = [p_{q1}, p_{q2}, \cdots, p_{qn}]$, where n represents the total number of principal cells in the population, and p_{mk}, p_{qk} indicate the binary firing state of the k-th neuron in each PV, I quantified their overlap using an asymmetric Jaccard coefficient:

$$I_{m,q} = \frac{|PV_m \cap PV_q|}{|PV_m|} \tag{5.3}$$

where $\cap$ represents the intersection of active cells between the two PVs (i.e. how many cells are active in both PVs), and $|\cdot|$ denotes the PV cardinality (i.e. the number of active cells in that PV). A value of 1 indicated that all cells active in PV_m were also active in PV_q, while a value of 0 indicated no shared activity. Additionally, I computed the reverse measure:

$$I_{q,m} = \frac{|PV_m \cap PV_q|}{|PV_q|}$$

Since this inclusion measure is inherently asymmetric, $I_{m,q}$ and $I_{q,m}$ were not necessarily equal, allowing for directional comparisons between PVs. To emphasise this directional nature, I adopted the notation $I_{m,q} = I_{m \to q}$.

5.2.8 Neural inclusion of LM$^{\text{sink}}$ into Rad$^{\text{sink}}$ ripples

For each sleep session, I quantified the inclusion among pairs of ripple-nested PVs (LM$^{\text{sink}}$ and Rad$^{\text{sink}}$) to examine how active neurons in LM$^{\text{sink}}$ ripples contributed to Rad$^{\text{sink}}$ ripples. This was done independently for CA1 and CA3 PVs. To account for potential biases due to class imbalance, I matched the number of PVs from each ripple type to the smaller of the two sets, denoted as D. The mean inclusion of LM$^{\text{sink}}$ into Rad$^{\text{sink}}$ ($I_{LM \to Rad}$) was computed as the average across these subsampled D PVs.

To mitigate potential biases introduced by specific subsampling choices, I used a resampling approach. In each iteration, I randomly selected D PVs from the larger class and recomputed $I_{LM \to Rad}$. This procedure was repeated 1,000 times to capture variability in the data. Additionally, to establish a chance-level baseline, I performed a shuffling control in which the active cell distribution within LM$^{\text{sink}}$ ripples was randomised while keeping Rad$^{\text{sink}}$ ripples unchanged (see Section 2.4.5). Specifically, for each LM$^{\text{sink}}$ ripple, active cells were reassigned randomly while maintaining both the original number of active cells (reflecting ripple sparsity) and the overall number of LM$^{\text{sink}}$ ripples in which each neuron was active. The observed inclusion values were then compared against the shuffled control distribution ($I_{LM_{\text{shuffled}} \to Rad}$) to determine whether the overlap in active neurons between LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ ripples was greater than expected by chance (Figure 5.17(B); $n = 208$ sleep/rest sessions for CA1 and 171 for CA3).

This approach ensured that the observed inclusion patterns were not merely a consequence of LM$^{\text{sink}}$ ripples being sparser than Rad$^{\text{sink}}$ ripples. Notably, direct comparisons between $I_{LM\rightarrow Rad}$ and $I_{Rad\rightarrow LM}$ were avoided, as the inherent sparsity differences between the two ripple types (i.e., the denominator in equation 5.3 being greater for Rad$^{\text{sink}}$ than for LM$^{\text{sink}}$; see Figures 5.14 and 5.15) could bias their relative inclusion measures.

5.2.9 Population-level dimensionality

To quantify the intrinsic dimensionality of CA1 and CA3 ripple-nested activity within a 50-ms window centred at the ripple peak, I applied the angle-based intrinsic dimensionality (ABID) method (Thordsen & Schubert, 2022). ABID is a non-linear measure, making it particularly suited for capturing complex neural activity patterns. It estimates dimensionality (D) by analysing the cosine similarity between each ripple PV and its k-nearest neighbours (k = 50). For each sleep/rest session with M ripples (computed separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$), the ripple-nested PVs were z-scored, and the dimensionality of each ripple m was computed as:

$$D_m = \frac{k^2}{\sum_{i=1}^{k}\sum_{j=1}^{k} S_{ij}^2} \tag{5.4}$$

where S_{ij} represents the cosine similarity between the normalised k-nearest PVs of ripple m. Dimensionality decreases when neighbouring PVs are tightly clustered, leading to higher summed similarities, while greater dispersion among PVs results in higher dimensionality. To control for class imbalance, I randomly sampled 100 PVs per ripple type, repeated this process 1,000 times, and computed the mean intrinsic dimensionality per session (Figure 5.18; n = 208 sleep/rest sessions for CA1 and 171 for CA3).

5.3 Firing responses of individual CA1 and CA3 neurons during Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples

In this section, I focus on single-cell responses across CA1 and CA3 principal cells and interneurons during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Having established in Chapter 4 that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples can be accurately discriminated based on their pyramidal LFP waveforms, I applied the classification model presented in Sections 4.2.3 and 4.3 to the tetrode dataset (Figure 5.2). As previously detailed, applying this model to the tetrode dataset yielded LFP features and ripple frequency distributions that were consistent with those obtained using ground-truth CSD classifications from silicon probe recordings (see Section 4.3.3).

The tetrode dataset used in this thesis, as detailed in Sections 2.1.3 and 2.1.6, comprises simultaneous multichannel recordings from CA1 and CA3 pyramidal layers (Figure 5.2). On average, each recording day included 42.9 (IQR: 27.0 - 62.8) principal cells and 9.0 (IQR: 6.0 - 11.0) interneurons, amounting to a total of 3,521 principal cells and 741 interneurons across 280 hours of sleep/rest from 13 mice. The dataset spans 244 sleep/rest sessions, with a mean session duration of 68.8 minutes (IQR: 48.0 - 90.2 minutes).

5.3.1 Neuronal responses during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

To investigate how individual hippocampal cells responded during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, I constructed peri-event time histograms (PETHs) for CA1 and CA3 principal cells and interneurons (see Section 2.4.2). CA1 principal cell activity transiently increased in both Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, but the mean peak firing rate was significantly higher in Rad$^{\text{sink}}$ ripples [Figure 5.3; mean peak rate (95% CI): Rad$^{\text{sink}}$, 21.88 (21.27 - 22.50) Hz; LM$^{\text{sink}}$, 13.12 (12.70 - 13.55) Hz; $p < 10^{-5}$, paired bootstrap test; $n = 2,196$ CA1 principal cells]. A similar pattern was observed in CA3, where principal cells also fired at higher rates during Rad$^{\text{sink}}$ ripples [Figure 5.3; mean peak rate (95% CI): Rad$^{\text{sink}}$, 8.82 (8.14 - 9.54) Hz; LM$^{\text{sink}}$, 5.18 (4.80 - 5.61) Hz; $p < 10^{-5}$, paired bootstrap test; $n = 1,325$ CA3 principal cells].

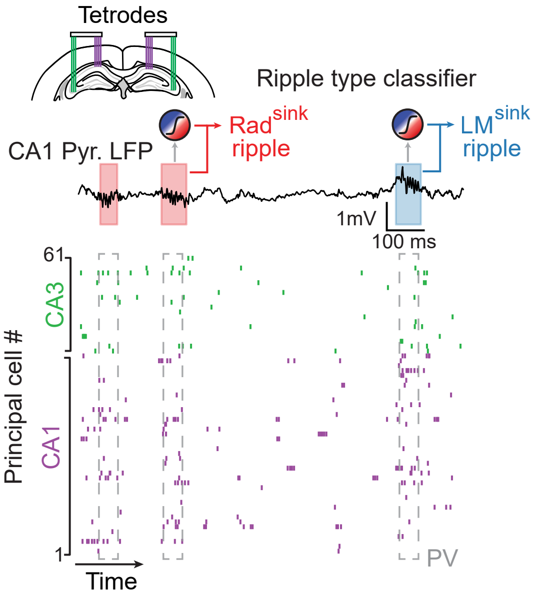

Figure 5.2: Raw dual-site tetrode recording distinguishing Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples
Example dual-site 14-tetrode ensemble recording of CA1 and CA3 principal cells (for clarity, only principal cells are shown, although interneurons were also recorded). Top trace: raw LFP signal from the CA1 pyramidal layer. Bottom raster plot: (colour-coded) spike trains of CA1 and CA3 principal cells, with each row representing a different cell (1-s sample shown for clarity). The population activity vector (PV) for individual ripples was extracted using a 50-ms window centred at the ripple envelope peak. Tetrode-recorded Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples were classified using the silicon probe-validated LDA model (see Section 4.3 and Figure 4.2).

Interestingly, while CA1 principal cells exhibited similar response patterns across both ripple types, CA3 principal cells displayed qualitatively distinct firing dynamics. During Rad$^{\text{sink}}$ ripples, CA3 principal cells showed a peak firing response centred around the ripple power peak. However, during LM$^{\text{sink}}$ ripples, their firing pattern was bimodal, with transient increases both at the ripple peak and approximately 100 ms before it (Figure 5.3). The timing of this earlier increase in CA3 activity coincided with a *stratum radiatum* sink (Figures 3.1 - 3.3).

Interneurons followed similar trends as principal cells, with higher overall firing rates in CA1 during Rad$^{\text{sink}}$ ripples and marked qualitative differences in CA3 across Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. CA1 interneurons exhibited stronger responses in Rad$^{\text{sink}}$ compared to LM$^{\text{sink}}$ ripples (Figure 5.4; mean peak rate (95% CI): Rad$^{\text{sink}}$, 53.04 (48.36 - 57.85) Hz; LM$^{\text{sink}}$, 31.57 (28.55 - 34.70) Hz; $p < 10^{-5}$, paired bootstrap test; $n = 408$ CA1 interneurons). Similarly, CA3 interneurons showed higher activity in Rad$^{\text{sink}}$ ripples (Figure 5.4; mean peak rate (95% CI): Rad$^{\text{sink}}$, 41.48 (37.02 - 46.08) Hz; LM$^{\text{sink}}$, 27.57 (24.65 - 30.57) Hz; $p < 10^{-5}$, paired bootstrap test; $n = 333$ CA3 interneurons).

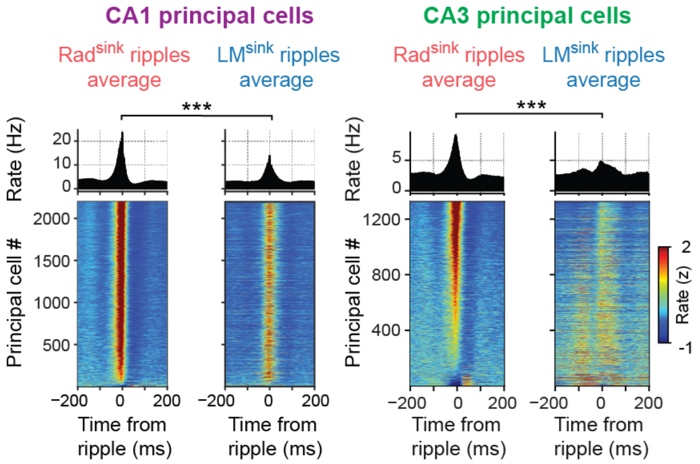

Figure 5.3: CA1 and CA3 principal cells firing responses during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
Triggered average firing response (peri-event time histograms, PETHs) of CA1 (left) and CA3 (right) principal cells aligned to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. *Top panels*: overall population-averaged responses. *Bottom panels*: z-scored firing responses of individual cells (relative to their mean and SD during Rad$^{\text{sink}}$), sorted by firing rate during Rad$^{\text{sink}}$ ripples. Both CA1 and CA3 principal cells exhibit significantly higher firing rates in Rad$^{\text{sink}}$ compared to LM$^{\text{sink}}$ ripples (paired bootstrap tests).
$n = 2,196$ CA1 and $1,325$ CA3 principal cells. $^{***}p < 0.001$.

Despite the significant overall reduction in interneuron firing during LM$^{\text{sink}}$ ripples, PETHs suggested that CA3 interneurons maintained a relatively strong response compared to CA3 principal cells, whereas CA1 interneurons exhibited a more uniform rate reduction across ripple types (Figures 5.3 and 5.4). To quantify this, I computed a ratio of interneuron activity normalised by the average concurrent principal cell activity (see Section 5.2.3). This metric revealed that the relative difference between CA3 interneuron and principal cell firing rates was significantly higher during LM$^{\text{sink}}$ ripples compared to Rad$^{\text{sink}}$ ripples (Figure 5.5). In contrast, CA1 interneurons maintained a similar relative firing rate across ripple types. These findings suggest that while CA1 interneuron activity scales proportionally with principal cell firing across ripple types, CA3 interneurons exhibit a relatively stronger

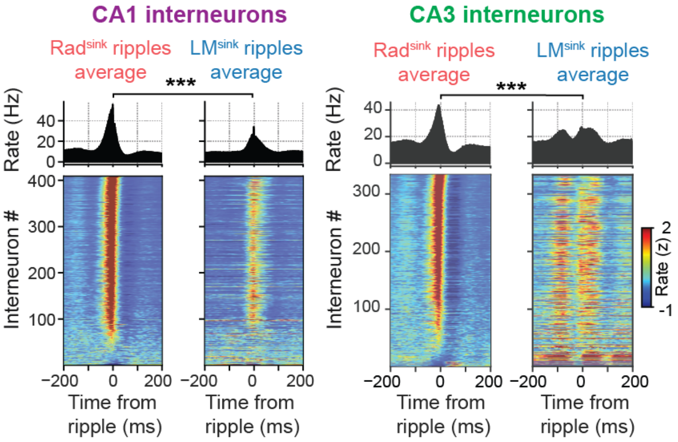

Figure 5.4: CA1 and CA3 interneurons firing responses during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
Triggered average firing response (peri-event time histograms, PETHs) of CA1 (left) and CA3 (right) interneurons aligned to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. *Top panels*: overall population-averaged responses. *Bottom panels*: z-scored firing responses of individual cells (relative to their mean and SD during Rad$^{\text{sink}}$), sorted by firing rate during Rad$^{\text{sink}}$ ripples. Both CA1 and CA3 interneurons exhibit significantly higher firing rates in Rad$^{\text{sink}}$ ripples compared to LM$^{\text{sink}}$ ripples (paired bootstrap tests).
$n = 408$ CA1 and 333 CA3 interneurons. $^{***}p < 0.001$.

response during LM$^{\text{sink}}$ ripples than would be expected based on the firing rate decrease observed in CA3 principal cells.

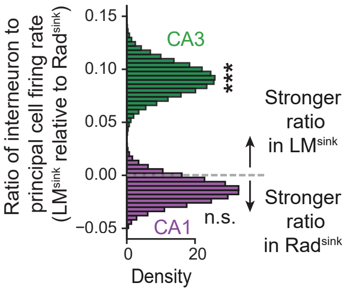

Figure 5.5: Interneuron-to-principal cell firing ratio in CA1 and CA3 during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
Mean difference in the interneuron-to-principal cell firing ratio between LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ ripples. While this ratio remains similar between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples in CA1 ($p = 0.48$, paired bootstrap test; $n = 408$ CA1 interneurons), CA3 exhibits a significantly higher ratio in LM$^{\text{sink}}$ ripples, indicating increased interneuron firing relative to principal cells compared to Rad$^{\text{sink}}$ ripples ($p < 10^{-5}$, paired bootstrap test; $n = 333$ CA3 interneurons).
$^{\text{n.s.}}p > 0.05$, $^{***}p < 0.001$.

Finally, I tested whether CA1 principal cells and interneurons coupled to different phases of the ripple oscillations across ripple types. To investigate this, I computed spike-phase probability distributions separately for CA1 principal cells and interneurons during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples (Figure 5.6), which were then used to measure their preferred ripple phase and coupling strength (i.e., spike-phase coherence) (see Section 5.2.1).

Overall, irrespective of the ripple type, CA1 neurons firing was coupled to the troughs of ripple oscillations, as previously reported (Buzsáki et al., 1992; Ylinen et al., 1995). Interestingly, CA1 principal cells fired at slightly later phases during Rad$^{\text{sink}}$ ripples compared to LM$^{\text{sink}}$ ripples, while CA1 interneurons exhibited the opposite trend, firing at earlier ripple phases during Rad$^{\text{sink}}$ ripples (Figure 5.7(A,B)). However, the strength of phase coupling remained similar across ripple types (Figure 5.7(C,D)).

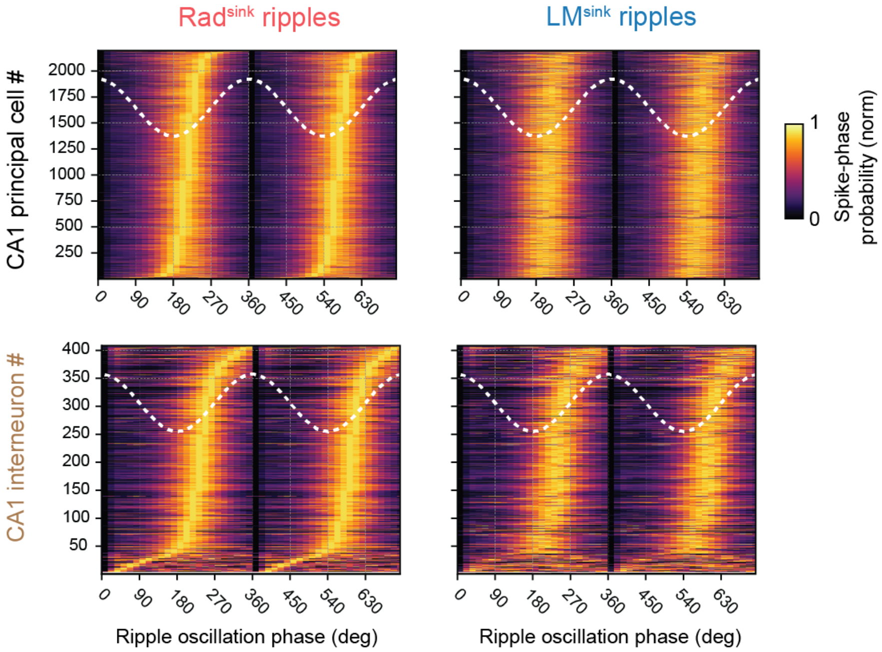

Figure 5.6: CA1 neurons spiking as a function of ripple oscillation phase. Heatmaps depicting spiking probabilities as a function of ripple phase for CA1 principal cells (top) and interneurons (bottom) during Rad$^{\text{sink}}$ (left) and LM$^{\text{sink}}$ (right) ripples. Spiking probabilities for each cell were max-normalised by dividing by the maximum probability within each ripple group. Since normalisation was performed independently for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, no direct assumptions can be made about phase coherence differences between the two (see Figure 5.7 for spikes-coherence analysis). Cells are sorted by their preferred phase during Rad$^{\text{sink}}$ ripples, from early to late phase (bottom to top rows), allowing a qualitative assessment of phase preferences across the population.

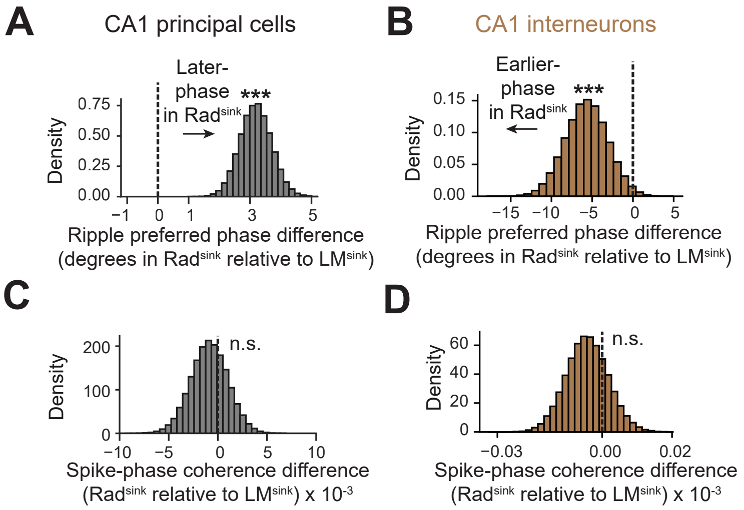

Figure 5.7: Phase-locking of CA1 neurons during Rad^{sink} and LM^{sink} ripples. (A-D) Mean difference between Rad^{sink} and LM^{sink} ripples (relative to LM^{sink}) in preferred ripple firing phase and spike-phase coherence for CA1 principal cells and interneurons.
(A,B) Preferred firing phase differences for CA1 principal cells (A) and interneurons (B). CA1 principal cells fired at later phases during Rad^{sink} ripples, whereas interneurons fired at earlier phases compared to LM^{sink} ripples (CA1 principal cells: $p < 10^{-5}$; interneurons: $p = 0.027$, paired bootstrap tests).
(C,D) Spike-phase coherence for CA1 principal cells (C) and interneurons (D), showing no significant differences between ripple types (principal cells: $p = 0.653$; interneurons: $p = 0.472$, paired bootstrap tests).
For (A),(C): $n = 2,196$ CA1 principal cells; for (B),(D): $n = 408$ CA1 interneurons. $^{\text{n.s.}}p > 0.05$, $^{***}p < 0.001$,

Overall, these results show that Rad^{sink} and LM^{sink} ripples differentially modulate CA1 and CA3 neurons, with a general increase in firing during Rad^{sink} ripples. Moreover, CA3 neurons exhibited qualitatively different firing patterns between ripple types, with a 'bimodal' response during LM^{sink} ripples. Finally, CA1 principal cells and interneurons fired at different ripple phases depending on the ripple type, further suggesting that Rad^{sink} and LM^{sink} ripples might originate differently within the hippocampus network.

5.3.2 Coactivity patterns within CA1 and CA3 during Rad^{sink} and LM^{sink} ripples

In this section, I explore whether specific CA1 or CA3 principal cells preferentially respond during either Rad^{sink} or LM^{sink} ripples. Specifically, I investigate whether certain clusters of cells are preferentially active in one ripple type but not the other, which does not appear to be the case based on peri-event time histograms (PETHs).

To gain insight into this, I examined the mean firing rate, or spiking probability, of principal cells in CA1 and CA3 during Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples for each recording day (Figure 5.8(A–B)). These plots revealed a high positive correlation between a cell's firing rate and spiking probability across ripple types [mean Pearson correlation (95% CI): CA1, 0.86 (0.85 - 0.88) r; CA3, 0.75 (0.72 - 0.78) r; see Section 5.2.2]. This strong cross-ripple response similarity indicated that principal cells with a high firing rate during one ripple type also exhibited high firing rates during the other, with this effect being more pronounced in CA1 than in CA3 (Figure 5.8(A–D)). These results suggest that cells modulate their firing proportionally across ripple types, although at different absolute firing rates, as shown in Section 5.3.1.

Next, I examined whether the differences observed between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples could be explained solely by overall firing rate differences, as described in Section 5.3.1, or whether coactivity patterns also played a role. In other words, beyond simply responding more or less strongly, does the specific combination of CA1 and CA3 neurons that are active together during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples also matter?

To address this, I tested whether Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples differed in their population-level coactivity patterns (see Section 5.2.4). I trained classifiers to distinguish ripple types based on the PVs containing the spiking activity for either CA1 or CA3 principal cells. To disentangle the contributions of firing rate differences and coactivity patterns to ripple classification, I generated surrogate PVs that preserved individual neuron firing rates (i.e., the number of spikes each cell discharged across ripples) and population-wide firing rates (i.e., the number of spikes discharged in each ripple) but shuffled the specific neurons that were transiently coactive in each ripple (see Section 2.4.5). These control models outperformed chance (50% accuracy ↔ 0 bits), confirming that some information about ripple types was derived from overall firing rate differences (Figures 5.9 and 5.3). However, when the

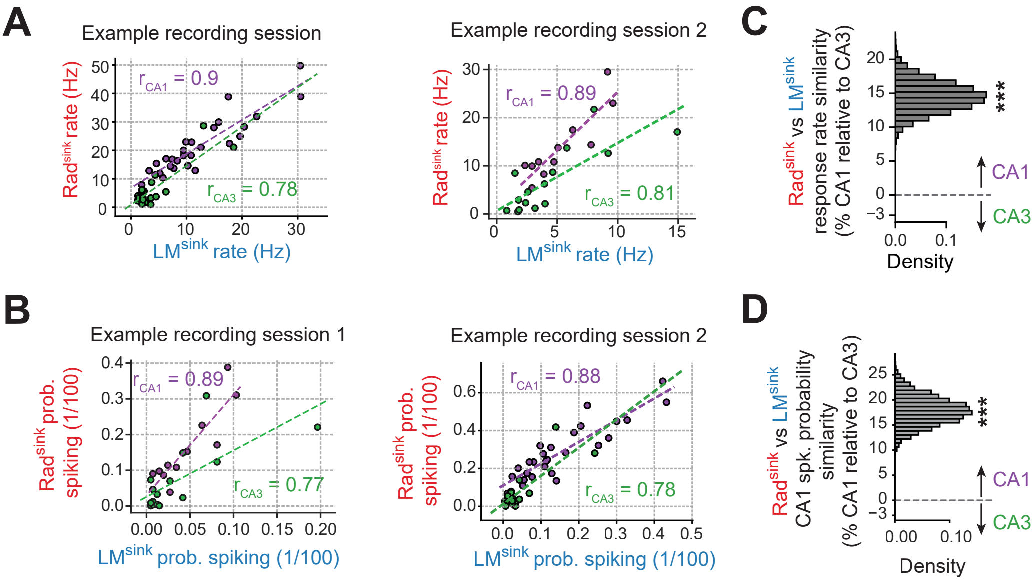

Figure 5.8: Activity similarity of CA1 and CA3 principal cells between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.

(A) Firing response similarity of individual principal cells to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples across two sleep sessions. Each point represents a CA1 (purple) or CA3 (green) cell. The Pearson correlation coefficient (r) quantifies the similarity in firing rate responses between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Dashed coloured lines indicate best-fit linear relationships.

(B) Same as (A), but showing the ripple-associated spiking probability instead of firing rate for the same sleep sessions.

(C–D) Mean difference in response rate similarity (C) and spiking probability similarity (D) between CA1 and CA3 principal cells, expressed as a percentage relative to CA3 response similarity (all $Ps < 10^{-5}$, bootstrap tests).

For (C),(D): $n = 208$ sleep sessions for CA1, $n = 171$ for CA3. $^{***}p < 0.001$.

true coactivity patterns were preserved (i.e., real data), the classifier significantly outperformed control models, suggesting that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples express distinct, non-redundant coactivity motifs (Figure 5.9 and Table 5.1).

Interestingly, CA1 exhibited more information about ripple type than CA3 in this model (Figure 5.9 and Table 5.1). This result was somewhat unexpected, as discussed in Section 5.3.1, since CA3 undergoes more qualitative changes between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, which could imply it contains more information for discriminating ripple types. However, this classification model only utilised activity at the ripple peak (i.e., PVs), whereas the qualitative differences in CA3 firing primarily arise from the overall temporal structure of the ripple response.

To account for this, I developed an alternative model that classified Rad$^{\text{sink}}$ and

LM$^{\text{sink}}$ ripples based on the entire spike train temporal structure surrounding each ripple event. Incorporating temporal information significantly improved classification accuracy, particularly within CA3, compared to models using only peak firing activity (Figure 5.10 and Table 5.1). Moreover, across different cell sample sizes, CA3 neurons provided more information for discriminating ripple types than CA1 neurons (Figure 5.10).

These findings demonstrate that while individual neurons exhibited high cross-ripple firing response similarity, the overall population and temporal patterns of coactive neurons differed significantly between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.

Together, the findings in this section highlight that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples differentially modulate hippocampal neurons, both at the level of single-cell responses and in terms of coactivity patterns. In the next section, I investigate these differences in cellpair coactivity in greater depth and examine how they relate to the overall structural organisation of the network.

Classifier input	**Ripple-peak population vectors**	**Entire temporal ripple response**
CA1 principal cells	9.25 (8.54 - 9.99)	13.64 (12.70 - 15.59)
CA3 principal cells	6.17 (4.81 - 7.62)	19.23 (17.70 - 20.81)

Table 5.1: Performances of models discriminating Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ from population activity.
Shown are the cross-validated performance [mean (95% CI); bits; $\times 10^{-2}$] of classifiers trained to predict Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripple types based on CA1 or CA3 activity, using either their population vectors (15 principal cells each; $n = 129$ sleep sessions for CA1, $n = 89$ for CA3) at the ripple peak or the entire temporal spiking profile around the ripple (see Section 5.2.4). Performance quantified as the mutual information (MI) between the true and predicted ripple types.

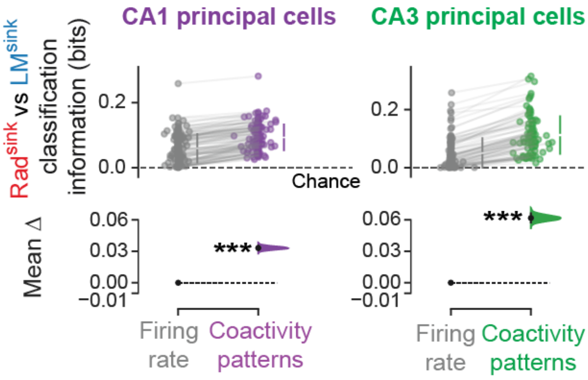

Figure 5.9: Discriminating Rad^{sink} and LM^{sink} ripples based on ripple-peak population activity.
Estimation plots showing the effect size of classifier accuracy (measured as mutual information) in distinguishing Rad^{sink} and LM^{sink} ripples using ripple-nested population vectors (PVs) from CA1 (left) and CA3 (right). Each distribution is compared to chance level (dashed lines) and a surrogate distribution that preserved individual cell firing rates and overall population statistics while shuffling coactivity patterns (gray distributions). Each dot represents a single sleep session with at least 15 CA1 or CA3 principal cells. For both CA1 and CA3, preserving population coactivity significantly improved classification accuracy compared to both controls (CA1: $p < 10^{-5}$; CA3: $p < 10^{-5}$; paired bootstrap tests; $n = 129$ sleep sessions for CA1, $n = 89$ for CA3).
$^{***}p < 0.001$.

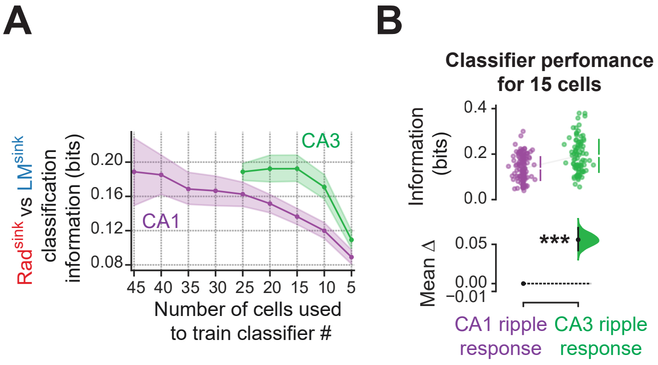

Figure 5.10: Discriminating Rad^{sink} and LM^{sink} ripples based on spiking responses around ripples.
(A,B) Classification accuracy of Rad^{sink} and LM^{sink} ripples in CA1 and CA3 using the method illustrated in Figure 5.1 (see Section 5.2.4). (A) Classifier accuracy as a function of the number of cells used during training, shown separately for CA1 and CA3. CA3 consistently carries more information than CA1 for discriminating between the two ripple types. (B) Corresponding estimation plot illustrating the effect size for the difference in classification accuracy between CA1 and CA3 classifiers (trained on 15 principal cells). CA3 classified Rad^{sink} and LM^{sink} ripples significantly more accurately than CA1 ($p < 10^{-5}$, bootstrap test; $n = 129$ sleep sessions for CA1, $n = 89$ for CA3; see also Table 5.1). This result aligns with the observation that CA3 principal cells exhibit greater differences in response between Rad^{sink} and LM^{sink} ripples compared to CA1 (Figure 5.3).
$^{***}p < 0.001$.

5.4 Levels of coactivity and structural stability during Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples

In Section 5.3.2, I demonstrated that beyond differences in firing rates, Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples are characterised by distinct coactivity patterns. This suggests that neuronal coactivity within CA1 and CA3 is not only different in magnitude between ripple types but may also be organised differently at the population level.

To explore this further, I first examined how pairs of CA1 and CA3 principal cells coordinate their activity during Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples by quantifying coactivity at the level of neuron pairs. Then, I analysed whether these coactivity motifs form stable, self-reinforcing relationships or, conversely, contain conflicting interactions that reduce coactivity cohesion.

5.4.1 Levels of principal cell coactivity

To characterise the coactivity motifs nested in Rad$^{\mathrm{sink}}$ versus LM$^{\mathrm{sink}}$ ripples, I applied the approach developed by Gava et al., 2024 (see Section 2.4.3). For each ripple type, I measured the coactivity of each cell pair (i, j) using the regression coefficient from the prediction of spike discharge in neuron i by that in neuron j, while controlling for the activity of the remaining population (Figure 5.11). Importantly, this method accounts for the coactivity of cell pairs while controlling for overall population firing rate differences between ripple types, which is crucial given the firing rate disparities observed in Section 5.3.1.

From this, for each recording day, I obtained an adjacency matrix representing the coactivity structure across all pairs of CA1 (or CA3) principal cells (Figure 5.11). Using these matrices, I estimated the average coactivity of single neurons relative to the rest of the population (see Section 2.4.3). At both the single recording-day resolution and across the entire dataset, I found that CA1 neurons exhibited higher coactivity during LM$^{\mathrm{sink}}$ ripples, whereas CA3 neurons displayed higher coactivity during Rad$^{\mathrm{sink}}$ ripples (Figure 5.12; all $Ps < 10^{-5}$, paired bootstrap test). These findings suggest that, after accounting for overall firing rate differences,

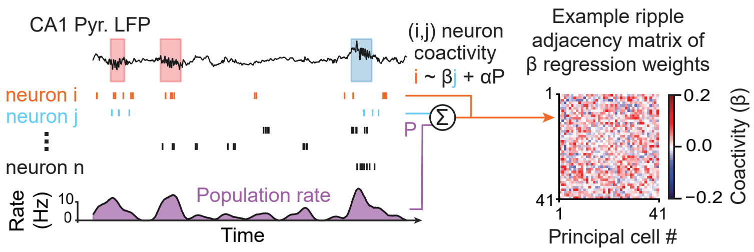

Figure 5.11: Measuring principal cell coactivity during ripples.
Schematic illustrating the quantification of ripple-nested principal cell coactivity. *Left*: Example LFP trace from the CA1 pyramidal layer alongside a raster plot of a subset of CA1 principal cells. The spiking activity of neuron i is predicted based on the spiking of neuron j, while controlling for the overall population firing rate (P), using a linear regression model. *Right*: Corresponding ripple adjacency matrix of beta coefficients, representing coactivity relationships between neurons.

CA1 principal cells exhibit higher average coactivity during LM$^{\text{sink}}$ ripples, whereas in Rad$^{\text{sink}}$ ripples, this coactivity is lower. Conversely, the opposite pattern is observed in CA3.

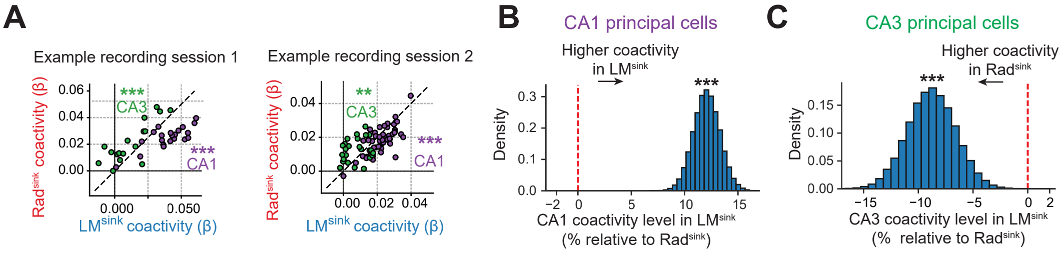

Figure 5.12: Differential CA1 and CA3 coactivity levels between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A) Scatter plots showing the mean coactivity of CA1 (purple) and CA3 (green) principal cells during LM$^{\text{sink}}$ versus Rad$^{\text{sink}}$ ripples from two example sleep sessions. Each point represents the average coactivity value of a CA1 or CA3 principal cell. The dashed black line indicates the diagonal of the first and third quadrants. CA1 principal cells exhibit higher coactivity during LM$^{\text{sink}}$ ripples, whereas CA3 principal cells show higher coactivity during Rad$^{\text{sink}}$ ripples (all $Ps < 0.01$, bootstrap tests; sleep session 1: $n = 21$ CA1 and $n = 17$ CA3 principal cells; sleep session 2: $n = 48$ CA1 and $n = 25$ CA3 principal cells). (B-C) Mean coactivity differences between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, expressed as a percentage change relative to Rad$^{\text{sink}}$ ripples. (B) CA1 principal cells show significantly higher coactivity during LM$^{\text{sink}}$ ripples than Rad$^{\text{sink}}$ ripples. (C) CA3 principal cells exhibit higher coactivity during Rad$^{\text{sink}}$ ripples compared to LM$^{\text{sink}}$ ripples. The vertical red dashed line represents the mean coactivity in Rad$^{\text{sink}}$ ripples. These results indicate, as highlighted by the black arrows, that CA1 principal cells are more coactive during LM$^{\text{sink}}$ ripples, whereas CA3 principal cells show the opposite pattern (CA1: $p < 10^{-5}$; CA3: $p = 2 \times 10^{-5}$; paired bootstrap tests; $n = 208$ sleep sessions for CA1 and $n = 171$ for CA3).
$^{**}p < 0.01$,$^{***}p < 0.001$.

5.4.2 Structural stability of CA1 and CA3 principal cell coactivity motifs

In the previous section, I showed that CA1 principal cells exhibit higher average coactivity during LM$^{\mathrm{sink}}$ ripples, whereas in Rad$^{\mathrm{sink}}$ ripples, this coactivity is lower. Conversely, the opposite pattern is observed in CA3. While this result indicates that neurons in CA1 are more coactive during LM$^{\mathrm{sink}}$ ripples, it does not reveal whether this coactivity forms internally stable interaction patterns or reflects a more variable and conflicting network structure.

To investigate whether LM$^{\mathrm{sink}}$ ripples promote more internally structured coactivity motifs, I applied structural balance theory from network topology to analyse the organisation of neuronal coactivity networks, as previously done by Gava et al., 2024. Graph-theoretic approaches have been shown to uncover the underlying architecture of hippocampal neuronal interactions (Gava et al., 2021, 2024; McHugh et al., 2024). Specifically, the arrangement of positively and negatively correlated spike trains determines whether a network exhibits structural balance, where interactions form internally stable, self-consistent motifs, or structural imbalance, where conflicting or opposing interactions may be less stable over time (Gava et al., 2024) (Figure 5.13(A)).

For instance, if neuron A tends to fire with both neuron B and neuron C, a balanced motif would be one where B and C are also coactive, reinforcing a stable coactivity structure (Figure 5.13(B)). Conversely, an imbalanced motif could occur if A fires with both B and C, but B and C rarely fire together or are even anti-correlated, creating an unstable configuration where their relationship is inconsistent across multiple events (Figure 5.13(B)).

By applying these mathematical tools, I aimed to determine whether the increased coactivity observed in LM$^{\mathrm{sink}}$ ripples corresponds to a more internally stable and structured coactivity network (Gava et al., 2024), as opposed to a more heterogeneous or conflicting organisation of neuronal interactions.

To investigate the stability of coactivity properties in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, I computed the structural balance of neuronal graphs embedding the measured coactivity relationships (Figure 5.13(A); see Section 2.4.3). In these graphs, each node represents a neuron, and each edge represents the coactivity strength of a cell pair (Figure 5.13(A)). I then identified balanced and unbalanced motifs of triadic relationships (Figure 5.13(B); see Section 5.2.5).

LM$^{\text{sink}}$ ripples contained a significantly higher proportion of balanced coactivity motifs in the CA1 population than Rad$^{\text{sink}}$ ripples, while structural stability in CA3 population activity remained similar across ripple types (Figures 5.13(C,D); CA1, $p = 6 \times 10^{-5}$; CA3, $p = 0.53$; paired bootstrap tests). These results indicate that LM$^{\text{sink}}$ ripples are associated with more stable CA1 coactivity motifs.

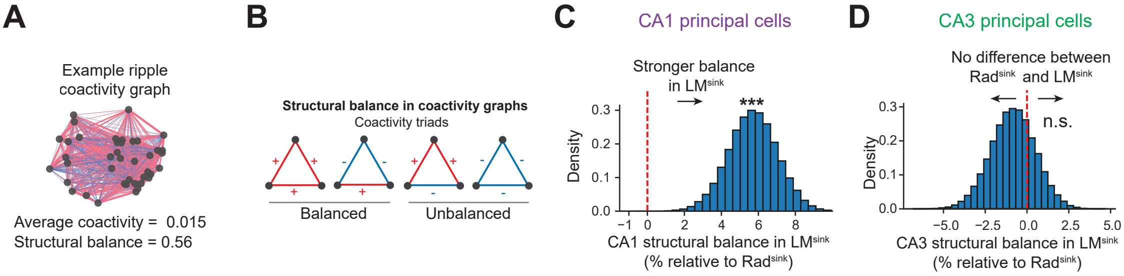

Figure 5.13: Structural balance of CA1 and CA3 coactivity motifs during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A) Example neuronal graph representing the ripple adjacency matrix shown in Figure 5.11, including its average coactivity and structural balance.
(B) Schematic illustration of balanced and unbalanced triads of coactivity used to assess structural balance in ripple-nested neuronal graphs.
(C,D) Mean difference in the structural balance of LM$^{\text{sink}}$ coactivity graphs relative to Rad$^{\text{sink}}$ graphs for CA1 (C) and CA3 (D) principal cells. The vertical red dashed line represents the mean structural balance in Rad$^{\text{sink}}$ ripples. (C) CA1 LM$^{\text{sink}}$ ripple coactivity graphs exhibit significantly higher structural balance than Rad$^{\text{sink}}$ ripple coactivity graphs ($p < 10^{-5}$; paired bootstrap tests; $n = 208$ sleep sessions). (D) In contrast, no significant difference in structural balance was observed between ripple types for CA3 coactivity graphs ($p = 0.53$, paired bootstrap test; $n = 171$ sleep sessions).
$^{\text{n.s.}}p > 0.05$, $^{***}p < 0.001$.

5.5 Population pattern similarities and differences in LM$^{\text{sink}}$ versus Rad$^{\text{sink}}$ ripples

In the previous section, I explored how Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples are associated with distinct coactivity patterns with different structural properties. Another way to examine the differences between these two ripple groups is to assess how individual neuronal responses and their organisation into coactive pairs influences the overall population activity. To explore this, I analysed the population vectors (PVs) associated with the ripple peak in CA1 and CA3 (Figure 5.14(A).

Similar to what was observed at the single-cell level, Rad$^{\text{sink}}$ ripples recruited significantly more CA1 principal cells at a given time compared to LM$^{\text{sink}}$ ripples [Figure 5.14(A,B); mean proportion of active CA1 principal cells (95% CI): Rad$^{\text{sink}}$, 41.4 (40.5 - 42.3) %; LM$^{\text{sink}}$, 28.0 (27.2 - 28.7) %; $p < 10^{-5}$, paired bootstrap test]. In other words, PVs corresponding to LM$^{\text{sink}}$ ripples were sparser (Figure 5.14(A,C); $p < 10^{-5}$, bootstrap test). CA3 PVs followed the same trend as CA1 PVs, being sparser in LM$^{\text{sink}}$ ripples and recruiting fewer principal cells (Figure 5.15; mean proportion of active CA3 principal cells (95% CI): Rad$^{\text{sink}}$, 24.3 (23.1 - 25.6) %; LM$^{\text{sink}}$, 17.7 (16.9 - 18.4); all $Ps < 10^{-5}$, paired bootstrap tests).

Placing these results in the context of previous findings, I found that LM$^{\text{sink}}$ ripples recruit fewer principal cells, forming smaller motifs, with CA1 motifs being structurally stable (Figures 5.14 and 5.13). Additionally, as shown in Section 5.3.2, transiently recruited CA1 and CA3 principal cells responded similarly across ripple types (Figure 5.8). In this sense, I hypothesised that the smaller motifs of active cells in LM$^{\text{sink}}$ ripples would be reappearing within the larger (denser) motifs of Rad$^{\text{sink}}$ ripples.

To test this, I measured the tendency for sets of neurons coactive in LM$^{\text{sink}}$ ripples to be contained within the coactive population observed in Rad$^{\text{sink}}$ ripples. I quantified this using an asymmetrical Jaccard coefficient, which evaluates, for each pair of ripples (m, n), the extent to which the set of active cells in ripple m is also active in ripple n (Figures 5.16 and 5.17(A); see Section 5.2.7). This

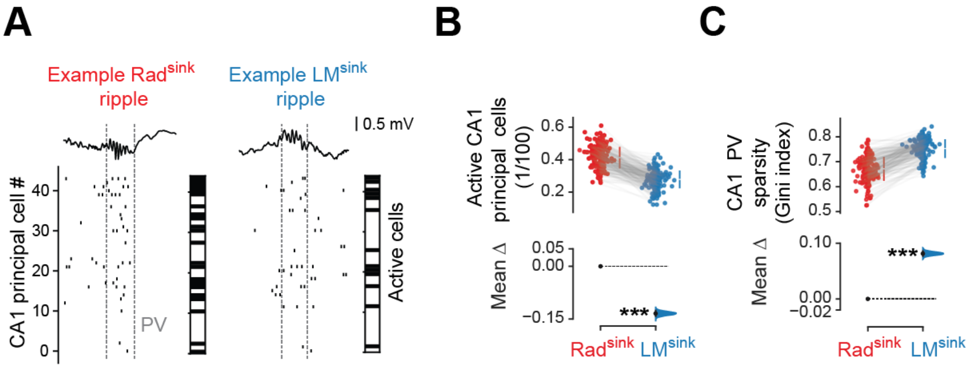

Figure 5.14: CA1 principal cells recruitment in Rad^{sink} versus LM^{sink} ripples. (A) Example *pyramidale* LFP trace and raster plot of CA1 principal cells for a Rad^{sink} ripple (left) and a LM^{sink} ripple (right), with corresponding binarised population vectors (PVs) representing cell recruitment (black squares: active cells; white squares: inactive cells).
(B) Estimation plot showing the effect size for the difference in the proportion of active CA1 principal cells within the population vectors nested in Rad^{sink} versus LM^{sink} ripples. Each point represents the mean proportion of active cells from a single sleep session.
(C) Same as (B), but showing the mean difference in sparsity (Gini index) of ripple population vectors.
For (B),(C): $n = 208$ sleep sessions. $^{***}p < 0.001$.

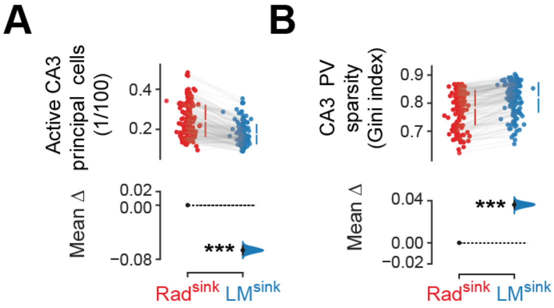

Figure 5.15: CA3 principal cells recruitment in Rad^{sink} versus LM^{sink} ripples. (A) Estimation plot showing the effect size for the difference in the proportion of active CA3 principal cells within the population vectors nested in Rad^{sink} versus LM^{sink} ripples. Each point represents the mean proportion of active cells from a single sleep session.
(B) Same as (A), but showing the mean difference in sparsity (Gini index) of ripple population vectors.
$n = 171$ sleep sessions. $^{***}p < 0.001$.

analysis allowed me to assess whether motifs of active neurons in LM^{sink} ripples also appeared within Rad^{sink} ripples at a single-ripple resolution.

To ensure that this metric was not simply a consequence of the difference in sparsity between LM^{sink} and Rad^{sink} PVs, I computed a control distribution by shuffling the identities of neurons jointly active in LM^{sink} ripples while preserving

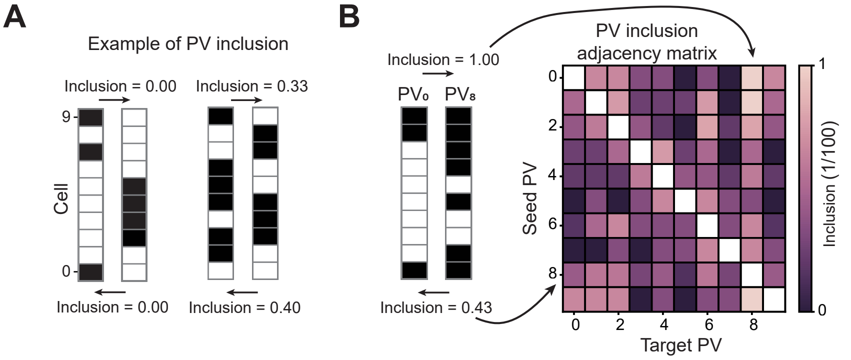

Figure 5.16: Schematic of active cell overlap measurement between population vectors.
(A,B) Inclusion analysis. Shown in (A) is a schematic illustrating the procedure for calculating the inclusion of active cells across two population vectors (PVs), where one serves as the seed vector and the other as the target vector. Each active cell in a PV is represented by a black square. Pairwise inclusion values are computed for each pair of PVs (m, n). Shown in (B) is an example adjacency matrix containing these inclusion values.

the structure of Rad$^{\text{sink}}$ PVs (see Section 5.2.8). This approach controlled for the possibility that a sparser population vector could overlap with a denser one by chance. In other words, this control allowed me to test whether the inclusion of LM$^{\text{sink}}$ motifs in Rad$^{\text{sink}}$ ripples depended on the specific identity of coactive neurons rather than arising from random sparsity effects.

CA1 coactivity motifs forming LM$^{\text{sink}}$ PVs were significantly more likely to appear within Rad$^{\text{sink}}$ PVs than expected by chance [Figure 5.17(B); mean CA1 inclusion LM→Rad (95% CI): z-scored against controls: 19.53 (18.37 - 20.70); $p = 5 \times 10^{-56}$, one-tailed t-test]. Similarly, CA3 coactivity motifs in LM$^{\text{sink}}$ ripples also appeared in Rad$^{\text{sink}}$ ripples, though less robustly than in CA1 (Figure 5.17(B)). This is expected given that CA3 responses are more diverse, as described in Sections 5.3.1 and 5.3.2 [mean CA3 inclusion LM→Rad (95% CI): z-scored against controls: 9.99 (9.32 - 10.67); $p = 7 \times 10^{-32}$, one-tailed t-test].

This result demonstrates that the smaller and stable motifs of principal cells, particularly within CA1, in LM$^{\text{sink}}$ ripples persist within Rad$^{\text{sink}}$ ripples, with additional cells recruited to form denser PVs. Importantly, the identity of the active neurons — that is, the specific coactivity motifs — matters, as demonstrated by the

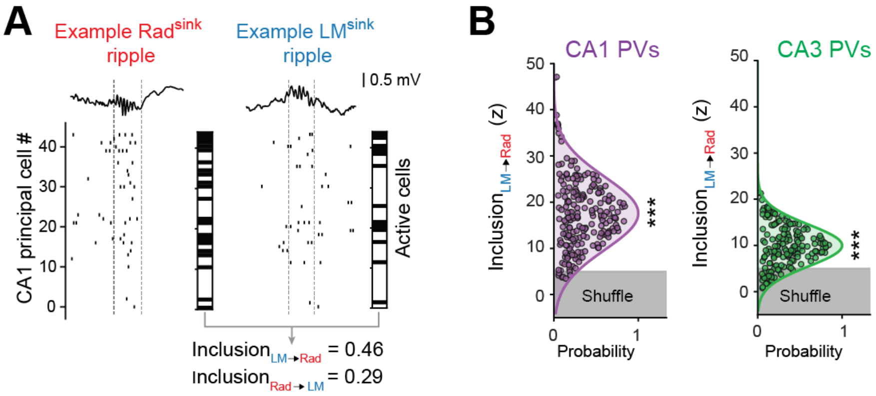

Figure 5.17: Overlap of active CA1 and CA3 principal cells between Rad^{sink} and LM^{sink} ripples.
(A) Same raw example *pyramidale* LFP trace and raster plot of CA1 principal cells for a Rad^{sink} ripple (left) and a LM^{sink} ripple (right) as shown in Figure 5.14(A), now with the addition of the inclusion measure. This measure quantifies the overlap of active cells between the two population vectors (PVs), capturing both directions (i.e. LM→Rad and Rad→LM).
(B) Distribution of the average inclusion values for motifs of active CA1 (left) and CA3 (right) principal cells in LM→Rad ripples. Each dot represents a sleep session. Inclusion values are z-scored relative to surrogate distributions generated by shuffling cell coactivity within LM^{sink} ripples. The gray shaded region represents ±5 standard deviations.
For (B): $n = 208$ sleep sessions for CA1 and $n = 171$ for CA3. $^{***}p < 0.001$.

control analysis. This suggests that the observed effect is not merely a consequence of differences in sparsity between the two ripple types. These findings suggest that LM^{sink} ripples host generic (lower-dimensional) population patterns with core motifs of coactivity, while Rad^{sink} ripples incorporate these motifs with additional cells to form larger, more complex (higher-dimensional) PVs (see Section 5.2.9).

Dimensionality, in this context, refers to the number of distinct population activity patterns expressed across multiple ripples. A lower-dimensional response indicates that neurons fire in highly coordinated and repetitive patterns across different events, while a higher-dimensional response reflects more independent activity across neurons, resulting in greater variability in population-level representations.

Consistently, LM^{sink} PVs (CA1 and CA3) exhibited lower intrinsic dimensionality (Thordsen & Schubert, 2022) [Figure 5.18; CA1 mean intrinsic dimensionality (95% CI): Rad^{sink}, 3.29 (3.24 - 3.35); LM^{sink}, 2.72 (2.63 - 2.81); CA3 mean intrinsic dimensionality (95% CI): Rad^{sink}, 2.22 (2.15 - 2.29); LM^{sink}, 1.83 (1.78 - 1.87); all

$Ps < 10^{-5}$, paired bootstrap tests]. Interestingly, the overall dimensionality was higher in CA1 PVs than in CA3. Thus, LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ population patterns differ not only in coactivity level but also in their structural organisation, with the lower-dimensional motifs featured in LM$^{\text{sink}}$ ripples reappearing within the higher-dimensional population vectors of Rad$^{\text{sink}}$ ripples.

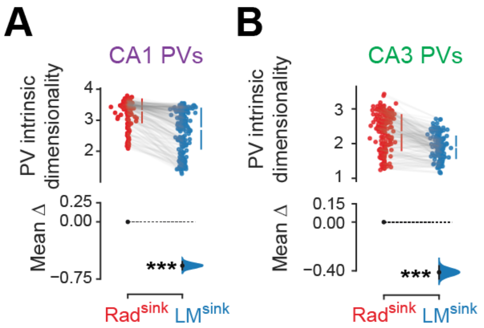

Figure 5.18: Dimensionality of CA1 and CA3 population vectors in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A,B) Estimation plots illustrating the effect size for the difference in intrinsic dimensionality of CA1 (A) and CA3 (B) population vectors during Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ ripples. Each point represents a single sleep session.
$n = 208$ sleep sessions for CA1 and $n = 171$ for CA3. $^{***}p < 0.001$.

5.6 Discussion

In this chapter, I have shown that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples recruit CA1 and CA3 neurons differently, from single-cell responses to the overall motifs shaping the population response.

I found that principal cells and interneurons in both CA1 and CA3 were recruited at different levels across the two ripple profiles. Rad$^{\text{sink}}$ ripples exhibited higher firing rates and denser recruitment of neurons compared to LM$^{\text{sink}}$ ripples. Furthermore, in CA3, beyond these quantitative differences in neuronal recruitment, I observed qualitatively distinct temporal firing responses between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. These findings further support the idea that these two ripple-associated CSD profiles emerge from distinct circuit dynamics in the hippocampus. The high recurrence in CA3 may enable strong excitatory currents converging to *stratum* *radiatum* during Rad$^{\text{sink}}$ ripples, activating a larger number of neurons at higher firing rates and generating higher-frequency ripples (see Section 3.3.2). In contrast, during LM$^{\text{sink}}$ ripples, CA3→CA1 inputs would be significantly weaker, either intrinsically or attenuated locally by interneurons within CA3 or CA1, leading to stronger differential current in *stratum lacunosum-moleculare*. These distal currents still recruit a notable population of neurons but at lower firing rates than in Rad$^{\text{sink}}$ ripples. This reduced recruitment may stem from weaker input currents from CA3 during LM$^{\text{sink}}$ ripples or from additional attenuation due to the longer distance these currents must travel to generate somatic spikes in CA1, compared to the shorter path from *stratum radiatum*. Regardless, these findings establish that within the continuum of CSD profiles, CA1 and CA3 principal cells are differentially modulated, supporting the hypothesis that ripple-associated recruitment is dynamically regulated.

Moreover, interneurons in CA1 and CA3 also responded differently to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, suggesting that the local circuits associated with these two ripple types are entrained differently. An interesting finding was that the ratio of interneuron activity to the mean principal cell firing rate was higher in CA3 during LM$^{\text{sink}}$ ripples but comparable across ripple types within CA1 (Figure 5.5). This suggests that although all neurons fire at higher rates during Rad$^{\text{sink}}$ ripples than during LM$^{\text{sink}}$ ripples, CA3 interneurons are comparatively more active during LM$^{\text{sink}}$ events relative to principal cells. In other words, CA3 shows a higher interneuron-to-pyramidal activity ratio in LM$^{\text{sink}}$ ripples, even after accounting for the overall firing-rate difference. In contrast, this ratio remained stable in CA1. One possible interpretation is that LM$^{\text{sink}}$ ripples might be associated with a higher level of inhibition within CA3, which could make CA1 more sensitive to other inputs arriving at *stratum lacunosum-moleculare*. This inhibitory effect in CA3 may originate from the region sending input currents to *stratum lacunosum-moleculare* during LM$^{\text{sink}}$ ripples, possibly the medial entorhinal cortex (MEC), which projects to this CA1 layer (Valero & de la Prida, 2018). A possible explanation for the increased inhibition within CA3 could therefore be that MEC exerts a dual role, sending currents to *stratum lacunosum-moleculare* to generate ripples via MEC3 while simultaneously, through MEC2 (Valero & de la Prida, 2018), targeting interneurons in CA3 to promote higher inhibition during LM$^{\text{sink}}$ ripples. By contrast, during Rad$^{\text{sink}}$ ripples, this inhibitory influence appears weaker, allowing strong excitatory drive from CA3 to propagate through the Schaffer collaterals to *stratum radiatum* and trigger Rad$^{\text{sink}}$ ripples — effectively forcing CA1 to 'listen' to CA3.

Another point of discussion arises from the response of CA3 during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. In particular, while Rad$^{\text{sink}}$ ripples exhibit a single peak response, LM$^{\text{sink}}$ ripples display a bimodal response. This suggests that the circuits associated with these ripples are entrained differently. Additionally, CA3 does not seem to be as active on average during a subset of LM$^{\text{sink}}$ ripples. Notably, CA3 neurons show increased activity approximately 100 ms before the ripple peak (Figure 5.3), coinciding with a current sink in *stratum radiatum* observed in the silicon probe dataset (Figures 3.1 - 3.3). This provides further validation of the accuracy of the LDA ripple-type classifier. Why is CA3 active before the peak of LM$^{\text{sink}}$ ripples? This may ensure that CA3, either intrinsically or through feedforward excitation from the dentate gyrus and MEC2, is biased to activate earlier, including interneurons, thereby reducing its excitatory impact at the ripple peak when strong currents drive CA1 principal cells in *stratum lacunosum-moleculare*. Future studies could investigate the role of CA3 and its upstream inputs, such as the dentate gyrus and MEC2, during this pre-ripple peak to assess how it influences subsequent ripple responses in CA1.

The differences between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples extend beyond modulation of spiking levels to distinct structural properties of population coactivity patterns. Specifically, I found that different coactivity motifs are recruited across the two ripple profiles (see Section 5.3.2), with more pronounced differences within CA3. Population vector analysis combined with topological analysis revealed that CA1 principal cells contributing to the sparser LM$^{\text{sink}}$ ripples formed core motifs of high and robust coactivity, which could reappear in Rad$^{\text{sink}}$ ripples but alongside additional cells. This suggests that although CA1 cell motifs are highly stable within LM$^{\text{sink}}$ ripples, this does not mean that LM$^{\text{sink}}$ ripples never change their coactivity patterns. Rather, it indicates that their structural coactivity patterns are more resistant to change and may require more time or broader network reorganisation to update. In other words, one can envision the small motifs of active cells in LM$^{\text{sink}}$ ripples as being confined to a fixed point — or stable state — while Rad$^{\text{sink}}$ ripples, through the recruitment of additional cells and the formation of denser motifs, enable transitions away from this fixed point, giving rise to more complex PVs. In this sense, I also showed that LM$^{\text{sink}}$ ripples nest lower-dimensional patterns, while Rad$^{\text{sink}}$ ripples nest composite, higher-dimensional population patterns. In short, these findings suggest that, despite potentially originating from parallel hippocampal circuits, CA1 principal cell responses to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples can influence each other.

Importantly, the difference in sparsity alone cannot fully account for the stability differences in coactivity motifs and overall coactivity levels. While both CA1 and CA3 population vectors were sparser during LM$^{\text{sink}}$ compared to Rad$^{\text{sink}}$ ripples, the CA3 coactivity analysis revealed distinct results from CA1, showing higher coactivity levels in Rad$^{\text{sink}}$ ripples but no significant difference in structural balance (see Figures 5.12 and 5.13). This suggests that, beyond differences in sparsity, intrinsic structural properties differentiate the network-level organisation of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples in CA3 compared to those in CA1. Additionally, the overall level of inclusion and dimensionality were significantly higher for CA1 PVs than in CA3 (see Section 5.5). This could be influenced by the greater number of active cells in CA1 compared to CA3 during ripples (Figures 5.14 and 5.15). Beyond this, the increased complexity and dimensionality observed in CA1 may also reflect the integration of CA3 assemblies with inputs from other upstream regions.

As outlined in Chapter 1, since CA1 assemblies are known to reactivate recent waking experiences, the differences described in this chapter raise important questions about how Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples influence the type of reactivation they mediate, particularly given their differences in dimensionality. In the next chapter, I explore this question by examining the content of these two ripple types.

... la città non dice il suo passato, lo contiene come le linee d'una mano

...the city does not speak its past, it contains it like the lines of a hand

— Italo Calvino's *Le città invisibili*

6

Reactivation of recent and prior motifs during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

Contents

6.1 Introduction

Hippocampal ripples have emerged as critical network events that facilitate the offline processing of neuronal population activity (Buzsáki, 2015). The high synchrony of CA1 principal cells during ripples plays a key role in the reactivation of coactivity patterns from waking experiences, enabling further offline processing during sleep and rest periods (Wilson & McNaughton, 1994). That is, changes in hippocampal network connectivity induced by exploration are also reflected in subsequent sleep. This process is crucial for memory consolidation, as disruption of post-exploration ripples has been associated with impaired memory performance in future tasks (Chang et al., 2025; Ego-Stengel & Wilson, 2010; Girardeau et al., 2009).

While ripple reactivation is a global phenomenon observed across CA1 principal cells, recent studies have highlighted significant heterogeneity between deep and superficial CA1 pyramidal neurons (Esparza et al., 2025; Geiller et al., 2017; Navas-Olive et al., 2020; Slomianka et al., 2011; Soltesz & Losonczy, 2018; Valero & de la Prida, 2018). During exploration, deep CA1 cells fire at higher rates and exhibit more rigid activation patterns, maintaining similar spatial tuning across different environments (Berndt et al., 2023; Gava et al., 2021, 2024; Harvey et al., 2023; Lopes-dos-Santos et al., 2023; Valero et al., 2015). In contrast, superficial cells show lower firing rates but more flexible, context-dependent recruitment, with activity patterns that vary across experiences and reflect environment-specific features (Berndt et al., 2023; Gava et al., 2021, 2024; Grosmark & Buzsáki, 2016; Harvey et al., 2023; Lopes-dos-Santos et al., 2023; Valero et al., 2015).

These functional differences extend to ripples, during which superficial cells fire at higher rates and play a dominant role in the reactivation of cell assemblies and the expression of offline replay (Harvey et al., 2023). This distinction likely arises from differential inputs to these two CA1 sublayers: deep CA1 cells receive stronger input from the medial entorhinal cortex layer III (MEC3), which projects to *stratum lacunosum-moleculare*, while superficial CA1 cells receive stronger input from CA3, which projects to *stratum radiatum* (Valero & de la Prida, 2018).

However, it remains unclear how this organisation influences ripple activity at the level of individual ripples. If deep and superficial cells exhibit distinct recruitment and functional roles, does this differentiation extend to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples? Furthermore, how does the diversity in laminar currents shaping ripples relate to these functional distinctions?

In Chapters 3 and 5, I showed that ripples are highly variable and can be categorised into two distinct profiles. Canonical Rad$^{\text{sink}}$ ripples are characterised by a sink in *radiatum*, elevated CA3 activity, and the formation of high-dimensional CA1 population motifs. In contrast, a subset of ripples — LM$^{\text{sink}}$ ripples — exhibit a prominent sink in *lacunosum-moleculare*, reduced CA3 activity, and the emergence of lower-dimensional, sparse CA1 motifs that exhibit high structural stability. Given the differences in inputs to deep and superficial CA1 cells and their respective contributions to reactivation, one would expect superficial cells to be more active in Rad$^{\text{sink}}$ than in LM$^{\text{sink}}$ ripples due to the enhanced CA3 activity associated with Rad$^{\text{sink}}$ ripples. This raises the question of whether the enhanced flexibility of superficial cells primarily contributes to the higher-dimensional population vectors observed in Rad$^{\text{sink}}$ ripples, whereas the rigidity of deep cells underlies the smaller, more stable motifs seen during LM$^{\text{sink}}$ ripples. Furthermore, how do these differences in neuronal activity affect the reactivation content? Do highly dimensional Rad$^{\text{sink}}$ ripples convey different information than the more rigid LM$^{\text{sink}}$ ripples?

To address these questions, in this chapter, I investigate how ripple diversity relates to neuronal recruitment and reactivation. I first examine within-ripple differences in the activity of deep and superficial CA1 principal cells, assessing how each sublayer responds to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Next, I relate these differences to the structural features of population motifs, as explored in Section 5.5. I then examine how these differences in neuronal recruitment influence the reactivation content of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Finally, I assess the temporal dynamics of reactivation within Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples over hour-long sleep sessions.

6.2 Methods

6.2.1 Classification of CA1 principal cells into deep and superficial

CA1 principal cells were classified as either deep or superficial following the methodology described in Lopes-dos-Santos et al., 2023. Briefly, LFP features were extracted from the tetrode on which each principal cell was recorded. These features were then projected onto a linearised trajectory, estimating the cell's depth within the pyramidal layer by determining whether the associated tetrode was closer to the *stratum radiatum* or the *stratum oriens*.

As illustrated in Figure 6.1(A), the depth distribution of CA1 principal cells along this trajectory exhibited a bimodal pattern. Based on this distribution, a threshold depth value of 6 was used to differentiate superficial and deep cells, yielding a total of 1,353 deep and 843 superficial cells. For analyses incorporating the criteria described in Section 2.1.6, these numbers were refined to 1,100 deep and 480 superficial cells.

6.2.2 Change in firing rate during ripples relative to baseline

To estimate how deep and superficial CA1 principal cells responded to Rad$^{\text{sink}}$ vs. LM$^{\text{sink}}$ ripples, I assessed how much each cell increased its rate during ripples from a baseline window (Figure 6.2(B); $n = 1,100$ deep and 480 superficial cells). To do so, I compared the change in firing rate of deep and superficial principal cells relative to a pre-ripple baseline. This analysis allowed for the comparison of cells with different firing rates during ripples (i.e., deep versus superficial CA1 principal cells).

To quantify this measure, I computed the mean instantaneous firing activity of each cell within a 50-ms window centred around the ripple peak ($\text{rate}_{\text{ripple}}$). To estimate the baseline firing rate ($\text{rate}_{\text{baseline}}$), I calculated the mean firing activity within a time window from -200 ms to -100 ms relative to the ripple peak. Then, independently for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, I quantified the firing rate increase as:

$$\Delta\text{rate} = \log_{10}\left(\frac{\text{rate}_{\text{ripple}}}{\text{rate}_{\text{baseline}}}\right) \tag{6.1}$$

where positive values of Δrate indicate an increase in firing rate relative to the pre-ripple baseline, while negative values indicate a decrease. To prevent contamination of the baseline rate, only isolated ripples were included in this analysis (see Section 2.2.2).

6.2.3 Additional cells active in Rad$^{\text{sink}}$ but not in LM$^{\text{sink}}$ ripples

To determine whether superficial cells were more likely than deep cells to be active preferentially in Rad$^{\text{sink}}$ ripples but not in LM$^{\text{sink}}$ ripples, I examined cells that exhibited selective firing during Rad$^{\text{sink}}$ ripples in Figure 6.2(C). First, I identified the cells active in each ripple using the approach described in Section 5.2.7. Then, for each cell, I computed the conditional probability of it being active in a Rad$^{\text{sink}}$ ripple given that it was inactive in a LM$^{\text{sink}}$ ripple.

To control for each cell's intrinsic firing probability across ripples, I estimated a chance level for each cell using a shuffling procedure. This procedure independently shuffled coactive neurons in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples while preserving both the total number of ripples in which each cell was active and the sparsity of each ripple (see Section 2.4.5). The shuffling process was repeated 200 times, and the resulting chance levels were used to z-score the observed conditional probabilities (Figure 6.2(C); $n = 1,100$ deep and 480 superficial cells).

6.2.4 Offline reactivation of waking coactivity patterns

To analyse offline reactivation of CA1 or CA3 principal cell spiking patterns during sleep/rest (Figure 6.3), I first trained a linear regression model using awake theta-cycle activity from the exploration session of that recording day. Specifically, the model was trained to predict the firing rate of a target cell based on the activity of four other cells, ensuring that none of the predictor cells were recorded on the same tetrode as the target cell. This criterion minimised biases in correlations that could arise from shared global tetrode activity.

Model accuracy was assessed using Pearson correlation between the actual firing rate of the target cell and the firing rate predicted by the model. To ensure robustness, I performed 100 bootstrap iterations, randomly selecting four predictor cells from the pool of cells recorded on the same day for each iteration. I then applied these cross-validated models to the z-scored ripple-nested firing activity of the target cell during pre-exploration and post-exploration sleep/rest sessions, using a 50-ms window centred around the ripple peak. A schematic illustrating this method — specifically when applied to deep and superficial CA1 cells separately (see Section 6.2.4), rather than to the general CA1 or CA3 population — is shown in Figure 6.4(A).

For each cell, the mean accuracy across the 100 cross-validated models was computed separately for pre-exploration and post-exploration sleep, yielding two model accuracy values. Reactivation was defined by comparing the overall accuracy across all cells in post-exploration sleep to the corresponding accuracy in pre-exploration sleep. Significant reactivation was identified if the post-exploration accuracy exceeded the pre-sleep accuracy (1-tailed paired bootstrap tests; Figure 6.4(B,C); $n = 1,580$ CA1 and $n = 866$ CA3 principal cells). This analysis was conducted independently for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples during both pre- and post-sleep periods.

Deep and superficial cell reactivation

In Figure 6.4, I computed the offline reactivation of awake spiking patterns for deep and superficial cells. This analysis followed the same methodology described in Section 6.2.4, with an additional constraint: predictor cells and the predicted cell were required to belong to the same CA1 pyramidal sublayer. Specifically, a set of four superficial cells was used to predict the activity of another superficial cell, and a set of four deep cells was used to predict the activity of another deep cell, as illustrated in Figure 6.4(A).

Together with the existing requirement that predictor cells and the predicted cell be recorded on different tetrodes (see Section 6.2.4), this sublayer-specific constraint reduced the pool of eligible cells to $n = 1,097$ deep and $n = 478$ superficial CA1 principal cells.

6.2.5 Coactivity pattern stability in LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ ripples

To assess whether the coactivity patterns nested in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples changed from pre-exploration to post-exploration sleep/rest, I first extracted ripple-nested population vectors (PVs) containing the instantaneous firing activity of CA1 principal cells within a 50-ms window centred around the ripple peak. This was done separately for pre-exploration and post-exploration sleep/rest sessions.

Using the z-scored PVs from pre-sleep, I trained a generalised linear model (GLM) to predict the firing response of each CA1 principal cell based on the activity of the remaining population during each ripple. This process was conducted independently for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, yielding two models per cell (one for LM$^{\text{sink}}$ and one for Rad$^{\text{sink}}$ ripples). Each model was cross-validated 20 times (80% training, 20% testing), and accuracy was quantified as the mean correlation between the predicted and actual firing activity in the test set for each iteration: $r_{\text{pre},i}$.

To estimate chance-level performance, I applied a shuffling procedure in which cell IDs within each PV were randomly permuted 500 times before training, thereby disrupting cell-specific correlations while preserving overall activity rates within each PV. To measure changes in coactivity from pre- to post-exploration sleep, I applied each cell's cross-validated model to the (z-scored) PVs from post-sleep and measured its accuracy ($r_{\text{post},i}$).

For both Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, coactivity stability for cell i was defined as:

$$\text{Stability}_i = 1 - (r_{\text{pre},i} - r_{\text{post},i}) \tag{6.2}$$

where $r_{\text{pre},i}$ and $r_{\text{post},i}$ are the model accuracies in pre- and post-sleep, normalised with respect to the chance level described above (Figure 6.5; $n = 140$ pre- and post-sleep pairs).

A stability score near 1 indicates minimal change in coactivity patterns between pre- and post-sleep, while a low stability score suggests significant alterations in coactivity across sessions.

6.2.6 Reactivation of prior and recent coactivity motifs

To determine whether Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples preferentially reactivated (i) recently formed neural coactivity motifs (i.e., neurons coactive during the open-field exploration session but not in pre-exploration sleep/rest) and/or (ii) prior neural motifs (i.e., coactivity patterns already present in pre-exploration sleep/rest), I first constructed coactivity graphs as described in Section 2.4.3.

For the pre-sleep coactivity graph, I sampled an equal number of Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples to establish a balanced coactivity baseline. Similarly, I constructed a coactivity graph for the awake session, using spiking activity within waking theta cycles instead of ripple events. Thus, for each recording day, I obtained two coactivity graphs: A_{pre} (pre-sleep) and A_θ (wakefulness).

To isolate coactivity motifs that emerged during wakefulness but were absent in pre-sleep, I modelled the changes in coactivity from from pre-sleep to wakefulness as a linear transformation:

$$A_\theta = W_{\text{recent}} \times A_{\text{pre}}$$

where the transformation matrix W_{recent} was computed as:

$$W_{\text{recent}} = A_\theta \times A_{\text{pre}}^{-1} \tag{6.3}$$

Similarly, to identify coactivity motifs present in pre-sleep but absent in the subsequent wakefulness, I modelled the transformation as:

$$A_{\text{pre}} = W_{\text{prior}} \times A_\theta$$

yielding the transformation matrix:

$$W_{\text{prior}} = A_{\text{pre}} \times A_\theta^{-1} \tag{6.4}$$

To assess the reactivation of prior and recent motifs in post-exploration sleep, I extracted post-sleep PVs, which contained the instantaneous firing activity of CA1 principal cells within a 50-ms window centred around the ripple peak. By applying the transformation matrices W_{recent} and W_{prior} to the post-sleep ripple PVs, I projected the original coactivity space onto subspaces representing either recent or prior motifs.

For each PV, I computed the dot product between the transformed PV and the original coactivity space, yielding a reactivation score for each ripple. A high reactivation score indicated strong alignment of the original coactivity space with either the recent or prior transformation space, while a score near zero reflected orthogonality to the transformation.

Reactivation scores were computed separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples in post-sleep, and the mean score within each ripple group was defined as the reactivation strength for that ripple type in the given sleep session.

Simulation of prior and recent coactivity motifs

In Figure 6.6, I validated the method described in Section 6.2.6 using *in silico* simulations. To achieve this, I first simulated a neuron ensemble's pre-sleep activity matrix containing two predefined patterns of coactive neurons (prior motifs). Next, I simulated an awake exploration phase by introducing additional coactivity patterns (recent motifs). Specifically, one pre-sleep motif remained unchanged, the second motif was modified by the integration of an additional neuron, and a third motif emerged during the simulation as a newly formed pattern.

Using these simulated data, I constructed two sets of post-sleep ripples. The first set (referred to as *rigid ripples*) contained only the neuronal motifs observed in pre-sleep. The second set (referred to as *plastic ripples*) included both these pre-existing motifs and those newly formed during awake exploration. I then applied the method described in Section 6.2.6 to this "ground truth" simulated data to assess its ability to accurately extract prior and recent motifs from the two post-sleep matrices. The strength of these motifs was computed for each post-sleep scenario to evaluate the method's performance.

6.2.7 Contribution of deep and superficial cells to reactivation of coactivity patterns

To assess whether deep and superficial cells differentially contributed to prior and recent motifs, I quantified their respective contributions to the reactivation of these motifs during post-exploration sleep (Figure 6.8; see Section 6.2.6). Specifically, I measured the strength of prior motifs during LM$^{\text{sink}}$ ripples and recent motifs during Rad$^{\text{sink}}$ ripples within groups of five deep or five superficial cells. To ensure robustness, I performed 500 permutations by randomly selecting five deep or five superficial cells from the pool of principal cells recorded that day. The mean across these permutations was used to define the reactivation strength for each sublayer in each ripple. Finally, for each pre- and post-sleep pair, I compared these reactivation values (prior and recent motifs) between deep and superficial cells.

To maintain a fair comparison and account for potential imbalances in the number of recorded cells (or the absence of cells from a specific sublayer) on a given recording day, for this analysis I required at least five deep and five superficial principal cells to be present (Figure 6.8; $n = 43$ pre- and post-sleep pairs from 6 mice).

6.2.8 Changes in reactivation over an hour-long timescale

To analyse how the reactivation of prior and recent motifs changed throughout sleep/rest, I quantified changes in the reactivation of prior and recent coactivity motifs over an hour-long timescale (Figures 6.9 and 6.11; see Section 6.2.6). For this analysis, I included only pre- and post-sleep pairs where the post-exploration sleep session lasted at least one hour ($n = 118$ pre- and post-sleep pairs from 13 mice).

For each session, I calculated the mean reactivation of prior and recent motifs during Rad$^{\text{sink}}$ or LM$^{\text{sink}}$ ripples, dividing the first hour of post-exploration sleep into 10-minute bins. To ensure comparability of reactivation values across days with different numbers of recorded cells, I normalised the reactivation scores by subtracting the global mean reactivation for each motif (i.e., prior or recent), regardless of ripple type (Rad$^{\text{sink}}$ or LM$^{\text{sink}}$). This allowed me to estimate the change from the overall mean across the sleep session. To determine whether reactivation significantly changed over time, I applied an exponential fit model (see Section *Exponential fit of reactivation over time*).

Exponential fit of reactivation over time

To assess whether reactivation changed exponentially over time (Figure 6.9), I modelled the temporal dynamics using an exponential decay function. First, data from all sleep/rest sessions were concatenated into two vectors: one representing the time of each bin and the other representing the corresponding reactivation values for prior motifs and recent motifs, either from Rad$^{\text{sink}}$ or LM$^{\text{sink}}$ ripples. Changes in reactivation over time were then modelled as:

$$\text{reactivation}(t) = a \cdot e^{-\frac{t}{\tau}} + c \tag{6.5}$$

where a represents the amplitude (scaling factor), τ is the time constant (decay parameter), and c is the offset or steady-state value. Note that the negative sign in a determines the direction of change: a positive a results in a negative leading to exponential increase towards c as $t \to \infty$, whereas a negative a results in exponential decrease to c as $t \to \infty$.

The function was fitted to the concatenated session data using the `optimize` submodule from the SciPy Python package. A hyperparameter search was conducted to determine the best initial estimates for a, τ, and c, optimising the fit between the model and the data. The goodness of fit was quantified using Pearson correlation (r) between the observed reactivation values and the model's predictions.

To estimate the mean fit score (r) and time constant (τ), I performed cross-validation by fitting the exponential function to bootstrapped subsets containing 80% of the pre- and post-sleep pairs. This procedure was repeated 10,000 times to generate distributions of r and τ. To assess the statistical significance of the fit, I conducted an additional 10,000 iterations in which the time values were shuffled, and the fitness score was computed for each shuffled dataset. The p-value was determined as the proportion of shuffled models achieving a fitness score equal to or greater than the actual score. A fit was considered significant if $p < 0.05$.

To minimise potential biases related to changes in sleep quality at the beginning of the session, the first 10 minutes of each sleep session were excluded from the analysis.

6.2.9 Reactivation over time controlling for ripple occurrence frequency and ripple population sparsity

To examine whether the observed gradual changes in reactivation over time during LM$^{\text{sink}}$ ripples (Figure 6.9) were genuinely time-dependent or instead driven by changes in ripple occurrence frequency (i.e., the number of ripples per time bin) or ripple population sparsity (i.e., the number of active neurons per ripple) (Figures 6.11, 6.10), I used a linear regression model.

Briefly, I computed the time-binned reactivation values of recently expressed motifs separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples for each session and then concatenated across sessions. A linear regression model was then trained to predict reactivation in each time bin (Figure 6.9(A); see Section 6.2.8). The predictors included ripple occurrence time (i.e., the time of the bin), ripple occurrence frequency (mean number of ripples per minute in each bin), and ripple population sparsity (mean proportion of active cells per ripple within each bin).

Given the non-linear relationship between time and reactivation (Figure 6.9), a $\log_{10}$ transformation was applied to ripple occurrence times to linearise this relationship. To minimise potential biases related to sleep quality, the first 10 minutes of each session were excluded from the analysis. The model was trained and cross-validated 20 times, with an 80% training and 20% testing split. Model accuracy was assessed as the Pearson correlation between the observed reactivation values and the model's predictions.

To establish a chance-level baseline, the reactivation values were shuffled 25,000 times, and the model's true accuracy was compared to these shuffled predictions (Figure 6.9(B)). Additionally, to evaluate the contribution of each predictor, I trained models in which individual features were shuffled 25,000 times, effectively removing the specific information carried by the shuffled feature.

The contribution of a given feature was quantified as the difference between the accuracy of the original model and that of the corresponding feature-shuffled model (Figure 6.9(C)). A high accuracy gain for a feature indicated strong predictive power, demonstrating that the feature played a significant role in predicting changes in reactivation. The significance of each feature was assessed by computing p-values as the proportion of shuffles with a gain less than zero, indicating that the feature significantly contributed to the model.

6.3 CA1 sublayer dynamics during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

6.3.1 Recruitment of CA1 deep and superficial cells during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

As mentioned in the introduction to this chapter, deep and superficial CA1 pyramidal cells — defined by the location of their somata closer to *stratum oriens* and *stratum radiatum*, respectively — have recently been shown to exhibit distinct functional correlates, shaped by differential input from MEC3 and CA3 (Berndt et al., 2023; Valero & de la Prida, 2018). Given that Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples are associated with different levels of CA3 activity (Chapters 3 and 5), it is plausible that deep and superficial CA1 cells are differentially recruited during the two ripple types.

In light of this, I hypothesised that superficial CA1 cells — receiving stronger input from CA3 — would be more actively recruited during Rad$^{\text{sink}}$ ripples, which are associated with heightened CA3 activity. In contrast, during LM$^{\text{sink}}$ ripples, where CA3 activity is significantly weaker, this recruitment of superficial cells would be reduced.

To test this, I examined whether LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ ripples differentially engaged deep and superficial CA1 cells recorded along the radial axis of the pyramidal layer (Figure 6.1(A)). First, I classified CA1 principal cells into deep ($n = 1,353$) and superficial ($n = 843$) categories (see Section 6.2.1). Consistent with previous reports (Harvey et al., 2023; Lopes-dos-Santos et al., 2023), classified deep and superficial cells exhibited differences in their mean firing rates during theta oscillations and ripples, as well as in their preferred theta phase (Figure 6.1(B-D)).

I then quantified their engagement in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples separately by computing the change in their firing rates relative to a pre-ripple baseline (see Section 6.1). This normalisation accounted for baseline firing rate differences between deep and superficial cells (Figure 6.1(C)), allowing for direct comparison.

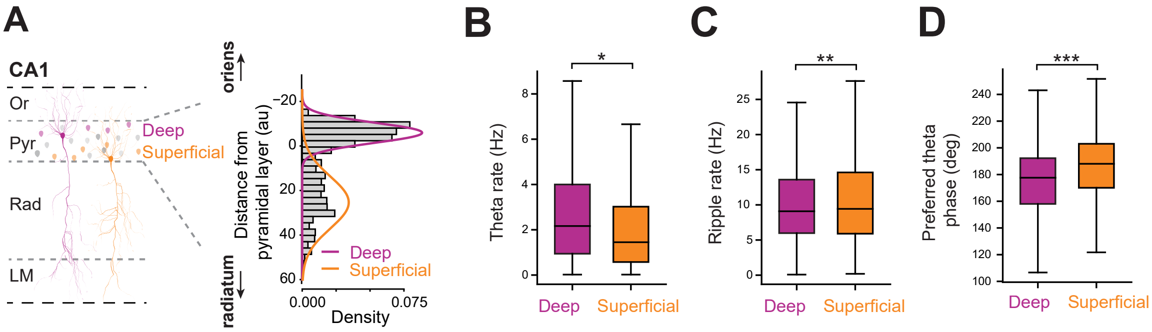

Figure 6.1: Firing properties of CA1 deep and superficial principal cells.
(A) Distribution of CA1 principal cells recorded in the deep (purple) and superficial (orange) pyramidal sublayers.
(B–D) Boxplots reporting for deep and superficial cells the mean firing rate during theta oscillations (B), across all ripples irrespective of ripple type (C), and the preferred theta phase (D). Deep cells fired at higher rates than superficial cells during theta oscillations [mean rate (95% CI): deep, 3.02 (2.82–3.23) Hz; superficial, 2.60 (2.30–2.93) Hz; $p = 0.017$, bootstrap test]. Conversely, superficial cells exhibited higher firing rates during ripples [mean rate (95% CI): deep, 10.60 (10.22–10.98) Hz; superficial, 11.55 (10.94–12.19) Hz; $p = 0.008$, bootstrap test]. Additionally, superficial cells fired later in the theta cycle than deep cells [mean preferred theta phase (95% CI): deep, 172.08 (169.84–174.30) degrees; superficial, 182.18 (179.34–184.96) degrees; $p < 10^{-5}$, bootstrap test].
For (A)-(D): $n = 1,353$ deep and $n = 843$ superficial cells. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

The analysis revealed that during Rad$^{\text{sink}}$ ripples, deep and superficial cells exhibited similar increases in firing rates from baseline (Figure 6.2(A,B)). In contrast, during LM$^{\text{sink}}$ ripples, deep cells showed a significantly greater firing rate increase than superficial cells, indicating preferential recruitment of deep cells. This indicated that superficial cells were preferentially recruited in Rad$^{\text{sink}}$ ripples when compared to their deep counterparts.

To further validate this finding, I conducted a second analysis. In Section 5.5, I demonstrated that LM$^{\text{sink}}$ ripples contain small coactivity motifs that are also present in Rad$^{\text{sink}}$ ripples, but with additional recruited cells. Thus, I assessed whether these 'additional cells' in Rad$^{\text{sink}}$ ripples were predominantly superficial. To do so, I computed the conditional probability of each CA1 principal cell being active in a Rad$^{\text{sink}}$ ripple given that it was inactive in a LM$^{\text{sink}}$ ripple (see Section 6.2.3). This probability was normalised against a shuffled distribution controlling for each cell's intrinsic firing probability across ripple events. The results confirmed that superficial cells were preferentially recruited in Rad$^{\text{sink}}$ ripples but not in LM$^{\text{sink}}$ ripples (Figure 6.2(A,C)).

These findings, together with those discussed in Chapter 5, support the hypothesis that LM$^{\text{sink}}$ ripples engage a lower-dimensional, stable motif dominated by deep CA1 cells, whereas Rad$^{\text{sink}}$ ripples integrate these stable motifs with additional superficial cells, forming a denser, higher-dimensional representation. Given the differential content and contributions to reactivation of deep and superficial cells (Gava et al., 2024; Harvey et al., 2023), this distinction may reflect a functional difference in how Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples support offline reactivation — a question I explore in the next section.

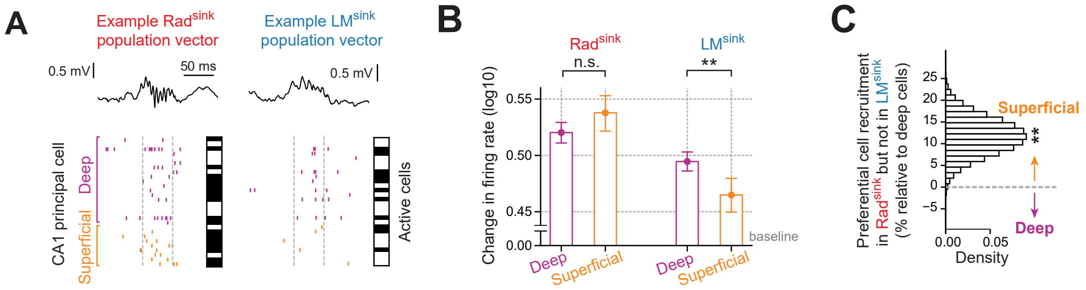

Figure 6.2: Firing and recruitment of CA1 deep and superficial cells in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A) Example raw LFP traces and raster plots showing the spiking activity (colour-coded) of deep and superficial cells during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, along with their corresponding binary population vectors indicating recruited cells.
(B) Change in firing response of deep and superficial cells. In Rad$^{\text{sink}}$ ripples, deep and superficial principal cells increase their firing rates similarly relative to their pre-ripple baseline. In contrast, during LM$^{\text{sink}}$ ripples, deep cells exhibit a significantly greater firing increase compared to superficial cells (mean ± 95% CI; Rad$^{\text{sink}}$: $p = 0.10$; LM$^{\text{sink}}$: $p = 3.2 \times 10^{-3}$; bootstrap tests).
(C) Preferential recruitment of superficial cells in Rad$^{\text{sink}}$ ripples. The conditional probability of cells being selectively active in Rad$^{\text{sink}}$ ripples (given their inactivity in LM$^{\text{sink}}$ ripples) was computed, normalised relative to chance by controlling for intrinsic cell firing probability in ripples, and expressed as a percentage relative to the mean of deep cells. Superficial cells were more likely than deep cells to be active in Rad$^{\text{sink}}$ but not in LM$^{\text{sink}}$ ripples ($p = 2.7 \times 10^{-3}$; bootstrap test). The dashed line represents chance level.
For (B),(C): $n = 1,100$ deep and $n = 480$ superficial cells. $^{\text{n.s.}}p > 0.05$, $^{**}p < 0.01$.

6.3.2 Reactivation of CA1 deep and superficial cells in Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples

Hippocampal ripples are well known to facilitate the offline reactivation of cell assemblies formed during wakefulness, a process crucial for memory consolidation (Chang et al., 2025; Ego-Stengel & Wilson, 2010; Girardeau et al., 2009; Roux et al., 2017; Wilson & McNaughton, 1994).

Thus, I investigated whether both Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples reactivated coactivity motifs from previous wakefulness exploration. To do so, I trained generalised linear models to predict the spiking activity of each principal cell based on the activity of its peers during theta oscillations in exploratory behaviour, using a leave-one-out approach (see Section 6.2.4). These models were then applied to ripple-nested activity during post-exploration sleep/rest, controlling for pre-exploration sleep/rest.

As expected from previous findings, CA1 principal cells were significantly reactivated during both Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples in post-exploration sleep/rest (Figure 6.3(A); $p < 10^{-5}$; one-tailed bootstrap tests). Interestingly, CA3 principal cells exhibited significant reactivation exclusively during Rad$^{\mathrm{sink}}$ ripples ($p < 10^{-5}$; one-tailed bootstrap test), with no significant reactivation during LM$^{\mathrm{sink}}$ ripples ($p = 0.39$; one-tailed bootstrap test; Figure 6.3(A)).

These results suggest that while CA1 principal cells seem to reactivate during both ripple types, CA3 reactivation might be restricted to Rad$^{\mathrm{sink}}$ ripples, consistent with the reduced involvement of CA3 during LM$^{\mathrm{sink}}$ ripples (see Chapter 5).

Having established that CA1 principal cells are reactivated during both Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples, I next asked whether this reactivation involved deep and superficial pyramidal cells equally. In the previous section, I demonstrated Rad$^{\mathrm{sink}}$ ripples robustly recruit both deep and superficial CA1 cells, whereas LM$^{\mathrm{sink}}$ ripples preferentially recruit deep cells (see Section 5.3.1). This led me to hypothesise that both deep and superficial cells would be reactivated during Rad$^{\mathrm{sink}}$ ripples, while the overall CA1 reactivation observed during LM$^{\mathrm{sink}}$ ripples would be primarily driven by deep cells, with minimal contribution from superficial cells.

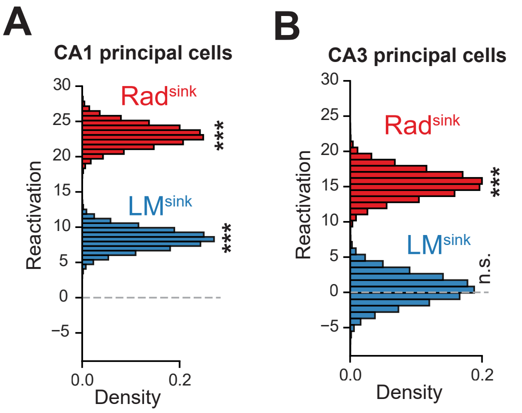

Figure 6.3: Reactivation of CA1 and CA3 principal cells in Rad$^{\mathrm{sink}}$ vs. LM$^{\mathrm{sink}}$ ripples.
(A, B) Offline reactivation of waking theta patterns in CA1 and CA3 principal cells. (A) Mean difference in the reactivation of CA1 principal cells during Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples, expressed as a percentage relative to pre-sleep (all $Ps < 10^{-5}$; one-tailed bootstrap tests; $n = 1,580$ CA1 principal cells). (B) Same as (A), but for CA3 principal cells, which show significant reactivation in Rad$^{\mathrm{sink}}$ ripples ($p < 10^{-5}$; one-tailed bootstrap test) but not in LM$^{\mathrm{sink}}$ ripples ($p = 0.39$; one-tailed bootstrap test; $n = 866$ CA3 principal cells). $^{\mathrm{n.s.}}p > 0.05$, $^{***}p < 0.001$.

To test this, I adapted the reactivation analytical framework described above to analyse the reactivation of deep and superficial cells separately (Figure 6.4(A); see Section 6.2.4). In line with my hypothesis, I found that Rad$^{\mathrm{sink}}$ ripples reactivated both cell types, whereas LM$^{\mathrm{sink}}$ ripples selectively reactivated deep cells (Figure 6.4(B,C); all $Ps < 10^{-5}$, except for superficial cells in LM$^{\mathrm{sink}}$ ripples: $p = 0.84$; one-tailed bootstrap test). This contrasts with the overall CA1 population reactivation observed during LM$^{\mathrm{sink}}$ ripples, which appears to be driven predominantly by deep cells (Figure 6.3(A)). Notably, the reactivation profile of CA1 superficial cells closely mirrored that of CA3 principal cells (Figure 6.3(B)), which are known to preferentially project to superficial rather than deep CA1 pyramidal neurons.

These findings indicate that Rad$^{\mathrm{sink}}$ ripples, associated with CA3 engagement, reactivate both deep and superficial CA1 cells, whereas LM$^{\mathrm{sink}}$ ripples predominantly reactivate deep cells, with no significant involvement of superficial or CA3 principal cells.

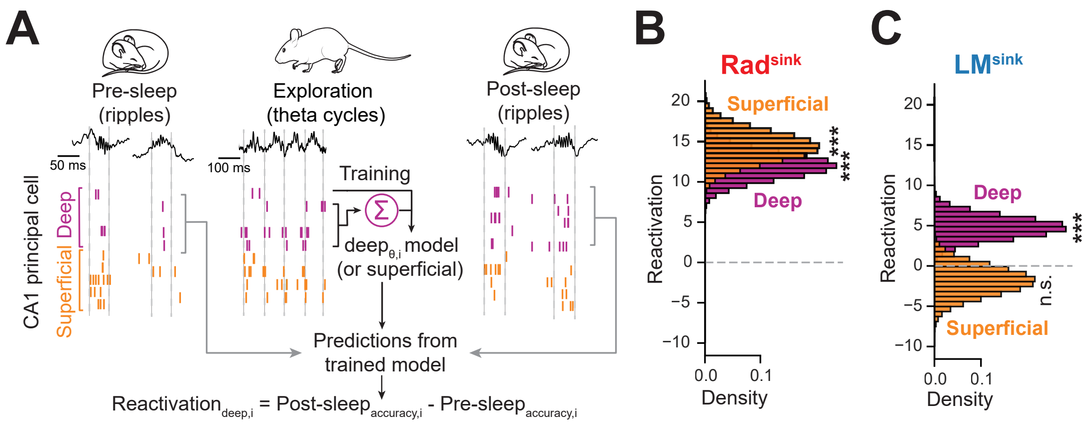

Figure 6.4: Reactivation of CA1 deep and superficial cells in Rad$^{\text{sink}}$ vs. LM$^{\text{sink}}$ ripples.
(A–C) Rad$^{\text{sink}}$ ripples reactivate waking theta coactivity of both deep and superficial cells, whereas LM$^{\text{sink}}$ ripples selectively reactivate deep cells. (A) Schematic of the method used to compute offline reactivation. A generalised linear model (GLM) was trained on theta cycles during active exploration to predict the activity of each principal cell based on the theta spiking activity of other cells in the same sublayer. Each trained theta model was then applied to predict the response of its target cell during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples in pre-exploration and post-exploration sleep. Reactivation was quantified as the change in model accuracy from pre- to post-exploration sleep. (B) Mean difference in reactivation during Rad$^{\text{sink}}$ ripples for deep and superficial cells. Both sublayers show significant reactivation during Rad$^{\text{sink}}$ ripples (all $Ps < 10^{-5}$; one-tailed bootstrap tests). (C) Same as (B), but for LM$^{\text{sink}}$ ripples. Notably, superficial cells do not exhibit significant reactivation during post-sleep LM$^{\text{sink}}$ ripples (deep cells: $p = 2.7 \times 10^{-4}$; superficial cells: $p = 0.84$; one-tailed bootstrap tests).
For (B),(C): $n = 1,097$ deep and $n = 478$ superficial cells. $^{\text{n.s.}}p > 0.05$, $^{***}p < 0.001$.

6.4 Composition of coactivity motifs in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

In the previous section, I demonstrated that Rad$^{\text{sink}}$ ripples contain higher-dimensional population patterns that reactivate both deep and superficial cells, forming dense coactivity motifs. In contrast, LM$^{\text{sink}}$ ripples host lower-dimensional population patterns that primarily engage deep cells in core motifs of coactivity.

Previous research has shown that deep and superficial CA1 principal cells respond differently to behavioural experiences. Specifically, during theta oscillations in wakeful exploration, deep cells tend to show similar tuning across different contexts, whereas superficial cells adapt their firing to discriminate between distinct contexts and experiences (Gava et al., 2021, 2024; Harvey et al., 2023). In other words, deep cells are more associated with stable and rigid tuning, while superficial cells exhibit more flexible, context-dependent responses (Gava et al., 2024; Harvey et al., 2023). Given these differences in adaptation, I hypothesised that the lower-dimensional motifs formed by deep CA1 cells in LM$^{\mathrm{sink}}$ ripples would reflect more stable and recurring activity patterns, whereas the higher-dimensional motifs involving both deep and superficial cells in Rad$^{\mathrm{sink}}$ ripples would integrate new information from wakefulness and exploration, reflecting greater flexibility and remapping potential.

To test this, I first examined whether the coactivity motifs expressed in LM$^{\mathrm{sink}}$ and Rad$^{\mathrm{sink}}$ ripples during post-exploration sleep were similar to or different from those in pre-exploration sleep. I quantified this using a measure of coactivity stability, comparing population coactivity patterns in pre- and post-exploration sleep (Figure 6.5; see Section 6.2.5).

Consistent with my hypothesis, I found that coactivity motifs in LM$^{\mathrm{sink}}$ ripples were significantly more similar to those in pre-sleep than in Rad$^{\mathrm{sink}}$ ripples, indicating that LM$^{\mathrm{sink}}$ ripples maintain stable coactivity patterns that persist from pre- to post-exploratory sleep. In contrast, Rad$^{\mathrm{sink}}$ ripples undergo greater changes in their coactivity patterns from pre- to post-exploration sleep, possibly reflecting the integration of new information encoded during wakefulness.

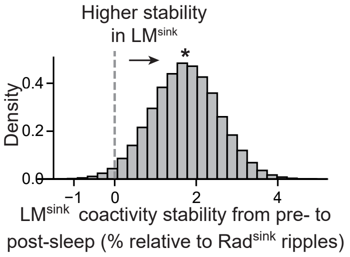

Figure 6.5: Coactivity stability in LM$^{\mathrm{sink}}$ ripples across sleep.
Mean difference in the stability of CA1 principal cell coactivity from pre-exploration to post-exploration sleep in Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples, expressed as a percentage relative to Rad$^{\mathrm{sink}}$ ripples. LM$^{\mathrm{sink}}$ motifs are more stable than Rad$^{\mathrm{sink}}$ motifs, showing greater consistency across the two sleep sessions surrounding exploration [mean stability in LM$^{\mathrm{sink}}$ ripples relative to Rad$^{\mathrm{sink}}$ ripples (95% CI): 1.74 (0.12–3.35)%; $p = 0.018$; 1-tailed paired bootstrap test; $n = 140$ pre-/post-sleep session pairs].
$^{*}p < 0.05$.

To further test whether this remapping in Rad$^{\mathrm{sink}}$ coactivity was due to the integration of recently acquired information, while LM$^{\mathrm{sink}}$ ripples reflected pre-existing coactivity patterns, I developed a method to distinguish between these two types of motifs (Figure 6.6). This method allowed me to isolate recent (awake-driven) versus prior (pre-existing) motifs and assess their relative contributions to Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples.

In the rest of this section, I first describe this method and validate it through simulations. I then show how Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples differentially engage recent and prior motifs.

6.4.1 Extraction and analysis of prior and recent coactivity motifs in simulated activity

To isolate the coactivity motifs formed during wakefulness from those that were already present in pre-sleep, I developed the method outlined in Section 6.2.6. A key assumption in this approach is that during pre-exploration sleep, a set of CA1 coactivity patterns exists. During subsequent exploration, these coactivity motifs may be strengthened or suppressed, acquire or lose neurons, or entirely new assemblies may emerge. Regardless of the specific changes, the assumption is that some transformation occurs between pre-exploration sleep and awake exploration, and this change can be captured using a linear transformation (see Equations 6.3, 6.4).

To validate that this method accurately captures the transformation from pre-sleep to awake exploration, I conducted *in silico* simulations, as described in Section 6.4.1.

I first simulated a pre-sleep population activity matrix containing two pre-defined motifs of coactive neurons, referred to as prior motifs. During the subsequent awake exploration, one of these prior motifs underwent re-organisation by incorporating an additional neuron, while a new motif of coactive neurons emerged, representing a recent motif. I then simulated two post-exploration sleep population activity matrices: one retained only the motifs previously observed in pre-exploration sleep, referred to as rigid ripples, while the other incorporated the newly acquired exploration-related motifs, referred to as plastic ripples. Figure 6.6(A) illustrates example activity matrices over time bins for each session, with the corresponding coactivity matrices in Figure 6.6(B).

Using the method detailed in Section 6.2.6, I extracted two subspaces from the pre-exploration sleep and exploration coactivity matrices (Figure 6.6(C)). The first subspace captured prior coactivity motifs that were already present in pre-sleep. The second subspace isolated recent motifs that emerged during exploration and were absent in pre-sleep. A qualitative assessment of these transformation matrices confirmed that they successfully highlighted both the pre-existing motifs and the changes introduced during wakefulness.

Finally, I examined the strength of recent and prior motifs in post-sleep rigid and plastic ripples. The analysis revealed that rigid ripples more strongly reactivated prior motifs, whereas plastic ripples predominantly reactivated recently acquired coactivity motifs (Figure 6.6(D)).

The results in this section demonstrate that the method outlined in Section 6.2.6 can reliably distinguish recent from prior motifs and track their strength in post-exploration sleep. In the next section, I apply this method to Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples to characterise the function of their coactivity motifs.

6.4.2 Recent and prior motifs in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

Given that the main hypothesis in this section aimed to test whether Rad$^{\text{sink}}$ coactivity was biased toward recently acquired information, while LM$^{\text{sink}}$ ripples reflected pre-existing coactivity patterns, I applied the simulation-validated method described above to neural data. This allowed me to first decompose neuronal coactivity motifs into *recent* motifs, defined as neurons selectively coactive during arena exploration but not during the preceding sleep/rest, and *prior* motifs, referring to neurons already coactive in pre-exploration sleep/rest. I then measured the

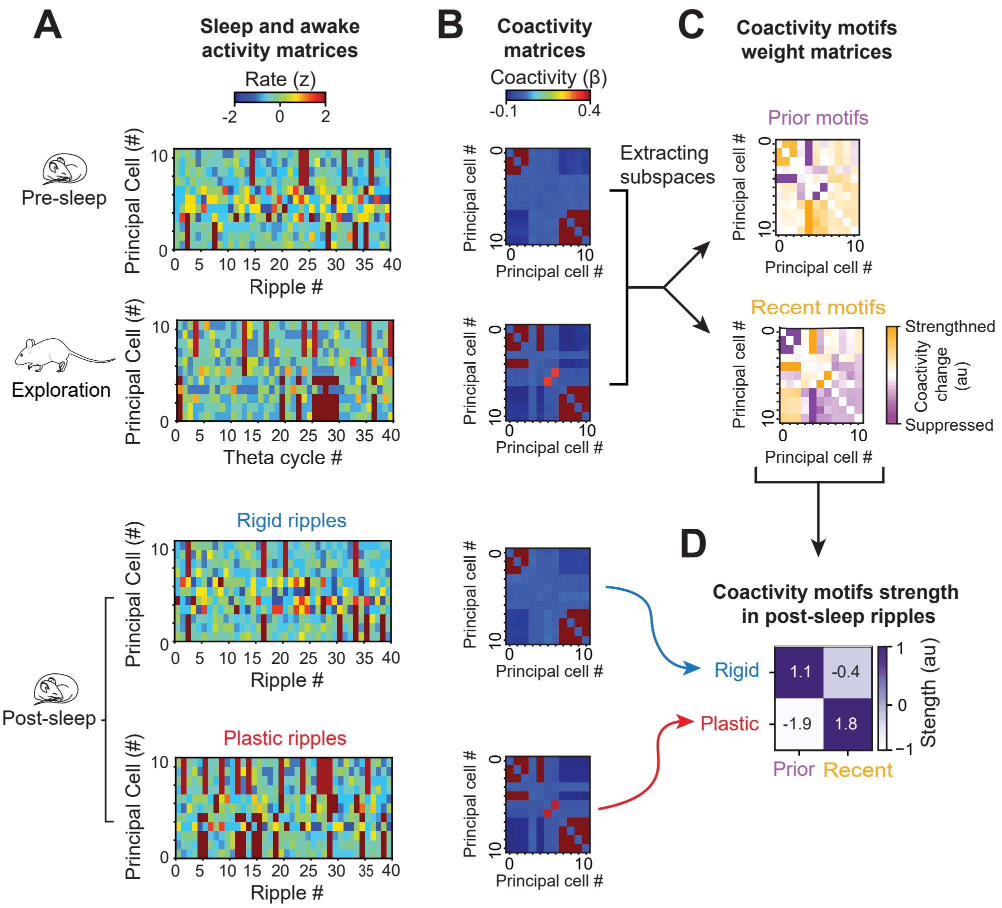

Figure 6.6: Extracting recent and prior coactivity motifs in simulated neural data.
(A–D) Validation of the method for isolating prior and recent coactivity motifs using in-silico simulations. Shown in (A) are example activity matrices over time bins for each session, including pre-sleep, awake exploration, and post-exploration sleep conditions (rigid and plastic ripples), with corresponding coactivity matrices in (B). From pre-sleep to awake exploration, coactivity motifs changed, and these changes were present in post-sleep plastic ripples but absent in rigid ripples. I then extracted two subspaces from the pre-exploration sleep and exploration coactivity matrices: prior coactivity motifs (present in pre-sleep) and recent motifs (emerging during exploration and absent in pre-sleep), as shown in (C). Transformation matrices representing these motifs highlight cell pairs composing these motifs, with strengthened connections indicated in orange (top panels for prior motifs, bottom panels for recent motifs). For instance, the recent motifs weight matrix highlights the newly formed ‘connections’ during wakefulness, effectively extracting a subspace orthogonal to the pre-sleep coactivity motifs. These transformation matrices were then applied to the simulated post-exploration sleep activity shown in (A) to evaluate the strength of each motif. The heatmap in (D) summarises the strength of prior and recent coactivity motifs in the two post-exploration sleep conditions from (A, B), showing that rigid ripples predominantly reactivated prior motifs, while plastic ripples primarily reactivated recently-expressed coactivity motifs.

reactivation strength of these recent versus prior motifs during post-exploration sleep/rest (Figure 6.7(A)). Strikingly, Rad$^{\text{sink}}$ ripples preferentially reactivated recent coactivity motifs, whereas LM$^{\text{sink}}$ ripples continued to express prior motifs (Figure 6.7(B); prior motifs: $p = 6 \times 10^{-4}$; recent motifs: $p < 10^{-5}$; bootstrap tests).

These results align with previous findings that LM$^{\text{sink}}$ coactivity motifs exhibit strong cross-sleep similarity (Figure 6.5), consistent with stronger structural balance (Figure 5.13) and the lower dimensionality of LM$^{\text{sink}}$ patterns (Figure 5.18).

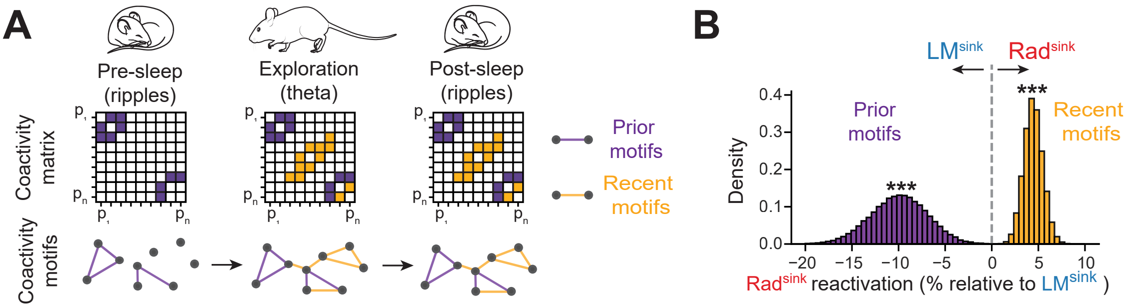

Figure 6.7: Recent and prior motifs in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.
(A) Schematic of the method used to identify prior and recent neuronal coactivity motifs (see also Figure 6.6). Coactivity matrices (top) and corresponding neuronal graph motifs (bottom) are shown across pre-sleep (ripples), exploration (theta cycles), and post-sleep (ripples). Purple squares denote cellpairs forming prior coactivity motifs (i.e., already present in pre-sleep), while orange squares indicate coactivity relationships that emerged selectively during exploration (i.e., absent in pre-sleep).
(B) Mean difference in reactivation of prior and recent motifs during post-sleep Rad$^{\text{sink}}$ ripples relative to LM$^{\text{sink}}$ ripples. Rad$^{\text{sink}}$ ripples selectively reactivated recently expressed theta coactivity motifs, whereas LM$^{\text{sink}}$ ripples preferentially reactivated prior motifs. For (B): $n = 140$ pre-/post-sleep session pairs. $^{***}p < 0.001$.

At the beginning of this chapter, I discussed how LM$^{\text{sink}}$ ripples are primarily composed of lower-dimensional motifs formed by stable deep cells, while Rad$^{\text{sink}}$ ripples integrate these stable motifs with plastic superficial cells into higher-dimensional PVs. Now, in line with this, I have shown that LM$^{\text{sink}}$ ripples express prior motifs, while the higher-dimensional motifs in Rad$^{\text{sink}}$ ripples incorporate recently formed motifs. While these findings suggest a complementary relationship between LM$^{\text{sink}}$ and Rad$^{\text{sink}}$ motifs, a key question remains: are superficial cells primarily responsible for driving the recently formed motifs in Rad$^{\text{sink}}$ ripples, while deep cells primarily define the prior, stable motifs in LM$^{\text{sink}}$ ripples?

To test this, I measured the contribution of deep and superficial cells to the reactivation of recent and prior motifs (see Section 6.2.7). In line with my hypothesis and previous findings, deep cells predominantly contributed to prior coactivity motifs (Figure 6.8(A); $p = 2 \times 10^{-5}$; bootstrap test), whereas superficial cells contributed more significantly to the reactivation of recent coactivity motifs (Figure 6.8(B); $p = 1.2 \times 10^{-4}$; bootstrap test).

To conclude, in this section I demonstrated that LM$^{\text{sink}}$ ripples host core motifs of prior neuronal coactivity that primarily engage deep cells and remain stable across sleep/rest epochs. In contrast, Rad$^{\text{sink}}$ ripples preferentially reactivate recent motifs of waking coactivity, incorporating superficial cells alongside deep cells to form high-dimensional population patterns.

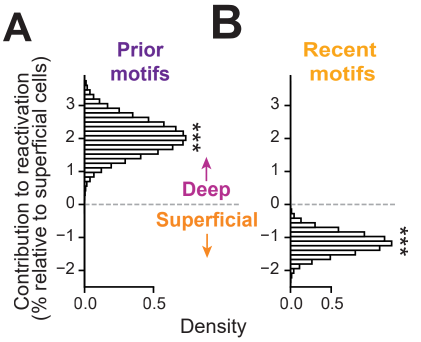

Figure 6.8: Contribution of CA1 deep and superficial to recent and prior motifs.
(A,B) Mean difference in the contribution of deep and superficial cells to the reactivation of prior (A) and recent (B) motifs (relative to the mean of superficial cells). Deep cells primarily contributed to prior motifs, whereas superficial cells predominantly contributed to recent motifs. .
$n = 43$ pre-/post-sleep session pairs. $^{***}p < 0.001$.

6.5 Temporal dynamics of reactivation in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples

In the previous section, I demonstrated that, during post-exploration sleep, LM$^{\text{sink}}$ ripples preferentially express motifs from prior sleep, while Rad$^{\text{sink}}$ ripples rapidly adapt to reactivate newly formed motifs through the coactivity of deep and superficial cells.

This raised an important question: does the strong coactivity of deep cells within LM$^{\text{sink}}$ ripples remain entirely stable and insensitive to recent information, or do these motifs gradually update over time? Given the stronger structural balance, the predominance of deep cells, and the lower dimensionality of LM$^{\text{sink}}$ ripples (see Chapter 5), it is possible that these motifs require more time to undergo changes. In other words, over a longer timescale of sleep, do LM$^{\text{sink}}$ ripples eventually drift or update their coactivity motifs?

To test this, I analysed the temporal dynamics of CA1 population patterns over hour-long sleep sessions [mean duration (IQR): 86.25 (76.33 - 91.62) minutes per sleep session]. Strikingly, in LM$^{\text{sink}}$ ripples, the strength of prior motifs (i.e., those already present in pre-exploration sleep) decreased exponentially during the first half-hour of post-exploration sleep/rest [Figure 6.9(A); time constant τ (95% CI): 13.12 (8.34 - 21.80) minutes; $p < 10^{-4}$; one-tailed bootstrap test], while the strength of recent motifs increased over the same period [Figure 6.9(B); time constant τ (95% CI): 13.77 (8.50 - 27.70) minutes; $p < 10^{-4}$; one-tailed bootstrap test]. In contrast, Rad$^{\text{sink}}$ ripples showed no significant changes in reactivation content throughout the sleep period [Figure 6.9; prior motifs, $p = 0.32$; recent motifs, $p = 0.23$; bootstrap tests]. To ensure that the observed exponential changes in LM$^{\text{sink}}$ ripples were not biased by changes in sleep quality at the beginning of the session, the first 10 minutes of each sleep session were excluded from this analysis.

These results show that while LM$^{\text{sink}}$ ripples preferentially express prior motifs on average, over longer timescales, they gradually update their content toward recently formed motifs. On the other hand, Rad$^{\text{sink}}$ motifs remain stable throughout sleep.

A critical question is whether this drift in LM$^{\text{sink}}$ motifs reflects a precise remapping process or whether it arises as a by-product of other factors. One possible explanation could be an increase in the number of LM$^{\text{sink}}$ ripples (i.e., ripple occurrence frequency), which might bias the observed changes in reactivation. Another possibility is that an increasing number of neurons become recruited into LM$^{\text{sink}}$ ripples over time (i.e., ripple population sparsity), and as a by-product, this

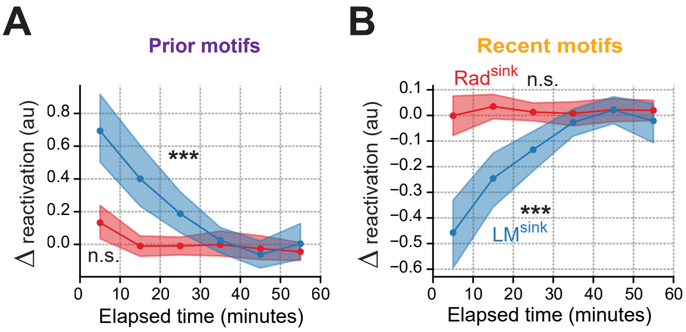

Figure 6.9: Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ reactivation dynamics of prior vs. recent motifs.
(A-B) Time course of the mean change in reactivation of prior (A) and recent (B) motifs in Rad$^{\mathrm{sink}}$ (red) and LM$^{\mathrm{sink}}$ (blue) ripples over the hour-long sleep/rest period. Reactivation values are normalised by subtracting the overall mean reactivation of each session. LM$^{\mathrm{sink}}$ ripples exhibited an exponential decay in the reactivation of prior motifs, whereas the reactivation of recent motifs showed an exponential increase. In contrast, Rad$^{\mathrm{sink}}$ ripples maintained temporally stable reactivation throughout sleep/rest, without significant temporal changes. Significance levels were determined based on the fitness score of the exponential function (see Section 6.2.8). Shaded areas represent the 95% confidence interval.
$n = 118$ pre-/post-sleep session pairs. $^{\mathrm{n.s.}}p > 0.05$, $^{***}p < 0.001$.

could lead to a gradual shift toward recent motifs. Alternatively, the observed temporal changes in LM$^{\mathrm{sink}}$ reactivation may be independent of these factors.

To explore this further, I examined the long-term dynamics of ripple occurrence frequency and ripple population sparsity over the hours-long sleep period (Figure 6.10). A preliminary inspection revealed that while ripple occurrence frequency showed some variability, ripple population sparsity remained relatively stable throughout sleep/rest. However, this qualitative assessment alone is not sufficient to determine whether the observed reactivation dynamics over sleep significantly (Figure 6.9) depend on these factors or whether they change strictly as a function of time. A more rigorous analysis was required to disentangle these contributions (see Section 6.2.9).

To determine whether the observed drift was intrinsic to time (i.e., ripple occurrence time) and not influenced by changes in ripple frequency or sparsity, I trained linear regression models to predict the reactivation strength of recent motifs over the sleep/rest period based on ripple occurrence time, frequency, and sparsity

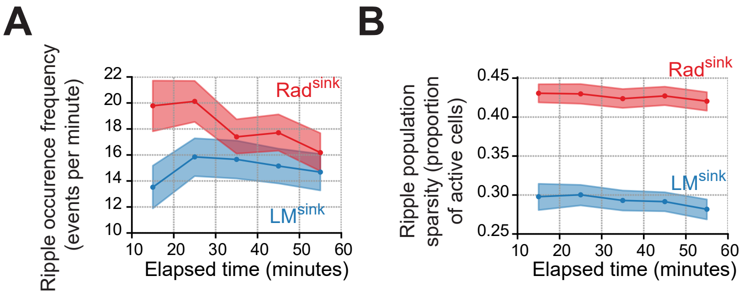

Figure 6.10: Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ dynamics of ripple occurrence frequency and sparsity.
(A–B) Time course of mean ripple occurrence frequency (A) and population sparsity (measured as the proportion of recruited cells) (B) in Rad$^{\text{sink}}$ (red) and LM$^{\text{sink}}$ (blue) ripples over the hour-long sleep/rest period. Shaded areas represent the 95% confidence interval ($n = 140$ pre-/post-sleep session pairs).

(Figure 6.11(A); see Section 6.2.9). I repeated this analysis separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Consistent with the stable reactivation dynamics observed throughout sleep, the Rad$^{\text{sink}}$ model did not significantly predict reactivation time course ($p = 0.068$; one-tailed bootstrap test). In contrast, reactivation strength in LM$^{\text{sink}}$ ripples was significantly predicted by the model ($p < 4 \times 10^{-5}$; one-tailed bootstrap test), consistent with the temporal drift observed in Figure 6.9.

To determine which variable most strongly influenced reactivation strength, I independently shuffled all predictors (i.e., ripple occurrence time, frequency, and sparsity) and assessed their contributions to overall model accuracy. For Rad$^{\text{sink}}$ ripples, only ripple population sparsity contributed significantly above chance (Figure 6.11(B,C); $p = 2.4 \times 10^{-3}$; one-tailed bootstrap test, Bonferroni corrected). In contrast, for LM$^{\text{sink}}$ ripples, only ripple occurrence time had a significant effect (Figure 6.11(B,C); $p < 4 \times 10^{-5}$; one-tailed bootstrap test, Bonferroni corrected).

These results suggest that the gradual drift in LM$^{\text{sink}}$ coactivity motifs from prior to recent is indeed time-dependent and does not arise from other changes in ripple dynamics (Figures 6.9 and 6.11(C)). More importantly, this shift occurs while maintaining a constant number of ripples (ripple occurrence frequency) and a stable number of recruited cells (ripple population sparsity) (Figure 6.10). For Rad$^{\text{sink}}$ ripples, reactivation strength was significantly correlated with ripple population sparsity, which likely explains the borderline non-significant model accuracy (Figure

6.11(B,C)). Nevertheless, this finding indicates that Rad$^{\text{sink}}$ reactivation does not depend on the elapsed time within the sleep session but rather on the number of recruited cells.

In summary, these findings demonstrate that Rad$^{\text{sink}}$ ripples consistently reactivate the same motifs throughout sleep. In contrast, LM$^{\text{sink}}$ ripples gradually shift from prior to recent motifs over time, all while maintaining a stable ripple occurrence rate and ripple population sparsity.

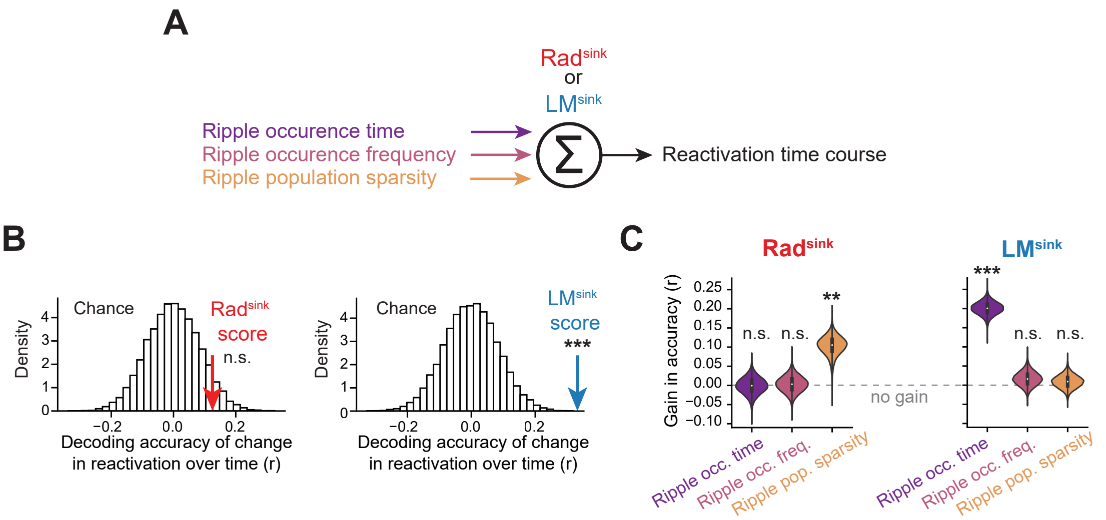

Figure 6.11: Factors influencing reactivation dynamics of prior vs. recent motifs during sleep.
(A–C) Factors explaining temporal changes in reactivation over an hour-long sleep period (see Section 6.2.9). Linear regression models were trained to predict the reactivation strength of recent motifs across sleep/rest (A). Separate models were trained for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, using three predictors: ripple occurrence time, ripple occurrence frequency, and ripple population sparsity (computed in 10-minute bins). (B) Model accuracy for Rad$^{\text{sink}}$ (left) and LM$^{\text{sink}}$ (right) ripples, with chance level estimated by shuffling reactivation values 25,000 times. Red/blue arrows indicate the observed model accuracy. The Rad$^{\text{sink}}$ model did not significantly predict the reactivation time course ($p = 0.068$, one-tailed bootstrap test), consistent with the stable reactivation observed throughout sleep/rest (Figure 6.9). In contrast, changes in reactivation over time during LM$^{\text{sink}}$ ripples were significantly predicted ($p < 4 \times 10^{-5}$, one-tailed bootstrap test). (C) Violin plots showing the contribution of each predictor to model accuracy. For Rad$^{\text{sink}}$ ripples (left), only ripple population sparsity significantly contributed ($p = 2.4 \times 10^{-3}$, one-tailed bootstrap test, Bonferroni corrected). For LM$^{\text{sink}}$ ripples (right), only ripple occurrence time significantly contributed ($p < 4 \times 10^{-5}$, one-tailed bootstrap test, Bonferroni corrected).
$^{\text{n.s.}}p > 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

6.6 Discussion

In this chapter, I have demonstrated that LM$^{\mathrm{sink}}$ ripples contain small motifs of stronger coactivity between principal cells in the deep CA1 pyramidal sublayer. Following waking experience, Rad$^{\mathrm{sink}}$ ripples append onto these core motifs the recently recruited superficial CA1 principal cells, yielding denser, higher-dimensional population patterns that consistently reactivate over hour-long sleep/rest. Meanwhile, the pre-existing coactivity backbone nested in LM$^{\mathrm{sink}}$ ripples, formed by sparser, lower-dimensional patterns, gradually drifted throughout sleep/rest, slowly updating its content to reflect recent wakefulness.

Recent studies have documented significant diversity among CA1 pyramidal cells, emphasising that their somatic location within the deep versus superficial sublayers of the stratum pyramidale predicts distinct contributions to hippocampal network dynamics (Cavalieri et al., 2021; Gava et al., 2024; Harvey et al., 2023; Mizuseki et al., 2011; Valero et al., 2015). Overall, during ripples, superficial CA1 principal cells exhibit stronger changes in firing activity compared to deep cells (Gava et al., 2024; Harvey et al., 2023; Lopes-dos-Santos et al., 2023; Valero et al., 2015). Here I show that while Rad$^{\mathrm{sink}}$ ripples reflect the canonical pattern of superficial cells being more active than deep cells, LM$^{\mathrm{sink}}$ ripples exhibit the opposite trend, with deep cells being predominantly more active than their superficial counterparts. Furthermore, I found that the additional CA1 principal cells recruited during Rad$^{\mathrm{sink}}$ ripples — but not during LM$^{\mathrm{sink}}$ ripples — are primarily superficial cells.

Importantly, while previous studies showed that the offline reactivation of waking patterns in CA1 is mainly driven by superficial cells (Harvey et al., 2023), I observed that while this seemed the case in the canonical Rad$^{\mathrm{sink}}$ ripples, this is not the case during LM$^{\mathrm{sink}}$ ripples. CA1 superficial cells and CA3 principal cells were not significantly reactivated during these ripples. These findings make sense given the differential anatomical inputs: CA1 deep cells receive stronger input from MEC3, which projects to *stratum lacunosum-moleculare*, while CA1 superficial cells are more strongly driven by CA3 inputs terminating in *stratum radiatum* (Fernández-Ruiz et al., 2017; Valero & de la Prida, 2018). Accordingly, I suggest that during Rad$^{\text{sink}}$ ripples, offline reactivation of waking patterns primarily engages CA3 ensembles, which in turn bias CA1 to reactivate the associated higher-dimensional response integrating superficial CA1 cells.

Deep and superficial CA1 principal cells have also been shown to respond differently to behavioural experiences. Notably, in theta cycles marking wakeful exploration, deep cells exhibit higher firing rates and more stable activity patterns, whereas superficial cells show lower firing rates and adapt flexibly to new experiences (Gava et al., 2024; Harvey et al., 2023; Lopes-dos-Santos et al., 2023). This difference could be partly attributed to stronger excitatory inputs that superficial cells receive from CA3, enabling them to respond to contextual changes (Berndt et al., 2023; Valero & de la Prida, 2018). Consistent with these findings, I show that superficial cells drive in Rad$^{\text{sink}}$ ripples of post-exploration sleep the expression of coactivity motifs expressed during the previous exploratory behaviour. In contrast, the more rigid deep cells instantiate pre-existing motifs of coactivity in LM$^{\text{sink}}$ ripples. Previous studies reported that CA1 superficial, but not deep, principal cells undergo postsynaptic potentiation in response to CA3 inputs after novel experience (Berndt et al., 2023). In line with this, I found that CA3 reactivation mirrored the behaviour of superficial cells CA1. These results suggest that CA3 might enable CA1 superficial cells to rapidly and flexibly reorganise coactivity patterns in Rad$^{\text{sink}}$ ripples following new waking experience. Conversely, the reduced activity of CA3 during LM$^{\text{sink}}$ ripples could favour the expression of deep cell motifs.

These findings are in line with a recent study by Chang et al., 2025, which showed that ripples associated with small pupil size — typically of higher ripple frequency — are more linked to novel or recent memories, while large pupil size ripples, of lower frequency, are associated with prior and familiar memories. In this sense, the small-pupil-size ripples share many characteristics with the Rad$^{\text{sink}}$ ripples characterised here, which are also biased toward reactivating recent motifs. However, it is slightly more complicated to directly associate LM$^{\text{sink}}$ ripples with large-pupil-size ripples. Although they share several features — such as lower ripple frequency, reduced CA1 cell recruitment, and weaker sinks in *stratum radiatum* — the study by Chang et al., 2025 only examined sharp-wave ripples, meaning CA1 ripples that are accompanied by a sharp-wave in *radiatum.* In this sense, as I discuss in more detail later in Chapter 7, the ripples analysed in that study are likely biased toward baseline and Rad$^{\text{sink}}$ ripples (see Chapter 3). It is possible that large-pupil-size ripples in that dataset overlap more with what I here define as baseline ripples rather than true LM$^{\text{sink}}$ ripples, which may explain why some of their features align only partially. Nevertheless, similarly to that study, in this thesis I demonstrate through various analytical and experimental methods that different ripples carry different types of content, either biased toward recently-formed or prior motifs.

The prior coactivity motifs nested in LM$^{\text{sink}}$ ripples are more stable and yet not stationary over the hour-long sleep following recent wakefulness. Population patterns in LM$^{\text{sink}}$ ripples, but not those in Rad$^{\text{sink}}$ ripples, gradually drift toward the coactivity space marking the recent waking activity. It remains for future studies to reveal the mechanistic implementation of this drift, occurring while the most recently activated coactivity motifs are quickly integrated into high-dimensional patterns of deep and superficial cells in Rad$^{\text{sink}}$ ripples. An interesting aspect concerns the input-output relationships established by superficial and deep CA1 principal cells, and the possibility of bidirectional influences. Indeed, CA1 superficial cells predominantly project to MEC (Fernández-Ruiz et al., 2017; Harvey et al., 2023; Soltesz & Losonczy, 2018; Valero & de la Prida, 2018); CA1 deep cells preferentially receive inputs from MEC3. This could actuate a feedback loop to CA1 deep cells. CA1 superficial cells could also modulate deep cells by recruiting local inhibitory circuits (e.g., parvalbumin-expressing basket cells (Lee et al., 2014; Valero et al., 2015)). By these means, CA1 superficial cells would influence both the hippocampal-entorhinal and the intra-hippocampal circuitries, shaping the core motifs of deep CA1 cell activity embedded in LM$^{\text{sink}}$ ripples. That is, during post-exploration sleep, motifs of superficial cells reflecting recent waking experience reactivate in Rad$^{\text{sink}}$ ripples. This could in turn bias MEC activity, which projects back to CA1 deep cells, gradually shifting the CA1 deep motifs toward the most recent activity state space. This interactive process could help internal linking of the recent experience with previous ones in the hippocampus, and allow network updating with the most recently formed information.

The temporal dynamics observed for the reactivation of recent and prior motifs within Rad$^{\mathrm{sink}}$ and LM$^{\mathrm{sink}}$ ripples depend on different factors, such as ripple occurrence frequency (i.e., how many ripples occurred) and ripple occurrence sparsity (i.e., how many cells were recruited on average). For LM$^{\mathrm{sink}}$ ripples, the exponential change in coactivity motifs expressed within CA1 principal cells appeared to be explained solely by the elapsed time in the sleep/rest session, rather than by ripple occurrence frequency or sparsity. This suggests that, over the course of hour-long sleep, the coactivity motifs within LM$^{\mathrm{sink}}$ ripples shift from expressing prior motifs to recent motifs without changing the number of LM$^{\mathrm{sink}}$ ripples or the number of recruited cells. This is particularly interesting because it implies that the sparse LM$^{\mathrm{sink}}$ ripples remain sparse throughout sleep but gradually remap their coactivity structure. In contrast, reactivation within Rad$^{\mathrm{sink}}$ ripples did not appear to change over time but instead seemed to depend primarily on the number of recruited cells. In other words, the more cells that were recruited in Rad$^{\mathrm{sink}}$ ripples — regardless of when the event occurred within the session or the total number of Rad$^{\mathrm{sink}}$ ripples — the reactivation was stronger. This could suggest that Rad$^{\mathrm{sink}}$ ripples facilitate higher reactivation than LM$^{\mathrm{sink}}$ ripples through a mechanism of recruiting more cells, potentially driven by stronger CA3 inputs (see Chapter 5). Rather than precisely mapping coactivity patterns, Rad$^{\mathrm{sink}}$ ripples may rely on large-scale CA3 activation to ensure that CA1 motifs associated with recent wakefulness experiences are consolidated. However, this remains a speculative hypothesis. Nevertheless, the data suggest that, throughout sleep, LM$^{\mathrm{sink}}$ ripples precisely remap their coactivity while preserving their sparsity, whereas Rad$^{\mathrm{sink}}$ reactivation remains stable and fluctuates randomly as a function of the number of recruited cells.

These findings suggest a model in which Rad$^{\mathrm{sink}}$ ripples enable rapid integration of new experiences through superficial CA1 cells, while LM$^{\mathrm{sink}}$ ripples scaffold these experiences onto pre-existing hippocampal representations through gradual motif updating.

7

General Discussion

In this thesis, I identified two ripple profiles based on their distinct laminar currents across the CA1 somato-dendritic axis, which differentially recruited hippocampal neurons during offline periods of sleep and rest. The first profile, characterised by stronger *stratum radiatum* currents and high CA3 activity, integrates superficial and deep CA1 principal cells into higher-dimensional composite patterns that undergo stable, hour-long reactivation of recently expressed waking motifs of neuronal coactivity. The second profile, marked by stronger *stratum lacunosum-moleculare* currents and significantly weaker CA3 spiking, contains core motifs of deep cells forming lower-dimensional patterns that exhibit time-varying reactivation, gradually drifting over sleep from prior to recent coactivity spaces. Collectively, these findings reveal a diversity of hippocampal ripple profiles linked to the offline organisation of neuronal population reactivation. By tuning differences in activity level, structural organisation, and reactivation content of population patterns (Figure 7.1), this ripple-by-ripple diversity may instantiate parallel reactivation channels that balance the stable integration of recent wakefulness with the flexible updating of prior internal representations.

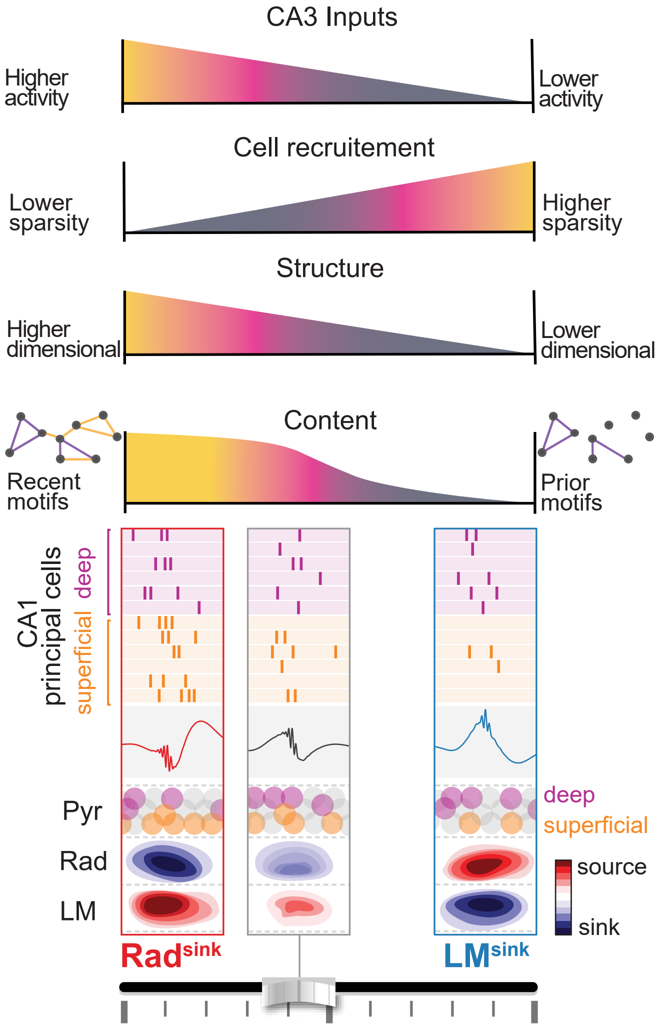

Figure 7.1: Summary schematic.
Two ripple profiles, identified by their radially organised currents, exhibit distinct CA1 population activity levels, structures, and contents. Rad$^{\text{sink}}$ ripples, characterised by stronger *radiatum* current sinks and higher CA3 activity, integrate recent motifs of waking coactivity, combining superficial and deep principal cells into higher-dimensional population patterns that undergo stable reactivation throughout hour-long sleep/rest. In contrast, LM$^{\text{sink}}$ ripples, marked by stronger *lacunosum-moleculare* current sinks and lower CA3 activity, contain small motifs of pre-structured coactivity, engaging deep cells into lower-dimensional patterns that exhibit reactivation drift, gradually updating prior motifs with recent activity content.

Ripple diversity as a continuum

I began this thesis by analysing the variability in CSD profiles of individual ripple events. These profiles form a continuum rather than discrete clusters, consistent with previous reports suggesting that ripple waveforms and features occupy a continuous parameter space (Sebastian et al., 2023). This continuum may enable dynamic modulation of hippocampal population activity on a ripple-by-ripple basis through transient changes in current flow within specific neural layers. The dominant contributors to this continuum are ripples associated with a stronger current sink in *stratum radiatum* (Rad^{sink} ripples), reflecting the canonical CSD profile of the grand average ripple (Buzsáki et al., 1983; Sullivan et al., 2011). Interspersed among these are ripples that instead show a stronger current sink in *stratum lacunosum-moleculare* (LM^{sink} ripples), which remain hidden when averaging across ripples. CA3 projections to CA1 *stratum radiatum* play a key role in sharp-wave/ripple generation (Buzsáki et al., 1983; Buzsáki, 1986; Nakashiba et al., 2009), consistent with the large average input current observed in this layer. Yet, both this study and prior work (Sebastian et al., 2023) reveal substantial variability in current profiles associated with individual ripples. Using an unsupervised approach, I identified a subset of ripples characterised by a stronger current sink in *stratum lacunosum-moleculare* and a slower ripple frequency. Given that *stratum radiatum* and *stratum lacunosum-moleculare* receive distinct inputs from CA3 and medial entorhinal cortex layer III (MEC3), respectively, this diversity in laminar sinks may reflect differences in upstream inputs modulating ripple generation at the level of individual ripple events. This transient, ripple-by-ripple tuning of laminar currents is reminiscent of the dynamic modulation observed across individual cycles of hippocampal theta oscillations during awake exploration (Lopes-dos-Santos et al., 2018). Similar to theta cycles, my findings demonstrate that the hippocampus dynamically tunes offline population activity on a ripple-by-ripple basis, supporting recent work suggesting that a temporal microstructure of sleep embeds the reactivation of co-existing hippocampal patterns (Chang et al., 2025).

The fact that ripples form a continuum rather than discrete clusters also prompts some follow-up discussion points. The idea that ripples are defined within a continuum, consistent with other studies, means that there are no clear-cut clusters of ripple types. Within this continuum, the currents in *stratum radiatum* and *stratum lacunosum-moleculare* are gradually tuned and define the extremes of this distribution (Figure 7.1). However, throughout this thesis — especially from Chapter 5 onwards — I focused on the two extremes, $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples, which may give the impression that ripples fall into two discrete clusters. As discussed in Section 3.4, the parameters used to focus on the extremes could indeed be optimised. Nevertheless, focusing on the extremes allowed me to highlight an important result in itself: that the laminar profiles of ripples are highly diverse, as also shown in previous work (Sebastian et al., 2023).

In a way, when the features defining the data are so variable, trying to look at everything together would make it impossible to see any clear trend, as everything would get diluted — especially when applying common ripple analysis methods such as averaging. In fact, although it is well established that ripples are associated with sinks in *stratum radiatum*, my analysis shows that a consistent subset of ripples has an intrinsically different profile, with a sink in *lacunosum-moleculare*, and this is something that would be completely hidden when averaging all events.

One could ask whether these distinct features might arise from biological noise or from how the extremes of the distribution were selected. I argue that this is unlikely, and throughout this thesis I show several pieces of evidence that do not support this hypothesis. First, while the thresholding used to separate $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples from the full distribution (see Chapter 3) could play a minor role in mixing and diluting features, this would not explain all the differences observed. Of course, the more extreme the events picked, the stronger the differences expected — but naturally, this also comes with fewer events, making it harder to achieve statistical significance. But more importantly, the differences I report are not just a matter of degree — for instance, they do not simply reflect a gradual increase or decrease in sink strength in *radiatum*. Instead, what is revealed at the extremes are qualitatively different profiles, where the sink switches from being in *radiatum* to *lacunosum-moleculare*. This matters because different regions project to these layers, and this supports the idea that these distinct ripple profiles are tuned by different input currents rather than being noise or artefacts from analysis. Second, I show throughout the thesis that these CSD profiles are consistently associated with other distinct features — features that are unrelated to how the ripples were split. These include differences in ripple frequency, the recruitment and modulation of CA1 and CA3 neurons, and reactivation content. If the two groups had simply arisen from noise or analysis artefacts, it would have been very unlikely to observe such consistent and functionally meaningful differences across so many independent measures. Third, in Chapters 3 and 4, I repeated many analyses across the full continuum — not only focusing on the extremes — by including baseline ripples. These baseline ripples were biased toward weak sinks in *radiatum* (see Figure 4.2E), and for all features used for cross-validation, baseline ripples always fell between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples, whether looking at ripple frequency, power, or waveform shape (see Figures 3.5 and 4.2). This strongly supports the idea that ripple current profiles exist along a continuum where features are gradually modulated, and it also validates my focus on the extremes. After all, if the middle ripples behave somewhere between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$, then studying the extremes helps to understand what sits on either end of this distribution.

In the future, it would be interesting to look more carefully at whether other features of baseline ripples also fall between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$. For example, are baseline ripples slightly biased to recruit superficial CA1 cells, but less so than Rad$^{\text{sink}}$ ripples? Is CA3 activity during baseline ripples stronger than in LM$^{\text{sink}}$, but not as strong as in Rad$^{\text{sink}}$? What about the joint activity of CA3 and superficial cells? This could help answer the question of whether a small current from CA3 into *radiatum* is enough to reactivate superficial CA1 and CA3 neurons, or whether stronger input is required to drive that reactivation. Exploring these questions would give further insight into how a small change in currents to CA1 can predict a substantial change in ripple dynamics within this continuum.

Biases in ripple detection: why LM$^{\text{sink}}$ ripples were previously overlooked?

Given the striking differences in ripple profiles and associated features, an immediate question arises: why have such distinct events, particularly LM$^{\text{sink}}$ ripples, remained largely undetected in previous work? In this section, I address how, on top of all the averaging issues discussed throughout this thesis, potential methodological biases may have led to overlooking these ripples.

First, LM$^{\text{sink}}$ ripples are the minority, while Rad$^{\text{sink}}$ ripples are much more common — hence why Rad$^{\text{sink}}$ ripples dominate the average profile. Given this, it is important to have long sleep sessions when designing the recording procedure; otherwise, short sleep sessions, simply due to sampling probability, are likely to be biased toward the more common ripple type, i.e., Rad$^{\text{sink}}$. To mitigate this bias in this thesis, all included rest/sleep sessions had a mean duration longer than 1 hour (see Section 2.2.2).

Another more subtle but very important point relates to how ripples are defined — particularly the distinction between ripples and sharp-wave ripples — and how studies focus on one or the other (See Section 1.2). In Chapter 3, I have shown and discussed the well-defined sharp-wave in *stratum radiatum* during Rad$^{\text{sink}}$ ripples and its absence during LM$^{\text{sink}}$ ripples. Many studies — including Sebastian et al., 2023, Chang et al., 2025, and Harvey et al., 2023 — focus specifically on sharp-wave ripples (SWRs), typically detecting them by requiring that putative CA1 ripple events are accompanied by a sharp-wave in *stratum radiatum* (see Sections 1.2 and 1.4). This naturally biases the detection towards Rad$^{\text{sink}}$ ripples, while missing LM$^{\text{sink}}$ ripples entirely.

Here, I avoided this bias by detecting ripples purely based on the LFP recorded from the CA1 pyramidal layer, without imposing any requirement on sharp-waves (see Section 2.2.2). Moreover, regarding the parameters used for ripple detection (see Section 2.2.2), if the frequency band chosen to detect ripples is too high — e.g., using 120 Hz or 130 Hz as the lower cut-off frequency, as in X. Liu et al., 2022 — this will bias detection toward high-frequency ripples, i.e., Rad$^{\text{sink}}$ ripples again. These aspects are critical to consider not only when trying to examine ripple variability, but also when placing the findings from this study in the context of previous work and understanding why LM$^{\text{sink}}$ ripples have been under-reported until now.

Interestingly, Sebastian et al. (2023) recently reported that a subset of ripples exhibited a current profile matching LM$^{\text{sink}}$ ripples, which was also associated with lower ripple frequency and distinct LFP waveforms. However, in that study, only 173 such events were detected — a relatively small number, which was probably biased by the detection method (see above). By contrast, in this thesis, the overall dataset comprised 942,331 ripples, with more than 200,000 LM$^{\text{sink}}$ ripples identified in the tetrode dataset alone (see Sections 2.2.2 and 4.3.3), providing a much more robust and comprehensive characterisation of distinct CA1 laminar profiles.

To summarise, it is possible that previous studies have (unintentionally) biased their experimental procedures and ripple detection methods toward Rad$^{\text{sink}}$ ripples, which may explain why LM$^{\text{sink}}$ ripples have been overlooked so far. The extensive characterisation of ripple diversity presented in this thesis could therefore not only put previous results into context but also serve as a useful reference for future studies aiming to minimise detection biases and preserve ripple variability.

Having established that LM$^{\text{sink}}$ ripples represent a distinct, yet previously under-recognised, set of hippocampal ripples, the next step is to explore which neural circuits might generate these distinct ripple profiles.

Potential hippocampal circuits for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples and future directions

In Sections 5.6 and 6.6, I have discussed potential hippocampal and entorhinal circuits involved in Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples. Summarising and bringing these ideas together, the canonical Rad$^{\text{sink}}$ ripples seem to be dominated by strong inputs from CA3 neurons that arrive in *stratum radiatum*, triggering large-scale recruitment of CA1 principal cells — both deep and superficial, especially superficial — firing at high rates and generating high-frequency ripples. In contrast, during LM$^{\text{sink}}$ ripples, strong currents arrive in *stratum lacunosum-moleculare*, and these events are associated with high inhibition within CA3, which likely prevents CA3 from firing strongly at the ripple peak or sending large sinks in *radiatum*. As a result, CA1 neurons may prioritise the *lacunosum-moleculare* currents which, in turn, trigger lower-frequency ripples formed by small coactivity motifs of mainly deep CA1 cells.

Thus, the widely accepted idea that CA3 generates ripples may primarily apply to a subset — namely, Rad$^{\text{sink}}$ ripples, which represent the majority of events. However, another subset, LM$^{\text{sink}}$ ripples, seems not to rely on CA3 as much, suggesting that they might originate elsewhere. In line with this, previous work showed that blocking CA3 outputs does not abolish ripple occurrence in CA1 but leads to slower-frequency ripples, resembling the LM$^{\text{sink}}$ profile (Nakashiba et al., 2009). Given that *stratum lacunosum-moleculare* receives prominent input from MEC3, these findings point to MEC3 as a likely candidate for driving LM$^{\text{sink}}$ ripples. Supporting this hypothesis, recent studies have demonstrated that entorhinal cortex inputs to CA1 can indeed modulate ripple characteristics (C. Liu et al., 2023; Yamamoto & Tonegawa, 2017; Zutshi & Buzsáki, 2023).

On top of the feed-forward input currents arriving in *stratum radiatum* or *stratum lacunosum-moleculare*, superficial interneurons within these layers might also play a critical role in gating information (Klausberger & Somogyi, 2008; Sakalar et al., 2022). One possible speculation is that, for instance, during Rad$^{\text{sink}}$ ripples, superficial interneurons in *stratum radiatum* are strongly recruited, while during LM$^{\text{sink}}$ ripples, interneurons within *stratum lacunosum-moleculare* might be more active. This local interneuron activity could mediate or bias CA1 pyramidal cells to be more 'attentive' to either Rad currents (i.e., CA3 input) or LM currents (e.g., MEC3 input). In this sense, it might be interesting — if not fundamental — for future works to examine the recruitment of these interneurons during Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples.

Surprisingly, no distinct ripple profiles with stronger current sinks in *stratum oriens* — potentially reflecting more CA2-driven ripples — emerged using our analytical framework (see Chapter 3). This might initially appear contradictory to previous studies reporting that CA2 contributes to ripple generation and shapes the ripple LFP waveform (Oliva et al., 2016; Sebastian et al., 2023). However, this is not unexpected when placed in the context of our experimental framework. Oliva et al. (2016), for instance, showed that CA2 contributes to ripple generation primarily during quiet wakefulness, whereas ripples during rest/sleep are mainly driven by CA3. All ripple-data analysed in this thesis — and perhaps a limitation of the study — come exclusively from sleep, not wakefulness (see Section 2). This could explain why no clear $\mathrm{Or}^{\mathrm{sink}}$-like profile emerged in our analyses; such events were likely merged into or diluted within the $\mathrm{Rad}^{\mathrm{sink}}$ cluster (see Chapter 3). A natural future direction for this work would be to explore CSD profiles during awake ripples and test whether a third, CA2-dependent profile emerges under those conditions. That said, it is important to highlight why, for a study focused on ripple diversity, using long sleep sessions was critical. To investigate variability meaningfully, a large sample of ripple events is required — otherwise, subsampling risks masking variability, as discussed earlier. Since ripple occurrence is significantly higher during nREM than wakefulness, this allowed me to avoid subsampling issues and establish a robust analytical framework to uncover ripple variability. Now that this framework has been established, future studies could apply it to awake ripples — without falling into the trap of subsampling biases when making claims about ripple diversity. In other words, we already know that ripples are diverse; the goal would now be to test how that diversity manifests specifically during awake states.

What about reactivation of waking coactivity patterns? A recent study has shown that CA1 superficial cells are the main contributors to offline reactivation and replay during ripples (Harvey et al., 2023). Again, this consensus view likely reflects $\mathrm{Rad}^{\mathrm{sink}}$ ripples, which indeed reactivate both deep and superficial CA1 cells. In contrast, $\mathrm{LM}^{\mathrm{sink}}$ ripples do not reactivate superficial CA1 cells nor CA3 principal cells. This is quite an interesting finding when considered alongside previous knowledge about CA3 and ripple-related reactivation. While the field has largely focused on the idea that ripples originate from CA3 and drive reactivation of superficial cells, I show here that a significant subset of ripples does not follow this mechanism. Although this thesis challenges some of the classical views on ripples, these findings should not be interpreted as saying that previous results are wrong. Rather, this work offers a way to think about ripple diversity that might help explain why certain features become diluted or masked when ripple variability is not taken into account.

An interesting future possibility for this line of work would be to assess causal involvement of these hippocampal circuits in generating $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples. There are multiple ways this could be tested, but it would likely focus on either CA3 (which is already shown to respond differently in the two ripple types and projects to *radiatum*) or on identifying who sends the input currents to *lacunosum-moleculare*. If focusing on CA3, one possible approach would be to optogenetically silence the CA3→*radiatum* pathway (i.e., Schaffer collaterals) during sleep and examine the effect on CA1 ripples. A possible experimental design would involve delivering light pulses to silence CA3 and then analysing whether the proportion of $\mathrm{Rad}^{\mathrm{sink}}$ ripples decreases compared to control sessions where the optogenetic virus was not expressed. As a control, it would also be important to examine whether the proportion of $\mathrm{LM}^{\mathrm{sink}}$ ripples changes. Based on the findings of this thesis, my hypothesis is that $\mathrm{Rad}^{\mathrm{sink}}$ ripples should be greatly reduced, while $\mathrm{LM}^{\mathrm{sink}}$ ripples would remain unaffected or only minimally impacted. This would directly test the role of CA3 in generating $\mathrm{Rad}^{\mathrm{sink}}$ ripples. If instead focusing on identifying the region driving $\mathrm{LM}^{\mathrm{sink}}$ ripples, a good candidate is MEC3. In that case, a similar experiment could be conducted to selectively silence MEC3 projections to CA1 *lacunosum-moleculare* and assess whether $\mathrm{LM}^{\mathrm{sink}}$ ripples are dependent on these inputs. Here, it would be critical to target MEC3 projections specifically, as unwanted spread of the virus to MEC2 could affect the trisynaptic pathway (MEC2→DG→CA3), potentially influencing $\mathrm{Rad}^{\mathrm{sink}}$ ripples as well. Another important point to consider is that, while MEC3 seems like the ideal candidate, other regions such as the nucleus reuniens of the thalamus also project to the *stratum lacunosum-moleculare* of CA1 (Wouterlood et al., 1990).

Above, I have described two open-loop intervention strategies to assess the circuits involved in generating Rad^{sink} and LM^{sink} ripples. Another possibility would be to perform closed-loop interventions to selectively silence one ripple type while preserving the other, and then assess the behavioural consequences in a memory task. This would require a real-time, closed-loop system capable of classifying ripples online as either Rad^{sink} or LM^{sink}. One possibility would be to take advantage of the very distinct LFP waveforms in *lacunosum-moleculare* between the two ripple types. For example, a tetrode or channel placed in LM could be used to detect a sharp-wave-like negative deflection characteristic of LM^{sink} ripples. The logic of a controller aimed at silencing LM^{sink} ripples would be as follows: if a ripple is detected in the CA1 pyramidal layer and a large negative deflection is simultaneously detected in LM, then CA1 principal cells would be silenced. Conversely, to silence Rad^{sink} ripples, the same system could be used but by inverting the detection logic — i.e., silencing if no sharp-wave-like deflection is detected in LM.

To assess behavioural effects, appropriate memory tasks should be used. For example, Chang et al. (2025) recently developed a closed-loop system that silences ripples based on pupil size and tested its effects on novel versus familiar memory using a chessboard-maze task. Along similar lines, a future study could adopt this approach, selectively silencing Rad^{sink} (which are biased toward reactivating recent motifs) or LM^{sink} ripples (biased toward prior motifs), and test whether this differentially affects novel versus familiar memory recall in a task like the chessboard maze. However, such a study would likely be substantial on its own, requiring significant development of closed-loop classification systems, extensive experimentation, and careful analysis to account for potential technical limitations or errors.

Ripple diversity as a mechanism for balancing stability and flexibility

Beyond understanding which circuits generate these ripple profiles, a fundamental question remains: why would the hippocampus engage such distinct forms of population activity during offline periods? In this final section, I propose that ripple diversity reflects a network mechanism for balancing the need to integrate new information while preserving stable internal representations.

Why does the hippocampus co-process offline rapid integration and gradual stabilisation of recent and prior population responses, respectively? While the prompt integration of new information is essential, maintaining a stable underlying population activity structure could yet provide a “backbone” upon which finer-grain information processing can be built (Chandra et al., 2025). Recent theoretical work suggests that the hippocampus integrates new sensory information using a pre-structured internal scaffold provided by MEC, enabling scalable, flexible, and efficient memory storage (Chandra et al., 2025). In line with this, the core motifs of deep cells associated here with $\mathrm{LM}^{\mathrm{sink}}$ ripples would define such a pre-structured scaffold. Recent waking experience would then flexibly recruit coactivity motifs of CA1 superficial cells within the network. Through CA3-driven activity, the network is transiently pushed out of the prior neural backbone to append new information. Combining deep and superficial cells would allow enhancing the specificity and distinctiveness of newly-acquired representations. The resulting composite population patterns would be later reactivated offline during $\mathrm{Rad}^{\mathrm{sink}}$ ripples. Meanwhile, $\mathrm{LM}^{\mathrm{sink}}$ ripples would continue expressing core activity motifs, which gradually drift to update the population backbone (Figure 7.2). Deep and superficial cells would therefore work synergistically towards the concomitant rapid integration and gradual refinement of hippocampal population representations.

Moreover, CA1 deep cells establish important projections to neocortical areas (prefrontal cortex) (Harvey et al., 2023). The gradual drift of the hippocampal backbone, evolving from prior to recent coactivity spaces over hour-long post-exploration sleep/rest, could exert important influence on systems consolidation of neuronal ensembles in downstream neocortical circuits (e.g., prefrontal cortex (Harvey et al., 2023)). In parallel, this gradual drift would also update the local neuronal priors within the hippocampus, which would then serve as the foundation for the next low-dimensional backbone in LM^{sink} ripples, while Rad^{sink} ripples continue to push the network out of this state through superficial cells, with high-dimensional patterns reporting recent wakefulness. Such a functional loop could support a network trade-off between stability and flexibility. Figure 7.2 illustrates this proposed dynamic between recent and prior motifs across sleep and wake.

Why might the hippocampus preserve and reactivate a backbone of prior population motifs, rather than focusing solely on consolidating recent experiences? Maintaining such a structured backbone could provide a stable scaffold that facilitates the integration of new information without overwriting existing representations. By reusing and gradually updating population patterns, the hippocampus may avoid the energetic cost of generating entirely new high-dimensional activity for each new experience. Organising distinct waking experiences within a shared *schema* could thus support flexible computations and reduce energy demand compared to a system that requires *de novo* creation of entirely new high-dimensional population patterns from the outset each time (Chandra et al., 2025). In this framework, LM^{sink} ripples and deep cells might promote an energy-efficient and flexible mechanism of memory consolidation, through which this *schema* is preserved and onto which new information is integrated by CA3 activity and superficial cells during Rad^{sink} ripples.

In conclusion, in this thesis I propose that ripple-by-ripple diversity in CA1 – achieved by tuning laminar current levels – allows the hippocampus to flexibly integrate new experiences during offline periods (sleep/rest) while preserving a stable internal scaffold. This mechanism, involving differences in CA1/CA3 network activity, structural organisation, and the neuronal composition of population patterns, allows the hippocampus to flexibly integrate new experiences offline

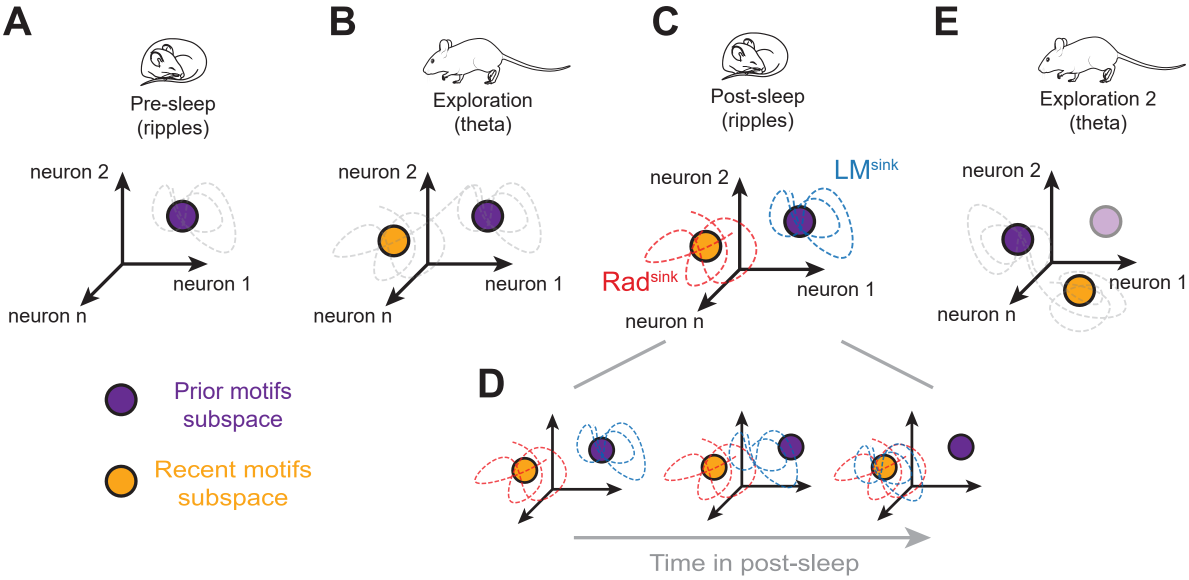

Figure 7.2: Schematic of recent and prior motif dynamics across sleep and wake.
(A–E) Schematic illustrating a possible interpretation of how recent and prior coactivity motifs may organise and shift within neural space across sleep and wake. Each schematic (e.g., A, B) represents the neural space, where each axis corresponds to the activity of an individual neuron (e.g., neurons 1, 2, n). (A) During pre-sleep (ripples), CA1 neurons express coactivity motifs whose neural trajectory (dashed line) gravitates around the 'backbone' subspace of prior motifs (purple circle). (B) During wakefulness (theta cycles), recent motifs are integrated into the network, pushing the neural trajectory toward a recent subspace (orange circle), while still intermittently sampling prior motifs. (C) In post-sleep (ripples), Rad^{sink} ripples predominantly express recent motifs, while LM^{sink} ripples initially express coactivity motifs aligned with the prior motif subspace, which then drift toward recent motifs (D). (E) In a subsequent exploration session, the previously recent subspace becomes the new 'backbone' (purple circle), and the most recent experience pushes the network toward a new 'recent' subspace. The old backbone is represented as a pale purple circle.

while preserving a stable internal scaffold, supporting a dynamic balance between memory stability and adaptability.

Appendices

A

Evaluating LDA ripple discrimination across CA1 depth

Contents

A.1 Introduction

In Chapter 4, I demonstrated that the low-frequency components (i.e., the sharp wave–like deflections) of the *pyramidale* LFP waveforms during ripples can be used to significantly distinguish between $\mathrm{Rad}^{\mathrm{sink}}$, baseline, and $\mathrm{LM}^{\mathrm{sink}}$ ripples. This was a key methodological advance in my thesis, and I leverage this classification model to examine how neurons respond to these different ripple profiles in tetrode datasets in Chapters 5 and 6.

However, as discussed in Chapter 4 and in previous studies (Lopes-dos-Santos et al., 2023), the LFP waveform during ripples — specifically its low-frequency component — also varies systematically with depth within the pyramidal layer (Figure A.1(A)). This depth-dependent variation is a well-known feature of CA1 LFPs and is one I also leverage to classify deep and superficial cells (see Section 6.2.1). In this context, one might ask whether the performance of the classifier is biased depending on the depth from which the ripple waveform is sampled. In other words, if the waveform is taken from a channel located closer to *oriens* rather than one closer to *radiatum*, does this increase the likelihood of false positives when classifying Rad^{sink}, baseline, and LM^{sink} ripples?

This consideration becomes particularly relevant when applying the model to tetrode datasets: which tetrode should be selected for classification? In the discussion of Chapter 4, I outlined a strategy to mitigate these depth-related biases. Specifically, I trained the model on the channel located at the centre of the pyramidal layer in the silicon probe recordings, and then, for tetrode data, selected the tetrodes whose ripple waveforms most closely resembled the template waveforms used to train the LDA classifier.

To further assess whether the ripple type classifier used throughout this thesis is robust to depth-dependent biases, in this chapter I examine how accurately and consistently the LDA model presented in Chapter 4 can distinguish ripple types across different depths in the silicon probe dataset.

A.2 Methods

To evaluate how well the ripple classifier could distinguish Rad^{sink}, baseline, and LM^{sink} ripples as a function of depth within the CA1 pyramidal layer, I used the pre-trained LDA model described in Section 4.2.3 and applied it to seven different channels within the pyramidal layer in the silicon probe dataset (38 sleep sessions from 5 mice; see Section 2.1.3). The central channel corresponded to the one used for training the model, and I selected three channels above and three channels below it for comparison (see Figure A.1(A)). Using the known spacing between channels in each probe (see Table 2.1), I estimated model performance as a function of distance from the pyramidal layer centre for each sleep session.

Since recordings from silicon probes allowed me to obtain ground truth ripple labels from the CSD profiles, I evaluated classification performance by applying the trained LDA model to ripple waveforms from each of the seven channels. For each mouse, I computed: (i) the overall classification accuracy across all ripple types (Figure A.1(B)), and (ii) the per-class accuracy, such as the proportion of correctly predicted Rad$^{\text{sink}}$ ripples (Figure A.1(C)).

Additionally, to assess whether the model could significantly distinguish between the two extreme ripple types — Rad$^{\text{sink}}$ versus LM$^{\text{sink}}$ — regardless of baseline ripples, I computed a discrimination index defined as:

$$\text{Discrimination Index} = \min\left(\frac{\text{True}_{LM^{\text{sink}}}}{\text{True}_{LM^{\text{sink}}} + \text{Misclassified}_{LM^{\text{sink}} \to Rad^{\text{sink}}}}, \frac{\text{True}_{Rad^{\text{sink}}}}{\text{True}_{Rad^{\text{sink}}} + \text{Misclassified}_{Rad^{\text{sink}} \to LM^{\text{sink}}}}\right) \tag{A.1}$$

where True_y denotes the number of correctly classified ripple events of type y, and $\text{Misclassified}_{y \to x}$ indicates the number of ripple events from class y that were misclassified as class x. The ratio $\frac{\text{True}_y}{\text{True}_y + \text{Misclassified}_{y \to x}}$ thus represents the model's accuracy in distinguishing ripple type y from the opposing class x. A value of 1 indicates perfect discrimination, with no confusion between the two classes, while a value near 0 indicates a complete failure to distinguish them. Because baseline ripples were excluded for this analysis, the chance level for this metric is 0.5. As shown in Equation A.1, I computed this ratio separately for Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples and, for each sleep session, defined the discrimination index as the minimum of the two class-specific ratios — that is, the one closer to chance. This index is shown in Figure A.1(D).

A.3 Depth-specific evaluation of the LDA ripple classifier

To assess whether the accuracy and predictions of the LDA ripple classifier — trained on the 'sharp-wave' component of CA1 *stratum pyramidale* ripple waveforms — were biased by the depth of the recording site, I extracted LFP signals from silicon probe channels located within a $\pm 60\,\mu$m range around the pyramidal layer centre (i.e., the reference channel used to train the model). As expected from previous studies, LFP waveforms showed systematic variations as a function of depth (Figure A.1(A); Lopes-dos-Santos et al., 2023).

To evaluate whether these waveform differences influenced classifier performance, I assessed the accuracy of the LDA model independently for each sleep session using individual channels within the depth window. For each depth, I measured how well the previously trained model could discriminate between $\mathrm{Rad}^{\mathrm{sink}}$, baseline, and $\mathrm{LM}^{\mathrm{sink}}$ ripples based on the LFP waveform from that specific channel, yielding an accuracy–depth profile for each session (Figure A.1(B)). Across channels and sessions, the LDA consistently classified ripple types significantly above chance. Importantly, classification accuracy did not significantly differ across depths, indicating that the model's performance was robust to LFP variability along the pyramidal layer.

Next, I investigated whether depth affected the classification accuracy of individual ripple types (Figure A.1(C)). All ripple types were classified significantly above chance across depths. A weak trend emerged whereby $\mathrm{LM}^{\mathrm{sink}}$ ripples were slightly better classified in channels closer to *stratum oriens*, while $\mathrm{Rad}^{\mathrm{sink}}$ ripples were marginally better predicted in channels closer to *stratum radiatum.* In contrast, baseline ripple accuracy remained stable across depths.

To examine whether the depth-dependent trends observed in $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ classification were driven by misclassification between the two ripple types (e.g., $\mathrm{Rad}^{\mathrm{sink}}$ classified as $\mathrm{LM}^{\mathrm{sink}}$ in channels closer to *oriens*), I computed the LDA discrimination ratio between $\mathrm{Rad}^{\mathrm{sink}}$ and $\mathrm{LM}^{\mathrm{sink}}$ ripples as a function of depth (Figure A.1(D)). Across all depths, the LDA significantly distinguished $\mathrm{Rad}^{\mathrm{sink}}$

from LM$^{\text{sink}}$ ripples, and no significant differences in discriminability were observed across channels. This suggests that the observed depth trends in Figure A.1(C) were primarily due to either ripple type being misclassified as baseline, rather than misclassification between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ themselves.

Together, these results confirm that the LDA classifier presented in Chapter 4 reliably distinguishes ripple types based on LFP waveforms across different depths of the pyramidal layer. This indicates that the classification is driven by consistent waveform differences across ripple types, and not confounded by depth-related LFP variability. Moreover, the classifier consistently discriminated between the two ripple extremes (Rad$^{\text{sink}}$ vs. LM$^{\text{sink}}$), validating its use in the analyses presented in Chapters 5 and 6.

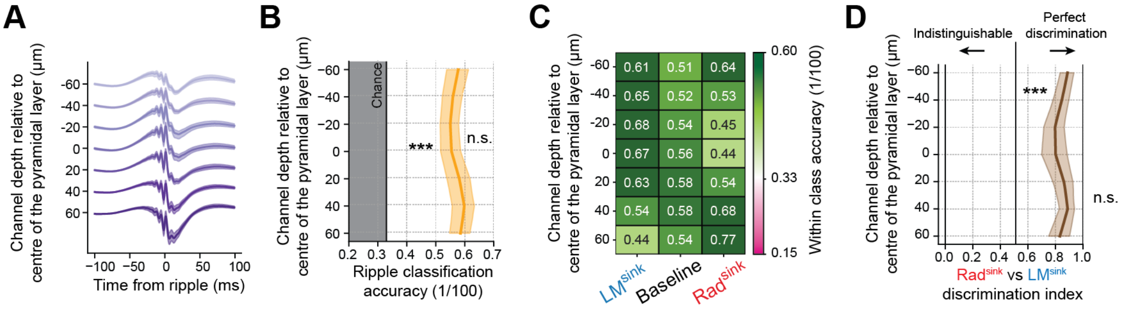

Figure A.1: Ripple classification using LDA across different depths of the pyramidal layer.
(A) Ripple LFP waveforms recorded at different depths relative to the centre of the pyramidal layer. *Shaded areas* around each trace represent the 95% confidence interval across sleep sessions.
(B) Overall LDA classification accuracy as a function of depth. The classifier predicted ripple types significantly above chance level at all depths ($p < 10^{-5}$; one-tailed bootstrap test; shaded grey area indicates chance level), with no significant difference in accuracy across depths ($p = 0.518$; one-way ANOVA with Tukey post-hoc test).
(C) Heatmap showing LDA classification accuracy for each ripple class (LM$^{\text{sink}}$, baseline, Rad$^{\text{sink}}$) as a function of depth from the centre of the pyramidal layer. Each entry reflects the mean accuracy across all sleep sessions. Green indicates above-chance performance (pink would indicate chance). Note the subtle depth-dependent trends: LM$^{\text{sink}}$ ripples were better predicted closer to *oriens*, while Rad$^{\text{sink}}$ ripples were better predicted near *radiatum*.
(D) Discrimination index between Rad$^{\text{sink}}$ and LM$^{\text{sink}}$ ripples as a function of depth. The LDA significantly distinguished between the two ripple types across depths ($p < 10^{-5}$; one-tailed bootstrap test), with no significant differences in performance across channels ($p = 0.446$; one-way ANOVA with Tukey post-hoc test).
For (A)-(D): $n = 54$ sleep sessions. $^{***}p < 0.001$.

A.4 Discussion

In this chapter, I demonstrated that the LDA ripple classifier used throughout this thesis reliably distinguished ripple types, regardless of the recording depth within the pyramidal layer. This suggests that, independent of the specific tetrode channel used for prediction, the model consistently discriminated between ripple groups.

Specifically, overall classification accuracy did not significantly differ across different depths within the pyramidal layer. Only a weak trend was observed, with Rad^{sink} ripples being better predicted closer to *stratum radiatum*, and LM^{sink} ripples closer to *stratum oriens*. However, this effect was primarily driven by events being misclassified as baseline ripples, rather than mislabelled as the opposing group. Importantly, the LDA consistently discriminated the extreme cases of Rad^{sink} and LM^{sink} ripples, which is particularly relevant for Chapters 5–6, where I focused exclusively on these two ripple types in the tetrode dataset. This reinforces that the two groups analysed there represent true positives.

As discussed in the Chapter 4 discussion, if systematic noise or bias in classification had been introduced due to tetrode positioning across days, one would expect considerable noise in the results. The fact that clear and consistent group differences still emerged suggests that the findings are robust and unlikely to have arisen in the presence of such noise.

It is also worth noting that, to analyse extreme cases, I included channels located $\pm 60\mu\text{m}$ from the centre of the pyramidal layer. Given that the overall thickness of the pyramidal layer is approximately $100\mu\text{m}$ (Lopes-dos-Santos et al., 2023), these extreme cases lie near the boundaries or slightly outside of the layer. Despite this, the classifier still performed well, indicating that ripple waveform features are preserved across different depths. This suggests a possible future direction: to train separate LDAs at different depths and use the classifier that matches the estimated depth of a given tetrode or channel for prediction.

Finally, while these results confirm the robustness of the classifier across depths, they also reinforce the importance of using the tetrode closest to the silicon probe template for ripple classification. In practice, this corresponds to using the tetrode nearest to the centre of the pyramidal layer, thereby minimising potential depth-related biases.

References

Barron, H. C., Reeve, H. M., Koolschijn, R. S., Perestenko, P. V., Shpektor, A., Nili, H., Rothaermel, R., Campo-Urriza, N., O'Reilly, J. X., & Bannerman, D. M. e. a. (2020). Neuronal computation underlying inferential reasoning in humans and mice. *Cell*, *183*, 228–243.e21.

Belluscio, M. A., Mizuseki, K., Schmidt, R., Kempter, R., & Buzsáki, G. (2012). Cross-frequency phase–phase coupling between theta and gamma oscillations in the hippocampus. *Journal of Neuroscience*, *32*, 423–435.

Berndt, M., Trusel, M., Roberts, T. F., Pfeiffer, B. E., & Volk, L. J. (2023). Bidirectional synaptic changes in deep and superficial hippocampal neurons following in vivo activity. *Neuron*, *111*, 2984–2994.e4.

Box, G. E. P. (1979). Robustness in the strategy of scientific model building. In R. L. Launer & G. N. Wilkinson (Eds.), *Robustness in statistics* (pp. 201–236). Academic Press.

Bragin, A., Jandó, G., Nádasdy, Z., van Landeghem, M., & Buzsáki, G. (1995). Dentate eeg spikes and associated interneuronal population bursts in the hippocampal hilar region of the rat. *Journal of Neurophysiology*, *73*(4), 1691–1705.

Bragin, A., Engel, J. J., Wilson, C. L., Fried, I., & Buzsáki, G. (1999). High-frequency oscillations in human brain. *Hippocampus*, *9*, 137–142.

Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, *25*, 1073–1188.

Buzsáki, G., Leung, L. W., & Vanderwolf, C. H. (1983). Cellular bases of hippocampal eeg in the behaving rat. *Brain Research*, *287*, 139–171.

Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Research*, *398*, 242–252.

Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, *31*(3), 551–570.

Buzsáki, G. (2006). *Rhythms of the brain*. Oxford University Press.

Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*(5679), 1926–1929.

Buzsáki, G., Horvath, Z., Urioste, R., Hetke, J., & Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science*, *256*, 1025–1027.

Buzsáki, G., Logothetis, N., & Singer, W. (2013). Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. *Neuron*, *80*, 751–764.

Buzsáki, G., & Moser, E. I. (2013). Memory, navigation and theta rhythm in the hippocampal-entorhinal system. *Nature Neuroscience*, *16*(2), 130–138.

Castelli, M. (2025). Mcastelli98/dentatespikeclassifier: Dentate spike classifier v2 [Software].

Castelli, M., Lopes-dos-Santos, V., Gava, G. P., Lambiotte, R., & Dupret, D. (2025). Hippocampal ripple diversity organises neuronal reactivation dynamics in the offline brain. *bioRxiv*.

Cavalieri, D., Angelova, A., Islah, A., Lopez, C., Bocchio, M., Bollmann, Y., Baude, A., & Cossart, R. (2021). Ca1 pyramidal cell diversity is rooted in the time of neurogenesis. *eLife*, *10*, e69270.

Chandra, S., Sharma, S., Chaudhuri, R., & Fiete, I. (2025). Episodic and associative memory from spatial scaffolds in the hippocampus. *Nature*, 1–13.

Chang, H., Tang, W., Wulf, A. M., Nyasulu, T., Wolf, M. E., Fernandez-Ruiz, A., & Oliva, A. (2025). Sleep microstructure organizes memory replay. *Nature*, 1–9.

Colgin, L. L. (2016). Rhythms of the hippocampal network. *Nature Reviews Neuroscience*, *17*, 239–249.

Colgin, L. L., Kubota, D., Jia, Y. S., Rex, C. S., & Lynch, G. (2004). Long-term potentiation is impaired in rat hippocampal slices that produce spontaneous sharp waves. *Journal of Physiology*, *558*, 953–961.

Deering, R., & Kaiser, J. F. (2005). The use of a masking signal to improve empirical mode decomposition. *ICASSP '05. IEEE International Conference on Acoustics, Speech, and Signal Processing*, 485–488.

Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*, 1241–1242.

Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance [Publisher: Nature Publishing Group]. *Nature Neuroscience*, *13*(8), 995–1002.

Dusek, J. A., & Eichenbaum, H. (1997). The hippocampus and memory for orderly stimulus relations [Publisher: Proceedings of the National Academy of Sciences]. *Proceedings of the National Academy of Sciences*, *94*(13), 7109–7114.

Dvorak, D., Chung, A., Park, E. H., & Fenton, A. A. (2021). Dentate spikes and external control of hippocampal function. *Cell Reports*, *36*(5), 109497.

Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, *20*, 1–10.

Eichenbaum, H., Dudchenko, P., Wood, E., Shapiro, M., & Tanila, H. (1999). The hippocampus, memory, and place cells: Is it spatial memory or a memory space? *Neuron*, *23*(2), 209–226.

English, D. F., Peyrache, A., Stark, E., Roux, L., Vallentin, D., Long, M. A., & Buzsáki, G. (2014). Excitation and Inhibition Compete to Control Spiking during Hippocampal Ripples: Intracellular Study in Behaving Mice [Publisher: Society for Neuroscience Section: Articles]. *Journal of Neuroscience*, *34*(49), 16509–16517.

Esparza, J., Quintanilla, J. P., Cid, E., Medeiros, A. C., Gallego, J. A., & Prida, L. M. d. l. (2025). Cell-type-specific manifold analysis discloses independent geometric transformations in the hippocampal spatial code [Publisher: Elsevier]. *Neuron*, *0*(0).

Farrell, J. S., Hwaun, E., Dudok, B., & Soltesz, I. (2024). Neural and behavioural state switching during hippocampal dentate spikes. *Nature*, *628*(8008), 590–595.

Fernandez-Ruiz, A., Sirota, A., Lopes-Dos-Santos, V., & Dupret, D. (2023). Over and above frequency: Gamma oscillations as units of neural circuit operations. *Neuron*, *111*(7), 936–953.

Fernández-Ruiz, A., Oliva, A., Nagy, G. A., Maurer, A. P., Berényi, A., & Buzsáki, G. (2017). Entorhinal-ca3 dual-input control of spike timing in the hippocampus by theta-gamma coupling. *Neuron*, *93*, 1213–1226.e5.

Fernández-Ruiz, A., Oliva, A., Oliveira, E. F. d., Rocha-Almeida, F., Tingley, D., & Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. *Science*, *364*, 1082–1086.

Fosso, O. B., & Molinas, M. (2018). Emd mode mixing separation of signals with close spectral proximity in smart grids. *2018 IEEE PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe)*, 1–6.

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

Gava, G. P., Lefèvre, L., Broadbelt, T., McHugh, S. B., Lopes-Dos-Santos, V., Brizee, D., Hartwich, K., Sjoberg, H., Perestenko, P. V., & Toth, R. e. a. (2024). Organizing the coactivity structure of the hippocampus from robust to flexible memory. *Science*, *385*, 1120–1127.

Gava, G. P., McHugh, S. B., Lefèvre, L., Lopes-dos-Santos, V., Trouche, S., El-Gaby, M., Schultz, S. R., & Dupret, D. (2021). Integrating new memories into the hippocampal network activity space. *Nature Neuroscience*, *24*, 326–330.

Geiller, T., Fattahi, M., Choi, J.-S., & Royer, S. (2017). Place cells are more strongly tied to landmarks in deep than in superficial ca1. *Nature Communications*, *8*, 14531.

Gini, C. (1921). Measurement of inequality of incomes. *The Economic Journal*, *31*, 124–126.

Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, *12*, 1222–1223.

Girardeau, G., & Lopes-dos-Santos, V. (2021). Brain neural patterns and the memory function of sleep. *Science*, *374*(6567), 560–564.

Grosmark, A. D., & Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*, *351*, 1440–1443.

Hagberg, A. A., Schult, D. A., & Swart, P. J. (2008). Exploring network structure, dynamics, and function using networkx. *Proceedings of the 7th Python in Science Conference*, 11–15.

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., & Smith, N. J. e. a. (2020). Array programming with numpy. *Nature*, *585*, 357–362.

Harvey, R. E., Robinson, H. L., Liu, C., Oliva, A., & Fernandez-Ruiz, A. (2023). Hippocampo-cortical circuits for selective memory encoding, routing, and replay. *Neuron*, *111*, 2076–2090.e9.

Huang, N. E., Shen, Z., Long, S. R., Wu, M. C., Shih, H. H., Zheng, Q., Yen, N.-C., Tung, C. C., & Liu, H. H. (1998). The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, *454*, 903–995.

Hunter, J. D. (2007). Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, *9*, 90–95.

Hurley, N., & Rickard, S. (2009). Comparing measures of sparsity. *IEEE Transactions on Information Theory*, *55*, 4723–4741.

Imbrosci, B., Nitzan, N., McKenzie, S., Donoso, J. R., Swaminathan, A., Böhm, C., Maier, N., & Schmitz, D. (2021). Subiculum as a generator of sharp wave-ripples in the rodent hippocampus. *Cell Reports*, *35*(3).

Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep [Publisher: Nature Publishing Group]. *Nature Neuroscience*, *10*(1), 100–107.

Joo, H. R., & Frank, L. M. (2018). The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation. *Nature reviews. Neuroscience*, *19*(12), 744–744.

Klausberger, T., & Somogyi, P. (2008). Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations. *Science*, *321*(5885), 53–57.

Klinzing, J., Niethard, N., & Born, J. (2019). Mechanisms of systems memory consolidation during sleep. *Nature Neuroscience*, *22*, 1598–1610.

Lee, S.-H., Marchionni, I., Bezaire, M., Varga, C., Danielson, N., Lovett-Barron, M., Losonczy, A., & Soltesz, I. (2014). Parvalbumin-positive basket cells differentiate among hippocampal pyramidal cells. *Neuron*, *82*, 1129–1144.

Liu, A. A., Henin, S., Abbaspoor, S., Bragin, A., Buffalo, E. A., Farrell, J. S., Foster, D. J., Frank, L. M., Gedankien, T., Gotman, J., Guidera, J. A., Hoffman, K. L., Jacobs, J., Kahana, M. J., Li, L., Liao, Z., Lin, J. J., Losonczy, A., Malach, R., … Buzsáki, G. (2022). A consensus statement on detection of hippocampal sharp wave ripples and differentiation from other fast oscillations. *Nature Communications 2022 13:1*, *13*(1), 1–14.

Liu, C., Todorova, R., Tang, W., Oliva, A., & Fernandez-Ruiz, A. (2023). Associative and predictive hippocampal codes support memory-guided behaviors. *Science*, *382*, eadi8237.

Liu, X., Terada, S., Ramezani, M., Kim, J. H., Lu, Y., Grosmark, A., Losonczy, A., & Kuzum, D. (2022). E-Cannula reveals anatomical diversity in sharp-wave ripples as a driver for the recruitment of distinct hippocampal assemblies. *Cell Reports*, *41*(1), 111453–111453.

Lopes-dos-Santos, V., Brizee, D., & Dupret, D. (2023). Spatio-temporal organization of network activity patterns in the hippocampus. *bioRxiv*.

Lopes-dos-Santos, V., van de Ven, G. M., Morley, A., Trouche, S., Campo-Urriza, N., & Dupret, D. (2018). Parsing hippocampal theta oscillations by nested spectral components during spatial exploration and memory-guided behavior. *Neuron*, *100*, 940–952.e7.

Magland, J., Jun, J. J., Lovero, E., Morley, A. J., Hurwitz, C. L., Buccino, A. P., Garcia, S., & Barnett, A. H. (2020). Spikeforest, reproducible web-facing ground-truth validation of automated neural spike sorters. *eLife*, *9*.

McHugh, S. B., Lopes-dos-Santos, V., Gava, G. P., Hartwich, K., Tam, S. K. E., Bannerman, D. M., & Dupret, D. (2022). Adult-born dentate granule cells promote hippocampal population sparsity. *Nature Neuroscience*, *25*, 1481–1491.

McHugh, S. B., Lopes-dos-Santos, V., Castelli, M., Gava, G. P., Thompson, S. E., Tam, S. K., & Dupret, D. (2024). Offline hippocampal reactivation during dentate spikes supports flexible memory. *Neuron*, *112*(22), 3768–3781.

McKinney, W. (2010). Data structures for statistical computing in python. *Proceedings of the 9th Python in Science Conference (SciPy)*, 51–56.

Mitzdorf, U. (1985). Current source-density method and application in cat cerebral cortex: Investigation of evoked potentials and eeg phenomena. *Physiological Reviews*, *65*, 37–100.

Mizuseki, K., Diba, K., Pastalkova, E., & Buzsáki, G. (2011). Hippocampal ca1 pyramidal cells form functionally distinct sublayers. *Nature Neuroscience*, *14*, 1174–1181.

Nádasdy, Z., Hirase, H., Czurko, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*, *19*, 9497–9507.

Nakashiba, T., Buhl, D. L., McHugh, T. J., & Tonegawa, S. (2009). Hippocampal ca3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron*, *62*, 781–787.

Navas-Olive, A., Valero, M., Jurado-Parras, T., de Salas-Quiroga, A., Averkin, R. G., Gambino, G., Cid, E., & de la Prida, L. M. (2020). Multimodal determinants of phase-locked dynamics across deep-superficial hippocampal sublayers during theta oscillations. *Nature Communications*, *11*, 2217.

Nitzan, N., Swanson, R., Schmitz, D., & Buzsáki, G. (2022). Brain-wide interactions during hippocampal sharp wave ripples. *Proceedings of the National Academy of Sciences*, *119*, e2200931119.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.

Oliva, A., Fernández-Ruiz, A., Buzsáki, G., & Berényi, A. (2016). Role of hippocampal ca2 region in triggering sharp-wave ripples. *Neuron*, *91*(6), 1342–1355.

O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, *11*, 209–215.

Pachitariu, M., Steinmetz, N. A., Kadir, S. N., Carandini, M., & Harris, K. D. (2016). Fast and accurate spike sorting of high-channel count probes with kilosort. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, & R. Garnett (Eds.), *Advances in neural information processing systems 29* (pp. 4448–4456). Curran Associates, Inc.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., & Dubourg, V. e. a. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Peyrache, A., Battaglia, F. P., & Destexhe, A. (2011). Inhibition recruitment in prefrontal cortex during sleep spindles and gating of hippocampal inputs. *Proceedings of the National Academy of Sciences USA*, *108*, 17207–17212.

Quinn, A., Lopes-dos-Santos, V., Dupret, D., Nobre, A., & Woolrich, M. (2021). Emd: Empirical mode decomposition and hilbert-huang spectral analyses in python. *Journal of Open Source Software*, *6*, 2977.

Quinn, A., Lopes-dos-Santos, V., Huang, N., Liang, W.-K., Juan, C.-H., Yeh, J.-R., Nobre, A. C., Dupret, D., & Woolrich, M. W. (2021). Within-cycle instantaneous frequency profiles report oscillatory waveform dynamics. *Journal of Neurophysiology*, *126*, 1190–1208.

Ramirez-Villegas, J. F., Logothetis, N. K., & Besserve, M. (2015). Diversity of sharp-wave–ripple lfp signatures reveals differentiated brain-wide dynamical events. *Proceedings of the National Academy of Sciences*, *112*, E6379–E6387.

Roux, L., Hu, B., Eichler, R., Stark, E., & Buzsáki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nature Neuroscience*, *20*, 845–853.

Sakalar, E., et al. (2022). Neurogliaform cells dynamically decouple neuronal synchrony between brain areas. *Science*, *377*, 324–328.

Scoville, W. B., & Milner, B. (2000). Loss of recent memory after bilateral hippocampal lesions. 1957. *Journal of Neuropsychiatry and Clinical Neurosciences*, *12*(1), 103–113.

Seabold, S., & Perktold, J. (2010). Statsmodels: Econometric and statistical modeling with python. *SciPy Conference*.

Sebastian, E. R., Esparza, J., & de la Prida, L. M. (2024). Quantifying the distribution of feature values over data represented in arbitrary dimensional spaces. *PLOS Computational Biology*, *20*, e1011768.

Sebastian, E. R., Quintanilla, J. P., Sánchez-Aguilera, A., Esparza, J., Cid, E., & de la Prida, L. M. (2023). Topological analysis of sharp-wave ripple waveforms reveals input mechanisms behind feature variations. *Nature Neuroscience*, *26*, 2171–2181.

Sirota, A., Csicsvari, J., Buhl, D. L., & Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences USA*, *100*, 2065–2069.

Slomianka, L., Amrein, I., Knuesel, I., Sørensen, J. C., & Wolfer, D. P. (2011). Hippocampal pyramidal cells: The reemergence of cortical lamination. *Brain Structure and Function*, *216*, 301–317.

Soltesz, I., & Losonczy, A. (2018). Ca1 pyramidal cell diversity enabling parallel information processing in the hippocampus. *Nature Neuroscience*, *21*, 484–493.

Stark, E., Roux, L., Eichler, R., Senzai, Y., Royer, S., & Buzsáki, G. (2014). Pyramidal cell-interneuron interactions underlie hippocampal ripple oscillations. *Neuron*, *83*, 467–480.

Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., & Buzsáki, G. (2011). Relationships between Hippocampal Sharp Waves, Ripples, and Fast Gamma Oscillation: Influence of Dentate and Entorhinal Cortical Activity. *Journal of Neuroscience*, *31*(23), 8605–8616.

Tavares, R. M., Mendelsohn, A., Grossman, Y., Williams, C. H., Shapiro, M., Trope, Y., & Schiller, D. (2015). A Map for Social Navigation in the Human Brain. *Neuron*, *87*(1), 231–243.

Thordsen, E., & Schubert, E. (2022). Abid: Angle based intrinsic dimensionality — theory and analysis. *Information Systems*, *108*, 101989.

Tingley, D., McClain, K., Kaya, E., Carpenter, J., & Buzsáki, G. (2021). A metabolic function of the hippocampal sharp wave-ripple. *Nature*, 1–5.

Valero, M., Averkin, R. G., Fernandez-Lamo, I., Aguilar, J., Lopez-Pigozzi, D., Brotons-Mas, J. R., Cid, E., Tamas, G., & Menendez de la Prida, L. (2017). Mechanisms for selective single-cell reactivation during offline sharp-wave ripples and their distortion by fast ripples. *Neuron*, *94*, 1234–1247.e7.

Valero, M., Cid, E., Averkin, R. G., Aguilar, J., Sanchez-Aguilera, A., Viney, T. J., Gomez-Dominguez, D., Bellistri, E., & de la Prida, L. M. (2015). Determinants of different deep and superficial ca1 pyramidal cell dynamics during sharp-wave ripples. *Nature Neuroscience*, *18*, 1281–1290.

Valero, M., & de la Prida, L. M. (2018). The hippocampus in depth: A sublayer-specific perspective of entorhinal–hippocampal function. *Current Opinion in Neurobiology*, *52*, 107–114.

van de Ven, G. M., Trouche, S., McNamara, C. G., Allen, K., & Dupret, D. (2016). Hippocampal offline reactivation consolidates recently formed cell assembly patterns during sharp wave-ripples. *Neuron*, *92*, 968–974.

Vandecasteele, M., Varga, V., Berényi, A., Papp, E., Bartho, P., Venance, L., Freund, T. F., & Buzsáki, G. (2014). Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus. *Proceedings of the National Academy of Sciences USA*, *111*, 13535–13540.

Viney, T. J., Lasztóczi, B., Katona, L., Crump, M. G., Tukker, J. J., Klausberger, T., & Somogyi, P. (2013). Network state-dependent inhibition of identified hippocampal ca3 axo-axonic cells in vivo. *Nature Neuroscience*, *16*, 1802–1811.

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., & Bright, J. e. a. (2020). Scipy 1.0: Fundamental algorithms for scientific computing in python. *Nature Methods*, *17*, 261–272.

Waskom, M. L. (2021). Seaborn: Statistical data visualization. *Journal of Open Source Software*, *6*, 3021.

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

Wouterlood, F. G., Saldana, E., & Witter, M. P. (1990). Projection from the nucleus reuniens thalami to the hippocampal region: Light and electron microscopic tracing study in the rat with the anterograde tracer phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*, *296*(2), 179–203.

Yamamoto, J., & Tonegawa, S. (2017). Direct Medial Entorhinal Cortex Input to Hippocampal CA1 Is Crucial for Extended Quiet Awake Replay. *Neuron*, *96*(1), 217–227.e4.

Ylinen, A., Bragin, A., Nádasdy, Z., Jandó, G., Szabo, I., Sik, A., & Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 hz) in the intact hippocampus: Network and intracellular mechanisms. *Journal of Neuroscience*, *15*, 30–46.

Zhang, Y., Cao, L., Varga, V., Jing, M., Karadas, M., Li, Y., & Buzsáki, G. (2021). Cholinergic suppression of hippocampal sharp-wave ripples impairs working memory. *Proceedings of the National Academy of Sciences*, *118*(15), e2016432118.

Zutshi, I., & Buzsáki, G. (2023). Hippocampal sharp-wave ripples and their spike assembly content are regulated by the medial entorhinal cortex. *Current Biology*, *33*(17), 3648–3659.e4.