

The hippocampal code - Towards understanding neural representations underlying spatial navigation

Archit Gupta



Electrical Engineering and Computer Sciences
University of California, Berkeley

Technical Report No. UCB/EECS-2023-25

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2023/EECS-2023-25.html>

May 1, 2023

Copyright © 2023, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

The hippocampal code - Towards understanding neural representations underlying spatial navigation

by

Archit Gupta

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Foster, Co-chair
Professor José Carmena, Co-chair
Professor Michel Maharbiz
Professor Frederic Theunissen

Summer 2021

The hippocampal code - Towards understanding neural representations underlying spatial navigation

Copyright 2021
by
Archit Gupta

Abstract

The hippocampal code - Towards understanding neural representations underlying spatial navigation

by

Archit Gupta

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Science

University of California, Berkeley

Professor David Foster, Co-chair

Professor José Carmena, Co-chair

Navigation is central to the survival of species. From a systems neuroscience perspective, the Hippocampus is one of the principal regions that enables mammalian navigation. Academic and medical literature have also implicated this part of the archicortex in the long-term storage of memories. The ability to record from hundreds of neurons at sub-millisecond time-scale from the hippocampus now provides a unique window into its inner workings, allowing us to test broader computational theories about the brain. In this dissertation, we explore the representation of space in the brain from multiple perspectives. First, we study oscillatory interactions in the brain - we combine a novel opto-genetic experimental paradigm with in-vivo electrophysiology to study inter-connected regions in the rodent brain. We then use cross-frequency analyses to study how two of the most prominent oscillations in the hippocampus, theta and gamma, interact. We show that the latter, gamma oscillations, are synchronized across a large extent of anatomical tissue and that this synchrony is driven by the CA3 sub-region of the Hippocampus. We then discuss functional representations of space in the brain that constitute a cognitive map. We demonstrate tools and algorithms that were developed as a part of our projects to perform Bayesian inference of neural activity in real time, as well as to detect known oscillatory phenomenon called Sharp-Wave Ripples (SWR) that have been associated with consolidation of memories in literature. We lay out experimental work to test the role of SWR in online spatial learning. Our work has been released as a publicly available open-sourced tool called ActiveLink. We discuss technological challenges in real-time inference of neural activity and offer a machine-learning solution based on neural networks to outperform existing methods for spatial inference.

To Papa, Mammi, Bhaiya, Bhabhi, Arnav, Ananya,
and Nehal.

Acknowledgments

I would like to thank my advisor, David Foster, not only for introducing me to the world of neuroscience, but also inculcating and nurturing scientific inquisitiveness and teaching me how to pursue scientific research. I have thoroughly enjoyed our many discussions designing experiments, interpreting our findings, and am thankful for the numerous pieces of wisdom that have greatly improved me as a researcher. I wish to thank José Carmena, my advisor in EECS, who lent me his support. The feedback and support I received from my committee members, Michel Maharbiz and Frederic Theunissen has greatly helped my research. I also owe a debt of gratitude to my former mentor Jaijeet Roychowdhury, and colleagues Tianshi Wang and Ahmet Gokcen Mahmutoglu, from whom I learnt a great deal as an engineer. That training gave me a unique perspective and ability to approach my current research. I am grateful to Linda Wilbrecht and Lance Kriegsfeld, who along with David were the first people to introduce me to neuroscience research and my discussions with them left a significant mark on me.

Every member of FosterLab has both inspired me to be a better researcher, and given me a hand at difficult points in research (which I learnt is rather frequent in experimental neuroscience). I am especially thankful to Heydar Davoudi for helping me build a research project early on, and who I could constantly turn to, to discuss ideas. John Widloski, Matthew Klienman and Alice Berners-Lee, thank you for inspiring me with your contagious passion for science and throughout my PhD, teaching me everything from handling a rat to performing surgery, and giving me invaluable feedback on my work. I am especially thankful to John for our late night conversations in the lab which ranged from troubleshooting spike-sorting programs to the meaning of life itself. Althea Cavanaugh and Charlie Walters, thank you for your support when I was new to the lab and you taught me the ways of drive building and electrophysiology.

My most enjoyable memories from my time at the lab have been with Caitlin Mallory and William Croughan, a.k.a “Team Stim”. All new to the lab, inquisitive and eager, we spent hours and hours trying to get our experiments running - adjusting electrodes, tugging on make-shift pulleys that were always falling apart, lugging rats around on makeshift platforms, bribing them with food to get them to calm down, mopping floors like maniacs - mostly doing things that I previously had not even imagined when I thought of science. It taught me that science is messy, and the true joy of being a scientist comes from sharing the experience with someone you respect. I will also miss out infinite ideation sessions trying to come up with better ideas to befuddle the rats, some of which I have tried to explain more scientifically in this dissertation. Over the last few years, I’ve learnt the most through our shared experiences and complementary skills.

Far away from home, I found a new family in my friends. Srimukh, Prateek, Niladri and Chirag, you were always there for me through good and bad times, looking out for me and nudging me towards the right direction when I couldn’t see it. Ishan, my longest standing friend, thank you for being a constant companion, near or far. Rishabh, Vidhi and Rishvi, thank you for providing me a home here and a loving family with whom I could lay back

and forget about the worries of everyday life. Deepali thank you for your zealous spirit and kind affection. Aakankha, life in the city would have been mundane if not for your company, thank you for introducing me to all the nooks and crannies.

Words aren't enough to thank my family, Mammi and Papa, bhaiya and bhabhi, who have constantly sought out my happiness, even over their own. Little ones, Arnav and Ananya, your innocence is the shining star of my life. Recently, I acquired a new family, Mammi, Papa and Priyal, who have brought with them enormous joy and happiness. Finally, my partner in life, Nehal. The last few years have been a roller coaster - getting married, living together, crossing the finish line, battered and bruised, at the end of AIDS LifeCycle, and then being apart. You have been the constant source of meaning and happiness in my life. From helping me get through harsh times in graduate school, to pushing me on quad-buster during AIDS LifeCycle, you have been with me every step of the way. I cannot imagine having accomplished this without you.

Contents

Contents	1
List of Figures	3
1 Introduction	12
1.1 Hippocampus and Space	14
1.2 Oscillations in the Hippocampus	20
1.3 Sharp-Wave Ripples	22
1.4 Hippocampal Replay	23
1.5 Closed loop experiments	26
2 Gamma Oscillations are synchronized in the Hippocampus	28
2.1 Materials and Methods	29
2.2 Theta-gamma interactions in the Hippocampus	33
2.3 Gamma oscillations are synchronized	35
2.4 Measurements along the dorsoventral axis	39
2.5 Role of CA3 in generating gamma oscillations	42
2.6 Role of CA3 in synchronizing gamma oscillations	49
2.7 Conclusion	53
3 ActiveLink: An open-source tool for neural targeting and closed-loop stimulation	58
3.1 Software Design	59
3.2 Quantified targeting of Pyramidal cell layer in CA1	62
3.3 Online analyses of neural activity	64
3.4 Application to closed-loop experiments	68
3.5 Discussion and Future Work	79
4 Neural-Network architectures for spatial inference using neural activity	82
4.1 Models	83
4.2 Results	88
4.3 Decoding Replays with Neural Networks	91
4.4 Down-sampling Performance	93

<i>CONTENTS</i>	2
4.5 Discussion	94
5 Conclusion	97
5.1 Gamma Oscillations: A synchronized marker for spatial representation . . .	97
5.2 The micro-structure of neural activity	98
5.3 A model for spatial navigation	99
Bibliography	105

List of Figures

1.1	Spatially tuned single units in the Hippocampus. (a.) Animals trajectory (blue) over the course of a 40min exploratory session in a $2m \times 2m$ open arena. Animal's physical location at time when action potentials were observed from a single cell have been marked in red. (b.) Spatial tuning curve constructed from (a.) by binning spikes and position into 4cm bins and calculating normalized firing rate in each position bin as described in Eq. (1.1).	14
1.2	Illustration of a Sharp-Wave Ripple complex detected in the rodent hippocampus.	22
1.3	Directional place fields recording during linear track traversal.	24
1.4	Replay (heatmap) in a $2m \times 2m$ open environment depicting a future path (red trace) to known goal location. Animal's past trajectory is shown in black.	25
1.5	Illustration of replay events suggestive of auto-associative dynamics (a.) A single replay event showing discrete locations where the replay event hovers before jump discretely. (b.) Total posterior over all replay events detected within a single session.	25
2.1	Gamma events are embedded within theta oscillations. (a.) Fast gamma event embedded within theta oscillations. (b.) Phase-Frequency Response (PFR) for a single recording site illustration theta phase locking for slow and fast gamma oscillations. (c.) Slow gamma event embedded within theta oscillations. (d.) Time-interval between consecutive fast-gamma events ($n = 147$ tetrodes with significant power in slow or fast gamma band, mean peak-lag 192ms). (e.) Time-interval between consecutive slow-gamma events ($n = 147$ tetrodes with significant power in slow or fast gamma band, mean peak-lag 171ms).	34
2.2	Statistics of slow and fast gamma events. (a.) Breakdown of theta cycles, $n_T = 273$ tetrodes, recorded $n_C = 744 \pm 147$ (mean \pm std.) theta cycles from individual tetrodes. (b.) Breakdown of recorded electrodes with significant theta locking ($n_T = 279$ tetrodes, $n_S = 178$ tetrodes have significant slow-gamma locking to theta, and $n_F = 201$ tetrodes have significant fast gamma locking to theta.)	34
2.3	Illustrating synchronized slow gamma events recorded from the Hippocampus.	35
2.4	Distribution of theta and gamma lags for recording sites. Measured lags for (a.) theta, (b.) slow and (c.) fast gamma. One-tailed, two-sample F-test statistics: slow gamma <i>vs.</i> theta $p = 7.8 \times 10^{-39}$, confidence interval (c.i.) for variance ratio $c.i. = [4.26, 6.90]$, fast gamma <i>vs.</i> theta $p = 7.3 \times 10^{-4}$, $c.i. = [1.19, 1.93]$	36

2.5 Construction of PFR and calculation of gamma-theta locking strength. (a.) Example phase frequency response from a one recording site over a single run session from Rat K. (b.) Average power in slow gamma frequency band as a function of theta phase. (c.) Slow gamma events extracted from the spectrogram (Rayleigh test for circular non-uniformity: $z = 28$, $p = 2.5 \times 10^{-13}$). (d.) Average power in fast gamma frequency band as a function of theta phase. (e.) Fast gamma events extracted from the spectrogram (Rayleigh test for circular non-uniformity: $z = 31$, $p = 1.2 \times 10^{-14}$). 37

2.6 Local *vs* Global theta reference - comparing PFR produced by local and global referencing for four different recording sites from one animal. (a.) T2 (Local reference), (b.) T22 (Local reference), (c.) T26 (Local reference), and (d.) T27 (Local reference). (e.) T2 (Global reference), (f.) T22 (Global reference), (g.) T26 (Global reference), and (h.) T27 (Global reference). 38

2.7 (a.) Illustration of synchronized fast gamma events from a recording session. (b.) Comparing local and global MRV Lengths over discrete gamma events in individual sessions (One-tailed, paired Wilcoxon sign-rank test comparing Local *vs*. Global Mean Resultant Vector (MRV) Lengths over $n = 13$ sessions: Slow Gamma $p = 0.02$, Fast Gamma $p = 0.01$). 39

2.8 Gamma events are referenced to a global theta phase (a.) Preferred phase of slow and fast gamma events referred to their local theta phase. (b.) Difference between global and local MRV difference for slow gamma. (c.) Preferred phase of slow and fast gamma events referred to a global theta phase definition. (d.) Difference between global and local MRV difference for fast gamma. 40

2.9 Local *vs*. Global referencing for individual sessions. Showing MRV constructed for slow (blue) and fast (red) gamma events for all the recorded sessions. For each session, locally reference gamma event MRV (left) and the same global event MRV (right) are shown. (a.) Rat I, S1, (b.) Rat I, S2, (c.) Rat I, S3, (d.) Rat K, S1, (e.) Rat K, S2, (f.) Rat K, S3, (g.) Rat G, S1. (h.) Rat J, S1. (i.) Rat J, S2. (j.) Rat J, S3. 41

2.10 Observed oscillation lags as a function of anatomical distance between pairs of recording electrodes. (a.) Measured theta lag for all recording sites relative to a single reference in a recording session. (b.) Measured slow gamma lag for all recording sites relative to a single reference in a recording session. (c.) Measured fast gamma lag for all recording sites relative to a single reference in a recording session. (d.) Theta lag *vs*. Anterior-Posterior (AP) distance (Pearson's correlation $r = 0.33$, $p = 4.5 \times 10^{-7}$). (e.) Slow gamma lag *vs*. AP distance (Pearson's correlation $r = 0.002$, $p = 0.98$). (f.) Fast gamma lag *vs* AP distance (Pearson's correlation $r = 0.07$, $p = 0.33$). 42

- 2.11 Relating lag and correlation between theta and gamma. Reporting paired Gamma-Theta measurements across $n = 3152$ pairs of recording sites. (a.) Slow gamma lag vs. Theta lag (Pearson's correlation $r = 0.005$, $p = 0.001$). (b.) Fast gamma lag vs. Theta lag (Pearson's correlation $r = 0.003$, $p = 0.85$). (c.) Slow gamma correlation coefficient vs. Theta correlation coefficient (Pearson's correlation $r = 0.22$, $p = 5.3 \times 10^{-37}$). (d.) Fast gamma correlation coefficient vs. Theta correlation coefficient (Pearson's correlation $r = 0.18$, $p = 1.2 \times 10^{-24}$). 43
- 2.12 Schematic of simultaneous recordings from dorsal and intermediate hippocampus. 43
- 2.13 Gamma synchrony is preserved along the longitudinal axis of the Hippocampus. (a.) Schematic showing simultaneous recordings of dorsal and intermediate hippocampus. (b.) Simultaneous recordings from electrodes shown in dorsal (blue) and intermediate (orange) CA1 with a slow gamma event. Top: Raw LFP recorded on the electrodes with highlighted gamma event. Bottom: Theta filtered LFP and slow gamma power envelope. (c.) Measured lags for electrode pairs within dorsal hippocampus, within intermediate hippocampus, and inter-regional pairs for theta filtered LFP, slow gamma power envelope and fast gamma power envelope. Wilcoxon signed-rank test for paired theta-gamma lags: Dorsal hippocampus, $n = 324$ electrode pairs, slow gamma *statistic* = 8.4×10^2 , $p = 2.5 \times 10^{-50}$, fast gamma *statistic* = 1.1×10^3 , $p = 4.3 \times 10^{-48}$. Intermediate hippocampus, $n = 244$ electrode pairs, slow gamma *statistic* = 8.2×10^2 , $p = 8.1 \times 10^{-35}$, Fast gamma *statistic* = 2.5×10^3 , $p = 5.8 \times 10^{-26}$. Inter-region lags: $n = 51$ electrode pairs, slow gamma *statistic* = 50, $p = 8.9 \times 10^{-9}$, Fast gamma *statistic* = 4.5, $p = 2.2 \times 10^{-9}$. (d.) Statistical comparison between theta lag and gamma lag within and across regions. Wilcoxon rank-sum test statistics: Slow gamma (dorsal vs. inter-region lags) *statistic* = -3.0 , $p = 0.002$, (intermediate vs. inter-region lag) *statistic* = -5.2 , $p = 1.9 \times 10^{-7}$. Fast gamma (dorsal vs. inter-region lags) *statistic* = 3.9, $p = 9.4 \times 10^{-5}$, (intermediate vs. inter-region lags) *statistic* = -6.6 , 5.7×10^{-11} 44
- 2.14 Schematic illustrating acute and selective opto-genetic suppression of CA3 input to CA1. 45

- 2.15 Suppressing CA3 input to CA1 - Changes in gamma power and events. (a.) Comparing change in average power in slow gamma frequency band during run between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -1.4$, $p_C = 0.08$, $n_E = 115$ tetrodes, $z_E = 1.2$, $p_E = 0.88$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -1.7$, $p = 0.09$). (b.) Measuring change in slow gamma power as a function of Sharp-Wave Ripples (SWR) suppression index (Pearson's correlation, $r = 0.11$, $p = 0.24$) in the experimental group. (c.) Comparing average power in fast gamma frequency band between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -1.7$, $p_C = 0.04$, $n_E = 115$ tetrodes, $z_E = -5.0$, $p = 3.1 \times 10^{-7}$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -2.0$, $p = 0.05$). (d.) Measuring change in fast gamma power as a function of SWR suppression index ($r = 0.47$, $p = 3.8 \times 10^{-7}$) in the experimental group. 46
- 2.16 Suppressing CA3 input to CA1 - Changes in gamma events. (a.) Comparing the change in rate of slow gamma events between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -2.1$, $p_C = 0.02$, $n_E = 115$ tetrodes, $z_E = -2.0$, $p_E = 0.02$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -0.88$, $p = 0.38$). (b.) Measuring change in rate of slow gamma events as a function of SWR suppression index (Pearson's correlation $r = 0.36$, $p = 3.0 \times 10^{-4}$) for the experimental group. (c.) Comparing the change in rate of fast gamma events Across control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -0.33$, $p_C = 0.37$, $n_E = 115$ tetrodes, $z_E = 1.3$, $p_E = 0.90$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -0.89$, $p = 0.37$), (d.) Measuring change in rate of fast gamma events as a function of SWR suppression index (Pearson's correlation $r = 0.04$, $p = 0.70$) for the experimental group. Comparing slow *vs.* fast gamma event incidence change within groups: *CON*, $z_C = -1.0$, $p_C = 0.32$, *EXP*, $z_E = -2.2$, $p_E = 0.03$ 47
- 2.17 Changes in PFR resulting from suppression of CA3 input to CA1. (left) PFR during running periods under Control condition (Light Off). (right) PFR during running periods under suppression of CA3 input (Light ON). 48
- 2.18 Comparing PFR four different recording sites from one animal in the *EXP* group across light conditions. (a.) T2 (Light OFF), (b.) T22 (Light OFF), (c.) T26 (Light OFF), and (d.) T27 (Light OFF). (e.) T2 (Light ON), (f.) T22 (Light ON), (g.) T26 (Light ON), and (h.) T27 (Light ON). 49
- 2.19 Changes in theta phase preference of gamma events. (a.) Theta phase distribution of slow gamma events during control (light off) condition. Rayleigh test statistics $z = 6.2$, $p = 0.002$. (b.) Theta phase distribution of slow gamma events during experimental (light on) condition. Rayleigh test statistics $z = 0.2$, $p = 0.82$. (c.) Theta phase distribution of fast gamma events during control (light off) condition. Rayleigh test statistics $z = 1.1 \times 10^2$, $p = 1.8 \times 10^{-54}$. (d.) Theta phase distribution of fast gamma events during experimental (light on) condition. Rayleigh test statistics $z = 90$, $p = 1.6 \times 10^{-45}$ 50

2.20 (a.) Change in MRV for slow gamma between control and experimental groups. One-tailed paired Wilcoxon signed-rank test statistics: $n_C = 38$ tetrodes, $z_C = -1.5$, $p_C = 0.06$. $n_E = 133$ tetrodes, $z_E = -3.0$, $p_E = 0.001$. Wilcoxon rank-sum test statistics across groups: $z = 0.03$, $p = 0.97$. (b.) Correlating change in slow gamma MRV in the experimental group as a function of SWR suppression during awake-rest (Pearson's Correlation: $r = 0.34$, $p = 1.4 \times 10^{-4}$). (c.) Change in MRV for fast gamma between control and experimental groups. One-tailed paired Wilcoxon signed-rank test comparing light conditions: $n_C = 38$ tetrodes, $z_C = -0.09$, $p_C = 0.46$. $n_E = 133$ tetrodes, $z_E = 0.59$, $p_E = 0.81$. Wilcoxon rank-sum test statistics across groups $z = 0.25$, $p = 0.80$. (d.) Correlating change in fast gamma MRV in the experimental group as a function of SWR suppression during awake-rest (Pearson's Correlation: $r = -0.07$, $p = 0.42$). One-tailed, paired Wilcoxon signed-rank test comparing slow gamma *vs.* fast gamma within *EXP* group: $z = -2.6$, $p = 0.004$ 51

2.21 The role of CA3 in locking gamma events to theta. (a.) Change in preferred phase of slow gamma events in control group. (b.) Change in preferred phase of slow gamma events in experimental group. Two-tailed Wilcoxon rank-sum test comparing phase shift between *EXP* and *CON* groups for slow gamma - $n_E = 131$ tetrodes, $n_C = 39$ tetrodes, $z = -1.8$, $p = 0.07$. (c.) Measuring change in preferred phase of slow gamma events against SWR suppression index for experimental group (Pearson's correlation for $n_E = 131$ tetrodes with both slow and fast gamma coherence over 0.1: $r = 0.1$, $p = 0.29$). (d.) Change in preferred phase of fast gamma events in control group. (e.) Change in preferred phase of fast gamma events in experimental group. Two-tailed Wilcoxon rank-sum test comparing phase shift between *EXP* and *CON* groups for slow gamma - $n_E = 131$ tetrodes, $n_C = 39$ tetrodes, $z = -3.8$, $p = 1.5 \times 10^{-4}$. (f.) Measuring change in preferred phase of fast gamma events against SWR suppression index for experimental group (Pearson's correlation for $n_E = 131$ tetrodes with both slow and fast gamma coherence over 0.1: $r = -0.02$, $p = 0.85$). 52

2.22 Schematic for calculation of a synchrony. 53

2.23 Measuring theta oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for LFP theta (Wilcoxon sign-rank test, OFF *vs.* ON: $n_C = 86$ electrodes, $z_C = 6.2$, $p_C = 7.5 \times 10^{-10}$, $n_E = 193$ electrodes, $z_E = -3.2$, $p_E = 0.002$, Wilcoxon rank-sum test, *EXP vs.* *CON* $z = 5.4$, $p = 8.1 \times 10^{-8}$). (b.) Change in mean lag for slow gamma (Wilcoxon sign-rank test, OFF *vs.* ON: $n_C = 86$ electrodes, $z_C = 3.3$, $p_C = 9.6 \times 10^{-4}$, $n_E = 193$ electrodes, $z_E = -4.7$, $p_E = 2.0 \times 10^{-6}$, Wilcoxon rank-sum test, *EXP vs.* *CON* $z = 5.4$, $p = 7.7 \times 10^{-8}$). (c.) Change in mean lag for LFP theta (Wilcoxon sign-rank test OFF *vs.* ON: $n_C = 86$ electrodes, $z_C = -0.99$, $p_C = 0.32$, $n_E = 193$ electrodes, $z_E = 1.9$, $p_E = 0.06$, Wilcoxon rank-sum test, *EXP vs.* *CON* $z = -1.9$, $p = 0.06$). (d.) Change in mean peak-correlation against SWR suppression for theta (Pearson's $r = -0.15$, $p = 0.05$). 54

2.24	Measuring slow gamma oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for slow gamma (Wilcoxon sign-rank test, OFF vs. ON: $n_C = 86$ electrodes, $z_C = -5.6$, $p_C = 1.6 \times 10^{-8}$, $n_E = 193$ electrodes, $z_E = 6.6$, $p_E = 4.2 \times 10^{-11}$. Wilcoxon rank-sum test, EXP vs. CON $z = -8.2$, $p = 2.6 \times 10^{-16}$). (b.) Change in mean peak correlation with SWR suppression for slow gamma (Pearson's correlation $r = 0.24$, $p = 0.001$). (c.) Change in mean lag for slow gamma. One-tailed, paired Wilcoxon sign-rank test, OFF vs. ON, $n_C = 86$ tetrodes, $z_C = 3.3$, $p_C = 9.6 \times 10^{-4}$, $n_E = 193$ tetrodes, $z_E = -4.7$, $p_E = 2.0 \times 10^{-6}$. Wilcoxon rank-sum test comparing EXP vs. CON, $z = 5.4$, $p = 7.7 \times 10^{-8}$. (d.) Change in mean lag with SWR suppression for slow gamma (Pearson's correlation $r = -0.17$, $p = 0.02$	55
2.25	Measuring fast gamma oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for fast gamma (Wilcoxon sign-rank test OFF vs. ON: $n_C = 86$ electrodes, $z_C = -0.68$, $p_C = 0.5$, $n_E = 193$ electrodes, $z_E = 5.8$, $p_E = 5.9 \times 10^{-9}$, Wilcoxon rank-sum test, EXP vs. CON $z = -4.4$, $p = 1.0 \times 10^{-5}$). (b.) Change in mean peak-correlation against SWR suppression for fast gamma (Pearson's correlation $r = 0.36$, $p = 1.2 \times 10^{-6}$). (c.) Change in mean lag for fast gamma (Wilcoxon sign-rank test between OFF vs. ON: $n_C = 86$ electrodes, $z_C = 2.7$, $p_C = 0.007$, $n_E = 193$ electrodes, $z_E = -3.9$, $p_E = 9.7 \times 10^{-5}$, Wilcoxon rank-sum test, EXP vs. CON $z = 4.4$, $p = 1.1 \times 10^{-5}$). (d.) Change in mean lag against SWR suppression for fast gamma (Pearson's correlation $r = -0.19$, $p = 0.01$).	56
3.1	Software architecture for real-time processing and close-loop feedback	59
3.2	SWR events occurring synchronously across multiple electrodes during a sleep session from a rat.	62
3.3	Statistics of SWR electrode adjustment in the brain. (a.) Mean power in the ripple frequency band. (b.) Standard-Deviation (over time) of power in the ripple frequency band. (c.) Correlation in the ripple band power between a recording site and a fixed reference. (d.) Average z-scored ripple power in a 500ms window surrounding SWR events detected on a reference electrode. (e.) Correlation in the Sharp-Wave envelope between a recording site and a fixed reference. (f.) Amplitude of the Sharp-Wave at a recording site.	63
3.4	Illustration of an SWR event detected online using ActiveLink.	65
3.5	(a.) Accuracy and (b.) Latency measures for detecting SWR events online. . . .	65
3.6	Clustering outcome after 3 minutes on a single tetrode bundle while the animals ran laps on a linear track. (a.) Electrode tips in neocortex, (b.) Electrode tips in Hippocampal CA1.	66

3.7	Real-time Bayesian Inference using 60ms decoding window with 20ms bin-shift between consecutive windows in ActiveLink. (a.) Bayesian decoding early less than 1minute into the session using hand-clustered single units in Trodes. (b.) Bayesian decoding ~ 5 minutes into the session using hand-clustered units. (c.) Illustration of a replay event decoded in real-time. (d.) Median decoding error over laps.	67
3.8	Decoding error over laps with online clustering.	68
3.9	Illustrating replay events detected using hand-clustered data and real-time online clusters built using ActiveLink (a.) Event 01, Hand-clustered, (b.) Event 01, clustered online, (c.) Event 02, Hand-clustered, and (d.) Event 02, Clustered online.	69
3.10	Effect of Ventral Hippocampal Commisure (VHC) stimulation on the activity of cells in CA1.	70
3.11	Block and trial structure for the learning task. Each training block consisted of 4 learning trials separated by an Inter-Trial Interval (ITI) as shown. Consecutive blocks were separated by 1 – 3 minutes. The maze and reward ports were wiped with a 70% Ethanol solution to clear odor-cues that might mark the current goal.	72
3.12	Configuration of the environment, reward locations and stimulation protocol for disruption blocks.	73
3.13	Spatial trajectories for a single learning block. Blue: Trajectory to the first approach to the goal location. Gray: Subsequent path taken after reaching the goal. Green Circle: Start location. Red Circle: Goal in the learning block. (a.)Trajectory from previous learning block. (b.)-(e.) Learning trajectories over 4 trials with the new goal location. (f.) Learning trajectory over the next block with a new goal location. Goal from the previous learning block has been marked in red.	74
3.14	(a.) Latency to goal over trials in a learning block compared between groups. (b.) Measuring savings in latency paired to Trial 1 of the same learning block.	74
3.15	Measuring spatial trajectories taken by the animal over learning trials. (a.) Average running speed during the time spent by the animal exploring the maze in search of food. (b.) Total distance covered by the animal in a trial before reaching the reward. (c.) Inefficiency in routing to the goal (total running distance normalized by the optimal path from the starting location) (d.) Improvement in routing after Trial 1.	75
3.16	Behavioral measures when the reward location was marked with a visual cue. (a.) Latency from start of a trial to the time the animal first reached the reward port when the port was not marked with a visual cue (Control) or otherwise. (b.) Total path length from the start of the to arrival at goal.	76
3.17	Measuring perseverance on the previous goal on Trial 1 of the subsequent block as a function of experimental condition (<i>CON vs. EXP</i>) in the current and previous block.	77

3.18	SWR rates during different behavioral states of the task. (a.) SWR rates from the start of a trial until the animal discovered the currently rewarded well. (b.) SWR rates during the time spent by the animal at the reward well. (c.) SWR rates during the rest period in the home platform following a trial on the maze. <i>T4</i> has been omitted since stimulation was stopped before placing the animal on the platform after the last trial of the block.	78
3.19	(a.) Experimental schematic for testing the role of replay in learning. (b.) Experimental schematic for testing the role of replay in planning.	80
4.1	Bayesian inference of spatial position in a 1-dimensional environment. (a.) Activity of a single place cell on a linear track. (b.) Posterior probability computed using a Naïve Bayes decoder using a $50ms$ decoding window.	84
4.2	Network architecture and data pre-processing. (a.) Illustration of a raw position trace. The dataset was broken down in train (70%), test (10%), and validate (20%) sets. (b.) Architecture of a 2-layer Neural Network used for directional decoding of population firing vectors. A hidden layer comprised of 256 units and output layer comprised of 114 directional place bins were used to decode animal's position from 185 single units. (c.) Sample training data (top) and labels (bottom) used for training the network.	86
4.3	Training the network. (a.) Kullback–Leible (KL) Divergence loss for training and validation sets over training epochs. (b.) Mean Absolute Error (MAE) measured for training and validation sets over the course of training. (c.) Illustration of probability distribution over position bins learned by the network on the test set.	88
4.4	Comparing decoding accuracy of Neural Networks with traditional Naive Bayes decoder. (a.) Decoding a single lap on a linear track using Naive Bayes decoder. (b.) Decoding the same lap as in (a.), but using a Neural Network (NN) decoder. (c.) Calculated Center of Mass (CoM) from posterior probability over the test set for NN decoder. (d.) Statistical comparison between Naive Bayes ($n_B = 1269$ decoded bins) and NN decoder ($n_N = 4962$ decoded bins, one-tailed Wilcoxon rank-sum test $statistic = 2.7 \times 10^6$, $p = 3.5 \times 10^{-15}$).	89
4.5	Hyper-parameter tuning and network optimization. (a.) KL divergence loss for training and validation sets when varying the number of NN layers. (b.) MAE for training and validation sets when varying the number of NN layers. (c.) KL divergence loss for training and validation sets when varying the width of the hidden layer in a 2-layer NN. (d.) MAE for training and validation sets when varying the width of the hidden layer in a 2-layer NN.	90
4.6	Hidden unit activation in the NN architecture.	91
4.7	Detection of replay event using NN architecture with different network depths (a.) 1-Layer NN, (b.) 3-Layer NN, (c.) 5-Layer NN.	91
4.8	Detection of replay event using 2-layer NN architecture with different network widths: (a.) 32 hidden units, (b.) 128 hidden units, (c.) 512 hidden units.	92

4.9	Agreement on replay metrics between Naïve Bayes and NN models over the six recorded sessions.	93
4.10	Comparing performance of tradition Naïve Bayes model again Neural-Networks with a down-sampled cell population. (a.) 50 Cells (b.) 75 Cells (c.) 100 Cells (d.) 125 Cells	94
5.1	Schematic of typical behavioral trajectories over the course of learning in a spatial navigation task. (a.) Typical trajectory before the goal location has been learned. (b.) Illustration of a behaviorial trajectory after the goal location has been learnt.	100
5.2	(a.) Model definition for an actor with uses a weighted combination of place-cell activity to produce directed movement in the maze. (b.) Model definition for a critic which combines weighted place cell activity with reward information to update the animal's navigational policy.	101
5.3	(a.) Trajectory for agent (blue circle marks start, red square marks trajectory end and green circle marks the reward) on Trial 2 in the current configuration. (b.) Spatial trajectory on the same configuration as in (a.), but for Trial 40. (c.) Value representation learnt by the critic model using a weighted combination of the firing rate maps described in Eq. (5.4).	102
5.4	Distance moved by an agent simulating reward search in an open environment using Value leaning. (a.) During learning trials. (b.) After learning.	103
5.5	Reward structure and inferred value function in spatial learning in a maze. (a.) Structure of reward information in the maze experienced by the learning agent. (b.) Value function learned as a superposition of weighted place cell firing rates by the	103

Chapter 1

Introduction

The ability to navigate space is fundamental to the survival of individuals. Virtually all species possess exceptional navigational skills. Some of the most remarkable examples of navigational capabilities are found in migratory birds. The Arctic tern, for an instance, covers 22,000 miles in its migratory pattern [109]. One of the early studies regarding the homing behavior of birds [158] found that they could home in to Key Island on Florida when released up to 1000 miles away.

Understanding the various mental facilities that enable intelligent navigation has drawn interest in psychology and neuroscience alike. Early research in psychology, including the work of Tolman [148], laid the foundation for navigational principles which can drive intelligent behavior. Very broadly, these navigational faculties can be divided into two groups - *taxon* navigation, or navigation using an egocentric perspective, and *locale* navigation, navigation from an allocentric perspective. Theories of taxon navigation are based on the idea that Stimulus-Response associations can be used to perform navigation. Locale navigation, on the other hand, presumes that animals have an inherent capability of forming a model of the world. Such a model, or a spatial *map* can then be inferred to perform intelligent navigation. Early evidence for locale navigation was presented in [148] where rats were placed in a maze comprised of a series of binary decisions, ultimately leading to an empty room. Later, when reward was introduced in the maze, animals who had previously experienced the environment in the absence of reward outperformed naïve animals suggesting some form of latent learning even in the absence of a reward which would lead to direct stimulus-response pairing.

Around the same time when cognitive theories of navigation were being debated, a landmark study [131] showed that patients with bilateral hippocampal lesions suffered from severe retrograde amnesia. While much of their reasoning ability and cognition had been spared, some of the patients had completely lost the ability to create new memories. These patients did not suffer deficits in recalling precise memories from before the surgical procedures. Subsequent studies in non-human primates [175] and humans [174] confirmed that Hippocampus was essential for forming new memories. Subsequent development of in-vivo electrophysiology in the last century offered a window into the working of the brain at the

resolution of individual neurons. More evidence for a cognitive map arose in [107, 106] when cells in the hippocampus were discovered to be tuned to space. We also received evidence for taxon navigation, with processing streams in Entorhinal Cortex (EC) depicting objects [33] and environmental boundaries [139] which could support a stimulus-response model of navigation in addition to the cognitive map which can co-exist in the same brain region [56]. In this dissertation, we primarily focus on the internal representation of space that is believed to constitute a cognitive map.

In this work, we delve into the structure and function of the Hippocampus in the context of navigating space. While early studies [99, 142, 111] have pointed out that Hippocampus is necessary for execution of spatial navigation, the precise neural phenomenon that enable spatial learning and planning are not very well understood. There is accumulating evidence that suggests that sequential structure of activity in the brain is vital for navigation. However, we do not understand how such sequences are constructed and utilized in navigation. Local Field Potential (LFP) is fundamental to this understanding, and a potential substrate for clinical diagnosis and therapies, using, for example, non-invasive Electroencephelogram (EEG) signals. We explore the structure and mutual-interaction of these signals in the brain and study their origin.

While in-vivo recordings were limited to a handful of single units in the early days of neuroscience, we are now able to record from hundreds [123], and in some cases thousands [124] of cells from the brain simultaneously. The increasing rate of information being extracted from the brain requires scale-able tools and algorithms for information processing. The statistical modeling paradigms from machine learning [78] provide ideas that can be borrowed into systems neuroscience to advance our capability to information neural signals at scale and with high fidelity.

The work in this dissertation explores three different directions to elevate our understanding of spatial navigation. 1) In Chapter 2, I assess the interaction between theta and gamma oscillations within the Hippocampus. This work was done in collaboration with Dr. Heydar Davoudi who collected the data involving opto-genetic suppression of CA3 input to CA1 and Althea Cavanaugh who collected simultaneous electrophysiological recordings from the dorsal and the intermediate CA1. 2) In Chapter 3, I describe an open-sourced tool to aid the study of spatial navigation in closed-loop experiments. I also discuss the application of the tool to an online spatial learning task for studying the effects of SWR disruption on learning. This work was done collaboration with Dr. Caitlin Mallory and William Croughan. 3) In Chapter 4, I describe machine-learning methods to improve our ability to infer the increasingly complex neural data being recorded today. The rest of this chapter is dedicated to the discussion of academic research that has shaped our understanding of how space is represented in the brain. Given the vast scope of the subject at hand, this review inevitably must be selective.

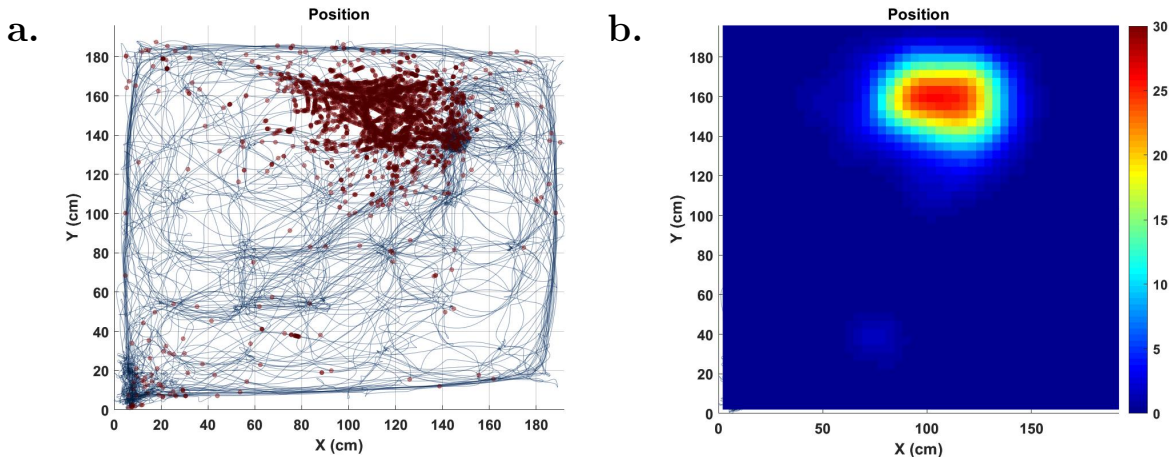


Figure 1.1: Spatially tuned single units in the Hippocampus. (a.) Animals trajectory (blue) over the course of a 40min exploratory session in a $2m \times 2m$ open arena. Animal's physical location at time when action potentials were observed from a single cell have been marked in red. (b.) Spatial tuning curve constructed from (a.) by binning spikes and position into 4cm bins and calculating normalized firing rate in each position bin as described in Eq. (1.1).

1.1 Hippocampus and Space

Early electrophysiology [107, 106] in the hippocampus showed that putative pyramidal cells are tuned to space. Fig. 1.1 illustrates the activity of a single cell recorded from the CA1 region of the hippocampus as the animal navigated a maze in search of reward delivered at one of thirty-six reward wells arranged in 6×6 grid (see [123] for details on the task). The blue trace in Fig. 1.1a. shows the animal's running trajectory during the session. The animal's location in the environment during every action potential emitted by the cell is reported in red circles. A firing rate map (Fig. 1.1b.) was constructed from the activity of the cell as

$$f(\mathbf{x}) = \frac{s(\mathbf{x}) + \epsilon}{o(\mathbf{x}) + \epsilon}, \quad (1.1)$$

where $f(\mathbf{x})$ is the firing rate map, $s(\mathbf{x})$ is the number of spikes or action potentials observed at the location \mathbf{x} and $o(\mathbf{x})$ is the amount of time spent by the animal at the location. A corrective parameter ϵ was used to for numerical reasons. This firing rate map is shown in Fig. 1.1b. and suggests that the firing of this cell is high selectivity to a specific region of the environment. Such cells in the hippocampus, which are tuned to space, have been called place cells.

Anatomically, the Hippocampus consists of well differentiated sub-regions, namely CA1, CA2, CA3 and Dentate Gyrus (DG) [3] and receives converging inputs from the visual stream via Entorhinal Cortex (EC). Place cells have been reported in all the sub-fields of

the hippocampus [69, 82]. The inherent recurrent connectivity of CA3 results in a robust spatial representation that shows higher invariance to local changes in the environment [80] while simultaneously driving place cell responses in CA1 [30].

Cells tuned to head direction have also been reported in the post-subiculum in freely moving rats [147]. Medial Entorhinal Cortex (MEC) is one of the principal inputs of the Hippocampus and receives converging inputs from the visual stream. Within the population of neuron in the MEC, it has been shown that a sub-population is tuned to an animal's running speed [76]. A study [56] found a population of cells in the MEC which is active in a triangular grid-like pattern during spatial navigation. Theoretical models suggest that a preconfigured attractor network [159] which receives inputs from cells tuned to the animal's velocity could produce a grid-like activation in the cell population. Further, combining inputs from grid cells could result in spatial tuning similar to what is observed in the hippocampus [138]. This is supported by experimental evidence suggesting that grid-properties of these 'grid cells' affects the size and stability of place cells in the Hippocampus [86]. A more functional view of place cells suggests that the slowness principle [162] when applied to the field-of-view results in tuning to space [48].

Experience-dependent changes

The hippocampal representation, interestingly, has also shown to be uniquely reconfigured for each different environment an animal is exposed to [104]. In the aforementioned experiment, for example, when animals were exposed to two different cylindrical environments, of the cells that were active in both the environments, 52% showed significantly different firing patterns. Furthermore, when changing the geometry of the environment (from circular to rectangular), the relative preferred firing locations of pairs of cells in one geometry could not be used to predict their relationship in the other. Rotating a single prominent visual cue in the environment often leads to a rotation of the firing rate maps by a consistent angle [11].

While place fields were typically viewed as aggregate activation, or long-term average spatial tuning curves of individual cells, [93] demonstrated that these fields change over the course of exploration in a novel environment. Over individual laps on a linear track, it was observed that the size of the place field increased with experience, and at the same time, directional fields shifted towards the running direction of the animal that the fields are tuned to. It has been hypothesized that this results from plasticity in CA1 - the order in which cells fire in a given running direction strengthens the synaptic connections following Hebbian learning, thereby pulling the fields towards the animal's running direction. A subsequent study [94] made similar findings in regard to the skewness of place fields over time. [133] made an interesting observation in mice that were exposed to multiple short (3-minute) trials on a linear track. They observed that multiple stable maps could co-exist for the same environment for over long periods of time. However, there is insufficient behavioral evidence to confirm if the animals identify the correspondence of multiple maps to the same environment.

Lesion Studies

Following a lesion of a majority of the medial temporal lobe, a group of patients reported severe impairments in episodic memory [174]. Subsequent studies in rodents and primates [175] suggested a prominent role for the Hippocampus in the formation of long-term episodic memories. While there was evidence for memory subsystems in the brain, the anatomical distribution and function of these memory systems has not been very well understood. A series of studies in the last century explored memory subsystems using electrolytic, radio-frequency or neurotoxic lesions to name a few of the experimental methods. Briefly, one of the brain regions believed to be responsible to creating, storing or retrieving memories was lesioned using the aforementioned methods. As a control, in some cases, a sham-surgery was performed wherein the procedure was performed at the same stereo-taxic coordinates but the pharmacological substance was not introduced, or replaced by another which didn't cause damage to the neural tissue. As an alternative to sham-surgery, in some cases another part of the brain which was similar in size to the one being investigated was damaged in a group of animals as a control.

Water Maze

After the discovery of place cells in the Hippocampus [107, 106], investigations of Hippocampus as the principal region involved in spatial navigation accelerated. In the early 1980s, a task called the "Water Maze" was developed to test spatial navigation [103]. A pool of water, made opaque by the addition of powdered milk consisted of an escape platform submerged under the surface of water. Rats were lowered into the water maze and swam until they were able to locate the escape platform. This process was repeated daily for a period of 8 days. Within each learning session, after the discovery of the platform, animals were taken out of the maze. The location of the platform was held constant across repeated exposures of the environment.

One of the experiment consisted of 4 groups of animals - no surgery ($n = 4$ animals), animals with superficial cortical surgery ($n = 13$ animals), sham surgery ($n = 4$ animals) and animals with the entire dorsal-ventral hippocampus lesioned [99]. Groups with either superficial cortical lesions or sham surgery had a stable latency at the end of learning ($< 8s$). Animals with hippocampal lesions, however, had profound difficulties in locating the hidden platform. After the specified period of training, the hidden platform was raised over the surface of the water and served as a visual cue. In this phase, all groups (including the animals with hippocampal lesions) were able to successfully escape the maze using the visible platform suggesting that they did not suffer from obvious motor or motivational deficits. However, when the platform was submerged once again, this time at the same location where it had previously been visible, the group with Hippocampal lesions had difficulty navigating to it.

Finally, the platform was removed altogether from the maze and animals were reintroduced. On measuring the amount of time the animals spent searching where the platform had previously been, it was found that the control groups spent significantly more time in

the correct quadrant of the maze where the platform had previously been. Animals with hippocampal lesions, on the other hand, did not show such a preference.

Multiple memory subsystems

While the water maze study established the hippocampus to be necessary for navigating space, the nature of this memory was not very well understood. [116] presented a double dissociation between the Hippocampus (using Fornix lesions) and the Caudate Nucleus using a variant of the water maze task. Similar to [99], the authors observed that animals with Fornix lesions performed poorly on the spatial water-maze task. However, animals with lesions of the Caudate Nucleus did not have difficulties in performing the task when compared to controls. In a visual discrimination variant of the task, where the water maze consisted of two visible cues, only one of which marked the location of a viable escape platform, they found that animals with Fornix lesions had no difficulty in the visual discrimination. Animals with lesions of Caudate Nucleus, however, had significantly worse performance than controls.

An extensive study of over 300 animals [90] attempted to further separate these memory subsystems by looking at 3 different kinds of memory tasks. Three principal regions - the Hippocampus, Lateral Amygdala and Dorsal Striatum were examined in this study. As opposed to the water maze used in previous studies, a radial 8-arm maze was used in this study.

Win-Shift (spatial) In the first of these experiments, all eight arms of the maze were baited with reward. A rat selected one of the eight arms and after exploring it, was forced to wait in the center of the maze for a period of 10 minutes. The trial ended when the animal had consumed the reward at all eight locations. Re-entering an arm that the animal had previously visited constituted an error. Among the three groups, animals with Fornix lesions performed significantly worse than controls. Animals with lesions to Lateral Amygdala or the Dorsal Striatum did not have significantly different performance from controls.

Conditioned Cue Preference in this version of the task, all but two arms of the maze were blocked. One of the arms was cued with a light source while the other one was kept dark. Animals spent 30 minutes in either of the two arms either with (paired) or without (unpaired) food. At the end of the session, animals were allowed to spend 20 minutes in the maze and their place-preference was measured as the total amount of time spent in the arm they were previously exposed to. The control animals spent significantly more time in the arm they were exposed to if it was paired with food. Animals with lesions to the Fornix or Dorsal Striatum showed a similar preference. However, animals with lesions to the Lateral Amygdala failed to develop a preference to the conditioning cue.

Win-Stay (Stimulus-Response) This followed a similar approach to the previous study [116] where a simple stimulus was paired with an immediate appetitive reward. Four of the eight arms in the maze were cued with a light signal. All the eight arms were baited with a reward. Entering any of the lit arms continued the experiment, wherein the arm was baited with a reward once again. However, if the animal entered one of the unlit arms, all lights were turned off and no more reward was introduced in the maze. In this

paradigm, the authors found that animals with Fornix lesions, in fact, performed better than controls. Animals with lesions to the Lateral Amygdala did not show a significantly different performance from controls. Animals with lesions to the Dorsal Striatum, however, had significantly worse performance in forming the simple stimulus-response association.

In a subsequent study [92], using the water-maze as the experimental substrate, another interesting finding was made. While animals with lesions to the dorsal striatum could indeed perform the water maze task, if the location of a visible platform was changed, animals with lesions to the dorsal striatum still preferentially checked the previous location where the visual platform had been. Combining the evidence presented in previous studies, the authors concluded that these animals, instead of learning the visual cue-escape relationship, had instead learnt a place-escape relationship.

Pharmacological Blockade of NMDA receptors

While lesion studies were highly informative about the role of hippocampus in spatial navigation, the neural mechanisms underlying this learning remained unclear. In order to shed light in the role of Long-Term Potentiation (LTP) mediated by N-Methyl-D-Aspartate (NMDA) receptors, multiple experiments were conducted in the presence of an NMDA receptor antagonist (AnimoPhosponovaleric Acid (AP5), (6)-3-(2-Carboxypiperazin-4-yl) Propyl-1-Phosponic Acid (CPP), MK-805 *etc.*) which prevented the induction of LTP.

[101, 100] studied learning in the water maze in the presence of AP5 which is known to prevent induction of LTP in the hippocampus. Naïve animals were implanted with micro-pumps for chronic intra-ventricular infusion of D/L AP5. Training started 4 days after the infusion surgery. One each training day, 3 trails, separated by a 4-hour Inter-Trial Interval (ITI) were provided. Animals with D/L AP5 infusions showed a degradation in their learning curves over the next 5 training days when compared to controls. Note that the degradation in performance was less severe than what had previously been observed in lesion studies [100]. These animals did not show a deterioration in a visual discrimination task involving a similar experimental setup. Importantly, if AP5 infusion was performed after the location of a fixed platform had learnt, animals did not show a loss in retention of the spatial information. These findings suggested that LTP in the hippocampus plays a significant role in the acquisition of the spatial information. Recent literature has challenged this view and several factors have gained prominence in assessing how spatial information is acquired and retained for navigation.

[15] presented the idea that the temporal requirements of a memory determined the need for LTP in accessing it. [142] addressed how time factors into spatial learning by systematically varying the ITI between consecutive trails in the water-maze experiment. Here, D-AP5 (the active enantiomer which inhibits LTP as opposed to the racemic D/L mixture previously used in [100]) was selectively infused in the Hippocampus of rats. On each training day, animals learnt a novel escape location in the water maze over 4 learning trials. The ITI between trials 1 and 2 was selected to be 15s, 20 minutes, or 2 house. All animal groups were pre-trained on the task and showed 1-trial learning of the novel escape

location. One of the animal groups had hippocampal lesions to compare the relative effects on spatial performance. Trials 3 and 4 were performed at a 15s ITI. When compared to animals with sham lesions, animals with Hippocampal lesions exhibited difficulty in learning the escape platform location irrespective of the $T1 - T2$ ITI. Interestingly, the performance of animals with D-AP5 infusion exhibited a dependence on the ITI. For short (15s) ITI, animals showed learning of the goal location in the previous trial as measured by savings in the escape latency. At a longer latency (20 minutes, 2 hours) however, the savings reduced significantly. These findings indicated that spatial information could be retained with an LTP independent mechanism for short duration. However, over longer time periods, the retention of information required LTP supported by NMDA receptors. In either of these scenarios, however, the Hippocampus was necessary for the acquisition of spatial information.

In [91], rats navigated a water maze with a fixed goal location (Phase 1). After 8 days of training on the ‘original’ platform location, animals were trained rapidly over the course of a single day to learn a new escape location in the maze (Phase 2). Prior to starting Phase 2, animals received CPP either directly into the hippocampus, or through system injections. In the control group, Artificial CerebroSpinal Fluid (ACSF) or saline was injected. Remarkably, animals were able to learn the location of a new goal location in the presence of systemic, as well as local NMDA antagonist. Animals were subsequently test 24 or 48 hours after this learning experience. During the recall experiment, the platform was moved back to its ‘original’ location. Animals who had CPP delivered to them either at the start or immediately after Phase 2 went directly to the original goal location, showing no memory of the second experience. The control animals, however, lingered at the ‘new’ goal location before heading to the original platform location.

[70] addressed three key questions regarding the stability and development of place maps in the absence of NMDA receptor activity. Using systemic CPP injections in rats, it was found that the representation of a familiar environment remained largely unaffected in the absence of NMDA activity, *i.e.*, place cells retained stable fields across days interspersed with NMDA inactivation. New spatial maps (remapping) was also observed in the same set of animals under NMDA inactivation, and remarkably, these maps remained stable over the course of hours as animals even when the spatial contexts changed between subsequent exposures. However, spatial maps developed under CPP were not stable over a longer-term. When animals were tested on the day following NMDA inactivation, the spatial representation for the newly experienced environment had changed significantly when compared to a familiar environment.

A recent study combined these ideas, assessing spatial learning in a task while also recording from place cells [37]. The spatial task comprised of a circular arena (120cm in diameter) with 177 reward wells. In any given learning session, 3 new reward wells were chosen and baited with food. Animals began a learning trial in a start-box and explored the arena until all the rewards were consumed. All three wells were baited and trials was repeated. One of the animals groups was administered intraperitoneal (R)-CPP injections. Both the control and CPP administered group acquired the new well locations over learning trials. On a subsequent (> 25 minutes after the learning session), animals were reintroduced

to the arena for a *probe* session during which, no food was delivered. It was observed that the control animals spent significantly more time near the rewarded locations from the learning session when compared to the CPP group. These findings also reinforce the idea that LTP is not necessarily required for short-term spatial learning, but plays an important role in the consolidation of spatial memories.

1.2 Oscillations in the Hippocampus

The Local Field Potential (LFP) recorded from the Hippocampus has been known to reflect a wide variety of cognitive and behavioral processes. In the early days of electrophysiology [152], it was already known that macroscopic electrical signals from the Hippocampus reflected states of arousal, such as walking, rearing, grasping *etc.* It has been hypothesized that oscillatory activity is reflective of synchronization processes in the brain [136]. As pointed out succinctly in [20], understanding the hippocampal LFP has two key merits. First, an understanding of these signals in the archicortex lays the foundation for a better understanding the structure and function of neocortical EEG. Second, understanding the structure of macroscopic activity in pathological models can translate to clinical applications using intra-cranial Electroencephelogram (EEG) as well as scalp EEG. This, for an instance today is being applied to understand human EEG for clinical diagnosis of brain anomalies like seizures [150] and stroke [67].

Theta oscillations

In rodents, theta oscillations (6–10 Hz) are primarily observed during alert behavioral states (running, rearing *etc.*) and Rapid Eye Movements (REM) sleep. Origin of theta oscillations lies in cholinergic inputs from the medial septum [5, 121, 6, 36]. Pyramidal neurons in hippocampal CA1 fire at progressively earlier theta phase over consecutive theta cycles [110]. More recently, it has also been shown that within a single theta cycle, neural activity is structured [46] sweeping from behind the animal (retrospective encoding) to locations further in front of the animal (prospective encoding) [154]. These ‘theta-sequences’ can be learnt rapidly in a novel environment [41]. From a functional prospective, theta sequences have been shown to reflect current and future choices in behavioral paradigms [160, 55]. Theta interactions of downstream regions have been implicated in learning in both spatial [134] and non-spatial [73] contexts.

A systematic anatomical study of theta phase across the longitudinal extent revealed that theta oscillations in fact travel along the the longitudinal (dorsoventral) axis of the hippocampus [85]. While it had originally been hypothesized that theta phase changes by 360° between the dorsal and ventral poles, recording along the entire longitudinal extent of the hippocampus revealed this phase spread to be 180° [119]. More recently, the traveling wave nature of theta has also been identified in humans [168]. It is interesting to know that theta oscillations are not universal to all species. In the Egyptian fruit bat, for an instance,

while the neurons exhibit place and grid responses similar to rodents, theta oscillations do not exist [165].

Gamma oscillations

Gamma oscillations (25 – 150 Hz) have long been studied in neocortex [157]. Recent studies [28, 97, 130, 42, 43] have identified 3 functionally and anatomically segregated sub-bands within the gamma frequency range. While we do not have a consensus of the exact frequency ranges, these sub-bands are typically defined as - slow gamma (30 – 60 Hz), mid gamma (60 – 90 Hz) and fast gamma (90 – 150 Hz). In [12], following bilateral lesions of Entorhinal Cortex (EC), slow-gamma (25 – 50 Hz) were observed in the CA1-CA3 network. Current Source Density (CSD) analysis suggested that there are two independent generators of gamma rhythms in the Hippocampus [29], namely CA3 and Dentate Gyrus (DG). A correlation study involving gamma oscillations [28] suggested that CA3 drives slow gamma oscillations in CA1 and Medial Entorhinal Cortex (MEC) drives fast gamma activity via DG.

A study of gamma oscillations during periods of restfulness suggested that increased gamma synchrony between CA1 and CA3 could result in more coherent spatial representation of space during Sharp-Wave Ripples (SWR) events [21]. In a recent study [43], it was shown that there is a differential task-specific increase in gamma synchrony between EC and DG. During a task with high spatial demand, it was shown that there is a selective increase in fast gamma synchrony between MEC and DG. On the other hand, during an object recognition task, within the same group of animals, there was an increased slow gamma synchrony between Lateral Entorhinal Cortex (LEC) and DG. Further, suppression of MEC and LEC led to behavioral deficits in the spatial and object-recognition tasks respectively.

Theta-Gamma interaction

[28, 27] noted that gamma power in slow and fast gamma bands was significantly modulated by theta phase. This further led to the discovery that during slow gamma events, the activity of CA1 and CA3 are highly synchronized. During fast gamma events, on the other hand, the activity of CA1 and MEC was highly synchronized. Combining this evidence with prior lesion studies [29] suggested CA3 to be the driver of slow gamma in CA1 and MEC to be the driver of fast gamma. Further, the phase-segregated oscillations provided distinct temporal windows to selectively drive the activity in CA1. Theoretical models of working memory [83, 84] suggest that the embedding of gamma oscillations with theta could serve as a general mechanism for sequence learning. This is supported further by phase-phase coupling between theta and gamma oscillations in the hippocampus [129].

1.3 Sharp-Wave Ripples

Within the Local Field Potential (LFP) recorded from the Hippocampus, periods of awake-rest and Slow-Wave Sleep (SWS) are characterized by fast (150 – 250Hz) oscillations called ripples riding on a slow oscillatory envelope called the sharp-wave (see Fig. 1.2). Sharp-Wave




Figure 1.2: Illustration of a Sharp-Wave Ripple complex detected in the rodent hippocampus.

Ripples (SWR) complexes arise from the excitatory recurrent system in the hippocampal CA3 [167, 16].

Early studies looking at memory [131, 175] had established that the hippocampus is essential for long-term storage of memories. Since SWR primarily occur during “offline” states, they have long been hypothesized to enable this long-term storage. Early studies showed temporally compressed reactivation of behavioral neural patterns accompanied with SWR events during sleep following behavioral experiences [161, 79]. We now know that this reactivation can be vivid in the extent to which it captures behavioral experiences [47, 123] and will be discussed in more detail in Section 1.4. The temporally compressed reactivation of neurons during SWR events reinforced the hypothesis of them mediating long-term storage of episodic memories. Simply put, a single experience can be remembered because the SWR event would compress time to a scale where synaptic plasticity induced changes can cement the experience in memory.

Further evidence comes from correlation studies measuring memory as a function of SWR events. In [7], human patients were implanted with macro-electrodes in the hippocampus and the rhinal cortex. Over two sessions, separated by a 60 minute sleep, patients were shown 80 images each. Subsequently, they were shown a mix of new and previously seen

images and asked to identify if they had previously seen an image. It was found that the performance on this task was highly correlated with the rate of ripples detected during the intermediate sleep in the rhinal cortex. In another study [4], involving reward modulation, animals ran trials on a linear track to receive equal amounts of reward on both ends. After a while, the reward on one of the ends was discrete increased to $4\times$ the reward on the other. It was found that the rate of SWR events increased dramatically on the end where the reward was increased. Interestingly, while the study of SWR events has largely been restricted to the offline or quiescent periods, in [37], it was found that exploratory-SWR were highly correlated with performance on a spatial navigation task.

In the context of SWR events, the dorsal-ventral axis of the hippocampus doesn't always undergo simultaneous oscillatory activity [118]. It has been hypothesized that these events are generated locally and then travel along the dorsal-ventral extent of the hippocampus. The functional role of this phenomenon has also been explored recently in a study [140], where it was found that dorsal and ventral SWR events can differentially activate downstream neurons (in this case in the Nucleus Accumbens (NAc)). Furthermore, [140] found that SWR in the dorsal hippocampus selectively activate NAc neurons encoding spatial trajectories.

1.4 Hippocampal Replay

It was first observed during Slow-Wave Sleep (SWS) that Sharp-Wave Ripples (SWR) were accompanied with pairwise reactivation of cells in an order that was consistent with cell activation during traversals of the environment [79]. It was observed that SWR during awake-rest are also accompanied with reactivation of cells in a manner consistent with the environment that the animal was currently experiencing [47]. These events have now been termed as “replay” events. However, the time scale at which the spatial trajectories were replayed was compressed, typically lasting 100 – 200ms to replay an environment that the animal would physically traverse over several seconds. It was found that in a linear environment, the two principal running directions are represented using different maps and replays could either involve cells activated in the same sequence that they were during traversal, called forward replays, or in the reverse order of activation during behavior, labelled reverse replays [34]. Forward and Reverse replays have been proposed to serve different functional roles in cognition. It has been reported that an unexpected increase in reward contingencies uniquely modulates the rate of reverse replay events [4], suggesting a role in reward prediction [61].

In [123], animal's alternated between a fixed ‘home’ well and a randomly selected ‘away’ well over alternate trails. Analysis of replay trajectories suggested that during stationary periods before animals navigated back to the fixed goal, the replay trajectories were more predictive of the upcoming or future paths that the animal's past path. The specificity of future-predictive replay trajectories during goal-oriented navigation suggested a role for replay in spatial planning. One such future-predictive replay trajectory is shown in Fig. 1.4. Here, the animal is currently at the location marked by the red-diamond. While the animal is

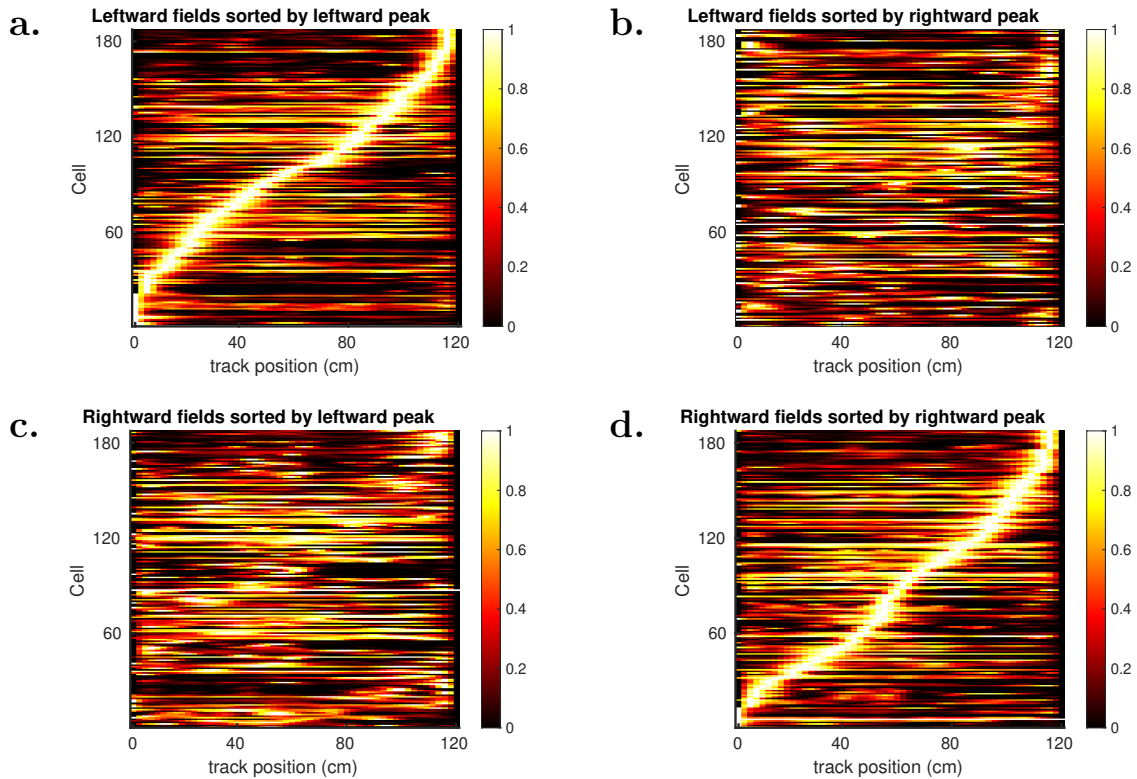


Figure 1.3: Directional place fields recording during linear track traversal.

at this location, the posterior probability for a 200ms decoded trajectory is shown using the heat-map. Golden dots mark the Center of Mass (CoM) of the decoded posterior probability. The animal's subsequent physical trajectory is shown in red - notice the alignment between the trajectory being mentally replayed and the subsequent path taken by the animal to a remembered goal location in the task. [164] made similar observations on a radial eight arm-maze, finding directional forward-replays to be predictive of future decisions and reverse-replays to be reflective of past choices.

Mechanistic models

The diversity of within and across region connectivity, combined with the heterogeneity of neural population in the Hippocampus has led to interest in computational and mechanistic modeling of its single unit activity. [122] suggests that the spatial representation exhibits auto-associative dynamics, in that the spatial representation, instead of representing space continuously, hovers at and discretely jumps between spatial representations. Fig. 1.5 illustrates replay dynamics on a linear track during stationary periods from a single recording session. In Fig. 1.5a., the posterior probability is shown for a single replay event as the animal consumed reward at $x = 0$ cm. Observe how the decoded location of the animal during




Figure 1.4: Replay (heatmap) in a $2m \times 2m$ open environment depicting a future path (red trace) to known goal location. Animal's past trajectory is shown in black.

the 200ms replay events hovers at distinct locations before jumping to the next (decoded using a Naïve Bayes decoder using 50ms windows shifted by 5cm). In Fig. 1.5b. we sum the posterior over all replay events and observe regions with high concentration of posterior probability.




Figure 1.5: Illustration of replay events suggestive of auto-associative dynamics (a.) A single replay event showing discrete locations where the replay event hovers before jump discretely. (b.) Total posterior over all replay events detected within a single session.

It has been also hypothesized that replays could arise from pre-configured cell assemblies in the Hippocampus [35]. However, these findings have been contradicted on statistical

grounds in subsequent studies [135]. [135] also establishes a role of N-Methyl-D-Aspartate (NMDA) receptors in enabling replay events, demonstrating that NMDA inactivation in a novel environment using (6)-3-(2-Carboxypiperazin-4-yl) Propyl-1-Phosphonic Acid (CPP) prevents the recent environment from being replayed in subsequent sleep. It is important to note that only a small fraction of SWR events are coherent spatial trajectories in the animal's current environment. One of the mechanistic views of replays, based on the diffusivity of spatial representation observed during stationary periods, suggests that Brownian motion in an abstract energy landscape [144] could lead to replay like observations.

1.5 Closed loop experiments

The advancements in recording and manipulation technology now allow us to infer neural activity in real-time and deliver closed-loop stimulation to make causal inferences for the underlying phenomenon. [50] applied this approach to study the role of Sharp-Wave Ripples (SWR) in spatial memory. Animals performed a spatial reference memory task on an radial eight-arm maze. Three of the eight arms were baited with reward and the animals were allowed 3 trials separated by 3×3 minute breaks in between. Following the behavioral task, animal spent 1-hour in a sleep box. During this period, they were subjected to online-detection and electrical disruption of SWR. In the control group, the stimulation was delivered after a random delay of 80 – 120ms. It was observed that over a periods of 15 days, the SWR-disrupted group exhibited a significantly worse learning curve than the groups that experienced no-stimulation, or the aforementioned delayed stimulation. A subsequent study [38] also found that behavioral performance over long-term learning of a spatial task deteriorated in animals who had SWR events disrupted in sleep.

More recently, awake ripples were also shown to play an important role in spatial learning. Naïve animals learnt to alternate between the exterior arms of a 'W'-shaped maze. Over the course of learning, it was observed that disrupting SWR events significantly deteriorated the outbound performance which required recalling the previous arm that the animal had visited [64]. Further, selective disruption of replay events corresponding to an environment in sleep led to degradation of spatial environment in that environment when compared to a control environment for which a delayed stimulation was delivered for replay events [53].

We discussed how the modulation of SWR is intertwined with behavioral and cognitive performance in Section 1.3, in [49], the after-effects for disruption of SWR events was studied as a function of cognitive load. Animals either performed a spatial reference-memory task in an eight-arm radial maze, or spent time in their home cage. After either of these sessions, animals were subjected to SWR disruption in a subsequent sleep session. It was observed that there was an overall increase in SWR events (as compared to a delayed stimulation control) for the animals who had previously performed a spatial task. The same increase was not observed in animals who had previously explored their home cages, suggesting a dependence of SWR events on cognitive demand. In [44], it has also been shown that optogenetic activation of CA1 cells can prolong existing SWR complexes. Additionally, in the

same task as [64], the authors found that prolonging the duration of SWR events improved performance in the outbound runs. In line with previous findings suggesting modulation of SWR events and associated replays with reward, [95] found that disruption of SWR events had differential effects on spatial learning for locations with different reward amounts, in that the performance deteriorated more for a highly rewarded location.

Chapter 2

Gamma Oscillations are synchronized in the Hippocampus¹

The hippocampus is one of the principal brain regions involved in spatial navigation [109, 102, 98, 123]. It is well known that hippocampal activity constitutes a cognitive map [109], allowing an animal's position to be decoded from its neural activity [13]. Hippocampal Local Field Potential (LFP) has long been studied to better understand circuit mechanisms that underlie the cognitive map [108, 18]. It is known that the hippocampus employs both rate and phase codes for memory and planning in the context of navigating space [108, 83, 137]. Theta oscillations (6-10 Hz) are strongest during ambulatory behavior and Rapid Eye Movements (REM) sleep. Single unit activity in the hippocampus is structured by theta phase such that each cell's spikes occur at progressively earlier theta phase as the animal runs through the cell's place field [137, 89]. There is a related, although dissociable progression at the ensemble level within theta cycles [41], in that the places decoded from the hippocampal population sweep from behind the animal towards future locations [46], also reflecting current and future choices [160, 55].

Gamma oscillations (25 – 150 Hz) are also prominent in the hippocampus in a variety of behavioral states [12, 29, 27]. In recent literature, gamma frequency oscillations have been classified into three overlapping frequency bands. Slow (25 – 60 Hz), mid (60 – 90 Hz) and fast (90 – 150 Hz) gamma frequency bands are locked to theta frequency, and occur at distinct phases of theta [28, 9, 42]. The theta phase-segregation of gamma frequency bands provides distinct windows for other brain areas to drive activity in CA1 [28, 97, 42], as well as for local circuit computations within the pyramidal cell layer [130]. Correlation studies have shown that during slow gamma events, CA1 activity is synchronized with CA3, and CA1 spiking during slow gamma events exhibits a functionally distinct representation of the animal's future locations [171]. Suppression of CA3 in a transgenic mouse [96] leads to disruption of gamma-associated temporal coding in CA1. In a recent study [30], acute silencing of CA3 demonstrated a vital role of CA3 in shaping place fields and SWR in CA1.

¹In collaboration with Dr. Heydar Davoudi and Althea Cavanaugh.

In this study, we aimed to examine the acute and causal contribution of CA3 input to slow and fast gamma oscillations in CA1.

It was recently observed that theta is not synchronized across CA1, but rather travels along the dorsoventral (longitudinal) axis, changing its polarity completely (a 180° phase shift) between the dorsal and ventral poles [85, 119]. The differential connectivity of the dorsal and ventral hippocampus suggests a functional role of the traveling wave [24]. It remains unknown whether gamma exhibits a similar anatomical profile, and how gamma and theta relate to each other at different anatomical positions along the dorsoventral axis.

Here we address the anatomical relationship between gamma and theta using (i) multi-tetrode recordings within dorsal CA1, (ii) simultaneous multi-tetrode recordings in both dorsal and intermediate CA1, and (iii) opto-genetic suppression of CA3 input to CA1. Our data suggest that, in contrast to the traveling wave of theta, both slow and fast gamma are synchronized in the hippocampus. We measure this synchrony first in the dorsal CA1 and then confirm that it is maintained as we move further along the dorsoventral axis into the intermediate hippocampus. Using optogenetics, we acutely and selectively suppressed CA3 input to CA1, in order to make a causal test of the role of CA3 with respect to gamma. We show that the suppression of CA3 input leads to de-synchronization of both slow and fast gamma events in dorsal CA1.

2.1 Materials and Methods

Part of the data that was used in this study was previously published in [30].

Experimental Procedures

Animals and Surgery

Procedures for studying the dorsal hippocampus were approved by Johns Hopkins Animal Care and Use Committee and followed US National Institutes of Health animal use guidelines. Prior to surgical procedures, animals were trained to run on a linear track for liquid chocolate reward while under a food restricted diet which allowed them to maintain over 90% of their free-feeding body weight. For viral injections and micro-drive implant, animals were anesthetized under 1 – 5% isoflurane in O_2 . Six male Long-Evans rats were used for studying dorsal CA1 and the effects of opto-genetic suppression of CA3 input to CA1. First, *AAV5 – CamKIIa – eArch3.0 – EYFP* (EXP) or *AAV5 – CamKIIa – EYFP* (CON) were injected at 6 different sites – Anterior-Posterior (AP), Medial-Lateral (ML), and Dorsal-Ventral (DV) coordinates (in mm) relative to Bregma: $(-3.1, \pm 3.5, -3.5)$, $(-4.0, \pm 4.3, -4)$, and $(-4.7, \pm 4.8, -4.8)$. 4 weeks after the viral injection, custom made micro-drive carrying up to 40 tetrode bundles and 2 fiber optic cables (one fiber per hemisphere) were implanted with fiber bundles situated at -4.0mm AP from Bregma and $\text{ML} \pm 2.85\text{mm}$.

Simultaneous recordings of the dorsal and intermediate hippocampus were performed under guidelines approved by Institutional Animal Care and Use Committee (IACUC) at

University of California, Berkeley. Custom micro-drives carrying 64 tetrode bundles were implanted into three male Long-Evans rats. Implants consisted of 2 cannulas, centered at -3.7mm AP and 2.3mm ML from Bregma for the dorsal site and -5.7mm AP and 5mm ML from Bregma for the intermediate site. Extremities of surgical coordinates for the implant cannula spanned 0.9mm ML, -3.0mm AP to 6.4mm ML, -6.6mm AP. Animals were allowed to rest and recover from the surgery for 1-3 weeks. During this period, tetrode bundles were gradually adjusted to target the hippocampal CA1. After clear Sharp-Wave Ripples (SWR) could be observed on a tetrode, small adjustments were made to target the pyramidal cell layer, characterized by bursts of high amplitude action potentials accompanying SWR. Tetrodes that had clear SWR during awake-rest (running speed $v < 5\text{cm/s}$) were selected for analysis.

Recording and Manipulation

For the dorsal hippocampus recordings, rats ran back and forth on a familiar 165cm long linear track for liquid chocolate reward. Raw electrode data was digitized at $32,000\text{Hz}$ and recorded using a Digital Lynx data acquisition system (Neuralynx). During traversals of the linear track, on alternate laps (starting at one end of the track and ending when the animal returned to it - the end at which stimulation started was randomized over recording sessions), green light (532nm , NI USB-6341) was delivered at 3.25mW .

For simultaneous recordings of the dorsal and intermediate hippocampus, animals ran either on a long (10m) linear track or a short (1m) linear track. Red and Green LED arrays were attached to the micro-drives to track the animal's position. Raw electrode data was digitized at $30,000\text{Hz}$ and recorded wirelessly using SpikeGadgets LLC data acquisition system and open-sourced recording software (Trodes). Camera stream was sampled at 29Hz . Animal's position was obtained by applying a threshold on the green and red color maps in the camera stream.

Data Analysis

LFP extraction and analysis

Local Field Potential (LFP) was extracted from one electrode from each tetrode bundle selected by visual inspection of the signal. For recordings from the dorsal hippocampus, the raw voltage data was sub-sampled to 3200 Hz and for the simultaneous dorsal and intermediate hippocampus recordings, voltage data was sub-sampled to 1500Hz . Theta was extracted from the LFP using a bandpass Butterworth filter of order 3 with frequency limits $5\text{-}10\text{Hz}$. Slow and fast gamma power envelopes were similarly calculated using $25\text{-}55\text{Hz}$ and $90\text{-}140\text{Hz}$ frequency bands. Instantaneous theta phase was extracted by applying Hilbert transform to the extracted theta waveform. Theta cycle boundaries were marked as time points where a phase roll-over could be seen in the Hilbert transform (corresponding to the

peaks in the waveform of a pure sinusoid). Consecutive theta boundaries that were within a theta period (100-200ms) were accepted.

Spectrogram analysis

A spectrogram was constructed for LFP data collected from each run segment (running speed $> 20\text{cm/s}$) using MATLAB's signal processing toolbox. We used a Hamming window of 320 samples (100ms) and a shift of 6 samples (1.875ms) between consecutive windows. Slow or fast gamma power within a theta cycle was calculated by averaging over all frequencies within the corresponding frequency band. Theta cycles with slow or fast gamma power 2 standard deviations over the mean were labelled as cycles with significant slow or fast gamma power respectively. For each gamma event (threshold crossing described above), the current theta cycle, later half of the previous cycle and first half of the next theta cycle were extracted. Because of the variability in theta period, the amplitude of the spectrogram was interpolated to fit a uniform sampling over two theta cycles (phase range $[-\pi, 3\pi]$) with 400 sample points corresponding to a period of 125ms.

Anatomical analysis of theta and gamma

Anatomical locations of recording electrodes within the dorsal CA1 were estimated using the implant designs. The implant cannulas were imaged. Within a hemisphere, the relative coordinates of all the recording sites were estimated with respect to the boundary of the implant. Comparisons were not made across hemispheres because of the variability in implant locations across hemispheres. For any pair of electrodes within the same hemisphere, the relative AP and ML coordinates were measured. Power envelopes were constructed by first filtering the raw LFP in the slow or fast gamma frequency band. The filtered signal was then squared and smoothed using a Gaussian filter with a 20ms standard deviation to obtain the power envelope. For each unique pair of electrodes, cross-correlograms were constructed for theta filtered LFP, slow gamma power envelope and fast gamma power envelope. The absolute value of lag at which a cross-correlogram attained its peak value was used for anatomical analysis. For a pair of electrodes to be included in the anatomical analysis, we required a peak theta-correlation cutoff of 0.6 and a peak correlation cutoff of 0.2 for both slow and fast gamma.

SWR detection

LFP signal was filtered in the SWR frequency band (150 – 250Hz). Ripple power envelope was obtained by smoothing the average power in the ripple band with either a 50ms moving average filter, or a Gaussian smoothing window with a 12.5ms standard deviation. 5 standard-deviation crossings over the mean ripple power were designated as ripple events. SWR suppression index was defined as $\left(\frac{\text{ON}-\text{OFF}}{\text{ON}+\text{OFF}}\right)$, where **ON** and **OFF** represent the SWR rates in the light on and light off conditions respectively.


Gamma event detection and suppression

Gamma events were defined as threshold crossings of the averaged gamma power spectrogram. Within the Phase-Frequency Response (PFR) measure, we first averaged power within gamma frequency band, then smoothed the resulting power envelope with a smoothing window of 20ms. Whenever the power envelope crossed 2 standard deviations over its mean value, the peak power time was denoted as a gamma event. The theta phase and absolute time at the event peak were recorded. A Mean Resultant Vector (MRV) was constructed with all the gamma events detected on a recording site for statistical comparison across light conditions and group data as follows

$$V = \sum_{k=1}^n p_k e^{i\phi_k}.$$

Here V is the resultant vector, p_k is the power of the k^{th} gamma event and ϕ_k is the theta phase of the k^{th} gamma event. This is a complex number with a magnitude which corresponds to the length of the MRV and a preferred theta phase.

Synchrony measurement



For any signal (LFP theta, slow or fast gamma power envelope), first, a cross-correlogram was constructed for lags ranging from $-60ms$ to $60ms$ for all the periods with running speed $> 20cm/s$ with a reference site r and a measurement site m . The peak value of the cross-correlogram and the lag at which this peak was detected were noted for the recording site pair. Let us call the peak $P(r, m)$ and absolute value of lag $L(r, m)$ for this pair of sites. For any reference r , the peak and lag were measured for all the other recording sites and the

average peak and average lag were reported as synchrony measures. Let's call these $\hat{P}(r)$ and $\hat{L}(r)$ respectively. These were defined as:

$$\hat{P}(r) = \sum_{m \neq r} P(r, m), \text{ and}$$

$$\hat{L}(r) = \sum_{m \neq r} L(r, m).$$

Statistical Analysis

Statistical tests were used to validate the differences in group data. One-tailed tests were used for signed comparisons between data groups (for example group A has a significantly smaller median than group B). Unless specified otherwise, non-parametric tests like Wilcoxon rank-sum test for testing difference in median values and Wilcoxon sign-rank test for testing if the median of a data group was significantly different from 0 were used. Pearson's correlation measure was used for testing if two continuous variables depended on each other. Further, in order to compare correlation coefficients, a Fisher z-transform was applied to the correlation coefficients and the resulting z-scores were compared for statistical significance. We used the circular statistics toolbox in MATLAB for testing dependence of variables on theta phase [10].

2.2 Theta-gamma interactions in the Hippocampus

In order to examine the gamma-theta relationship in CA1 and the contribution of CA3 to CA1 gamma activity, we bilaterally implanted six male Long-Evans rats with custom micro-drives each carrying 20 – 40 independently adjustable tetrode bundles and two optical fibers [30]. 4 – 6 weeks prior to the micro-drive implant surgery, the dorsal and intermediate CA3 sub-regions of four rats were bilaterally injected with light-activated proton pump *AA5 – CamKIIa – eArch3.0 – EYFP* (EXP) [166, 52]. Two rats were injected with *AAV5 – CamKIIa – EYFP* (CON).S Fiber-optic cables were located at the center of the implant cannula to target the pyramidal cell layer in CA1, so that light power could reach Stratum Radiatum (SR) below. Tetrodes were gradually lowered to the cell layer starting the day after the surgery and over a period of 2 – 4 weeks. Animals were maintained on a food-restricted diet to keep them over 90% of their free-feeding body weight. Prior to recording, animals received 5 – 7 consecutive days of pre-training where they ran 30 – 40 laps on a linear track for liquid chocolate reward. On each recording day, animals ran a single 15 – 30min long session on a highly familiar, 165cm long linear track for liquid chocolate reward (0.2ml, Carnation). Periods when the animals ran over 20cm/s were separated and used for subsequent analysis.

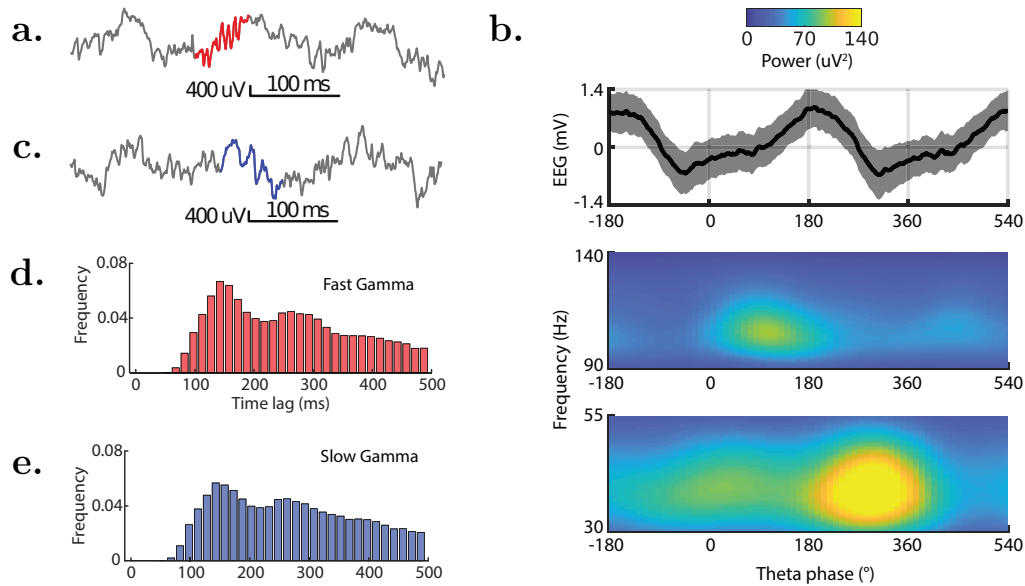


Figure 2.1: Gamma events are embedded within theta oscillations. (a.) Fast gamma event embedded within theta oscillations. (b.) Phase-Frequency Response (PFR) for a single recording site illustration theta phase locking for slow and fast gamma oscillations. (c.) Slow gamma event embedded within theta oscillations. (d.) Time-interval between consecutive fast-gamma events ($n = 147$ tetrodes with significant power in slow or fast gamma band, mean peak-lag 192ms). (e.) Time-interval between consecutive slow-gamma events ($n = 147$ tetrodes with significant power in slow or fast gamma band, mean peak-lag 171ms).




Figure 2.2: Statistics of slow and fast gamma events. (a.) Breakdown of theta cycles, $n_T = 273$ tetrodes, recorded $n_C = 744 \pm 147$ (mean \pm std.) theta cycles from individual tetrodes. (b.) Breakdown of recorded electrodes with significant theta locking ($n_T = 279$ tetrodes, $n_S = 178$ tetrodes have significant slow-gamma locking to theta, and $n_F = 201$ tetrodes have significant fast gamma locking to theta.)

Slow and fast gamma events were detected on recording electrodes (Figs. 2.1a. and 2.1c.). As reported previously [28], we found that both slow and fast gamma oscillations were locked to theta (Figs. 2.1d., 2.1e. and 2.2). As shown in Fig. 2.2a., roughly 25 – 30% of all theta cycles had prominent slow or fast gamma events. Of all the recording electrodes, a majority had significant gamma power-theta phase locking in the context of both slow and fast gamma oscillations (Fig. 2.2b.)². In over 40% of the recording electrodes, the distribution of slow and fast gamma events over theta phase had statistically significant distributions.

Moreover, slow gamma events occurred on the falling edge Fig. 2.1c. and fast gamma oscillations on the rising edge Fig. 2.1a. of theta waveform. We computed a spectrogram of the Local Field Potential (LFP) in the gamma frequency bands of interest. The spectrogram was clipped and averaged over two theta cycles. The measure PFR shown in Fig. 2.1b. also suggested that peaks in the gamma power envelope were synchronized with theta as previously suggested.

2.3 Gamma oscillations are synchronized




Figure 2.3: Illustrating synchronized slow gamma events recorded from the Hippocampus.

Simultaneous recording from sites distributed over 2mm in the dorsal CA1 revealed that there was a measurable spread in theta phase over the sites, consistent with spatially distributed samples from a traveling wave [85, 119]. Gamma events, in contrast, occurred synchronously. Fig. 2.3 illustrates simultaneous recordings from 4 electrodes in the dorsal CA1, highlighting the contrast between peaks in theta waveform (black traces, squares mark

²Rayleigh test for circular non-uniformity, $p < 0.05$

peaks) and gamma power envelopes (blue traces, circles mark peaks). A similar observation for fast gamma events is shown later in Fig. 2.7a..

We found that both slow and fast gamma power envelopes had a smaller distribution of lags when compared to theta as shown in Fig. 2.4 (One-tailed two-sample F-test, slow-gamma *vs.* theta $p = 7.8 \times 10^{-39}$, fast-gamma *vs.* theta $p = 7.3 \times 10^{-4}$). In order to reconcile




Figure 2.4: Distribution of theta and gamma lags for recording sites. Measured lags for (a.) theta, (b.) slow and (c.) fast gamma. One-tailed, two-sample F-test statistics: slow gamma *vs.* theta $p = 7.8 \times 10^{-39}$, confidence interval (c.i.) for variance ratio $c.i. = [4.26, 6.90]$, fast gamma *vs.* theta $p = 7.3 \times 10^{-4}$, $c.i. = [1.19, 1.93]$.

the phase-amplitude relationship between theta and gamma previously established with the synchronized gamma events illustrated here, we hypothesized that gamma events are locked to a global theta phase definition which is coherent in the hippocampus. We compared the locking of slow and fast gamma events to their local theta phase against a global theta phase defined by selecting one of the recording electrodes with strong theta oscillations to define theta phase. Fig. 2.5 illustrates the computation of the preferred theta phase and gamma-theta synchronization strength at any electrode. First a reference electrode was used to measure theta phase - this could be the same as the electrode on which gamma events were measured. The gamma-amplitude theta-phase PFR was constructed using the method previously described in Section 2.1 and shown here in Fig. 2.5a.. Averaging the PFR in the appropriate frequency band resulted in the mean amplitude for every phase value (Fig. 2.5b. - slow gamma, Fig. 2.5d. - fast gamma). Peaks detected on individual theta cycles were aggregated in theta phase and reported in Figs. 2.5b. and 2.5e. for slow and fast gamma respectively. These peaks have been referred to as gamma events in this study. Circular mean for gamma response as well as events was computed as described in Section 2.1 and referred to as the Mean Resultant Vector (MRV).

Fig. 2.6 shows the PFR measure for 4 sites recorded simultaneously from the hippocampus. Figs. 2.6a. to 2.6d. report the PFR when theta phase was measured locally. Notice the large variation in the phase at which peak slow and fast gamma power are observed. When we used a fixed theta phase definition for all the electrodes Figs. 2.18e. to 2.18h., as evident in the local theta waveform, the preferred theta phase of both slow and fast gamma power was more consistent.

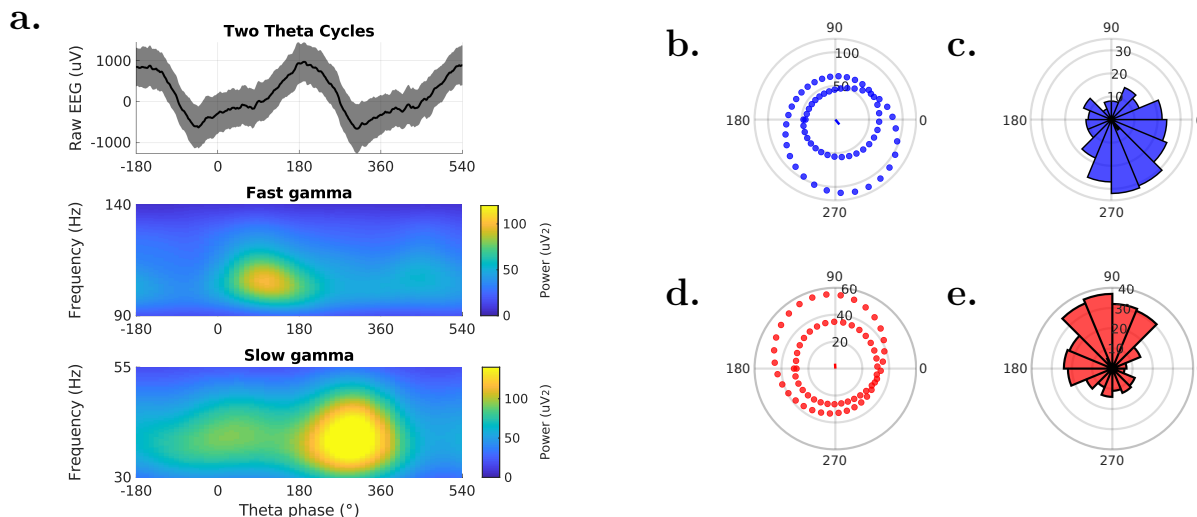


Figure 2.5: Construction of Phase-Frequency Response (PFR) and calculation of gamma-theta locking strength. (a.) Example phase frequency response from a one recording site over a single run session from Rat K. (b.) Average power in slow gamma frequency band as a function of theta phase. (c.) Slow gamma events extracted from the spectrogram (Rayleigh test for circular non-uniformity: $z = 28$, $p = 2.5 \times 10^{-13}$). (d.) Average power in fast gamma frequency band as a function of theta phase. (e.) Fast gamma events extracted from the spectrogram (Rayleigh test for circular non-uniformity: $z = 31$, $p = 1.2 \times 10^{-14}$).

Fig. 2.8 illustrates the locking strength of slow (blue) and fast (red) gamma events detected on different recording sites to either their local theta phase (Fig. 2.8a.), or a global theta phase (Fig. 2.8c.). The overall locking strength was measured by computing a MRV over all the recording sites for slow and fast gamma as shown in Fig. 2.5. We found that a global theta phase definition resulted in stronger locking than local theta phase definition for gamma events (Fig. 2.7b.). MRV measures for individual recording sessions are shown in ???. To ensure that the heightened locking using a global theta phase definition was not dependent on the electrode selected for phase definition, we defined theta phase using every recording site and compared the distribution of MRVs thus obtained against the corresponding MRV obtained using local theta phase from the same recording session, with the same result (One-tailed Wilcoxon signed-rank test statistics comparing the difference between global MRV score and local MRV score against a null distribution centered at 0 for $n = 279$ electrodes: slow gamma, Fig. 2.8b., $p = 3.7 \times 10^{-44}$, fast gamma, Fig. 2.8d., $p = 3.4 \times 10^{-13}$).

In order to verify that the shifts in theta we observed could be attributed to the traveling wave and not variation in recording depth, we calculated the lag in theta between pairs of recording sites and measured this lag against the anatomical distance between the sites. Figs. 2.10a. to 2.10c. illustrate cross-correlograms measured for theta filtered Local Field Potential (LFP), slow gamma power and fast gamma power envelopes at all the recording

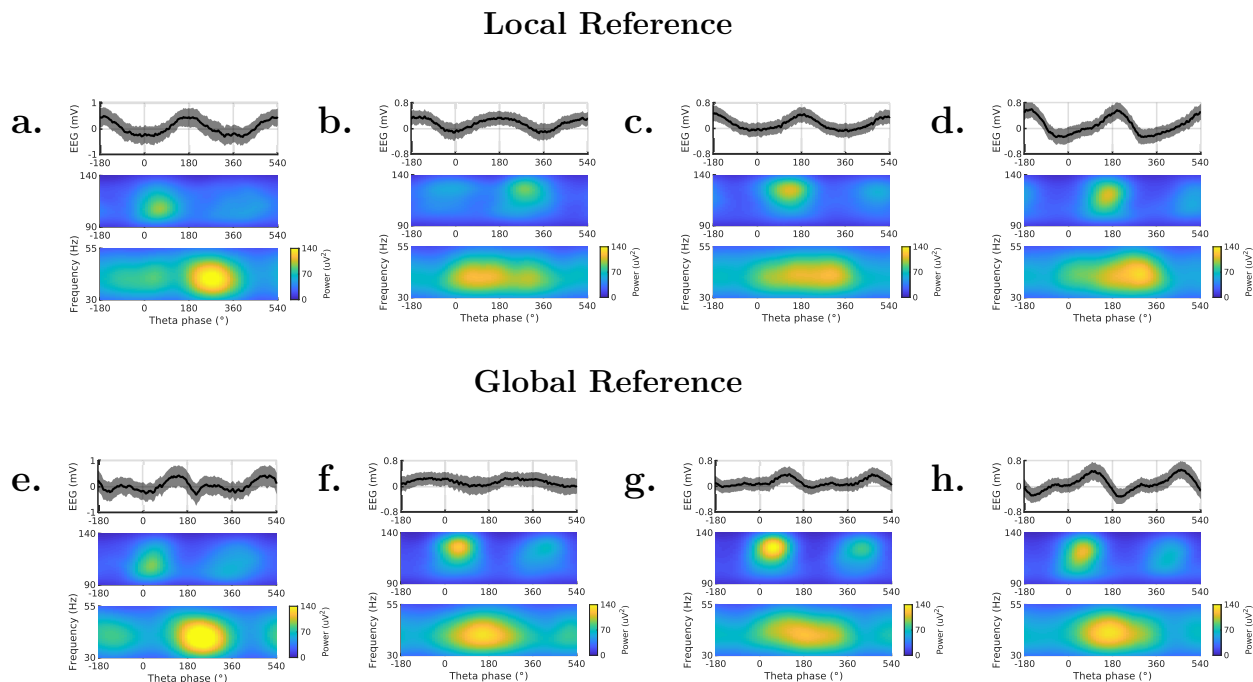


Figure 2.6: Local *vs* Global theta reference - comparing PFR produced by local and global referencing for four different recording sites from one animal. (a.) T2 (Local reference), (b.) T22 (Local reference), (c.) T26 (Local reference), and (d.) T27 (Local reference). (e.) T2 (Global reference), (f.) T22 (Global reference), (g.) T26 (Global reference), and (h.) T27 (Global reference).

sites relative to a single reference. These cross-correlograms were used to calculate the aforementioned lag. For pairs of sites located within the same hemisphere, we found that theta-lag and anatomical distance were significantly correlated (Fig. 2.10d., $n = 223$ electrode pairs with peak cross-correlation above 0.6, $r = 0.33$, $p = 4.5 \times 10^{-7}$). When comparing the anatomical profile observed for theta against that for slow and fast gamma power envelopes, we observed that such a relationship could not be seen for slow or fast gamma lags (Pearson correlations: Slow gamma, Fig. 2.10e., $r = -1.5 \times 10^{-4}$, $p=0.98$. Fast gamma: Fig. 2.10f., $r = 0.07$, $p = 0.33$). On comparing the correlation measures shown in Figs. 2.10d. to 2.10f., we found that the correlation for theta-lag with anatomical distance was significantly larger than that for slow and fast gamma power envelopes (One-tailed comparison of Fisher z -transformed correlation coefficients - theta *vs.* slow gamma $p = 1.1 \times 10^{-4}$, test-statistic 3.69, theta *vs.* fast gamma $p = 2.7 \times 10^{-3}$, test-statistic 2.78).

Fig. 2.11 shows comparisons for theta and gamma lags paired by recording sites. While there was no significant relationship between slow gamma lag and theta lag (Fig. 2.11a.) or fast gamma lag and theta lag (Fig. 2.11b.), we observed a significant positive correlation between peak slow gamma and peak theta correlation (Fig. 2.11c., $r = 0.22$, $p = 5.3 \times 10^{-37}$) and peak fast gamma and peak theta correlation (Fig. 2.11d., $r = 0.18$, $p = 1.2 \times 10^{-24}$).

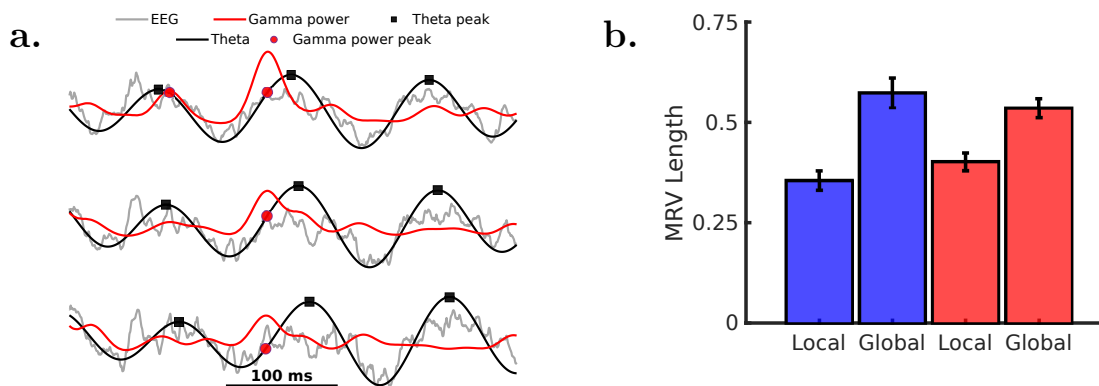


Figure 2.7: (a.) Illustration of synchronized fast gamma events from a recording session. (b.) Comparing local and global MRV Lengths over discrete gamma events in individual sessions (One-tailed, paired Wilcoxon sign-rank test comparing Local *vs.* Global MRV Lengths over $n = 13$ sessions: Slow Gamma $p = 0.02$, Fast Gamma $p = 0.01$).

Our findings suggested that within the dorsal CA1, there is variability in theta phase which can be attributed to the anatomical traveling wave nature of theta. In contrast, slow and fast gamma power envelopes are synchronized. While slow and fast gamma events are locked to theta frequency, they occur synchronously, and are locked to an abstract global theta phase definition.

2.4 Measurements along the dorsoventral axis

Synchronized gamma events observed within dorsal CA1 led us to hypothesize that this synchrony may be preserved along the dorsoventral (longitudinal) axis. In order to test this hypothesis, we simultaneously recorded from the dorsal and intermediate hippocampus while rats ran trials along a linear track. Hippocampal recordings spanned roughly half of the dorsoventral axis of the hippocampus. Three male Long-Evans rats (9 recording sessions) were unilaterally implanted with custom micro-drives carrying 64 independently adjustable tetrode bundles as shown in Fig. 2.12. 32 tetrode bundles targeted the dorsal CA1, whereas the remaining 32 tetrode bundles targeted the intermediate CA1. Tetrode bundles were independently adjusted to target the pyramidal cell layer in the respective regions. Subsequently, tetrodes that had clear Sharp-Wave Ripples (SWR) during awake-rest on the track (running speed $v < 5\text{cm/s}$) were selected for analysis.

We observed that gamma events were synchronized across the recording sites in dorsal and intermediate CA1. Fig. 2.13a. illustrates a synchronized slow gamma event across the dorsal (blue) and intermediate (orange) CA1. Theta filtered Local Field Potential (LFP) and both slow gamma power envelope for the same event are shown in Fig. 2.13b. In order to quantify the synchrony of gamma events and measure it against theta, we categorized

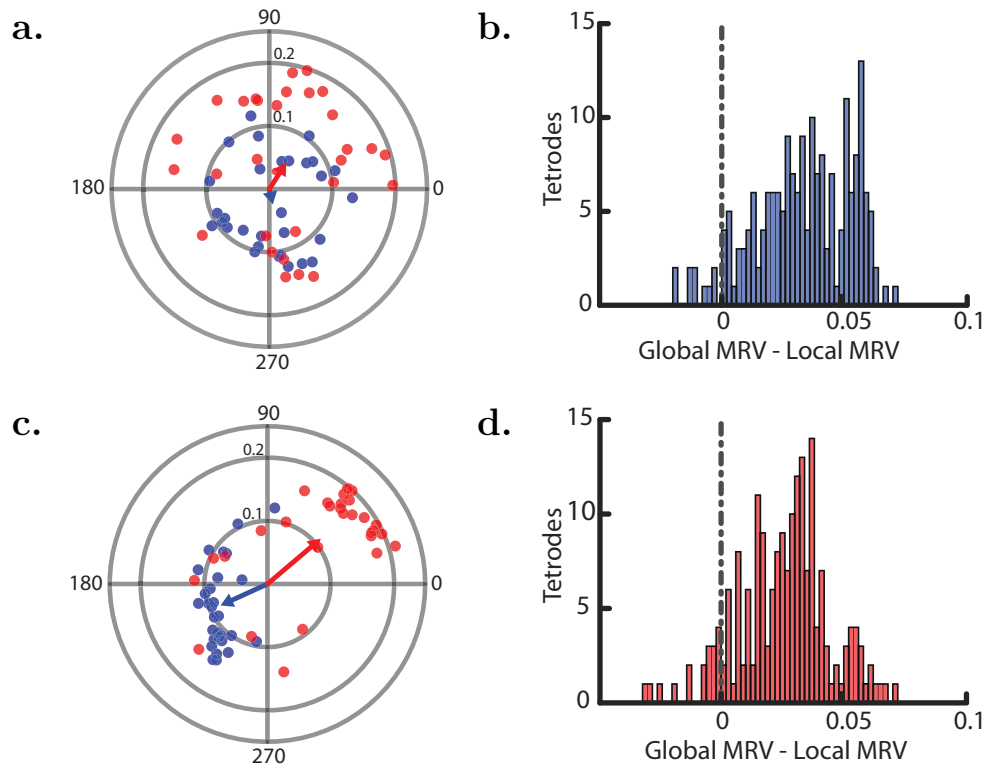


Figure 2.8: Gamma events are referenced to a global theta phase (a.) Preferred phase of slow and fast gamma events referred to their local theta phase. (b.) Difference between global and local MRV difference for slow gamma. (c.) Preferred phase of slow and fast gamma events referred to a global theta phase definition. (d.) Difference between global and local MRV difference for fast gamma.

any pair of electrodes into one of three categories - 1) both electrodes in the pair were in the dorsal CA1, 2) both electrodes in the pair were in the intermediate CA1, and 3) one of the electrodes was in the dorsal CA1 while the other was in intermediate CA1. We found that for all three groups, the lag in theta was significantly larger than that for slow and fast gamma power envelopes (Fig. 2.13c., Paired Wilcoxon signed-rank test statistics considering electrode pairs with peak theta-correlation > 0.6 and peak gamma-correlation > 0.2 : Within the dorsal CA1, $n = 324$ electrode pairs - theta *vs.* slow gamma $p = 2.5 \times 10^{-50}$, theta *vs.* fast gamma $p = 4.3 \times 10^{-48}$. Within the intermediate CA1, $n = 244$ electrode pairs - theta *vs.* slow gamma $p = 8.1 \times 10^{-35}$, theta *vs.* fast gamma $p = 2.2 \times 10^{-9}$).

If gamma events were synchronized across the regions but theta phase had a systematic offset, we predicted a significantly larger difference between theta and gamma lags when comparing across regions than when measured within an individual region. We compared the differences of the paired lags in Fig. 2.13d. and found that when compared to slow and fast gamma, the difference in theta lags was larger across the regions than within either

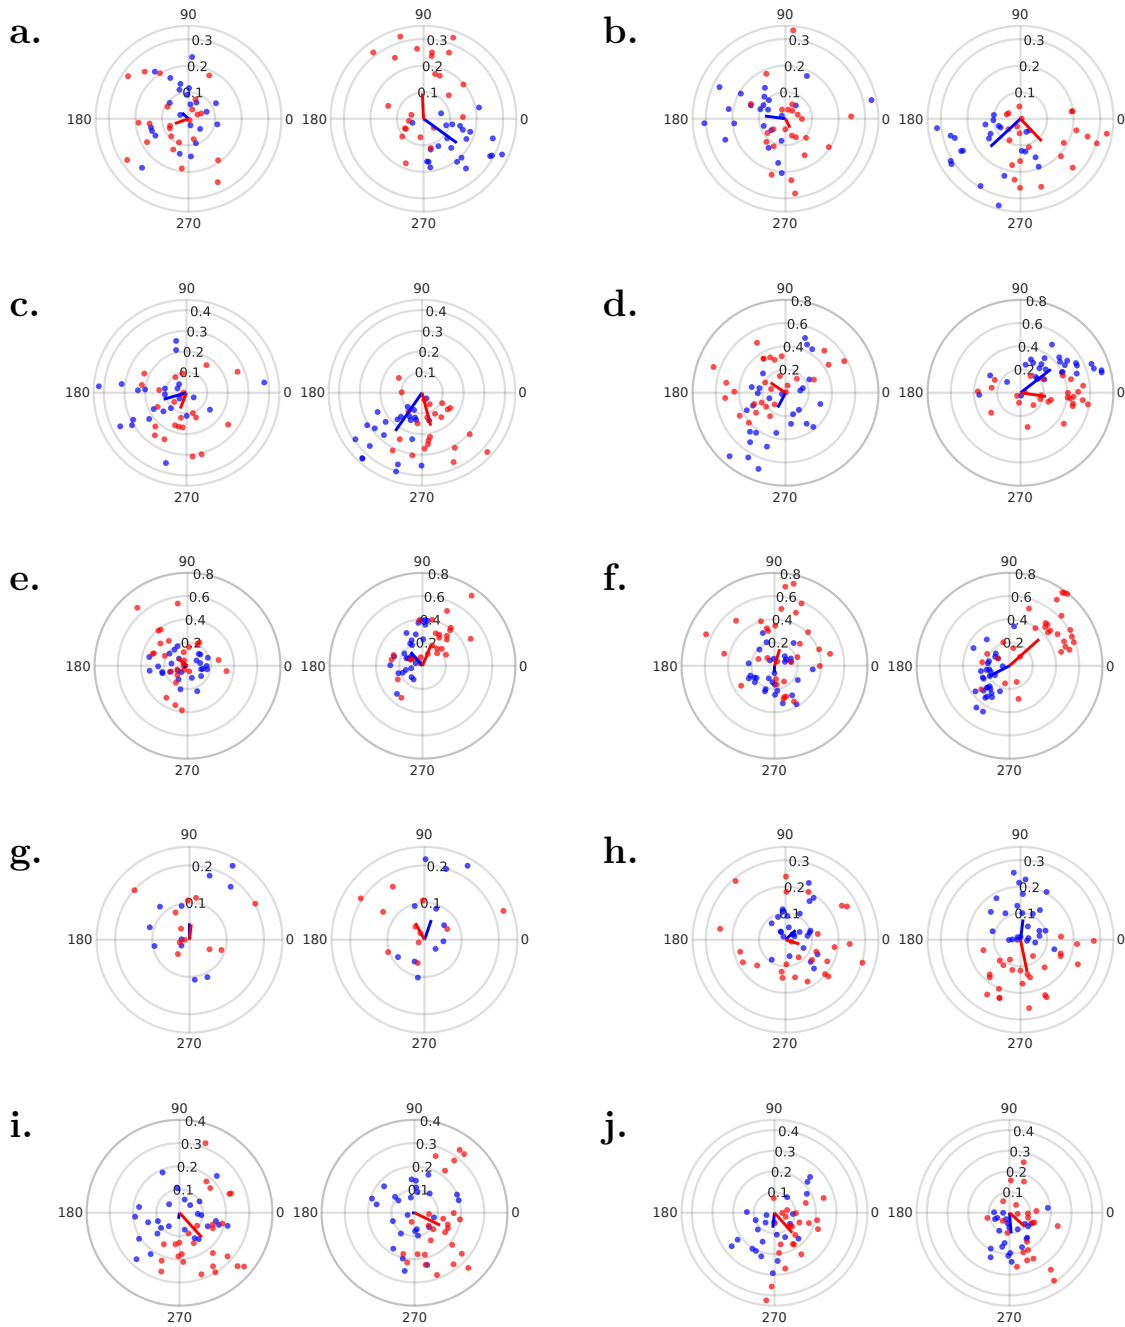


Figure 2.9: Local *vs.* Global referencing for individual sessions. Showing MRV constructed for slow (blue) and fast (red) gamma events for all the recorded sessions. For each session, locally reference gamma event MRV (left) and the same global event MRV (right) are shown. (a.) Rat I, S1, (b.) Rat I, S2, (c.) Rat I, S3, (d.) Rat K, S1, (e.) Rat K, S2, (f.) Rat K, S3, (g.) Rat G, S1. (h.) Rat J, S1. (i.) Rat J, S2. (j.) Rat J, S3.

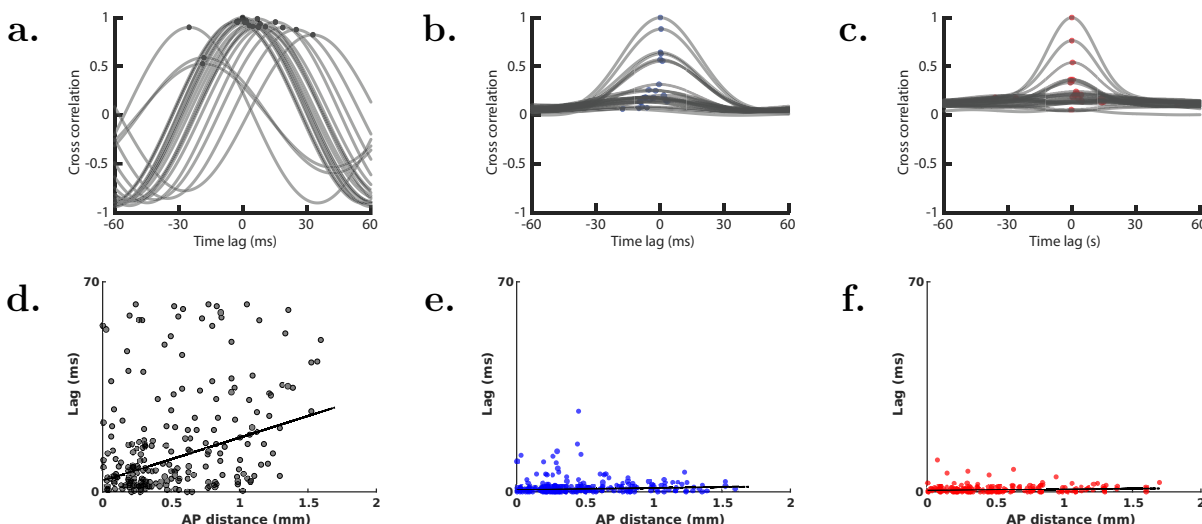


Figure 2.10: Observed oscillation lags as a function of anatomical distance between pairs of recording electrodes. (a.) Measured theta lag for all recording sites relative to a single reference in a recording session. (b.) Measured slow gamma lag for all recording sites relative to a single reference in a recording session. (c.) Measured fast gamma lag for all recording sites relative to a single reference in a recording session. (d.) Theta lag *vs.* Anterior-Posterior (AP) distance (Pearson's correlation $r = 0.33$, $p = 4.5 \times 10^{-7}$). (e.) Slow gamma lag *vs.* AP distance (Pearson's correlation $r = 0.002$, $p = 0.98$). (f.) Fast gamma lag *vs.* AP distance (Pearson's correlation $r = 0.07$, $p = 0.33$).

individual region (Wilcoxon rank-sum test statistics: Theta - slow gamma lag, inter-region *vs.* within dorsal CA1 $p = 0.002$, inter-region *vs.* within intermediate CA1 $p = 1.9 \times 10^{-7}$. Theta - fast gamma lag, inter-region *vs.* within dorsal CA1 $p = 9.4 \times 10^{-5}$, inter-region *vs.* within intermediate CA1 $p = 5.7 \times 10^{-11}$).

These findings suggested that as the theta wave travels along the dorsoventral axis, it creates a larger separation in the instantaneous theta phase between the dorsal and intermediate CA1 than within any individual region. Slow and fast gamma do not match this instantaneous phase separation, and instead, appear to have far smaller timing difference across regions, consistent with being synchronized.

2.5 Role of CA3 in generating gamma oscillations

We next examined the role of CA3 in generating slow gamma events, as hypothesized in the literature [29, 28]. In order to study the immediate role of CA3 in driving the gamma-theta relationship in the hippocampus, we used selective opto-genetic suppression of CA3 input onto CA1 during alternate laps of the linear track traversal. Fig. 2.14 illustrates a schematic of the suppression protocol. After the proton-pump inhibitor eArch3.0 was expressed in the

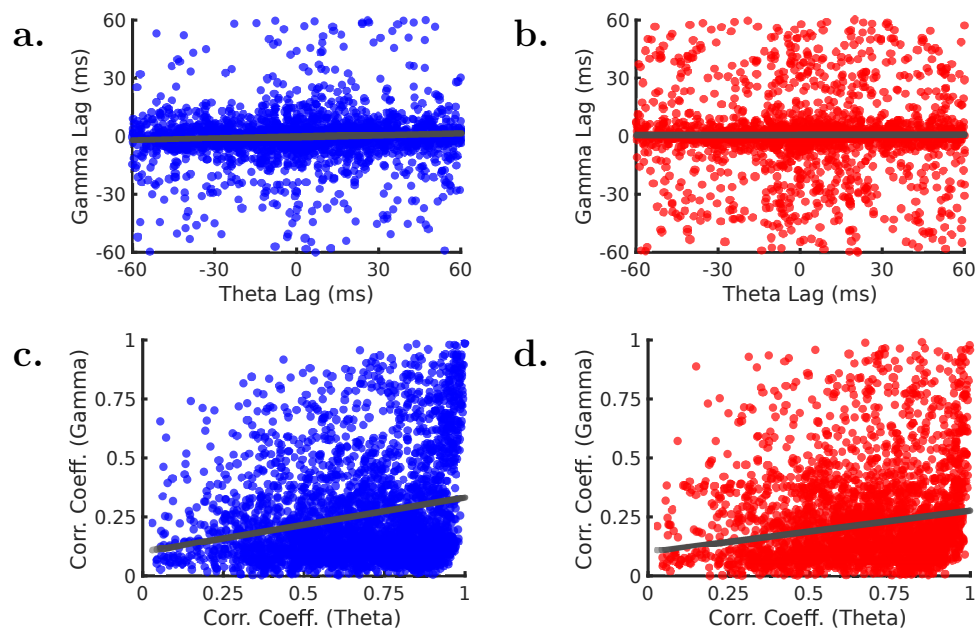


Figure 2.11: Relating lag and correlation between theta and gamma. Reporting paired Gamma-Theta measurements across $n = 3152$ pairs of recording sites. (a.) Slow gamma lag vs. Theta lag (Pearson's correlation $r = 0.005$, $p = 0.001$). (b.) Fast gamma lag vs. Theta lag (Pearson's correlation $r = 0.003$, $p = 0.85$). (c.) Slow gamma correlation coefficient vs. Theta correlation coefficient (Pearson's correlation $r = 0.22$, $p = 5.3 \times 10^{-37}$). (d.) Fast gamma correlation coefficient vs. Theta correlation coefficient (Pearson's correlation $r = 0.18$, $p = 1.2 \times 10^{-24}$).




Figure 2.12: Schematic of simultaneous recordings from dorsal and intermediate hippocampus.

cell bodies and axons of CA3 neurons, shining light in the CA1 region selectively suppressed neurotransmitter release in CA1.

While prior literature suggested a dominant role of CA3 specifically in slow gamma oscillations, we found that slow gamma oscillations largely persisted after CA3 input was suppressed. We observed a differential effect of the suppression of CA3 input on overall

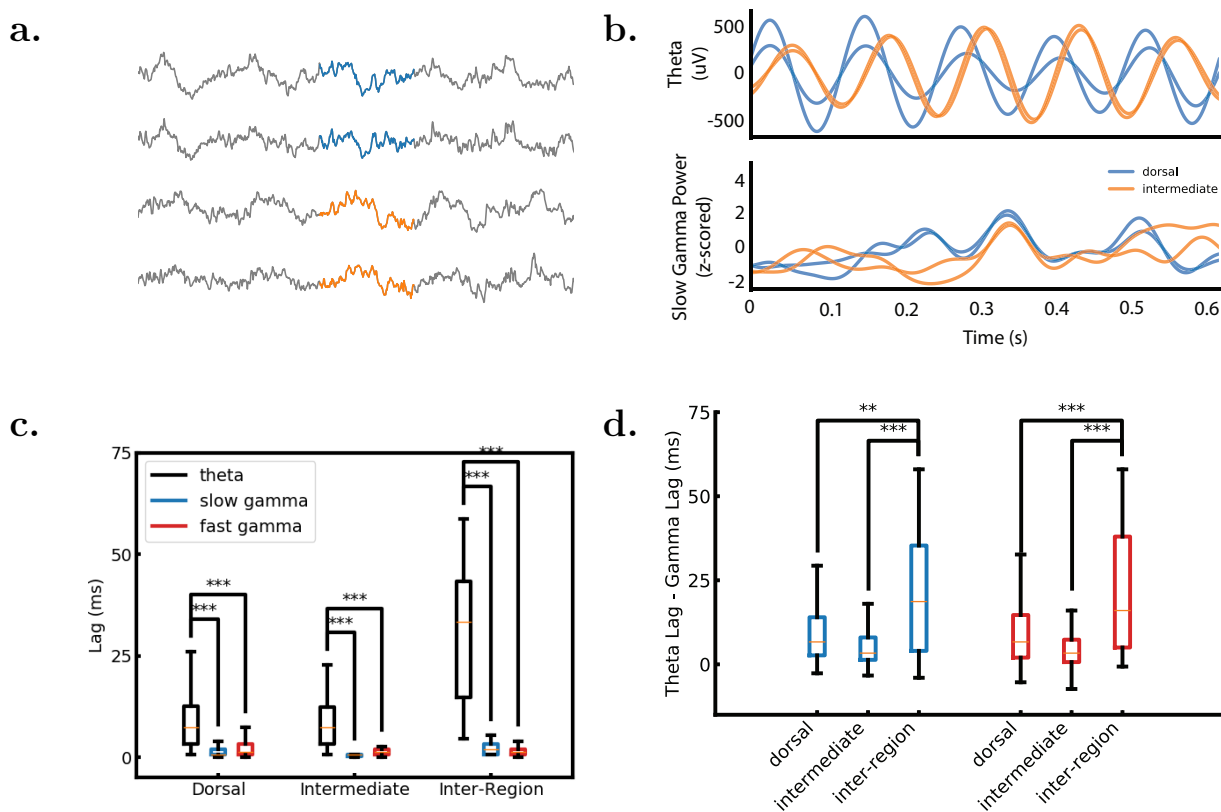


Figure 2.13: Gamma synchrony is preserved along the longitudinal axis of the Hippocampus. (a.) Schematic showing simultaneous recordings of dorsal and intermediate hippocampus. (b.) Simultaneous recordings from electrodes shown in dorsal (blue) and intermediate (orange) CA1 with a slow gamma event. Top: Raw LFP recorded on the electrodes with highlighted gamma event. Bottom: Theta filtered LFP and slow gamma power envelope. (c.) Measured lags for electrode pairs within dorsal hippocampus, within intermediate hippocampus, and inter-regional pairs for theta filtered LFP, slow gamma power envelope and fast gamma power envelope. Wilcoxon signed-rank test for paired theta-gamma lags: Dorsal hippocampus, $n = 324$ electrode pairs, slow gamma $statistic = 8.4 \times 10^2$, $p = 2.5 \times 10^{-50}$, fast gamma $statistic = 1.1 \times 10^3$, $p = 4.3 \times 10^{-48}$. Intermediate hippocampus, $n = 244$ electrode pairs, slow gamma $statistic = 8.2 \times 10^2$, $p = 8.1 \times 10^{-35}$, Fast gamma $statistic = 2.5 \times 10^3$, $p = 5.8 \times 10^{-26}$. Inter-region lags: $n = 51$ electrode pairs, slow gamma $statistic = 50$, $p = 8.9 \times 10^{-9}$, Fast gamma $statistic = 4.5$, $p = 2.2 \times 10^{-9}$. (d.) Statistical comparison between theta lag and gamma lag within and across regions. Wilcoxon rank-sum test statistics: Slow gamma (dorsal vs. inter-region lags) $statistic = -3.0$, $p = 0.002$, (intermediate vs. inter-region lag) $statistic = -5.2$, $p = 1.9 \times 10^{-7}$. Fast gamma (dorsal vs. inter-region lags) $statistic = 3.9$, $p = 9.4 \times 10^{-5}$, (intermediate vs. inter-region lags) $statistic = -6.6$, 5.7×10^{-11} .




Figure 2.14: Schematic illustrating acute and selective opto-genetic suppression of CA3 input to CA1.

gamma power and the incidence of gamma events in dorsal CA1 (Figs. 2.15 and 2.16). There was a significant decrease in fast gamma power during running periods in the EXP animals across light conditions (Fig. 2.15c.). We did not observe significant changes in slow gamma power in either groups across light conditions (Fig. 2.15a.). Within the experimental group, the reduction in fast gamma power was significantly larger than that in slow gamma power (One-tailed Wilcoxon rank-sum test, $p = 1.4 \times 10^{-6}$). Furthermore, in the experimental group, reduction in fast gamma power during running periods was correlated with a reduction in SWR rate during periods of awake-rest which in turn has been shown to be correlated with the efficacy of light delivery [30] (Fig. 2.15d., Pearson's correlation $r = 0.47$, $p = 3.8 \times 10^{-7}$.) We did not find such a correlation for slow gamma power (Fig. 2.15b., Pearson's correlation $r = 0.11$, $p = 0.24$).

We did not observe an overall decrease in the rate of slow (Fig. 2.16a.), or fast (Fig. 2.16c.) gamma events. However, within the experimental group, there was a significantly higher decrease in slow gamma events when compared to fast gamma events (One-tailed Wilcoxon rank-sum test $p = 0.03$).

Since CA3 and CA1 Local Field Potential (LFP) are strongly correlated during slow gamma events [28], and yet an acute suppression of CA3 activity did not result in suppression of slow gamma events, we hypothesized that CA3 input is required for the precise locking of slow gamma events to theta phase. Fig. 2.17 shows the Phase-Frequency Response (PFR) for a recording site in the presence (Fig. 2.17 - left) and absence (Fig. 2.17 - right) of CA3 input.

Computing the theta phase for gamma events extracted from the PFR on the recording site, we observed that in the presence of CA3 input, both slow (Fig. 2.19a.) and fast (Fig. 2.19c.) gamma events were locked to theta phase (Rayleigh test statistics: Slow gamma $z = 6.2$, $p = 0.002$, Fast gamma $z = 1.1 \times 10^2$, $p = 1.8 \times 10^{-54}$). For the same recording site, in the absence of CA3 input, slow gamma events (Fig. 2.19b.) were no longer locked to theta phase, however, fast gamma events (Fig. 2.19d.) remained locked (Rayleigh test statistics:

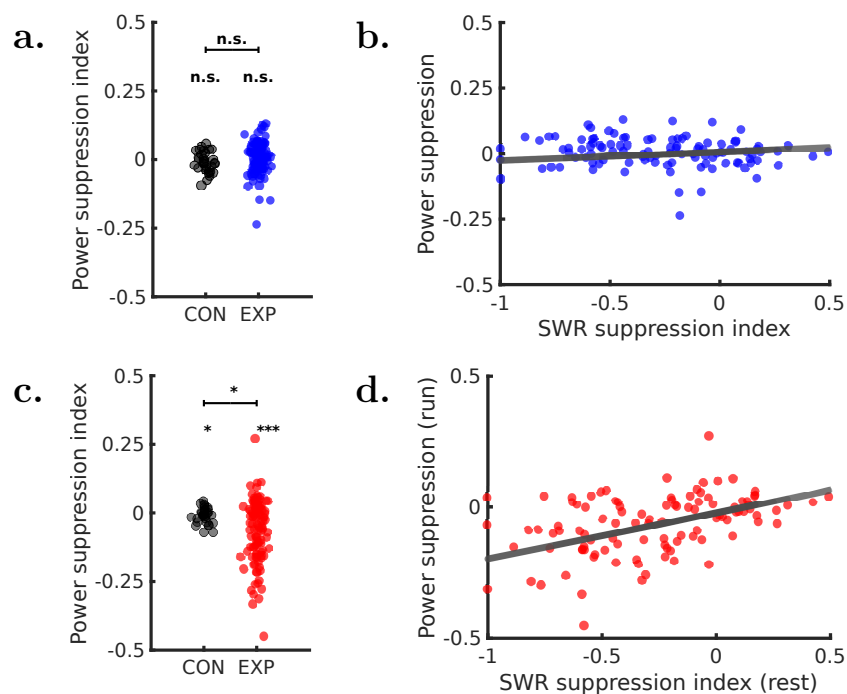


Figure 2.15: Suppressing CA3 input to CA1 - Changes in gamma power and events. (a.) Comparing change in average power in slow gamma frequency band during run between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -1.4$, $p_C = 0.08$, $n_E = 115$ tetrodes, $z_E = 1.2$, $p_E = 0.88$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -1.7$, $p = 0.09$). (b.) Measuring change in slow gamma power as a function of Sharp-Wave Ripples (SWR) suppression index (Pearson’s correlation, $r = 0.11$, $p = 0.24$) in the experimental group. (c.) Comparing average power in fast gamma frequency band between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -1.7$, $p_C = 0.04$, $n_E = 115$ tetrodes, $z_E = -5.0$, $p = 3.1 \times 10^{-7}$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -2.0$, $p = 0.05$). (d.) Measuring change in fast gamma power as a function of SWR suppression index ($r = 0.47$, $p = 3.8 \times 10^{-7}$) in the experimental group.

Slow gamma $z = 0.2$, $p = 0.82$, Fast gamma $z = 90$, $p = 1.6 \times 10^{-45}$). In order to quantify the locking of gamma events statistically, we computed an Mean Resultant Vector (MRV) for gamma events detected on each recording site and evaluated if the strength of gamma-theta locking, as measured by the MRV length degraded with suppression of CA3 input. We found that there was a selective decrease in the MRV length for slow gamma events in the *EXP* group (Fig. 2.20a. One-tailed Wilcoxon signed-rank test between light conditions $p = 0.001$). Comparing the change in MRV for slow and fast gamma within the *EXP* group, the MRV for slow gamma events showed more degradation than fast gamma events (One-tailed Wilcoxon signed-rank test, $p = 0.004$). We also correlated the change in MRV

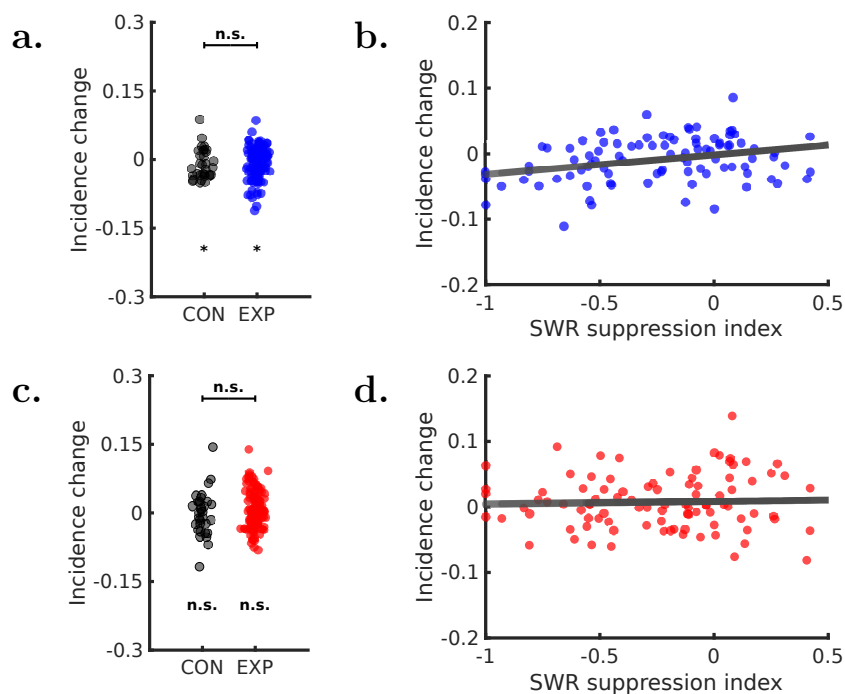


Figure 2.16: Suppressing CA3 input to CA1 - Changes in gamma events. (a.) Comparing the change in rate of slow gamma events between control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -2.1$, $p_C = 0.02$, $n_E = 115$ tetrodes, $z_E = -2.0$, $p_E = 0.02$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -0.88$, $p = 0.38$). (b.) Measuring change in rate of slow gamma events as a function of SWR suppression index (Pearson's correlation $r = 0.36$, $p = 3.0 \times 10^{-4}$) for the experimental group. (c.) Comparing the change in rate of fast gamma events Across control and experimental groups (One-tailed Wilcoxon signed-rank test, $n_C = 32$ tetrodes, $z_C = -0.33$, $p_C = 0.37$, $n_E = 115$ tetrodes, $z_E = 1.3$, $p_E = 0.90$, Wilcoxon rank-sum test *EXP vs. CON*, $z = -0.89$, $p = 0.37$), (d.) Measuring change in rate of fast gamma events as a function of SWR suppression index (Pearson's correlation $r = 0.04$, $p = 0.70$.) for the experimental group. Comparing slow *vs.* fast gamma event incidence change within groups: *CON*, $z_C = -1.0$, $p_C = 0.32$, *EXP*, $z_E = -2.2$, $p_E = 0.03$.

length for gamma events during run with the suppression of SWR events during awake rest. Here we found that for slow gamma oscillations, there was a significant correlation between the reduction in MRV length and SWR suppression (Fig. 2.20b. Pearson's correlation: $r = 0.28$, $p = 1.9 \times 10^{-3}$). The same relationship was not statistically significant for fast gamma oscillations (Fig. 2.20d. Pearson's correlation: $r = -0.03$, $p = 0.76$). A paired comparison of correlation values suggested that the change in slow gamma MRV was more strongly correlated with SWR suppression than the change in fast gamma MRV (One-tailed comparison of Fisher z-transformed correlation coefficients - $p = 3.0 \times 10^{-4}$, test-statistic

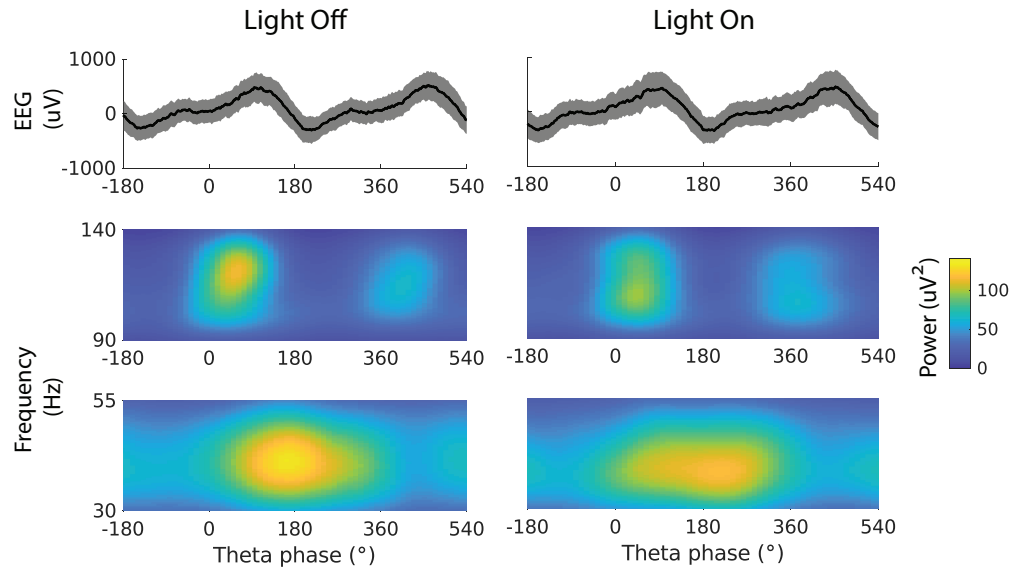


Figure 2.17: Changes in PFR resulting from suppression of CA3 input to CA1. (left) PFR during running periods under Control condition (Light Off). (right) PFR during running periods under suppression of CA3 input (Light ON).

3.53). While there were differences in the *EXP* and *CON* groups in theta phase offsets for slow and fast gamma (Figs. 2.21a., 2.21b., 2.21d. and 2.21e.), we did not find a relationship between this phase offset and extent of CA3 suppression (Figs. 2.21c. and 2.21f.).

These findings suggested that the CA3 input modulates both slow and fast gamma oscillations in the hippocampus. The nature of these modulations is different for slow and fast gamma frequency bands. For fast gamma, we observed a suppression dependent reduction in overall power. On the other hand, for slow gamma, we observed that suppression of CA3 input degraded the precision of gamma-theta locking at a recording site and the extent of this degradation was commensurate with the extent to which CA3 input to the site was suppressed.

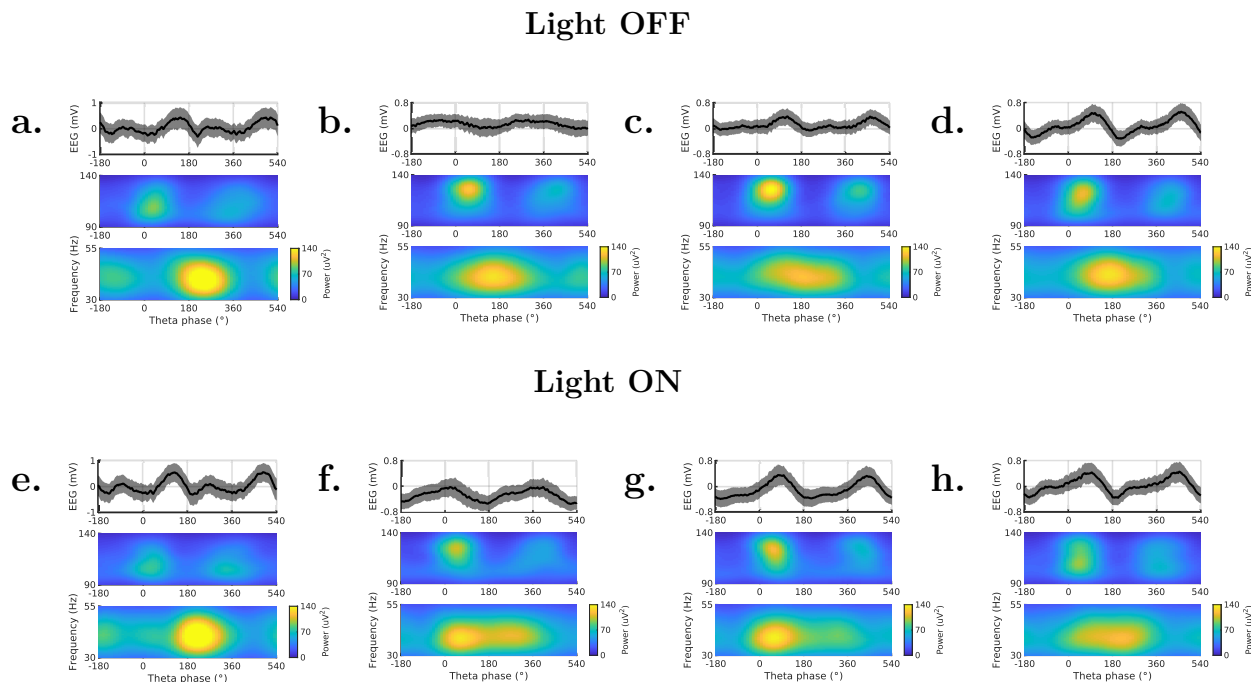


Figure 2.18: Comparing PFR four different recording sites from one animal in the *EXP* group across light conditions. (a.) T2 (Light OFF), (b.) T22 (Light OFF), (c.) T26 (Light OFF), and (d.) T27 (Light OFF). (e.) T2 (Light ON), (f.) T22 (Light ON), (g.) T26 (Light ON), and (h.) T27 (Light ON).

2.6 Role of CA3 in synchronizing gamma oscillations

Noting the dependence of gamma-theta locking on CA3 input, we hypothesized that the synchrony of gamma events within CA1 should, in turn, exhibit a dependence on CA3 input. We first defined measures of gamma synchrony in the dorsal CA1, and thereafter, studied their response to suppression of CA3 input, which was quantified as previously by the degree of Sharp-Wave Ripples (SWR) suppression during periods of awake-rest. Fig. 2.22 illustrates two such measures of synchrony: (i) the mean absolute lag of a recording site from all other sites; and (ii) the mean peak cross-correlation of a recording site with all the other sites. For the sites that were synchronized, we expected a small mean lag and a large mean peak cross-correlation. Furthermore, to the extent that this synchronization depended upon CA3 input, we expected the mean lag to increase and the mean peak cross-correlation to decrease as CA3 input was suppressed. For both slow and fast gamma oscillations, we observed that the mean absolute lag increased in the experimental group when compared to controls. Paired Wilcoxon signed-rank test statistics comparing light conditions - $n_E = 193$ electrodes in the experimental group and $n_C = 86$ electrodes in the control group: (Fig. 2.24c.), slow gamma, $p_E = 0.01$, $p_C = 0.01$, (Fig. 2.25c.) fast gamma $p_E = 0.08$, $p_C = 0.01$. Wilcoxon rank-sum test comparing experimental and control groups: slow gamma $p_s = 7.7 \times 10^{-8}$, fast gamma

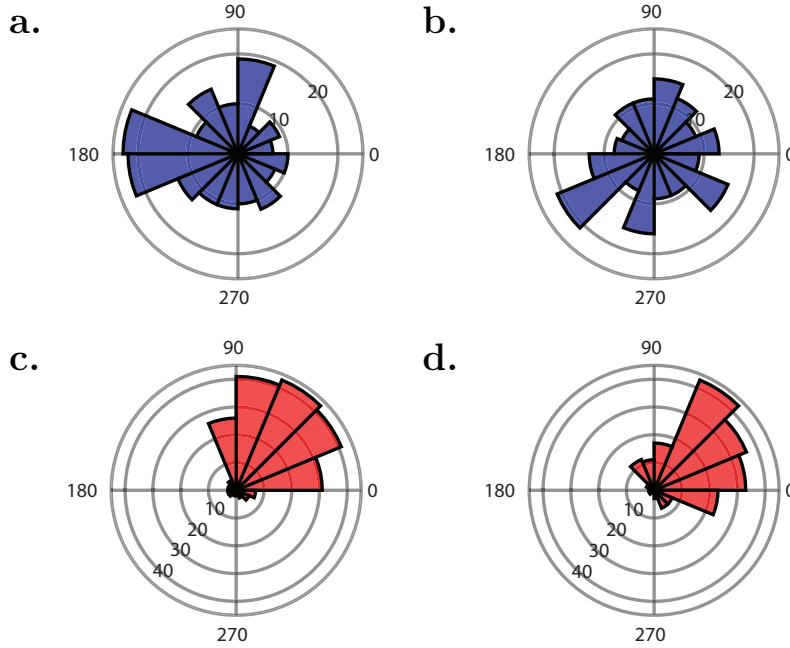


Figure 2.19: Changes in theta phase preference of gamma events. (a.) Theta phase distribution of slow gamma events during control (light off) condition. Rayleigh test statistics $z = 6.2$, $p = 0.002$. (b.) Theta phase distribution of slow gamma events during experimental (light on) condition. Rayleigh test statistics $z = 0.2$, $p = 0.82$. (c.) Theta phase distribution of fast gamma events during control (light off) condition. Rayleigh test statistics $z = 1.1 \times 10^2$, $p = 1.8 \times 10^{-54}$. (d.) Theta phase distribution of fast gamma events during experimental (light on) condition. Rayleigh test statistics $z = 90$, $p = 1.6 \times 10^{-45}$.

$p_f = 1.1 \times 10^{-5}$. Simultaneously, there was a significant decrease in the mean peak cross-correlation across the groups for both slow and fast gamma. Paired Wilcoxon signed-rank test statistics comparing light conditions - $n_E = 193$ electrodes in the experimental group and $n_C = 86$ electrodes in the control group: (Fig. 2.24a.), slow gamma $p_E = 2.0 \times 10^{-8}$, $p_C = 6.7 \times 10^{-7}$, (Fig. 2.25a.), fast gamma $p_E = 1.6 \times 10^{-4}$, $p_C = 0.6$. Wilcoxon rank-sum test comparing experimental and control groups: slow gamma $p_s = 2.6 \times 10^{-16}$, fast gamma $p_f = 1.0 \times 10^{-4}$.

We also tested if the loss of synchrony at a particular recording site, as measured by the decrease in mean peak cross-correlation, depended on the extent to which CA3 input to the site had been suppressed. By correlating mean peak cross-correlation change with SWR suppression, we found a strong relationship between the desynchronization of a recording site and the extent to which CA3 input to it had been suppressed for both slow and fast gamma - Pearson's correlation measures: Slow gamma (Fig. 2.24b.) $r = 0.24$, $p = 0.001$, Fast gamma (Fig. 2.25b.) $r = 0.36$, $p = 1.2 \times 10^{-6}$. In the experimental group we found that the increase in mean absolute lag was also correlation with SWR suppression. Pearson's correlations:




Figure 2.20: (a.) Change in MRV for slow gamma between control and experimental groups. One-tailed paired Wilcoxon signed-rank test statistics: $n_C = 38$ tetrodes, $z_C = -1.5$, $p_C = 0.06$. $n_E = 133$ tetrodes, $z_E = -3.0$, $p_E = 0.001$. Wilcoxon rank-sum test statistics across groups: $z = 0.03$, $p = 0.97$. (b.) Correlating change in slow gamma MRV in the experimental group as a function of SWR suppression during awake-rest (Pearson's Correlation: $r = 0.34$, $p = 1.4 \times 10^{-4}$). (c.) Change in MRV for fast gamma between control and experimental groups. One-tailed paired Wilcoxon signed-rank test comparing light conditions: $n_C = 38$ tetrodes, $z_C = -0.09$, $p_C = 0.46$. $n_E = 133$ tetrodes, $z_E = 0.59$, $p_E = 0.81$. Wilcoxon rank-sum test statistics across groups $z = 0.25$, $p = 0.80$. (d.) Correlating change in fast gamma MRV in the experimental group as a function of SWR suppression during awake-rest (Pearson's Correlation: $r = -0.07$, $p = 0.42$). One-tailed, paired Wilcoxon signed-rank test comparing slow gamma *vs.* fast gamma within *EXP* group: $z = -2.6$, $p = 0.004$.

Slow gamma (Fig. 2.24d.) $r = -0.17$, $p = 0.02$, fast gamma $r = -0.19$, $p = 0.01$.

In contrast, we found an opposite effect on the synchronization of theta. We did not see a significant change in the mean lag measure for theta (Fig. 2.23c.). As opposed to slow and fast gamma, when measuring mean peak cross-correlation, we observed an overall increase in the synchrony of theta across recording sites when CA3 input was suppressed (Fig. 2.23a., Paired Wilcoxon signed-rank test, $n_E = 193$ electrodes in experimental animals, $p_E = 0.06$, $n_C = 86$ electrodes in control animals, $p_C = 1.3 \times 10^{-14}$. Wilcoxon rank-sum test comparing groups: $p = 8.1 \times 10^{-8}$). We measured this increase in synchrony with the extent of CA3

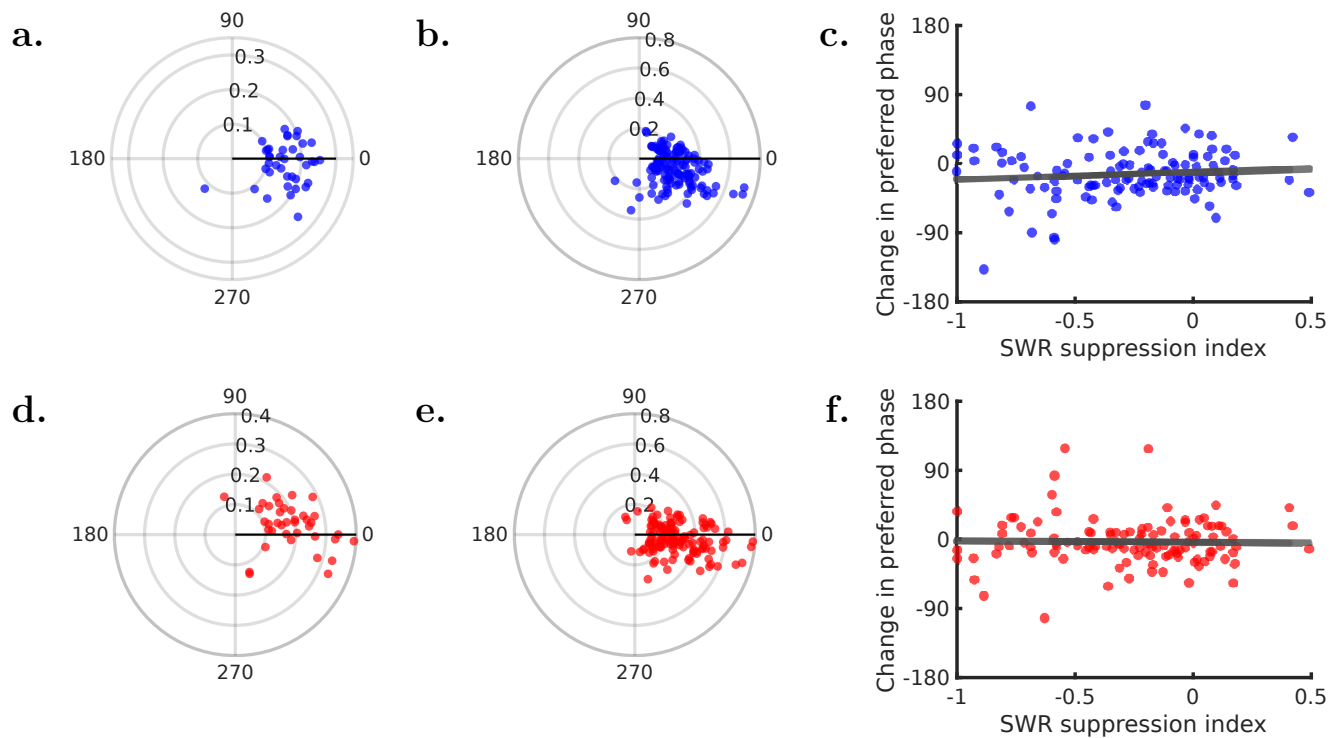


Figure 2.21: The role of CA3 in locking gamma events to theta. (a.) Change in preferred phase of slow gamma events in control group. (b.) Change in preferred phase of slow gamma events in experimental group. Two-tailed Wilcoxon rank-sum test comparing phase shift between *EXP* and *CON* groups for slow gamma - $n_E = 131$ tetrodes, $n_C = 39$ tetrodes, $z = -1.8$, $p = 0.07$. (c.) Measuring change in preferred phase of slow gamma events against SWR suppression index for experimental group (Pearson's correlation for $n_E = 131$ tetrodes with both slow and fast gamma coherence over 0.1: $r = 0.1$, $p = 0.29$). (d.) Change in preferred phase of fast gamma events in control group. (e.) Change in preferred phase of fast gamma events in experimental group. Two-tailed Wilcoxon rank-sum test comparing phase shift between *EXP* and *CON* groups for slow gamma - $n_E = 131$ tetrodes, $n_C = 39$ tetrodes, $z = -3.8$, $p = 1.5 \times 10^{-4}$. (f.) Measuring change in preferred phase of fast gamma events against SWR suppression index for experimental group (Pearson's correlation for $n_E = 131$ tetrodes with both slow and fast gamma coherence over 0.1: $r = -0.02$, $p = 0.85$).

suppression and found those to be negatively correlated (Fig. 2.23b., Pearson's correlation $r = -0.15$, $p = 0.05$).

The data recorded from dorsal CA1, thus, suggested that the synchrony of both slow and fast gamma events is regulated by the instantaneous CA3 input. Suppression of CA3 input has a desynchronizing effect on gamma events that depends on the localized strength of suppression.

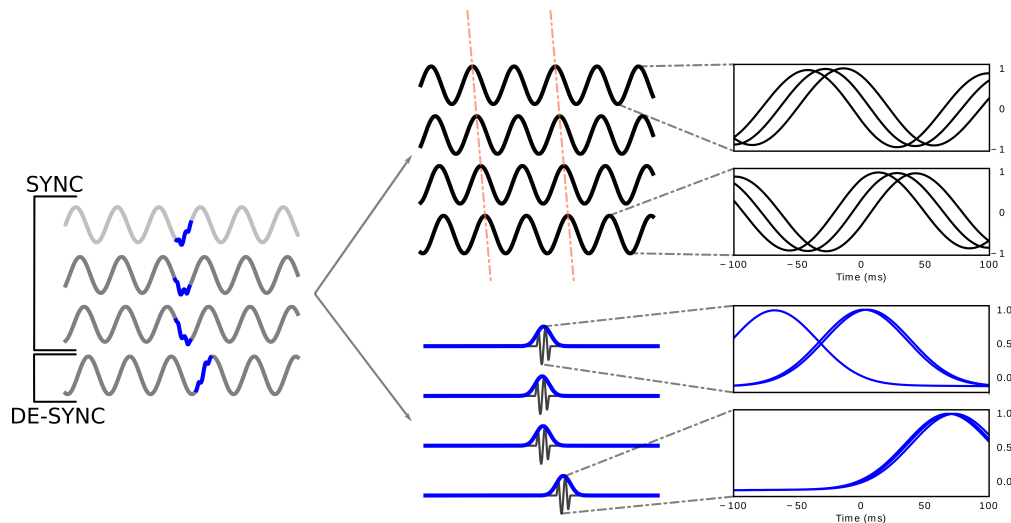


Figure 2.22: Schematic for calculation of a synchrony.

2.7 Conclusion

Our findings suggest that: 1) While theta travels along the septo-temporal axis of the hippocampus, both slow and fast gamma are locked to theta frequency, slow and fast gamma events occur synchronously across CA1. Furthermore, slow and fast gamma events are locked to an abstract, global theta phase definition. 2) CA2 plays a dominant role in shaping slow gamma events in the hippocampus. While slow gamma events persist even when CA3 input to CA1 is suppressed, the locking of slow gamma events to global theta phase deteriorates in the absence of CA3 input. 3) CA3 input to CA1 plays a role in synchronizing both slow and fast gamma events in the hippocampus.

The Local Field Potential (LFP), arising due to the accumulation of synaptic discharges, ionic conductance and other physiologic processes from multiple neurons [19], presents a unique window into the macroscopic properties of a cellular network. Historically, oscillations in the LFP have been of interest because of their relation to behavioral and cognitive processes. It has been hypothesized that oscillations in the LFP reflect synchronized processes in the cellular networks [136]. Here, we have analyzed two major LFP oscillations – theta and gamma. Theta oscillations exhibit large amplitudes and are associated with ambulatory behavior as well as Rapid Eye Movements (REM) sleep. Gamma oscillations, while lower in amplitude when compared to theta and Sharp-Wave Ripples (SWR), have been observed in a variety of behavioral states in both neocortical [157] and subcortical regions, in particular, in the hippocampus [12]. Theoretical models have suggested that an embedding of gamma cycles within theta could serve as a mechanism for sequence learning [83]. This theory has gained experimental support from both human [58] and rodent studies [84]. There is further evidence that the strength of theta phase-gamma amplitude locking is modulated by learning both within the hippocampus [149, 125] as well as between the hippocampus

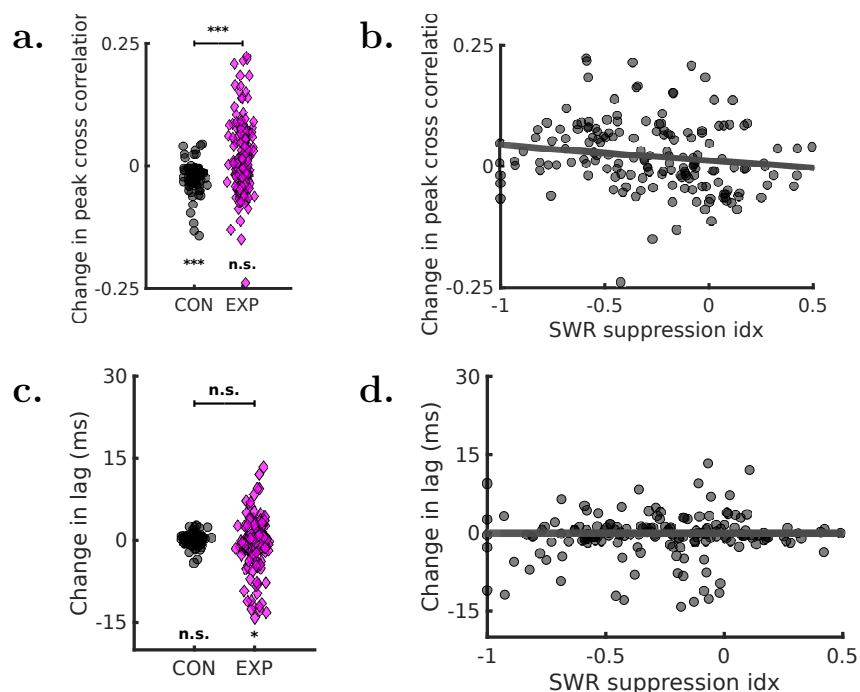


Figure 2.23: Measuring theta oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for LFP theta (Wilcoxon sign-rank test, OFF vs. ON: $n_C = 86$ electrodes, $z_C = 6.2$, $p_C = 7.5 \times 10^{-10}$, $n_E = 193$ electrodes, $z_E = -3.2$, $p_E = 0.002$, Wilcoxon rank-sum test, EXP vs. CON $z = 5.4$, $p = 8.1 \times 10^{-8}$). (b.) Change in mean lag for slow gamma (Wilcoxon sign-rank test, OFF vs. ON: $n_C = 86$ electrodes, $z_C = 3.3$, $p_C = 9.6 \times 10^{-4}$, $n_E = 193$ electrodes, $z_E = -4.7$, $p_E = 2.0 \times 10^{-6}$, Wilcoxon rank-sum test, EXP vs. CON $z = 5.4$, $p = 7.7 \times 10^{-8}$). (c.) Change in mean lag for LFP theta (Wilcoxon sign-rank test OFF vs. ON: $n_C = 86$ electrodes, $z_C = -0.99$, $p_C = 0.32$, $n_E = 193$ electrodes, $z_E = 1.9$, $p_E = 0.06$, Wilcoxon rank-sum test, EXP vs. CON $z = -1.9$, $p = 0.06$). (d.) Change in mean peak-correlation against SWR suppression for theta (Pearson's $r = -0.15$, $p = 0.05$).

and cortical areas [146]. Theta phase-gamma amplitude locking has been shown causally to correct spatial memory deficits in a mouse model for Alzheimer's disease [40]. In this study, opto-genetic stimulation of the medial septum at slow gamma frequency rescued spatial memory deficits.

In the context of theta-gamma locking the notion of theta phase, however, is ill-defined. Since theta exhibits traveling wave nature, the instantaneous phase of theta as well as the instantaneous local firing rate of cell populations is different for different parts of the hippocampal cell layer [85, 119]. Downstream regions rely on LFP information to tune into the spatial information contained in the firing of these cell populations [172]. In the absence of a synchronized signal, it remains unclear how downstream regions can selectively list to

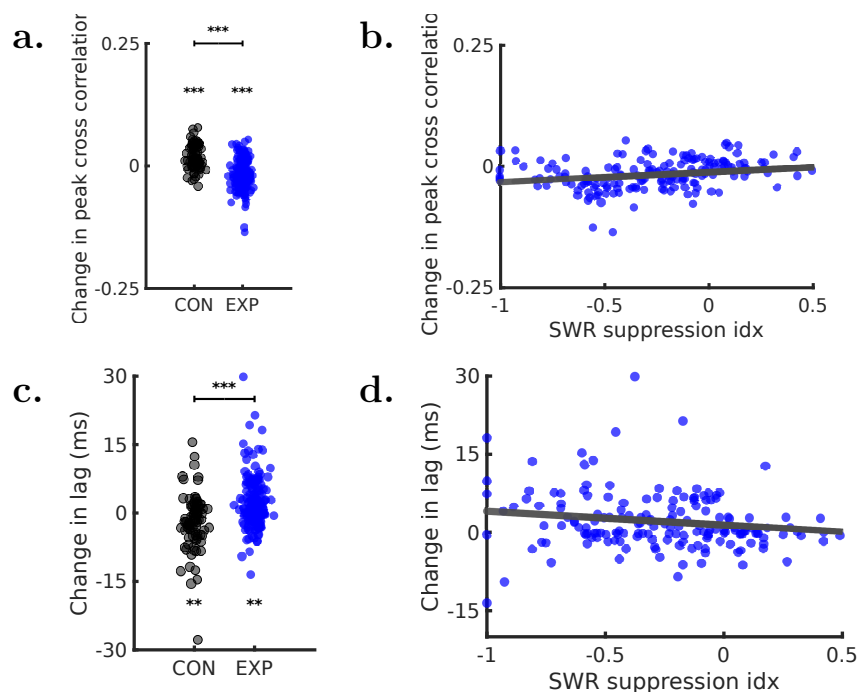


Figure 2.24: Measuring slow gamma oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for slow gamma (Wilcoxon sign-rank test, OFF vs. ON: $n_C = 86$ electrodes, $z_C = -5.6$, $p_C = 1.6 \times 10^{-8}$, $n_E = 193$ electrodes, $z_E = 6.6$, $p_E = 4.2 \times 10^{-11}$. Wilcoxon rank-sum test, EXP vs. CON $z = -8.2$, $p = 2.6 \times 10^{-16}$). (b.) Change in mean peak correlation with SWR suppression for slow gamma (Pearson's correlation $r = 0.24$, $p = 0.001$). (c.) Change in mean lag for slow gamma. One-tailed, paired Wilcoxon sign-rank test, OFF vs. ON, $n_C = 86$ electrodes, $z_C = 3.3$, $p_C = 9.6 \times 10^{-4}$, $n_E = 193$ electrodes, $z_E = -4.7$, $p_E = 2.0 \times 10^{-6}$. Wilcoxon rank-sum test comparing EXP vs. CON, $z = 5.4$, $p = 7.7 \times 10^{-8}$. (d.) Change in mean lag with SWR suppression for slow gamma (Pearson's correlation $r = -0.17$, $p = 0.02$).

animal's local or non-local position represented, for example, during theta sequences [173, 154]. The synchronization of gamma oscillations offers a resolution towards a coherent LFP signature that can be used by downstream areas to reliably interpret the spatial code in the hippocampus. By matching the global theta frequency, gamma events offer an alternate global reference which can be used to study theta-gamma locking.

In [85], the authors suggested that the spatial representation along the septo-temporal axis of the hippocampus varies as a function of the theta phase. However, direct evidence has not been established that different points along the dorsoventral axis of the hippocampus show differences in encoding the animal's position. While putative pyramidal cells increase in place field size along the septo-temporal axis [72], large place fields do not necessarily imply a coarser spatial code at the population level, nor is the population representation




Figure 2.25: Measuring fast gamma oscillation synchrony while suppressing CA3 input to CA1. (a.) Change in mean peak-correlation for fast gamma (Wilcoxon sign-rank test OFF vs. ON: $n_C = 86$ electrodes, $z_C = -0.68$, $p_C = 0.5$, $n_E = 193$ electrodes, $z_E = 5.8$, $p_E = 5.9 \times 10^{-9}$, Wilcoxon rank-sum test, EXP vs. CON $z = -4.4$, $p = 1.0 \times 10^{-5}$). (b.) Change in mean peak-correlation against SWR suppression for fast gamma (Pearson’s correlation $r = 0.36$, $p = 1.2 \times 10^{-6}$). (c.) Change in mean lag for fast gamma (Wilcoxon sign-rank test between OFF vs. ON: $n_C = 86$ electrodes, $z_C = 2.7$, $p_C = 0.007$, $n_E = 193$ electrodes, $z_E = -3.9$, $p_E = 9.7 \times 10^{-5}$, Wilcoxon rank-sum test, EXP vs. CON $z = 4.4$, $p = 1.1 \times 10^{-5}$). (d.) Change in mean lag against SWR suppression for fast gamma (Pearson’s correlation $r = -0.19$, $p = 0.01$).

of position necessarily offset between septal and temporal hippocampal poles. The spatial representation from different anatomical position in the hippocampus can be synchronized using gamma oscillations as a reference which is coherent. Despite the differences in spatial and temporal properties of cells along the septo-temporal axis, the emergence of a coherent spatial representation mediated by gamma synchrony can be viewed as an instance of the binding problem [127].

Slow gamma has previously been implied in synchronization of spatial representation in the hippocampus. We have evidence that within SWR occurring during periods awake-rest and quiescence, slow gamma coherence between CA1 and CA3 sub-regions results in a more coherent reactivation of animal’s current or recently experienced environment [22]. This suggests that slow gamma can indeed be a synchronizing signal, or a clocking mechanism which

unifies the network state along the transverse (CA1-CA3) axis. This is further supported by Bayesian decoding of replay events which exhibit auto-associative dynamics with discrete jumps in the animal's decoded position locked to peaks slow gamma [122]. Our findings add to these ideas and suggest that gamma synchrony could be a generalized mechanism for synchronizing spatial representation. These findings can be bound together into a single notion – gamma events clock the hippocampus across behavioral states.

In our previous work [30], we studied the role of CA3 in shaping single unit responses and associated LFP features including theta and SWR in CA1. We observed that suppression of CA3 input to CA1 led to macroscopic changes in the CA1 activity across behavioral states including sleep, ambulatory behavior on a linear track, as well as periods of awake-rest accompanying reward consumption. In experimental animals, overall firing rates as well as SWR rates were significantly suppressed. Place cell attributes, including firing rates, field center of mass and field sizes were significantly affected by suppression of CA3 input. We also found that the fidelity of awake hippocampal replay deteriorated in the experimental group. In this paper, we expand on the role of CA3 in shaping CA1 activity at a mechanistic level in the context of gamma synchrony. The lamellar routing of information via the different gamma frequency bands in the hippocampus is an active area of research [130, 77, 42, 112]. Different gamma bands have been shown to have different physiological and functional properties. It is believed that slow gamma reflects CA3 input to Stratum Radiatum (SR) [29]. Our work provides causal evidence for this hypothesis showing that while slow gamma events persist in the absence of CA3 input (in agreement with prior work using a genetic mouse line [96]), their relationship with theta is disrupted. Higher frequency gamma bands, on the other hand, are believed to reflect converging input from Entorhinal Cortex (EC) [97], as well as local computation in the pyramidal cell layer [130]. We find that CA3 input also affects higher frequency bands. While the relationship of fast gamma events remained largely unaffected by our perturbation of CA3 activity, we observed an overall decrease in fast gamma power.

Finally, we bring together the two aspects of this study – synchronization of gamma events and the mechanistic role of CA3 in shaping gamma events. The unique recurrent connectivity of the CA3 sub-region [3] makes it an ideal candidate to synchronize these events. On quantifying synchrony, we found that gamma events were desynchronized upon suppressing CA3 input. We observed this desynchronization using two different measures. On suppressing CA3 input, average lag or jitter in gamma increased across sites and at the same time, the peak synchronization decreased. Furthermore, we observed this desynchronization for both slow and fast gamma, further reinforcing the idea that the influence of CA3 is not limited to the slow gamma frequency band. The synchronization of slow gamma events, which could potentially make them a clocking mechanism throughout the hippocampus, and the role of CA3 in mediating this synchronization, also suggest a role of CA3 in clocking and gating information within CA1.

Chapter 3

ActiveLink: An open-source tool for neural targeting and closed-loop stimulation¹

High density electrophysiological recordings, offering high spatial and temporal resolution, are key to understanding the structure and function of the brain. Over the last few decades, we have packed more and more recording sites into probes allowing a large number of cells to be simultaneously recorded from the brain [143]. The neural code is often sparse, and the ability to record large populations of neurons is now revealing phenomenon at the ensemble level that were not known before. For instance, sparse populations of hippocampal neurons participate in replaying spatial trajectories that have been implicated in spatial planning and navigation [47, 123].

Targeting cells in different brain regions is difficult, often requiring years of experience and specificity for the region being studied and recording technique being employed. While high density probes can potentially offer high yields, presently, cell targeting primarily employs stereo-taxic coordinates and qualitative feedback from the recording stream. In Hippocampus research, for an instance, individually adjustable tetrode bundles have been used to target the thin sheet of pyramidal cells in the CA1 which contains place cells for over half a century [107]. Targeting this cell layer in rodents and bats for example, takes 2 – 4 weeks for experts with years of training. Additionally, upon reaching the target region with electrodes, the instability results a consistent cell yield over 1 – 2 weeks in a typical experiment, making the experimental procedure cumbersome. The current standard for estimating the relative coordinate of electrode tips to the cell layer is auditory feedback received by playing the raw signal on a speaker, and visual inspection of the putative action potentials. Low-latency Application Processing Interface (API) are pervasively used to stream and process neural data in real-time in many of these experiments, especially for closed-loop protocols. With the advent of faster, custom hardware and software API for real-time processing of neural

¹In collaboration with Dr. Caitlin Mallory and William Croughan.

data, we have an opportunity to build an automated solution to the problem of navigating to neural targets. Moreover, publicly available open-source solutions do not exist for a wide-variety of commonly used sub-tasks in such closed-loop experiments.

In this study, we describe an open source tool that quantifies neural measurements that can aid targeting cells in different parts of the brain. We talk in depth about targeting cells in the hippocampus, but the procedure can be easily adapted to other brain regions. We attempt to provide a software architecture that can also be used for some of the common pre-processing tasks such as LFP filtering, spike-sorting, and interfacing with peripheral environmental and stimulation devices. We also introduce a novel spike-sorting algorithm which can cluster streaming data in real-time, requiring as little as 60 seconds from the start of a recording session, to produce cluster definitions which can be then used for real-time processing of single unit activity. We demonstrate this ability using recordings from the hippocampus as animals ran on a linear track for 15 minute long sessions.

We also demonstrate the implementation a real-time Bayesian inference algorithm in which single-units clustered by hand are used to build place fields for cells in real-time and compute a posterior distribution of the animal's position from its neural activity. We show that high-accuracy can be achieved for such tasks in a very short amount of time, opening up the way for a large family of closed-loop experiments. Finally, we use our program in a closed-loop experiment to study the role of SWR events in spatial-learning as well as reproduce evidence for homeostatic changes to maintain SWR activity upon electrical manipulation in this context [49].

3.1 Software Design

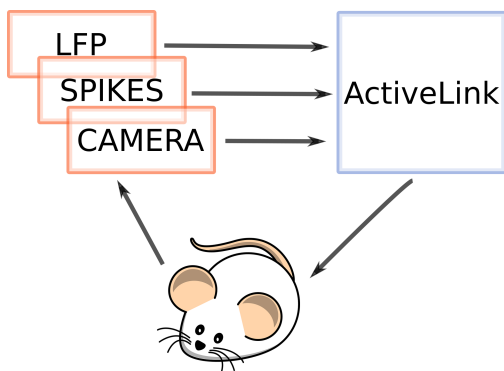


Figure 3.1: Software architecture for real-time processing and close-loop feedback

In this section, we describe the software architecture for our program in detail. The program was written in Python and the source-code has been provided at <https://github.com/architgupta93/activelink>. Fig. 3.1 details the overall architecture of the program. Our program is currently compatible with Trodes (Version 2.0), an open-sourced recording software freely available from SpikeGadgets LLC. The modular design of ActiveLink allows easy adaptation for other recording schemes. The developer can re-write the primary data streams (spike, Local Field Potential (LFP) and position) as needed and the rest of the

program will be able to interact with these streams.

Streaming Data

Procuring data from the streaming software (Trodes) largely consists of waiting for the producer stream to generate data packets and transmitting them to the receiving port. This task is efficiently achieved by parallel processing threads which can be allocated to the same processing core. ActiveLink executes three primary data fetching threads, *LFPFetcher*, *SpikeFetcher*, and *PositionEstimator*. In the source code, these are supplied in individual python files that contain the respective classes.

LFP

After LFP packets are received, they are filtered in the ripple frequency band (150 – 250 Hz) using a Butterworth filter of order 3, using cascaded Second-Order Sections (sos) filter for numerical stability. Real-time, causal smoothing is applied using an implicit exponential decay filter that implements

$$\mathbf{y}_i = \lambda F(\mathbf{x}_i) + (1 - \lambda) \mathbf{y}_{i-1}, \quad (3.1)$$

where \mathbf{y}_i is the power computed at time-step i , \mathbf{x}_i is the raw input at time-step i , F is the filter operation and λ is the smoothing factor. In practice, the data suffers from correlated noise which is corrected by subtracting the ripple power on a baseline electrode positioned in the white-matter. The raw ripple power computed in Eq. (3.1) is z-scored by computing the running mean ($\mathbf{m}(y)$) and standard deviation ($\mathbf{s}(y)$) of the baseline-subtracted ripple power. After receiving the n^{th} data-point, these were updated over as follows:

$$\mathbf{m}(y_n) = \mathbf{m}(y_{n-1}) + \frac{(y_n - \mathbf{m}(y_{n-1}))}{n} \quad (3.2)$$

$$\mathbf{s}(y_n) = \sqrt{\mathbf{s}(y_{n-1})^2 + \frac{(y_n - \mathbf{m}(y_n))(y_n - \mathbf{m}(y_n))}{n}} \quad (3.3)$$

The computer code for detection of Sharp-Wave Ripples (SWR) events is supplied in *RippleAnalysis.py* and the parameters for detection are defined in *RippleDefinitions.py*.

Single Unit Activity

Voltage waveforms simultaneously recorded from multiple electrodes can be clustered to identify the underlying single-unit activities. Several approaches have been described to clustered voltage waveforms [26, 14]. While algorithmic approaches have also been described for real-time spike sorting [151, 115], much of the literature is devoted to ‘offline’ clustering or sorting of voltage waveforms. In this regard, we differentiate our approach to be an ‘online’ clustering algorithm. Consider a set of bundled recording sites – these could represent 4 recording tips in a tetrode bundle, or more generally n sites on a recording probe or array. The recording software triggers when the raw voltage waveform crosses a fixed threshold (in

this case $50\mu V$). A fixed window around the triggered time-point constitutes the voltage waveform - Let's call the waveform to be \mathbf{w} .

First, estimate the Principal Components (PCs) from our sampled voltage waveforms $[\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_p]$. Projecting the voltage waveforms onto PCs has been shown to yield better clustering results than in the space of peak-amplitudes on the individual channels, for examples, in [26]. We call the projection coefficients of the waveform vector \mathbf{w} onto the major PCs \mathbf{c} . However, because of the prevalence of clustering visualization in the said 'amplitude' space, we visualize our clustering both in the amplitude space and in the space of projection onto PCs. Following an initial phase of estimating the PCs, we built a decision tree for clustering using the Balanced Iterative Reducing and Clustering using Hierarchies (BIRCH) algorithm in python [170, 120]. BIRCH clustering resulting in fragmented clusters which were merged in the amplitude space by back projecting from the Principal Component (PC) basis to the original basis. Agglomerative clustering was using with an average-distance metric between clusters to merge the sub-clusters [120]. Computer code for real-time spike-sorting is provided in *SpikeAnalysis.py* as a Python class *SpikeSorter*.

Cluster Metrics

Quantitative metrics help ascertain the quality of a single-unit cluster [59, 26], allowing end users a systematic approach to select clusters for subsequent processing. We have incorporated some of the commonly used metrics for single-unit clusters into ActiveLink, namely, firing rate, number of spikes, inter-spike interval, cluster centroid, minimum and maximum spike amplitudes. Clusters failing to meet used defined criteria on these metrics are labelled 'noise' clusters and not used for subsequent analysis. In the supplied computer code, cluster metrics are defined in *SpikeAnalysis.py* using the python class *ClusterMetrics*.

Real-Time decoding

As an illustration of a real-time application of the analysis of LFP and single-unit activity described thus far, in this section, we discuss real-time Bayesian decoding of spatial position. A candidate action-potential detected on a set of electrodes is first passed through the clustering pipeline to assign a cluster identity. Provided the cluster identity is not that of a 'noise' cluster, the identity is received. Because of the asynchronous nature of real-time processing, we maintain a First-In First-Out (FIFO) queue of all the time bins that are currently being processed. Decoding bins are uniformly spaced making it easy to match the incoming spike. When a spike is received past the current times in the FIFO queue, a new decoding bin is initialized with the posterior $P(x)$ corresponding to 0 spikes received from each neuron currently being analyzed. This is computed as

$$P(x) \Big|_{t=0} = \exp \left(W \cdot \sum_{i=1}^N f_i(x) \right). \quad (3.4)$$

Here, W is the size of the decoding window in seconds, $f_i(x)$ is the firing rate map of cell i over either directional or euclidean position. Say a spike is received at time t from neuron k , the posterior probability is updated as follows:

$$P(x) \Big|_{t=t} = P(x) * f_k(x). \quad (3.5)$$

Here $P(x)$ is the posterior for the time-bin corresponding to the spike. It will be computationally efficient to maintain log-probabilities require addition operations instead of multiplication used here. Computer code for real-time Bayesian inference is supplied in *PositionDecoding.py* as a Python class *BayesianEstimator*.

3.2 Quantified targeting of Pyramidal cell layer in CA1

Local Field Potential (LFP) features in the brain provide strong landmarks for targeting different brain regions. In this study, we focus on targeting cells in the hippocampal CA1 sub-region. CA1 consists of a thin layer ($\sim 100\mu\text{m}$ thick) of densely packed pyramidal cells. It is marked by several strong LFP features including Sharp-Wave Ripples (SWR) complexes (150 – 250 Hz) during awake rest and Slow-Wave Sleep (SWS) states, 6 – 10 Hz theta oscillations during ambulatory behavior (running, rearing, whisking) and Rapid Eye Movements (REM) sleep, and also oscillations in the gamma band.

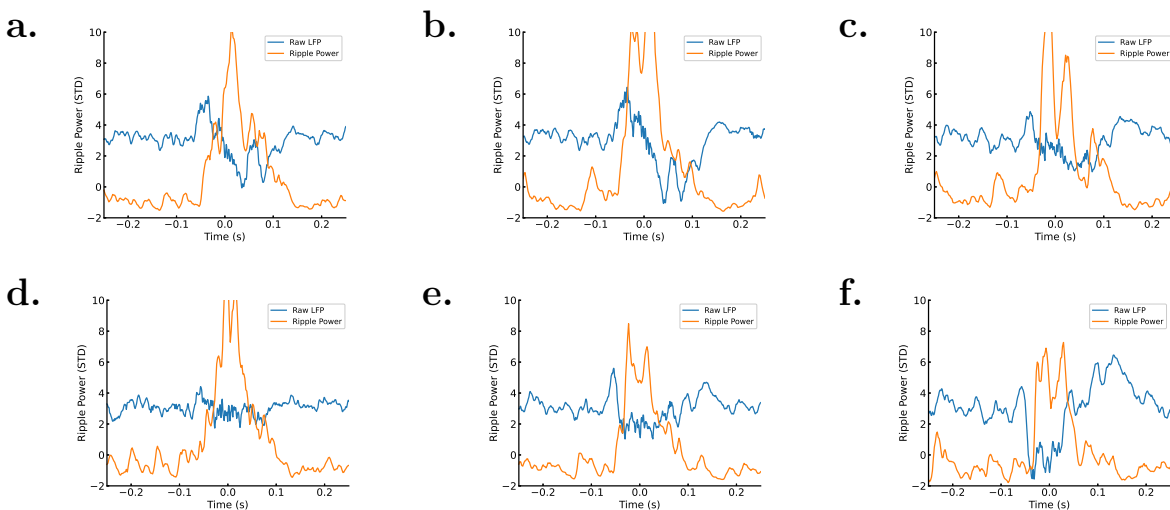


Figure 3.2: SWR events occurring synchronously across multiple electrodes during a sleep session from a rat.

Conventionally, recording electrodes start in the cortex after surgery. After confirming that electrodes are in the cortex (with the help of cortical single units or spindles), elec-

trodes are driven into white matter over 1 – 2 weeks. Arrival at white matter is marked by the lack of action potentials and high-frequency LFP activity. The primary marker for this adjustment has historically been cortical spiking activity. Our findings, adjusting electrodes in rodents suggest that quantitative measures like SWR envelope size can be effectively used for navigating to white matter as well (See Sharp-Wave amplitude over time in Fig. 3.3f.). As recording electrodes cross the white matter and enter the hippocampal formation, conventionally, SWR complexes are heard and seen over the recording equipment and then used to guide cell targeting until electrodes are very close to the cell layer. By quantifying SWR complexes and detecting them in real time, not only were we able to quantify SWR complexes in the hippocampal formation, we can also detect them much earlier, at smaller amplitudes.

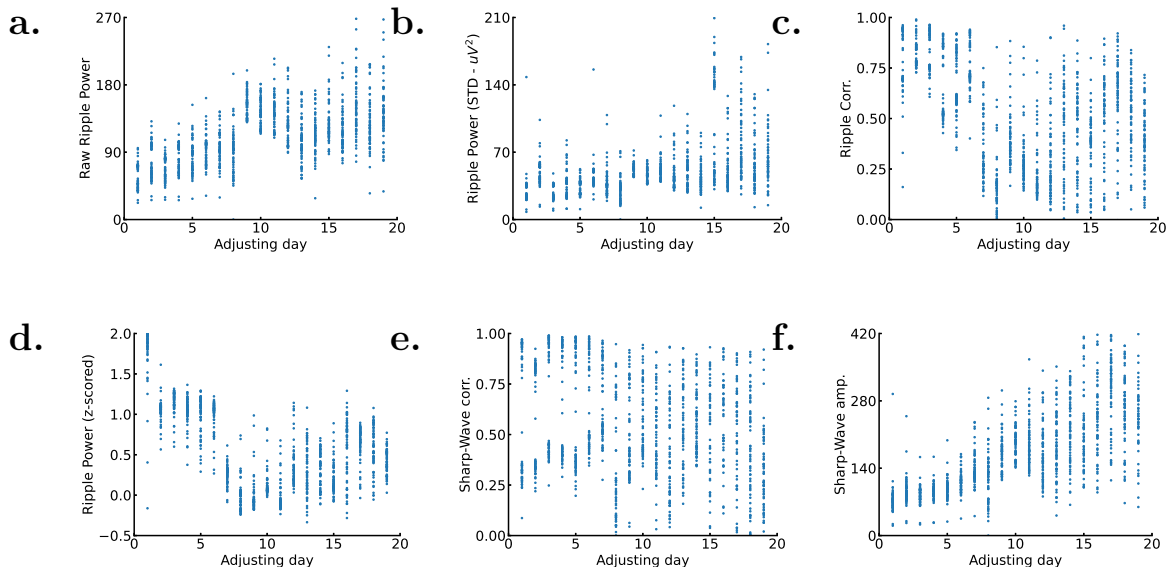


Figure 3.3: Statistics of SWR electrode adjustment in the brain. (a.) Mean power in the ripple frequency band. (b.) Standard-Deviation (over time) of power in the ripple frequency band. (c.) Correlation in the ripple band power between a recording site and a fixed reference. (d.) Average z-scored ripple power in a 500ms window surrounding SWR events detected on a reference electrode. (e.) Correlation in the Sharp-Wave envelope between a recording site and a fixed reference. (f.) Amplitude of the Sharp-Wave at a recording site.

Fig. 3.3 highlights the some of the SWR parameters that can be measured in real-time. The first three parameters, shown in Figs. 3.3a. to 3.3c., measure bulk properties of activity in the ripple frequency band across recording electrodes. In Fig. 3.3a., we report the mean ripple power, or the mean power in the ripple frequency band during an extended period of time over the course of electrode adjustment. Fig. 3.3b. reports the standard deviation of ripple power from its mean value/ Fig. 3.3c. measures the correlation of the ripple power

envelope between a recording electrode and a designated reference. The remainder of the parameters, shown in Figs. 3.3d. to 3.3f. report measurements in a 250ms window surrounding SWR events detected on the reference recording electrode. Fig. 3.3d. reports the average z-scored power in the ripple frequency band in a 250ms window surrounding a ripple event detected on a reference electrode. Fig. 3.3f. measures the amplitude of the sharp-wave envelope, computed by filtering the raw voltage waveform in the sharp-wave frequency band (10–50Hz) across the recording sites when accompanied with a ripple event. Fig. 3.3e. measures the correlation of the sharp-wave envelopes for different recording electrodes against a reference on which ripple events were detected.

It is important to note that some of these features, the mean ripple power and sharp-wave amplitude in particular, increase monotonically with recording depth, and as a corollary in Fig. 3.3, the electrode adjustment time. These can also be used, therefore, as markers for estimating the depth on an electrode relative to the hippocampal cell layer where ripple power is known to be maximum, or Stratum Radiatum (SR), where the sharp-wave is known to have the maximum amplitude. Fig. 3.2 illustrates an SWR event on six simultaneously recorded electrodes (blue trace depicts raw LFP). Note that while power in the ripple frequency band (orange trace) crosses the threshold for an SWR event on all electrodes, the amplitude and polarity of Sharp-Wave is different for the electrodes. This measure is quantified by the amplitude and correlation of the Sharp-Wave and can be used to target the cell layer. A negative deflection (Fig. 3.2f. for an instance) is indicative of the electrode tip having passed the the Pyramidal cell layer - This electrode can be moved backwards to reach the cell layer. Similarly, a large positive deflection (Fig. 3.2a. for an instance) is indicative of the electrode tip being in Stratum Oriens and can be used to move the electrode tip forward towards the cell layer in CA1.

3.3 Online analyses of neural activity

Fig. 3.4 shows an SWR event detected in real-time using ActiveLink. In Fig. 3.5, we report the efficacy of detecting SWR from a single session online. As shown in Fig. 3.5a., when compared to offline Local Field Potential (LFP) processing, we were able to detect over 97% of the ripples. The fairly large number of false positives result from an inclination towards high sensitivity for the experimental requirements. Further, we were able to detect over 50% of the SWR events in under 30ms as shown in Fig. 3.5b..

Real-time clustering of Single Unit Activity

Here, we present results combining the BIRCH clustering algorithm [170, 169] with hierarchical agglomerative clustering [120] to build clusters in real-time. Cluster centers are created and updated for tetrode bundles as individual spikes are received. Having seen a preset number of spikes, a cluster merge step is run which merges the leaf nodes of the BIRCH tree using distance between the cluster centers in the peak amplitude space.




Figure 3.4: Illustration of an Sharp-Wave Ripples (SWR) event detected online using ActiveLink.





Figure 3.5: (a.) Accuracy and (b.) Latency measures for detecting SWR events online.

In addition to LFP features, shape of the single unit waveform is typically employed for identifying brain regions that a recording electrode is presently in. Fig. 3.6 shows the clustering results for a single tetrode bundle while the animal ran trials on a linear-track for chocolate reward on both ends. Fig. 3.6a. shoes clustering outcome when the electrode tips were located in the neocortex above the Hippocampus. Contrast the shape of the clusters with Fig. 3.6b. which shows real-time clustering results when the electrode tips were located in the Hippocampal CA1 (marked by distinct SWR complexes). We also report decoding accuracy during running periods (Fig. 3.8) and illustrate replay events decoded using real-time clustering results (Fig. 3.9).

a.



b.




Figure 3.6: Clustering outcome after 3 minutes on a single tetrode bundle while the animals ran laps on a linear track. (a.) Electrode tips in neocortex, (b.) Electrode tips in Hippocampal CA1.

Real-time Bayesian inference

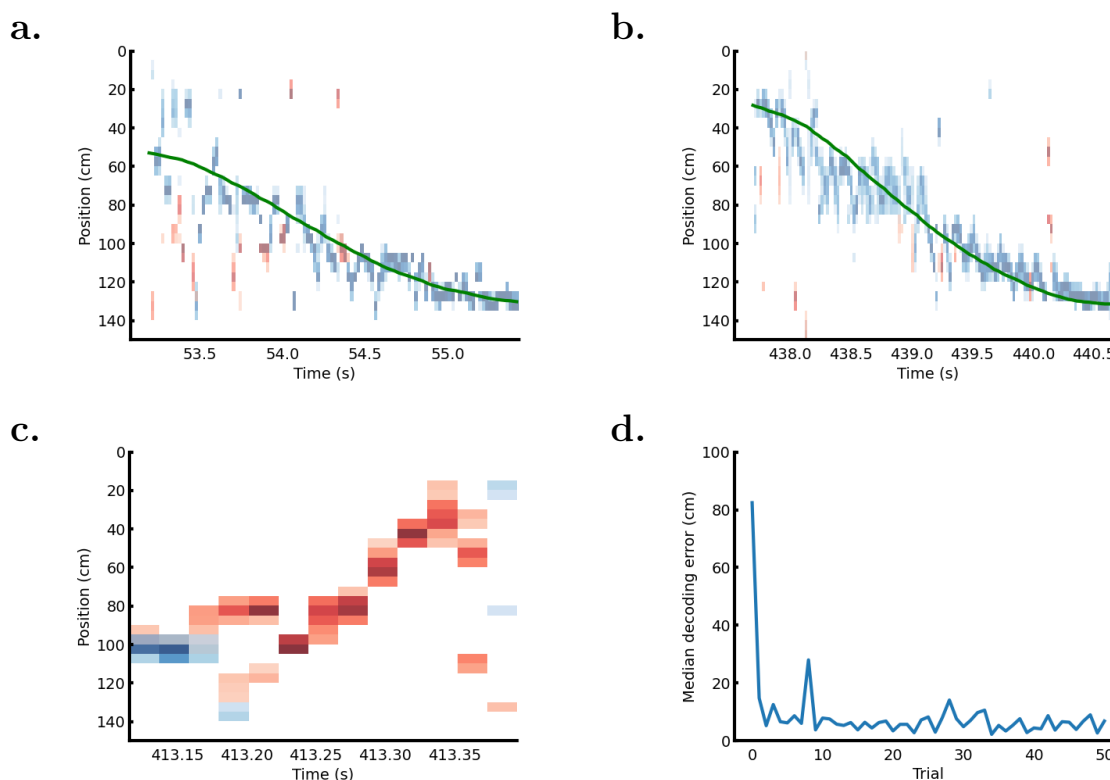


Figure 3.7: Real-time Bayesian Inference using 60ms decoding window with 20ms bin-shift between consecutive windows in ActiveLink. (a.) Bayesian decoding early less than 1 minute into the session using hand-clustered single units in Trodes. (b.) Bayesian decoding ~ 5 minutes into the session using hand-clustered units. (c.) Illustration of a replay event decoded in real-time. (d.) Median decoding error over laps.

Hippocampal activity is feature rich, providing a cognitive map of an animal’s environment [109]. It is possible to decode the animal’s current position using its neural activity in the hippocampus [13]. Fig. 3.7 highlights the performance of the online Bayesian decoding algorithm on a linear track environment. On a linear track, it is known that the running directions are represented with different contextual maps [34] and as a result, in addition to the animal’s position on the track, the contextual map being used to represent the animal’s position can also be decoded. Here the two running directions, running *up* a track (position increasing) is represented using the blue map and running *down* the track (position decreasing) is represented using the red map. As early as 1 minute into the session, from place fields for individual cells being constructed in real-time, both the animal’s contextual map and position can be decoded with high-fidelity Fig. 3.7a.. The decoding accuracy improves over time, both qualitatively as shown in Fig. 3.7b., and numerically Fig. 3.7d. when mea-

sured as the error between the Center of Mass (CoM) of the posterior probability and the animal’s true position. We also found that awake replay events, and the contextual map being replayed in the event at a compressed time scale can be effectively decoded using the real-time Bayesian inference algorithm Fig. 3.7c..




Figure 3.8: Decoding error over laps with online clustering.

Finally, we merged the online-clustering with real-time Bayesian inference to create an end-to-end processing pipeline which can process raw action potentials to produce a posterior probability distribution over space. Quantitative results for decoding accuracy with the end-to-end pipeline implemented in ActiveLink are shown in Fig. 3.8. The accuracy is lower than what we observed in Fig. 3.7d. with hand-clustered data and further quantification is needed to isolate and improve on the sources of error. Further, we were also able to observe replay event with the fully automated pipeline. Two such events, detected both with the hand-clustered data (Figs. 3.9a. and 3.9c.) as well as data clustered in real-time (Figs. 3.9b. and 3.9d.) are shown in Fig. 3.9. Note the agreement in both the directional map on the linear track (determined by posterior color red *vs.* blue) and the replay direction across the methods for the events shown here.

3.4 Application to closed-loop experiments

With the emergence of rapid recording and neural manipulation technologies, it is now possible to manipulate neural activity in real-time [38, 64, 53]. Closed-loop feedback is typically provided by triggering an external stimulus (laser, electrical pulse, reward delivery *etc.*) contingent on neural or behavioral states. With ActiveLink, we provide a way to trigger external stimuli based on such states. In this section, we discuss the application of ActiveLink to the real-time detection and disruption of Sharp-Wave Ripples (SWR) ripples in a novel spatial-learning task. The data in the subsequent study are derived from a single




Figure 3.9: Illustrating replay events detected using hand-clustered data and real-time online clusters built using ActiveLink (a.) Event 01, Hand-clustered, (b.) Event 01, clustered online, (c.) Event 02, Hand-clustered, and (d.) Event 02, Clustered online.

Long-Evans rat. Two more male Long-Evans rat was used for the track data presented in the previous sections.

Recording and Manipulation

All procedures and protocols were approved by the Institutional Animal Care and Use Committee (IACUC) at UC Berkeley. Animals were housed on a 12-hour non-inverted light schedule. Custom micro-drives carrying 64 tetrode bundles were designed in-house using SolidWorks and 3D printed using FormLabs2.0 and PreForm software. 32 tetrodes targeted the left hippocampus (cannula center coordinates -2.3mm Medial-Lateral (ML) and 4.2mm Anterior-Posterior (AP) from Bregma) while the remaining tetrodes targeted the right hippocampus (cannula center coordinates 2.3mm ML and 4.2mm AP from Bregma). A concentric bipolar electrode (FHC Inc. CBAPC10) was located at 0.8mm ML and -1.3mm AP from Bregma to target VHC. Surgical procedures were performed under $0.5 - 2\%$ isoflurane in O_2 . Animals' breathing rates were maintained at $60/\text{min}$ and body temperatures were maintained in $35.9^\circ\text{C} - 37.5^\circ\text{C}$ using an electric heating pad. 0.2ml Lidocaine was injected




Figure 3.10: Effect of Ventral Hippocampal Commisure (VHC) stimulation on the activity of cells in CA1.

at the incision site, and following a sterilization using Betadiene and Ethanol, an incision was made into the scalp. The skull was levelled, cleaned and dried. 6-8 titanium screws were bolted to the skull and reinforced with dental cement. Two stainless steel screws were bolted to the skull over the Cerebellum and anterior of Bregma, and served as the ground connection. Three craniotomies were made in total, one in each hemisphere to target the hippocampus and the third to target VHC. After lowering the implant, the craniotomies were covered with sterile bone wax and the drive was secured using dental cement. Tetrodes was lowered to a depth of 1mm from the bottom of the cannula into the brain during the surgery. The stimulating electrode was adjusted to a depth of 3.6mm at the end of the surgery.

Data acquisition and processing

Over the subsequent 2 – 4 weeks, electrodes were adjusted to target the Hippocampal CA1 using the methods previously discussed in Section 3.2. Raw signal was pre-amplified using SpikeGadgets LLC 256 channel head-stage and discretized at 30,000Hz. All recordings were done using Trodes2.0 Beta software with a real-time communication Application Processing Interface (API). One tetrode was positioned in the white matter based on ripple metrics described above in Section 3.2 in each hemisphere and used as a reference for Local Field Potential (LFP). To get the LFP, one of the 4 channels of the tetrode was selected, the

reference voltage from the electrode in the white matter in the same hemisphere was subtracted and the signal was down-sampled to 1500Hz. Spikes were detected using a $50\mu V$ threshold detection on any of the 4 channels on a tetrode. A 40 sample section (starting 12 time-samples before the threshold cross) was isolated on all the channels for waveform analysis and clustering.

Electrical Stimulation

ActiveLink was used to calibrate the electrical stimulation delivered to VHC to silence neural activity. Prior to starting an experiment, the animal was placed in a sleep box. Two different protocols were used to select the amplitude of stimulation current. ActiveLink communicated with an Arduino to trigger electrical stimulation at $0.2Hz$. Each stimulation consisted of a biphasic electrical pulse, $400\mu s$ long with a 50% duty cycle. The amplitude of the pulse was started at $40\mu A$ and we checked for a visible depression in the neural activity. If no depression was observed, the current amplitude was increased incrementally to $200\mu A$. However, if none of the current amplitudes produced a visible depression in neural activity, the stimulating electrode was adjusted in depth, and the calibration process was repeated. The immediate effect of electrical stimulation in the VHC is illustrated in Fig. 3.10, showing the activity of a population of neurons in the Hippocampus. Here, immediately after electrical current was delivered into the VHC (marked by black triangles at the top and bottom of the figure), no spikes were observed from virtually all cells for $\sim 100ms$.

Experiment design

Our experimental setup consists of a $180cm \times 180cm$ arena surrounded by 3ft high enclosure as illustrated in Figs. 3.11 and 3.12. Animals were food-restricted to maintain over 85% of their free-feeding body weight received extensive pre-training on the task. Early in training, multiple reward ports were baited in each quadrant of the maze to encourage exploration. The number of reward ports simultaneously baited was reduced to 1 over multiple days to shape the behavior. An extended platform was used to transfer the animal in and out of the maze. As a part of pre-training, the animal was placed on and rewarded with food pellets on the platform.

Each experimental *block* consisted of 4 entries into the maze - for each of these entries, a single reward port was baited with reward (see Fig. 3.11 - red circle marks the current goal location). The rewarded port was held the same for all 4 trials within a block, and a new port was pseudo-randomly selected for every new experimental block. An ITI of 20 minutes was enforced between Trial 1 (first exposure to a new reward location) and Trial 2 (subsequent re-discovery of the reward at the same location) based on delay-dependent learning effects reported previously in the literature [142]. Trials 2, 3, and 4 were separated by a 3 minute ITI. The starting location animal on each trial was pseudo-randomly selected to be in one of the 3 quadrants that the reward port was not located in (Fig. 3.11 - green squares mark the starting location for the trial within the block). During the ITI, the animal




Figure 3.11: Block and trial structure for the learning task. Each training block consisted of 4 learning trials separated by an Inter-Trial Interval (ITI) as shown. Consecutive blocks were separated by 1 – 3 minutes. The maze and reward ports were wiped with a 70% Ethanol solution to clear odor-cues that might mark the current goal.

was removed from the maze and placed on a sleep platform outside the maze. The maze and currently selected reward port were cleaned with a 70% ethanol solution to discourage the use of an odor-based strategy for navigating to the goal. Blocks were pseudo-randomly chosen to be *CON* (no stimulation) or *SWR* (detection and disruption of SWR events). On SWR disruption blocks, all SWR events were disrupted starting at the beginning of Trial 1 and until the animal was removed from the maze at the end of Trial 4. Since multiple training blocks were performed on the same day, we also analyzed Trial 0 (last trial of the preceding block, unless a block was the first in the day) and Probe (First trial on the next block in the day, if any). The analysis of the probe block was limited to the first time that the animal reached the goal location. Notably, on virtually all blocks, animals checked the port that had been baited in the previous block even though there was no reward present. Consecutive blocks were separated by a short (1 – 3 minute) interval to clean the maze and reconfigure the stimulation protocol.

Behavioral Results

Fig. 3.13 illustrates a typical experimental block in the aforementioned task. In the last trial of the previous block (*T0*, Fig. 3.13a.), the animal received a reward at the location marked by the red circle (top-left). Subsequently, the reward location was changed at the start of the new block (Figs. 3.13b. to 3.13e.). Notice that on entering the maze, the animal qualitatively spent a substantial amount of time searching near the previous goal location before finally stumbling on the newly baited reward port. On subsequent trials (Trial 2 - Fig. 3.13c., Trial 3 - Fig. 3.13d., and Trial 4 - Fig. 3.13e.), the animal took more directed routes to the




Figure 3.12: Configuration of the environment, reward locations and stimulation protocol for disruption blocks.

goal location, exhibiting single trial learning. After Trial 4, the reward port was switched again, and the animals was re-introduced into the maze at a pseudo-randomly chosen starting location (Fig. 3.13f.). The time and path to the first approach to the previously rewarded port were noted (as shown in Fig. 3.13f.).

In order to quantify the learning rate over trials, we measured the latency from the start of the trial to reach the reward port for all trials. Fig. 3.14a. shows this latency over control ($n_C = 12$ blocks) and experimental ($n_E = 15$ blocks) conditions. Using 2-way ANalysis Of VAriance (ANOVA), we found a significant effect of Trial ($F_4 = 10.5$, $p = 2.4 \times 10^{-7}$) on latency to reward. There was no significant effect of the group on latency ($F_1 = 2.6 \times 10^{-4}$, $p = 0.99$). The interaction between animal group and trial was also not significant ($F_4 = 0.36$, $p = 0.84$). Latency to the rewarded location on Trial 2 was not significantly different across groups (Two-tailed Wilcoxon rank-sum test *statistic* = 208, $p = 0.23$). Prior literature has also shown that latency saving, as paired difference between the latency in a trial compared to trail 1 from the same learning block, is an effective marker of learning in spatial navigation tasks [142]. We measured the savings in latency over our experimental blocks (Fig. 3.14b.). 2-way ANOVA revealed no significant effect of trial ($F_3 = 1.8 \times 10^{-3}$, $p = 0.99$), group ($F_1 = 0.14$, $p = 0.71$) or interaction between the two ($F_3 = 0.14$, $p = 0.94$). Two-tailed Wilcoxon rank-sum test between latency savings on Trial 2 between control and experimental groups show no significant difference (*statistic* = 225, $p = 0.37$).

Since raw latency to the goal did not reveal a significant effect of SWR disruption, in order to better quantify behavior, we measured other metrics for performance on the task. Animal's running speed was measured over trials as shown in Fig. 3.15a. 2-Way ANOVA with

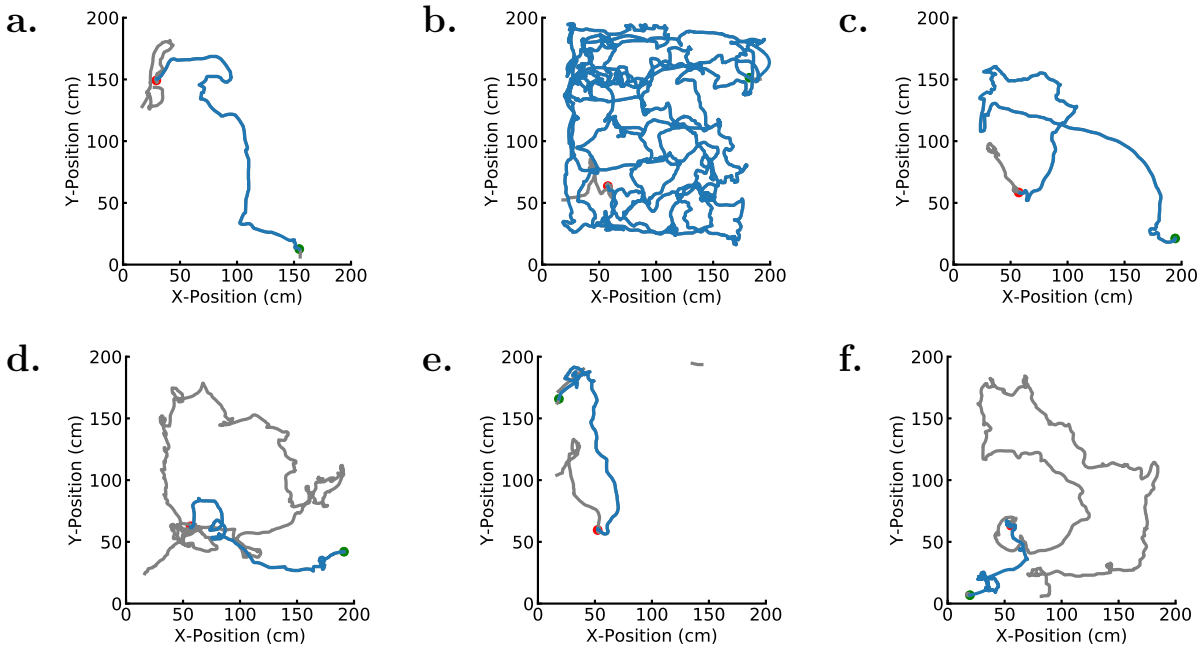


Figure 3.13: Spatial trajectories for a single learning block. Blue: Trajectory to the first approach to the goal location. Gray: Subsequent path taken after reaching the goal. Green Circle: Start location. Red Circle: Goal in the learning block. (a.) Trajectory from previous learning block. (b.)-(e.) Learning trajectories over 4 trials with the new goal location. (f.) Learning trajectory over the next block with a new goal location. Goal from the previous learning block has been marked in red.




Figure 3.14: (a.) Latency to goal over trials in a learning block compared between groups. (b.) Measuring savings in latency paired to Trial 1 of the same learning block.

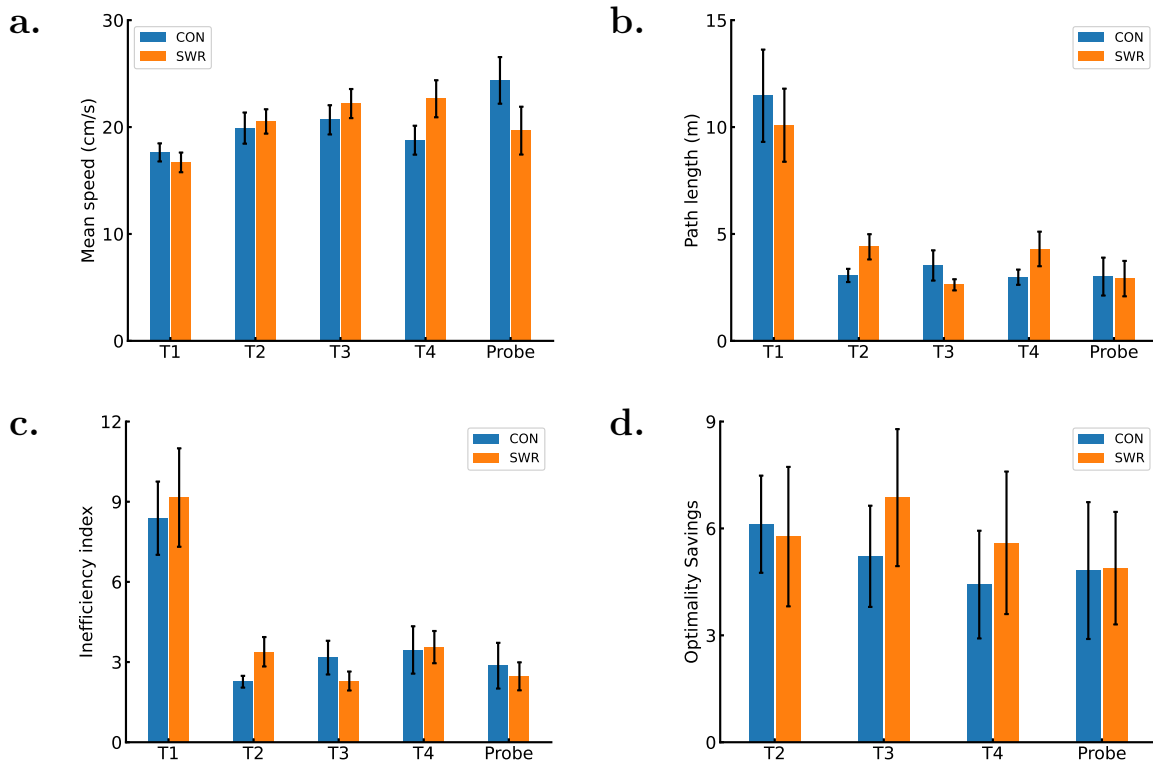


Figure 3.15: Measuring spatial trajectories taken by the animal over learning trials. (a.) Average running speed during the time spent by the animal exploring the maze in search of food. (b.) Total distance covered by the animal in a trial before reaching the reward. (c.) Inefficiency in routing to the goal (total running distance normalized by the optimal path from the starting location) (d.) Improvement in routing after Trial 1.

$n_C = 12$ control blocks and $n_E = 15$ experimental blocks showed no significant differences in running speed between groups ($F_1 = 2.7 \times 10^{-2}$, $p = 0.87$). There was a trend in that the running speed increased over trials, however the effect of trials on running speed was not statistically significant ($F_4 = 2.0$, $p = 0.09$). We did not observe a significant interaction between group and trials ($F_4 = 1.7$, $p = 0.15$).

The total path length has been shown in Fig. 3.15b.. Similar to the analysis of latency, 2-Way ANOVA showed a strong effect of trial on the total path length ($F_4 = 12.7$, $p = 1.1 \times 10^{-8}$). However, the effect of group was not significant ($F_1 = 3.3 \times 10^{-2}$, $p = 0.86$). The interaction between group and trial was not significant either ($F_4 = 0.66$, $p = 0.62$). We did, however, observe a trend in the total path length to be higher in the experimental blocks when compared to controls on Trial 2 (One-tailed Wilcoxon rank-sum test *statistic* = 178, $p = 7.4 \times 10^{-2}$). This trend was not observed on Trial 1 when the animals had not learnt the reward location (*statistic* = 240, $p = 0.50$), and on Trial 3 after the animals had been re-exposed to reward well in Trial 2 (*statistic* = 280, $p = 0.83$).




Figure 3.16: Behavioral measures when the reward location was marked with a visual cue. (a.) Latency from start of a trial to the time the animal first reached the reward port when the port was not marked with a visual cue (Control) or otherwise. (b.) Total path length from the start of the to arrival at goal.

Because of the variability in the starting location, we computed the optimality of the path taken by the animal as the ratio of the path length and the optimal path to the goal from the animal’s starting location in any given trial. This measure of path optimality has been reported in Fig. 3.15c.. Statistical measures of path-optimality were similar to those observed for total path length. A 2-way ANOVA suggested a strong dependence of path-optimality on trial ($F_4 = 8.3$, $p = 6.0 \times 10^{-6}$). Group did not have a significant effect on path-optimality ($F_1 = 7.1 \times 10^{-2}$, $p = 0.79$), and the interaction between group and trials was not significant ($F_4 = 0.46$, $p = 0.77$). Statistical comparisons within trials showed a trend in lower path-optimality in the experimental group when compared to controls for Trial 2 (One-tailed Wilcoxon rank-sum test *statistic* = 176, $p = 6.7 \times 10^{-2}$). Similar to the statistical measures for total path length, we did not see this trend on Trial 1 (One-tailed Wilcoxon rank-sum test *statistic* = 239, $p = 0.50$) and Trial 3 (One-tailed Wilcoxon rank-sum test *statistic* = 205, $p = 0.21$). Within a block, the change in path optimality over trial (Trial n compared to Trial 1) is reported in Fig. 3.15d. 2-Way ANOVA show no significant effect of group ($F_1 = 0.17$, $p = 0.68$) or Trials ($F_3 = 3.1 \times 10^{-2}$, $p = 0.91$). The interaction between animal group and trial was not significant ($F_3 = 0.18$, $p = 0.91$). On Trial 2, the difference in optimality savings was not significant (One-tailed Wilcoxon rank-sum test *statistic* = 219, $p = 0.31$).

In order to ensure that the animals did not employ a search strategy based on odor or visual cues, we performed a visually cued version of the task described previously. In a cued block, for all 4 learning trials, an glass bottle ($15\text{cm} \times 6\text{cm} \times 6\text{cm}$) filled with a fluorescent fabric was placed next to the reward well. Behavioral measures for this version of the task have been reported in Fig. 3.16. Fig. 3.16a. reports a comparison of latency to goal between the control and cued versions of the task. Latency to goal was significantly shorter in Trial 1 for the cued task when compared to the control task (One-tailed Wilcoxon rank-sum test

$n_C = 20$ control trials, $n_U = 6$ cued trials, $statistic = 12$, $p = 0.002$). Total path length to goal on Trial 1 (Fig. 3.16b.) was also significantly lower in the cued version of the task in Trial 1 (One-tailed Wilcoxon rank-sum test $n_C = 20$ control trials, $n_U = 6$ cues trials, $statistic = 100$, $p = 0.008$).

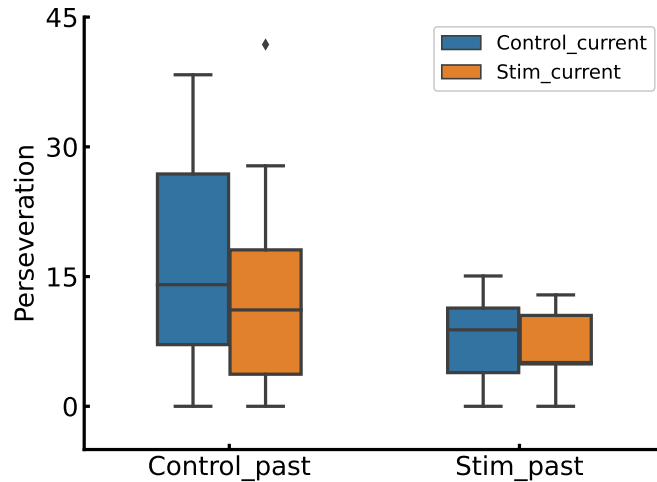


Figure 3.17: Measuring perseverance on the previous goal on Trial 1 of the subsequent block as a function of experimental condition (*CON* vs. *EXP*) in the current and previous block.

One of the observations in behavioral trajectories (Fig. 3.13) was that on Trial 1 of a learning block, animals spent time searching near the previous goal location. In Fig. 3.17, we measured this perseverance as a function of the experimental condition in the current and previous block. Perseverance was measured as the percentage of total time in Trial 1 spent in a 20cm radius around the previous goal location. A 2-Way ANOVA of this measure revealed a significant effect of the previous block type ($F_1 = 4.6$, $p = 3.9 \times 10^{-2}$). The current block type did not have a significant effect on perseverance ($F_1 = 0.77$, $p = 0.39$). The interaction between current and past block type was not significant ($F_1 = 0.13$, $p = 0.72$).

We found trends in our data suggesting that online SWR events aid spatial learning. While our findings pertaining online-learning were not statistically significant, the trend in spatial behavior could be seen across multiple behavioral metrics (path length and path-optimality). Furthermore, these trends were observed on Trial 2 which was the first time the animal recalled the learnt location. The same trends were not seen on Trial 1 and Trial 3. Significant differences in behavioral performance on a visually cued version of the task suggested that the animals did not simply rely on a trivial visual strategy to solve the task. We also found that on subsequent learning blocks, when the reward location had been changed, animals spent significantly more time near the previously rewarded well if SWR events were intact in the previous block. If SWR events were disrupted in the previous

block, animals showed a weaker memory of being rewarded at the well by spend lesser time near it. The significant effect on subsequent recall could also imply that perseverance is a stronger behavioral measure of spatial memory when compared to latency or path length. Behavior is complex and the animal may have employed intelligent strategies that do not rely on SWR based mechanisms for learning the goal location. However, the observations made here make a strong case for the necessity of online SWR events learning spatial goals in a familiar environment.

SWR Rebound

It has previously been reported that the disruption of SWR events during sleep leads to an overall elevation in the rate of these events selectively following a spatial task [49]. Here, were disrupted SWR during spatial exploration as opposed to previous work [49, 38, 50, 53] where SWR events were disrupted in subsequent or interleaved sleep. We hypothesized that a disruption of SWR events during the task would also lead to a rebound effect, leading to an overall higher ripple rate in the disrupted condition when compared to controls. It is important to note that in previous studies, the increase was measured with respect to a delayed stimulation control [49], which we did not perform. The control condition in our results corresponds to no electrical stimulation.




Figure 3.18: SWR rates during different behavioral states of the task. (a.) SWR rates from the start of a trial until the animal discovered the currently rewarded well. (b.) SWR rates during the time spent by the animal at the reward well. (c.) SWR rates during the rest period in the home platform following a trial on the maze. $T4$ has been omitted since stimulation was stopped before placing the animal on the platform after the last trial of the block.

In Fig. 3.18 we report the SWR rates in the three different behavioral states in our experiment. Fig. 3.18a. reports SWR rates during the exploratory phase of the task in each trial - starting at maze entry and until the animal reached the reward port for $n_C = 16$ control blocks and $n_E = 20$ experimental blocks with SWR disruption. These could potentially include exploratory SWR events that have previously been reported [37], SWR events during grooming, stopping periods, and stimulation delivered on movement or chewing artifacts. A 75ms cool-off period was imposed after each stimulation which successfully avoided stimulation based on electrical artifacts resulting from current delivery. Using 2-Way

ANOVA, we found a significant effect of experimental group on the SWR rates ($F_3 = 12.0$, $p = 7.1 \times 10^{-4}$). There was not effect of trial ($F_3 = 1.0$, $p = 0.37$) and the interaction between group and trials was not significant ($F_3 = 1.4$, $p = 0.24$). Fig. 3.18b. reports the SWR rates specifically during the times when the animal had stopped at the reward location for $n_C = 16$ control blocks and $n_E = 20$ experimental blocks with SWR disruption. During reward consumption, a 2-way ANOVA revealed a significant effect of group on the SWR rates ($F_1 = 24.5$, $p = 2.0 \times 10^{-6}$). While the effect of trial was not significant, we saw an increase in SWR rates over trials ($F_3 = 2.3$, $p = 0.08$). The interaction between group and trial was not significant ($F_3 = 0.17$, $p = 0.91$). Fig. 3.18c. reports the SWR rates during the ITI when the animal was taken out of the maze and placed on a sleep platform for $n_C = 18$ control blocks and $n_E = 20$ experimental blocks with SWR disruption. Note that SWR disruption was active during the ITI. Using a 2-way ANOVA we found a significant effect of experimental group on the SWR rates ($F_1 = 50.7$, $p = 1.2 \times 10^{-10}$). The effect of trials was not significant ($F_2 = 5.3 \times 10^{-2}$, $p = 0.95$) and the interaction between trial and group was not significant ($F_2 = 0.58$, $p = 0.56$).

Despite the modest behavioral effects observed in Section 3.4, we observed a robust increase in the rate of SWR events on the disruption blocks across behavioral states. These findings are in line with [49], where authors observed an overall increase in SWR rates following disruption in a sleep session following spatial learning. We add to that by assessing SWR rates following disruption during the course of navigation. We found that online SWR rates also increased following disruption on the task. These findings point to the importance of SWR events in navigation, wherein their disruption causes the underlying processes to produce more ripples than the baseline rate potentially as a compensation for the lack of spatial information.

3.5 Discussion and Future Work

Our findings in Section 3.4 did not suggest a significant degradation in behavioral measures during the rewarded trials. We saw trends in the data suggesting a degradation in performance but the results were not statistically significant. In contrast, when we measured performance on the next learning block when the animals spent significantly longer amount of time exploring the maze, the animal spent more time near the goal following a control block when compared to exploration following a Sharp-Wave Ripples (SWR) disruption block. There could be multiple reasons for our observations - Multiple navigational systems exist in the brain and are employed in varied strategies for navigation. Early parts of the SWR events that were not disrupted by electrical stimulation could be sufficient for providing spatial information over a short duration. However, it is possible that when the confidence of the animal is tested over a longer period of time, the loss of spatial memory driven by SWR disruption drives the differences in behavior. In Section 3.4, we found a strong effect of SWR disruption on the rate of SWR events. Disruption of SWR events elevated their overall rate, suggesting a fast time-scale compensatory mechanism which possibly preserves

the replay of spatial trajectories in the Hippocampus associated with SWR events.

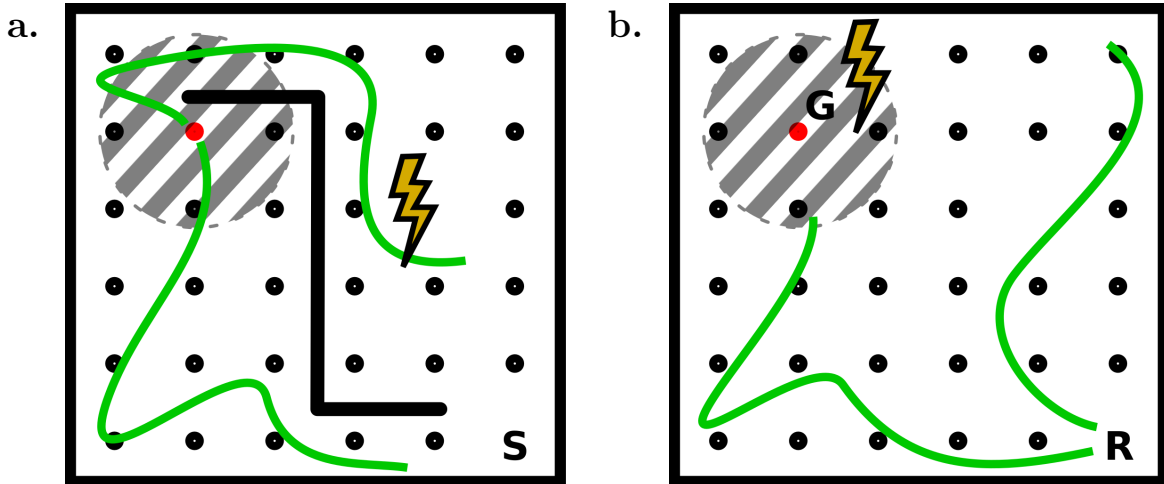


Figure 3.19: (a.) Experimental schematic for testing the role of replay in learning. (b.) Experimental schematic for testing the role of replay in planning.

Our experimental procedures lay the foundation for ongoing and future research that could delve deeper into specific role of replays in learning and planning. With the ability to decode replay events in real-time using techniques described in Section 3.1, and also expanded later in Section 4.3 using machine-learning methods, we can now design experiment where selective SWR events are disrupted contingent on the underlying spatial trajectory being replayed. Fig. 3.19 illustrates two such experimental ideas which can inform us about the role of replay trajectories in spatial planning and learning. Either of these paradigms follow the experimental approach described in Section 3.4. Briefly, in each session, a novel goal location is baited with reward. Animals are provided 4 learning trials - each trial starts at a pseudo-randomly selected location in the environment and after reward consumption, animals are taken out of the maze to discourage a path-integration based strategy. In Fig. 3.19a., we consider the situation when the animal is at a reward location (red circle) and replaying trajectories in an environment separated into two segments by a barrier. We hypothesize that replays propagate the value of the current goal specifically along the trajectory being replayed and interfering with replays disrupts this value propagation². In this experiment, we aim to selectively disrupt replays that traverse the right half of the environment. We hypothesize that when the animal starts a new trial at the location ‘S’, the animal will have worse performance when navigating the right half of the maze to reach the goal when compared to the left half of the maze. In contrast, in Fig. 3.19b., we describe an experiment to test the role of replays in spatial planning. Supposing the animal is stationary at a location ‘R’ which is not the current goal, we will selectively disrupt replays that approach a region

²The ideas around value representation in spatial navigation using a place code have been discussed in more detail in Section 5.3.

around the goal. If replays are involved in spatial planning, we expect this manipulation to degrade the animals' ability to navigate to the goal location.

Chapter 4

Neural-Network architectures for spatial inference using neural activity

Modern machine learning methods have made significant advancements in domains like computer vision, natural language processing, genomics *etc.* With the consistent growth in compute power delivered by Moore's law [128], we are witnessing previously unseen improvements in non-linear classification [74] and regression tasks [78, 57]. Specialized architectures can make use of the computation performed in large families of these models to accelerate performance at scale [68]. Multiple open-sourced libraries [1, 120, 65, 117] now support a broad range of machine-learning algorithms allowing users to deploy complex, optimized algorithms and models using a handful of lines of computer code.

Understanding the structure and function of brain is one of the most challenging problems in science. While much of neuroscience research in the last century focused on isolating cognitive function to individual brain regions and subsequently circuits, there is a need for complexity in our models of cognition [66, 87]. Fundamental to systems neuroscience, and therefore our understanding of the brain, is the inference of behavioral correlates from neural activity. In the field of Brain Machine Interfaces (BMI), for example, activity of motor neurons is used to predict motor actions [132], and subsequently, the actions of these motor neurons can be used to for manipulate objects in the external world. Research in spatial navigation, especially involving Entorhinal Cortex (EC) and the hippocampus relies on inferring spatial position from neural activity [13]. In [8], neural activity from the medial PreFrontal Cortex (mPFC) was used to predict previous and future goal choices on a delayed spatial alternation task. The problem of inferring abstract entities from the neural code is ubiquitous and central to our understanding of the brain as a computational unit.

An extensive study [51] discussed the applications and advantages of applying modern machine learning methods to neural decoding in different domains. Here, we narrow down the application of machine learning to neural networks which outperformed all other methods in the previous study. We propose an end-to-end training methodology which can make use of neural network architectures to address questions pertaining neural decoding in Hippocampal research. We show that the networks can be trained at different temporal resolutions

and that training the network at different resolutions allows the models to capture different physiological phenomenon previously not shown using neural network models. Lastly, we show that at the time-scale of decoding behavioral correlates, neural networks significantly outperform traditional Naïve Bayes model used for decoding spatial position from hippocampal activity. Recently, neural networks have also been brought into the fold of hippocampal research and used for detection of SWR events that were discussed in detail in Chapter 3. In [126], authors propose using neural networks as a substrate understanding the function of the brain. For example, training an hierarchical Artificial Neural Network (ANN) on a visual discrimination task and using the activation of the individual layers to understand mechanism of visual processing. While prior models exist along these lines ([114] for instance uses a sparse coding model to explain coding in V1), neural networks offer an end-to-end model that can be trained using behaviorally relevant objective functions. Incorporating Neural Network (NN) models into neuroscience, therefore, has advantages that surpass that of a computational black-box to ease transforming the neural code into a behavioral or cognitive measure of interest.

Modern machine-learning architectures offer a host of advantages - better infrastructure, robust performance, scalability, and the ability to interpret latent representations which are difficult to get from a single model. Here, we explore some of these aspects in the context of spatial representation in the Hippocampus and demonstrate the viability of neural network architectures as a method for inferring neural activity. In Section 4.1 we define the network architecture as well as the Naïve Bayes model traditionally used in the field. In Section 4.2, we compare the performance of the NN models against the traditional model. Subsequently, we show that these models, while having a superior decoding performance at a behavioral time-scale, can also decode replay events. Finally we show that NN models outperform the traditional Naïve Bayes model for a wide range of recording qualities as measured by total number of simultaneously recorded single units. The flexibility offered by NN models will allow us to expand on the current results and offers robust solutions to problems like cluster-less decoding [32], or even decoding spatial position from mesoscopic Local Field Potential (LFP).

4.1 Models

First we discuss Naïve Bayes, the most commonly used models in Hippocampal research for spatial inference. Then, we describe the neural network architectures that perform spatial inference and methods for training and optimizing these networks.

Naïve Bayes

The Naïve Bayes model for spatial inference as previously described in [13] has been illustrated for a linear environment in Fig. 4.1. For a linear track, the firing rate maps for a single cell recorded from the hippocampus are computed for the two different running directions

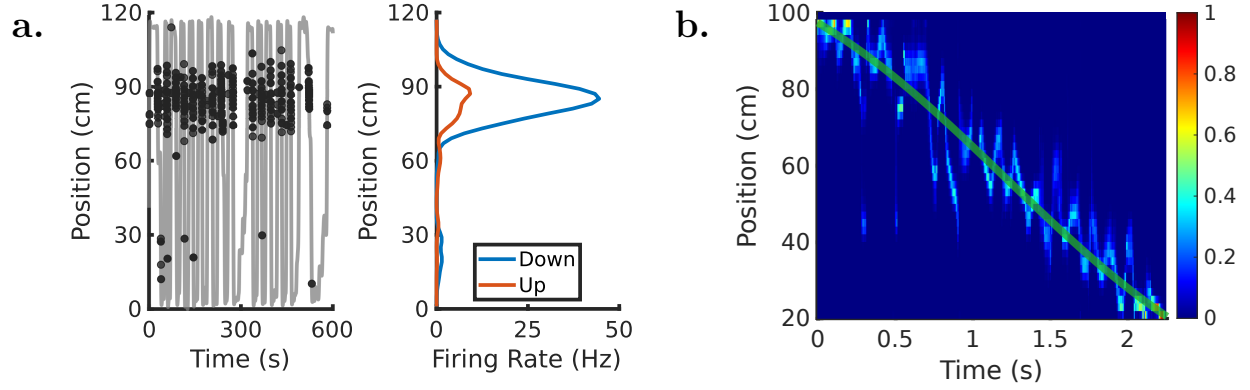


Figure 4.1: Bayesian inference of spatial position in a 1-dimensional environment. (a.) Activity of a single place cell on a linear track. (b.) Posterior probability computed using a Naïve Bayes decoder using a 50ms decoding window.

as shown in Fig. 4.1a.. We further assume that the number of action potentials emitted by a neuron in a given window of time follow a Poisson distribution and the parameter λ_i derived from the firing rate maps previously computed. For cell i , the probability of observing k_i spikes from the cell, given the animal's current position is given by

$$p_i(k_i|x) = \frac{\lambda_i(x)^{k_i} e^{-\lambda_i(x)}}{k_i!}. \quad (4.1)$$

The function $\lambda(x)$ in Eq. (4.1) is the cell's firing rate map, an example of which is shown in Fig. 4.1a.. Further for a population of n simultaneously recorded cells, assuming that the firing distributions for the cells are independent, we can write

$$p(\vec{k}|x) = \prod_{i=1}^n p_i(k_i|x), \quad (4.2)$$

where \vec{k} is the current population firing rate vector. In order to decode position from the current firing rate vector, we apply the Bayes' rule and evaluate

$$p(x|\vec{k}) = \frac{p(\vec{k}|x)p(x)}{p(\vec{k})}. \quad (4.3)$$

Combining Eqs. (4.1) to (4.3), we can build a Bayesian decoder for the animal's position from firing rates for simultaneously recorded cells.

$$p(x|\vec{k}) = \prod_{i=1}^n \left(\frac{\lambda_i(x)^{k_i} e^{-\lambda_i(x)}}{k_i!} \right) \frac{p(x)}{p(\vec{k})}. \quad (4.4)$$

In the modeling equation described in Eq. (4.4), $p(x)$ is the prior on the probability distribution. We implemented this model in MATLAB with a uniform prior to estimate the animal’s position from its population firing rate. Fig. 4.1b. shows the decoded posterior for a single running bout in the recorded session using a 50ms decoding window¹.

Neural Network Architecture

Neural Networks comprise of hierarchical layer of computation units (see Fig. 4.2b. for an example). The hidden units perform a weighted summation of the inputs with a bias term.

$$\mathbf{z} = W\mathbf{x} + b \tag{4.5}$$

Each hidden layer is followed by a non-linear Rectified Linear Unit (ReLU).

$$\mathbf{a} = ReLU(\mathbf{z}) \tag{4.6}$$

$$= \max(\mathbf{z}, 0) \tag{4.7}$$

The dense layers are followed by dropout regularization [141] to enforce a robust representation in the hidden layers. To avoid over-fitting, we incorporate weight-decay (L2-regularization) [75] to the hidden layers of the network. The output layer computes a softmax over the input variables.

$$s(\mathbf{a}_i) = \frac{e^{\mathbf{a}_i}}{\sum_i e^{\mathbf{a}_i}} \tag{4.8}$$

We trained the network to minimize the Kullback–Leible (KL) divergence loss between the predicted and prescribed probabilities over space. The KL divergence loss is defined for discrete probability distributions \mathbf{P} and \mathbf{Q} defined, in this case, over the space of position bins ξ as

$$L(\mathbf{P}|\mathbf{Q}) = \sum_{\mathbf{x} \in \xi} \mathbf{P}(\mathbf{x}) (\log(\mathbf{P}(\mathbf{x})) - \log(\mathbf{Q}(\mathbf{x}))). \tag{4.9}$$

This represents the directed divergence between the two probability distributions \mathbf{P} and \mathbf{Q} . Optimization over the loss function is performed using the Adam optimizer [71].

We implemented the neural network model in Python using TensorFlow2.0 [1] and Keras [25]. Fig. 4.2 highlights the overall architecture of the neural network used for decoding spatial position from neural activity. A typical 15 minute long recording session (typical behavior from a session shown in Fig. 4.2a.) was divided up in 70% train, 20% validation and 10% test sets which were non-overlapping in time. Notably, because place fields of individual cells are known to change over experience [93], later sections of the session were also used for training (first 40% and last 30% of the data constituted the training set). Since neural activity is known to represent non-local positions during periods of awake-rest [47, 123], time periods when the animal ran over 10cm/s were considered for decoding analysis. The train, test and validation sets were then split into 100 or 200ms time-bins for behavioral decoding.

¹The firing rate map used to reconstruct this posterior omitted the decoded lap.

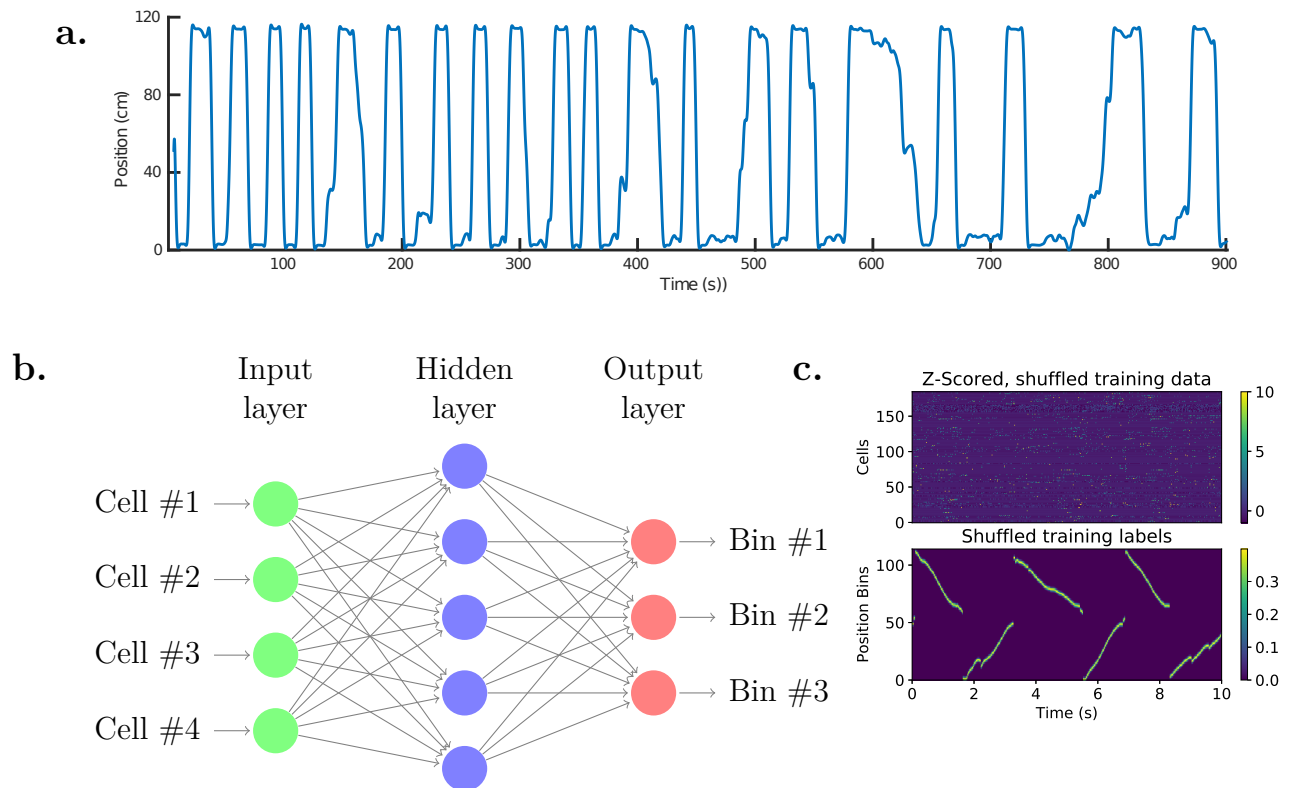


Figure 4.2: Network architecture and data pre-processing. (a.) Illustration of a raw position trace. The dataset was broken down in train (70%), test (10%), and validate (20%) sets. (b.) Architecture of a 2-layer Neural Network used for directional decoding of population firing vectors. A hidden layer comprised of 256 units and output layer comprised of 114 directional place bins were used to decode animal’s position from 185 single units. (c.) Sample training data (top) and labels (bottom) used for training the network.

Consecutive time-bins were offset by 25ms for data-augmentation. Note that overlapping bins were expressly not split across train, test and validation sets by first splitting the raw data into these set and subsequently binning the data into time bins. Firing rates were z-scored to normalize aggregate firing rate differences across cells (Fig. 4.2c., top, illustrates the z-scored firing rates of 185 simultaneously recorded cells in the session).

It is known that place cells exhibit directionality on a linear track [34]. In order to capture this directional sensitivity in the model, the training labels capture the directional position. Raw position was first binned into 2cm bins. For a 120cm track, this would produce 60 position bins. Directionality was incorporated by first doubling the number of position bins and allocating half of the bins for the ‘up’ running direction and the other half for the ‘down’ running direction. Since place cells typically span a larger portion of the environment than the position bin used here [107, 93], firing rates of cells, and in turn, the

entire population, are correlated at nearby locations. We incorporated this aspect into the model by smoothing training labels over position bins, finally training the model to compute a probability distribution over the directional space as shown in Fig. 4.2c., bottom.

We trained fully-connected, 2-layer neural networks to predict a probability distribution over space from the population firing rates as shown in Fig. 4.2b.² The first hidden layer consisted of 256 hidden units, which performed a weighted summation of the z-scored firing rates at the input layer with bias followed by ReLU activation described in Eq. (4.7). Networks with more than 2-layers had a 20% reduction the number of hidden units with every subsequent layer. The output layer’s size was determined by the total number of position bins being decoded. This layer also performed a weighted summation of the previous layer’s activation with bias, and followed it with softmax activation described in Eq. (4.8). Each dense layer was followed by a Dropout layer as a means for regularization [141] when training the network. Further, the learnt weights in each layer were also subject to an $L2$ -regularization loss [75], a method also known as *weight-decay* in the literature.

Fig. 4.3 illustrates the metrics collected when training the model over a single recording session. In Fig. 4.3a., the loss function has been illustrated on the training and validation sets over the course of training the network on a single session over 30 epochs. Since the loss function in this case doesn’t capture the accuracy in prediction animal’s position directly, we calculate that as the MAE as follows:

$$\text{MAE}(t) = \sum_{x \in \xi} |x \mathbf{Q}_m(x) - x_{\text{true}}|. \quad (4.10)$$

Here, $\mathbf{Q}_m(x)$ is the posterior probability distribution over the map m , which is chosen to the the map with higher total posterior using

$$m = \arg \max_d \sum_{x \in \xi} \mathbf{Q}_d(x), \quad (4.11)$$

Therefore the map with a higher total probability is selected first, and then the Center of Mass (CoM) is calculated over the probability distribution in that map as described in Eq. (4.10) The MAE over network training is shown in Fig. 4.3b. Note that even though the network is not optimizing the decoded location of the animal’s position directly, it decreases monotonically over the course of training for both the training and validation sets. Finally, we show the decoded probability averaged across the two maps in Fig. 4.3c. The probability distribution calculated by the model in this case closely follows the animal’s true position. We next quantify this accuracy and compare it against the accuracy of traditional Naïve Bayes model.

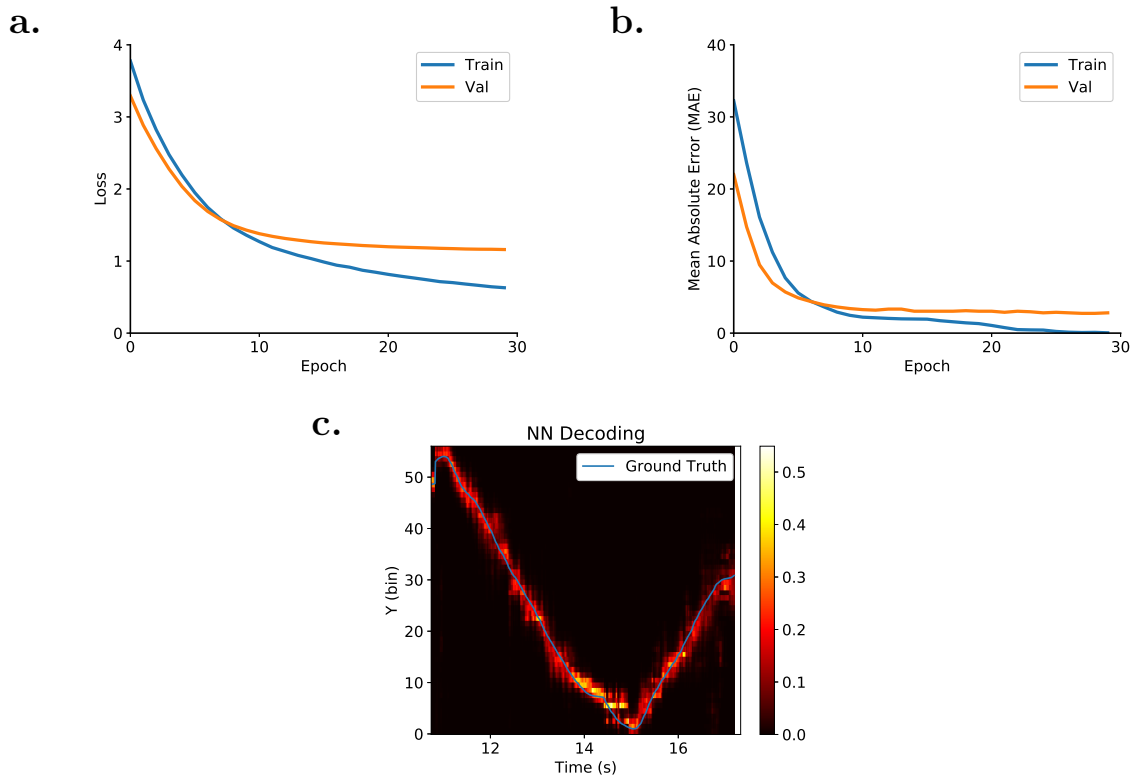


Figure 4.3: Training the network. (a.) KL Divergence loss for training and validation sets over training epochs. (b.) Mean Absolute Error (MAE) measured for training and validation sets over the course of training. (c.) Illustration of probability distribution over position bins learned by the network on the test set.

4.2 Results

Fig. 4.4 shows decoding performance of the NN model (Fig. 4.4b. showing decoding for a single run from one end of the track to the other, decoding window $T = 200ms$, window slide $w = 2ms$) and compares it against Naïve Bayes (Fig. 4.4a. decoding for the same lap as shown in Fig. 4.4b.). In both Figs. 4.4a. and 4.4b., the colorbar shows the posterior probability. Qualitatively, the Naïve Bayes models is more confident and often incorrect in predicting the animal’s position. The NN model on the other hand, effectively captures the variability in the position because of the smoothed training labels. The point estimate computed as the CoM for the NN model is show in Fig. 4.4c. The prediction error between the CoM and average true position in the decoding bin is shown in Fig. 4.4d.. Comparing the decoding error, we found it to be significantly lower for the NN model (One-tailed Wilcoxon rank-sum test $n_B = 1269$ decoded bins with Naïve Bayes, $n_N = 4962$ decoded bins with the

²The neural-network sketch was adapted from <https://tex.stackexchange.com/questions/365404/tikz-neural-network-draw-notation>.

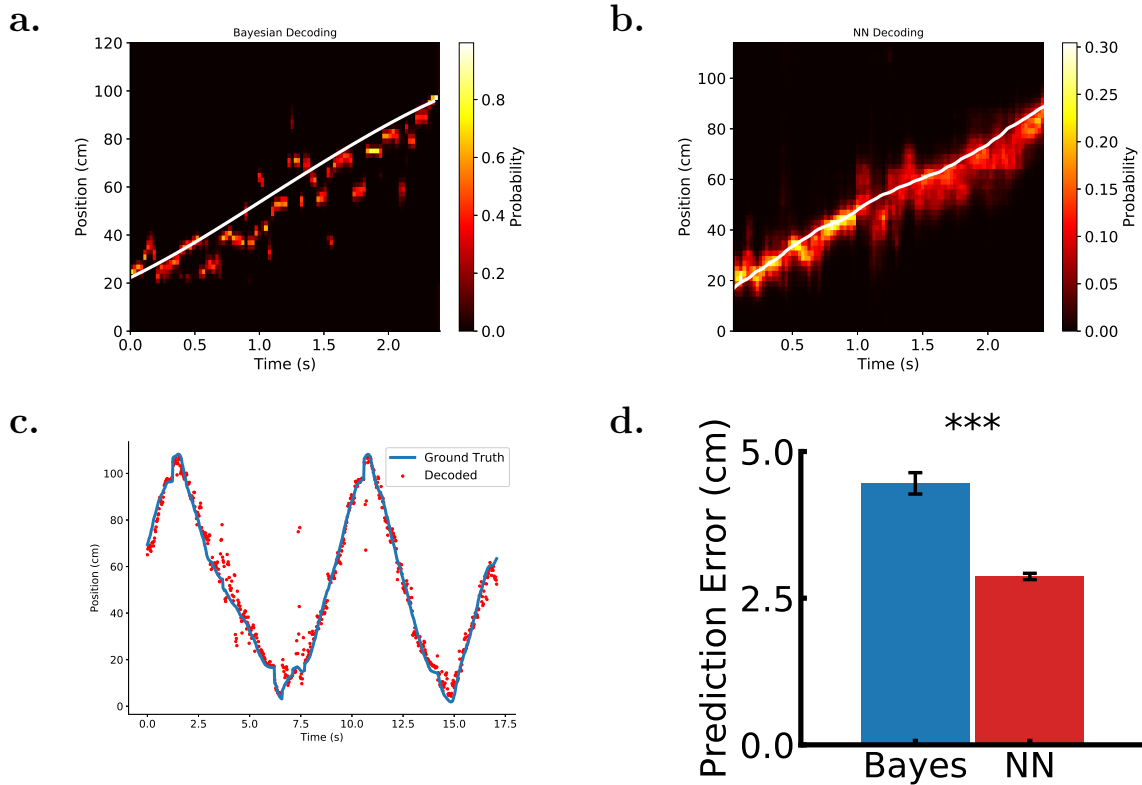


Figure 4.4: Comparing decoding accuracy of Neural Networks with traditional Naive Bayes decoder. (a.) Decoding a single lap on a linear track using Naive Bayes decoder. (b.) Decoding the same lap as in (a.), but using a Neural Network (NN) decoder. (c.) Calculated Center of Mass (CoM) from posterior probability over the test set for NN decoder. (d.) Statistical comparison between Naive Bayes ($n_B = 1269$ decoded bins) and NN decoder ($n_N = 4962$ decoded bins, one-tailed Wilcoxon rank-sum test *statistic* = 2.7×10^6 , $p = 3.5 \times 10^{-15}$).

NN model, $p = 3.5 \times 10^{-15}$). The difference in number of decoded bins largely comes from the way the data was split into train, test and validation sets.

In Fig. 4.5, we discuss the tuning of hyper-parameters for the network architecture. The depth of the network (number of hidden layers with 256 units in the first hidden layer - Figs. 4.5a. and 4.5b.) and the width of the network (number of hidden units in a 2-layer NN - Figs. 4.5c. and 4.5d.) were assessed. The number of epochs were fixed to 25 epochs for training the network on each session³. The validation set was used for measuring performance for hyper-parameter tuning. Networks were pseudo-randomly initialized and trained 25 times for each training session and aggregate performance was reported in Fig. 4.5. Note that the

³A more fair approach for measuring generalization accuracy would be to train the networks to the same training loss, however, we opted to allow similar amounts of training time instead here.

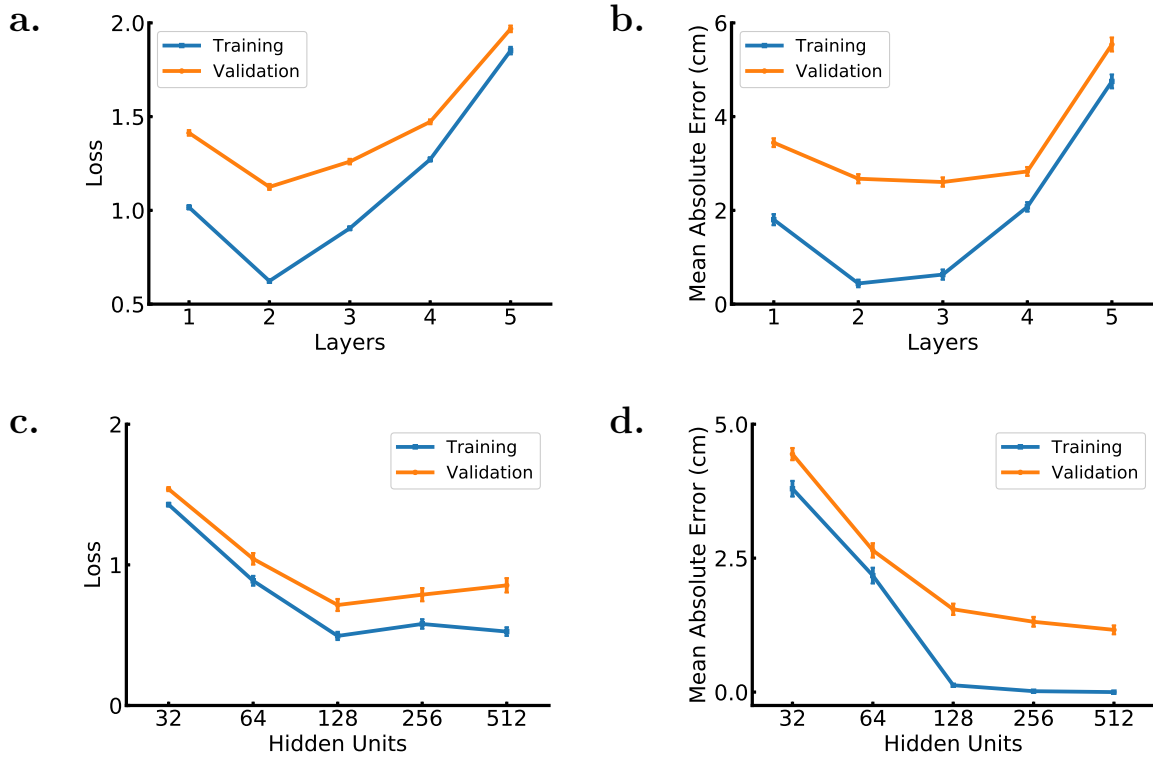


Figure 4.5: Hyper-parameter tuning and network optimization. (a.) Kullback–Leible (KL) divergence loss for training and validation sets when varying the number of NN layers. (b.) Mean Absolute Error (MAE) for training and validation sets when varying the number of NN layers. (c.) KL divergence loss for training and validation sets when varying the width of the hidden layer in a 2-layer NN. (d.) MAE for training and validation sets when varying the width of the hidden layer in a 2-layer NN.

networks were trained to optimize the loss Figs. 4.5a. and 4.5c.. The improvement in *MAE* was simply a measurement here and not used in the training procedure.

Introducing a hidden layer (2-layer *vs.* 1-layer network which performs multi-class logistic-regression - Fig. 4.5d.) led to a significant reduction in the MAE for both the training set (One-tailed Wilcoxon rank-sum test $n = 125$ samples, $statistic = 3.2 \times 10^3$, $p = 1.2 \times 10^{-28}$) and the validation set (One-tailed Wilcoxon rank-sum test $n = 125$ samples, $statistic = 7.1 \times 10^3$, $p = 1.8 \times 10^{-8}$). Increasing the network depth further did not improve the network performance. Allowing greater number of epochs for deeper networks lowered both the training and validation errors but has not been quantified. Tuning the number of hidden units in a 2-layer NN model, we found a near-monotonic improvement in network performance as more and more units were added to the hidden layer. However, the returns diminished after ~ 128 hidden units for the population recorded in our data. These network sizes may vary depending on the neural recordings that the networks are trained on.




Figure 4.6: Hidden unit activation in the NN architecture.

Fig. 4.6 shows the activation of 16 hidden units on the validation data set from a single session. We observed that the hidden units were highly tuned to the directional position of the animal on the track. The spatial and directional selectivity of the hidden units seen in Fig. 4.6 is qualitatively higher than what is observed for a typical place cells in the Hippocampus. However, further quantification is required to understand what hidden units in the network learn and how, if at all, it is different from place coding typically seen in hippocampal cells.

4.3 Decoding Replays with Neural Networks




Figure 4.7: Detection of replay event using Neural Network (NN) architecture with different network depths (a.) 1-Layer NN, (b.) 3-Layer NN, (c.) 5-Layer NN.

Unlike traditional Bayesian inference for spatial decoding, where place fields constructed at a behavioral time-scale can directly be used for decoding replay events by re-scaling the firing rates to the appropriate decoding window, the neural-networks need to be re-trained with the decoding window of interest. For decoding replays, the run-time data was binned into 20ms decoding windows shifted by 2ms. Time periods designated as replays detected




Figure 4.8: Detection of replay event using 2-layer NN architecture with different network widths: (a.) 32 hidden units, (b.) 128 hidden units, (c.) 512 hidden units.

using the Naïve Bayes decoder were isolated and analyzed in the replay posterior⁴. Figs. 4.7 and 4.8 show a single replay event decoding using a wide range of hyper-parameters. While there were qualitative changes in the posterior-probability, the event map and directionality could be correctly identified across the board.

We measured the agreement between replays detected using the Naïve Bayes model and the NN model. For both models, decoding was done using 50ms decoding windows shifted by 5ms. Two metrics were used for comparing detection of replays using the models - (i) The map with overall higher total posterior over the course of the replay (see Eq. (4.11)); and (ii) Sign of weighted-correlation of the total posterior. Weighted correlation was calculated as previously described in [163]:

$$m(\mathbf{x}; \mathbf{w}) = \frac{\sum_{i=1}^M \sum_{j=1}^N w_{ij} x_i}{\sum_{i=1}^M \sum_{j=1}^N w_{ij}}, \quad (4.12)$$

$$cov(\mathbf{x}, \mathbf{y}; \mathbf{w}) = \frac{\sum_{i=1}^M \sum_{j=1}^N w_{ij} (x_i - m(\mathbf{x}; \mathbf{w})) (y_j - m(\mathbf{y}; \mathbf{w}))}{\sum_{i=1}^M \sum_{j=1}^N w_{ij}}, \quad (4.13)$$

$$corr(\mathbf{x}, \mathbf{y}; \mathbf{w}) = \frac{cov(\mathbf{x}, \mathbf{y}; \mathbf{w})}{\sqrt{cov(\mathbf{x}, \mathbf{x}; \mathbf{w}) \cdot cov(\mathbf{y}, \mathbf{y}; \mathbf{w})}}. \quad (4.14)$$

In the computation of Weighted Correlation (WC) of a replay event described in Eq. (4.14), x_i is the i^{th} time bin out of a total of M bins and y_j is the j^{th} position bin out of a total of N position bins. w_{ij} is the posterior probability at the j^{th} position bin at the i^{th} time points in the replay. A replay traversing the map in the ‘up’ direction (for example Fig. 4.8a.) will have a positive WC, whereas a replay traversing a map in ‘down’ direction will have a negative WC. Both the map being replayed and the sign of WC deduced from the replay

⁴Events were identified as peaks in Sharp-Wave Ripples (SWR) power and event end-points were determined by identifying time points where ripple power fell to 1-STD over baseline on either sides of peak power.




Figure 4.9: Agreement on replay metrics between Naïve Bayes and NN models over the six recorded sessions.

event showed high agreement between the NN model and Naïve Bayes as shown in Fig. 4.9 - Map accuracy (mean±std) $81 \pm 11\%$, Sign accuracy $75 \pm 7\%$.

4.4 Down-sampling Performance

In order to further stress the models, we down-sampled the cells simultaneously recorded within a single session and used the down-sampled population for spatial inference. For each of the recorded sessions and a down-sampled population target say p , p cells were randomly selected from the population of recorded cells (ranging in 150 – 200 simultaneously cells in individual sessions). The selected cells were used to build the place maps for the Naïve Bayes decoder and to train the neural network for the Neural Network (NN) model. The held out data (test-set) was used to measure decoding performance for either models. This procedure was repeated 25 times and the aggregate performance measures were reported in Fig. 4.10.

Statistical tests comparing the performance of the two models have been reported in Table 4.1. For cell counts as low as 50 (which amounts of $\sim 25\%$ of the recorded population for some sessions), the NN model outperformed Naïve Bayes. When the cell counts were lowered to 25, the difference in performance between the two models was not statistically significant. These findings suggest that NN models are more robust for decoding spatial position from neural activity from place cells, and as the number of simultaneously recorded cells increases, these models will continue to outperform the existing methods for analysis of neural activity.




Figure 4.10: Comparing performance of tradition Naïve Bayes model again Neural-Networks with a down-sampled cell population. (a.) 50 Cells (b.) 75 Cells (c.) 100 Cells (d.) 125 Cells

Cell Count	Test-Statistic	P-Value
25	1.1×10^4	0.99
50	6.1×10^3	1.5×10^{-3}
75	4.6×10^3	1.5×10^{-8}
100	5.4×10^3	9.2×10^{-6}
125	3.8×10^3	7.4×10^{-13}

Table 4.1: One-tailed Mann-Whitney U test statistics comparing the performance of Naïve Bayes decoder against neural network for different cell counts $n_C = 25$ random cell population selection over $n_S = 5$ sessions for a total of $n = 125$ data-samples per comparison.

4.5 Discussion

Here, we have described a novel approach for designing and training neural networks for inferring spatial position from neural data that was electrically recorded from rodents. At longer decoding time-scales, we showed that Neural Network (NN) models outperform the traditional Naïve Bayes model used for spatial decoding. Further, we showed that on training the network at a shorter time-scale on behavioral (run) data, it generalizes to time-compressed replay events that occur during restfulness. Despite the time-compression of neural activity during replay events, networks generalized to those representations that they had never been trained or optimized for. We showed that these events can be effectively decoded for a wide range of hyper-parameters, suggesting that fine grained network tuning is not required for the adoption of neural networks into neural decoding in this context. Our findings suggest that

neural network models have exciting potential for inferring high resolution recordings from the brain, and that as the scale of neural recordings expands, modern machine learning models might excel at learning underlying representations in the data better than conventional models that are presently used.

While it is possible that the improved performance of NN model for behavioral time-scale decoding discussed here results from a difference in model size, models with fewer parameters (Fig. 4.5) also showed good decoding performance. Moreover, there is a fundamental difference between the Naïve Bayes decoder and the NN decoder which could potentially account for the difference in performance. It is well known that firing patterns of different cells in the hippocampus vary in the nature of spatial information contained within them. For example, there is a systematic variation in the spatial tuning curves along the longitudinal axis of the hippocampus [69, 72]. Furthermore, electrical recordings can include both pyramidal cells and inter-neurons. While inter-neuron firing contains spatial information that could be useful for decoding, the firing patterns of inter-neurons are characteristically different from Pyramidal cells [39]. Within the Naïve Bayes model, the Poisson distribution assumption is highly sensitive to the baseline firing rate of a cell, allocating higher weights to cells with higher firing rates. The model doesn't have a way of allocating weights to individual cells based on their ability to capture the animals' position. The NN model on the other hand, utilizes back-propagation to allocate different weights to different cell activities based on their ability to predict the animal's position.

Another fundamental question that remains to be answered pertains to the meaning of a *better* decoder in the context of hippocampal replay. In the absence of a ground truth for a replay, often arbitrary thresholds are used to enforce continuity and directionality among other constraints to establish an event as a replay. Here, we have been able to ascertain that there is an agreement between the Naïve Bayes decoder and the NN model on these properties. However, if replay were to exist on a continuum of these properties, as illustrated by analysis of all spatial trajectories observed during immobility periods [144], it is possible that the flexibility offered by the NN model makes it a better decoder for studying replays.

Abstract models that do not rely on an explicit behavioral model of neural activity allow analyses to expand into informative signals that are not well structured. While the rate code is more commonly believed to be the way information is represented in brain, there is evidence of oscillatory phase being recruited to represent information in the brain. Recently, it was shown that unfolding the theta phase at different points in the hippocampus can yield a spatial decoded that relies solely on the local oscillation phase [2]. It is possible to borrow architectures from computer vision [74] and natural language processing [60] to construct an architecture that can learn oscillatory features directly from the Local Field Potential (LFP) signal. LFP-based decoding models also raise the possibility of learning representation that can generalize across different sessions within the same animal because of lower variability in the LFP, or also perhaps, across different animals if we can learn general spectro-temporal patterns that generalize well across individuals. Neural networks provide a more suitable framework for learning such patterns. Finally, they offer a structure for a multi-modal decoder incorporating spiking, LFP, and imaging among a wide range of

signals for decoding neural activity.

Chapter 5

Conclusion

In this dissertation, we've touched on three different aspects of the representation of space in brain. We first assessed the anatomical profile of theta and gamma oscillations in Chapter 2, finding that gamma oscillations are synchronized in the hippocampus and can potentially act as a synchronizing signal across different spatial scales of representation. Then, in Chapter 3, we described the tools and methods that are necessary to collect data from large populations of place cells in the Hippocampus. We built a pipeline from the ground up for processing and visualizing neural data in real-time and delivering closed-loop stimulation. We applied the closed-loop stimulation protocol to a novel spatial-learning task and found trends in learning deficits. Finally, in Chapter 4, we proposed machine-learning solutions for inferring neural activity at scale and showed that these methods outperform the techniques that are used today while simultaneously generalizing to latent phenomenon that the models had never been trained on.

5.1 Gamma Oscillations: A synchronized marker for spatial representation

We showed that in our recordings, we could observe the anatomical traveling-wave nature of theta within the dorsal hippocampus. While we observed that both slow and fast gamma oscillations were locked to theta frequency, across recording sites, there was a greater agreement in gamma amplitude-theta phase locking if we defined an abstract global theta phase, than when compared to the local theta phase which was used to compute the gamma amplitude. This led us to believe that gamma events were globally synchronized in the Hippocampus and the gamma power envelope did not exhibit traveling wave characteristics despite being strongly modulated by, and locked to theta. We further tested our hypothesis by simultaneously recording from dorsal and intermediate hippocampus which spanned roughly half of the dorsoventral extent of the hippocampus. We found our hypothesis to be true - theta exhibited a much larger lag across the regions than both slow and fast gamma. Finally, we suppressed CA3 input to CA1 to causally test the role of CA3 in driving gamma os-

cillations in CA1. Contrary to the existing belief that CA3 selectively drives slow gamma oscillations in CA1, we surprisingly found that both slow and fast gamma persisted in the absence of CA3 input. Moreover, both oscillations saw graded changes when CA3 input was suppressed commensurate with the extent of input suppression. Finally we also observed that the global synchrony of both slow and fast gamma events degraded as CA3 input to CA1 was suppressed.

These findings suggest interesting directions to advance the study of spatial representation in the Hippocampus. It is becoming increasingly clear that cross-frequency coupling is a robust means of communication within and across brain regions. In this context, there is prior evidence suggesting gamma synchrony during periods of awake-rest after important for spatial memories [22, 122]. We reinforce and generalize this idea to the alert and ambulatory periods and propose that gamma oscillations in the Hippocampus provide a coherent clock that the spatial representations can be referenced to. This finding has deeper implications about how spatial memories are inferred. While sequential activity is often studied in the context of either theta oscillations or Sharp-Wave Ripples (SWR) events, neither of these phenomena are synchronized across the hippocampus [119, 118]. In order to assess a sequential activation of spatial trajectory from the entire population of place cells in hippocampal CA1, we propose the use of gamma events - both during active (theta sequences) and inactive periods (replays). This brings us to the study of these sequential events in a closed-loop experiment.

5.2 The micro-structure of neural activity

Static tuning curves have been pivotal in our understanding of how the brain forms representations of the world. Tuning curves in the visual cortex [62], for an instance, have allowed researchers to build a computational framework for the early layers of the visual cortex relying on robust sparse codes [113]. Increasing evidence suggests that such static tuning curves do not provide a complete picture of how information is represented in the brain. Information, in fact, is often represented using dynamic codes or sequences of individual cells or ensembles. This is not restricted to sequences in the Euclidean space that we delved into earlier in this dissertation - In humans for examples, it was shown that memorization of ordered word pairings depends on sequential codes [153]. The origin, structure and function of such sequential codes is not well understood and stands to deliver deep insights into information processing in the brain.

Interestingly, these sequential codes are often accompanied with oscillations in the Local Field Potential (LFP). Replay of word-pairing codes in the aforementioned study, for example, were accompanied with Sharp-Wave Ripples (SWR) events. In rodents, this micro-structure is seen both during restful periods as replays accompanied with SWR events [47], and during running bouts as theta sequences [46]. Importantly, oscillations frequently do not generalize across species. In bats, for example, theta oscillations are not observed [165]. SWR, on the other hand, appear invariant across species and have been reported in ro-

dents [17, 47, 64, 30], non-human primates [81, 63] and humans [105]. Irrespective of the species specific oscillatory traits, they have one major advantage over the direct study of single-unit activity - Oscillatory activity is easier to acquire and assess. Designing closed-loop experiments around oscillatory phenomena allows us to tap into the underlying sequential codes without having to explicitly record them. It also provides a rare opportunity to translate our understanding into clinical applications. Our behavioral study in Chapter 3 suggests a role for SWR events in online spatial-learning. While some of our results were not conclusive statistically, and the results came from a single animal, we observed trends in the behavioral measures suggesting that SWR events are needed for learning a novel spatial goal in a familiar environment. In the future, we propose combining our prior findings and designing closed-loop experiments centered on gamma-activity to be better able to target events that are synchronized across the entire hippocampus.

As our ability to gather information from the brain improves, we will reach the state where intelligent algorithms for extracting information from neural recordings become a necessity. Modern machine learning methods have had major success in previously arduous problems like computer vision, natural language processing, and generative modeling to name a few [78, 57]. Within this context, neural networks have been particularly successful. Further, existing open-source machine-learning frameworks are both efficient and scalable [1, 117, 65]. We have shown that neural-networks can be deployed in neuroscience to improve the quantitative performance of decoding models in Chapter 4. The ability of neural networks to learn patterns from data open the gateway for models that take in raw data and produce behavioral correlates needed for analytics and closed-loop processes like Brain Machine Interfaces (BMI).

5.3 A model for spatial navigation

Here we discuss a model-free Reinforcement Learning (RL) framework to understand how place cells could be used for making decisions in the spatial tasks previously discussed in Chapter 3. The model is principally based on learning a value representation in an environment based on reward contingencies. While this is discussed in detail in subsequent sections, the key idea here is that spatial tuning curves of place cells can serve as basis functions, which can then be weighed to learn a value function over space.

Value Learning

Consider a state-transition system with a set of states \mathbf{S} . The value function is defined over the state variables as the discounted, expected long-term reward for any state that the agent is currently in. Mathematically, the value function at the current state $\mathbf{s}(t)$ is given by

$$V^\pi(\mathbf{s}) = r(\mathbf{s}(t)) + \sum_{i=0}^{\infty} \left(\gamma^i r(\mathbf{s}(t+i)) \right). \quad (5.1)$$




Figure 5.1: Schematic of typical behavioral trajectories over the course of learning in a spatial navigation task. (a.) Typical trajectory before the goal location has been learned. (b.) Illustration of a behavioral trajectory after the goal location has been learnt.

Here, γ is the discount factor and $r(\mathbf{s})$ is the reward assigned to the state \mathbf{s} . The state transitions for the agent depend on the policy π that the agent applies. One of the popular methods for learning the state-value function V^π under the policy π is Temporal-Difference (TD) Learning. Under TD Learning, the estimated state-value function \hat{V}^π is updated in an iterative manner using

$$V(\mathbf{s}) \leftarrow V(\mathbf{s}) + \alpha (r(\mathbf{s}) + \gamma V(\mathbf{s}') - V(\mathbf{s})), \quad (5.2)$$

assuming that as a result of the current action, the agent transitioned from state \mathbf{s} to \mathbf{s}' and earned a reward $r(\mathbf{s})$. Within this framework, a learning-rate hyperparameter α is also used.

$$Q(\mathbf{s}_t, a_t) \leftarrow Q(\mathbf{s}_t, a_t) + \alpha \left(r_t + \gamma \max_a Q(\mathbf{s}_{t+1}, a) - Q(\mathbf{s}_t, a) \right) \quad (5.3)$$

As opposed to learning the value function defined in Section 5.3 explicitly, we instead work with the q-learning framework. The policy ($Q(\mathbf{s}_t, a_t)$), which is a function over state-action pairs is learnt as the agent navigates the environment..

Actor-Critic Model

We now explore an actor-critic model for goal-oriented spatial navigation, originally described in [45]. Computer code for this work is publicly available at <https://github.com/architgupta93/RL.Hippocampus>. Our agent explores a 2-dimensional maze in search of a single reward location as shown in Fig. 5.5a.. Synthetic place cells reflect the animal's current position using a firing rate map that follows a Gaussian distribution around a preferred

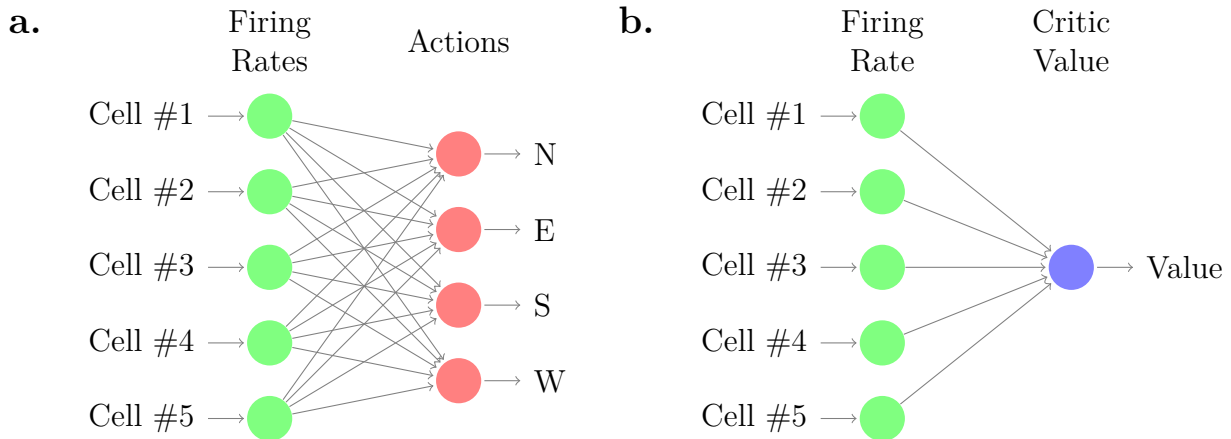


Figure 5.2: (a.) Model definition for an actor which uses a weighted combination of place-cell activity to produce directed movement in the maze. (b.) Model definition for a critic which combines weighted place cell activity with reward information to update the animal’s navigational policy.

firing location. Consider cell i with a preferred firing location p_i . The instantaneous firing rate of this cell at animal’s current location (or state) \mathbf{s} is given by

$$f_i(\mathbf{s}) = m_{fr} \exp\left(-\frac{\|\mathbf{s}-p_i\|^2}{2\sigma^2}\right). \quad (5.4)$$

In Eq. (5.4), m_{fr} represents the mean firing rate of a cell and was set to 0.5Hz, the variance σ was set such that place fields covered $\sim 5\%$ of the environment. The agent’s current location fully defined its state \mathbf{s} , and at any such state, the agent chose between 4 possible directional actions, namely (N, S, E, W) using a weighted combination of the firing rates as shown in Fig. 5.2a..

$$a_j(\mathbf{s}) = \sum_i z_{ji} f_i(\mathbf{s}). \quad (5.5)$$

Here a_j represents the activation for the four directional actions (N, E, S, W), and z_{ji} represents the contribution of cell i to action j . In order to incorporate the observed momentum in behavioral trajectories, the most recently selected action was up-weighted by a *momentum* factor set in (1.1, 2.0). Actions were selected using a softmax function from the activation a_j such that the probability p_j of selecting the action j is given by

$$p_j = \frac{\exp(2a_j)}{\sum_m \exp(2a_m)} \quad (5.6)$$

Near the edge of the maze, or an obstructing wall, the set of actions is restricted to the ones that the agent can legally take. The critic model, shown in Fig. 5.2b., similarly computed a weighted summation of the firing rates to estimate the *value* of the current state. Actor




Figure 5.3: (a.) Trajectory for agent (blue circle marks start, red square marks trajectory end and green circle marks the reward) on Trial 2 in the current configuration. (b.) Spatial trajectory on the same configuration as in (a.), but for Trial 40. (c.) Value representation learnt by the critic model using a weighted combination of the firing rate maps described in Eq. (5.4).

and Critic were simultaneously learnt using a variation of TD learning described in detail in [45]. The agent traded-off a greedy exploitation policy and an exploration policy using a probabilistic selection of the actions weighed by the expected discounted long-term reward for selecting the action. In Fig. 5.3 we illustrate how our agent learnt to traverse an environment with barriers to a fixed goal location over learning trials. This was a square $6m \times 6m$ environment with two barriers bisecting the arena separated by a small opening. Fig. 5.3a. illustrates the trajectory of the agent early in learning a novel goal configuration (Trial 2). Notice that the trajectory is not directed to the goal and similar to our experimental observations in Section 3.4, the agent explored unrewarded parts of the environment before reaching the goal location. Later in learning, as shown in Fig. 5.3b., from a similar starting location, the agent traced a direct path to the goal. Fig. 5.3c. shows the value representation learnt by the critic model inside the maze which captures both the the reward structure (high value near the goal) as well as the environmental structure (low value near the barriers). In order to further quantify the agent’s behavior and compare these data with the experimental observations in Section 3.4, we measured the agent’s latency to goal in Fig. 5.4.

In Fig. 5.4a. we report the excess distance moved by the agent compared to an optimal trajectory over the first 40 learning trials. After the learning trials have expired, the agent was reintroduced to the maze with the same reward contingency, however, now the Q-learning algorithm has stopped and the state-action values had been frozen. The navigational performance in this phase is shown in Fig. 5.4b. where the agent consistently took near-optimal routes to the previously learned goal location.

Fig. 5.5b. shows the value function learnt by the agent in another maze configuration over the learning procedure. Here, reward structure is shown in Fig. 5.5a. and the value representation learnt by the critic model is shown in Fig. 5.5b.. Again, we observe that the learnt value function captures both the reward structure and environmental contingencies




Figure 5.4: Distance moved by an agent simulating reward search in an open environment using Value learning. (a.) During learning trials. (b.) After learning.

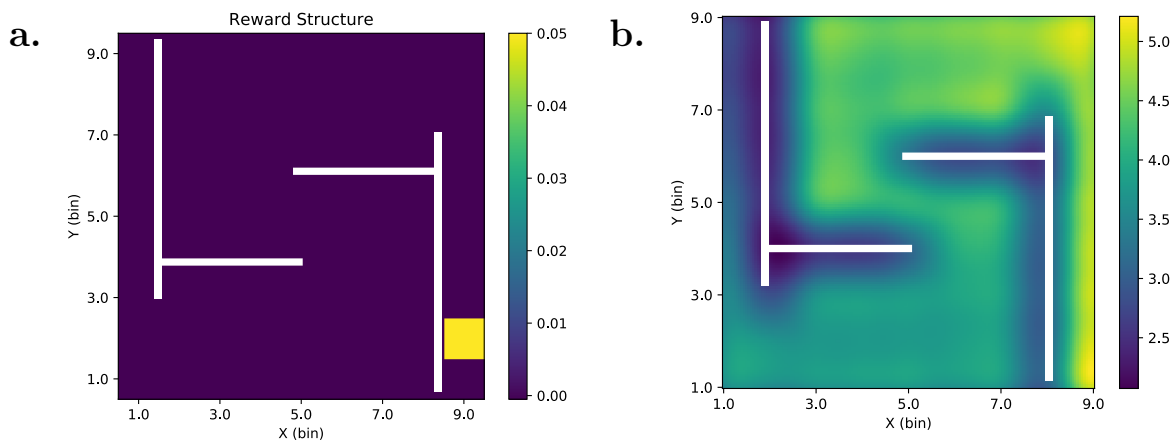


Figure 5.5: Reward structure and inferred value function in spatial learning in a maze. (a.) Structure of reward information in the maze experienced by the learning agent. (b.) Value function learned as a superposition of weighted place cell firing rates by the

by accounting for the distance the agent would have to travel from any initial location to the reward.

While this models serves as a foundation for experimental design, a lot questions remained to be answered about it biological plausibility as well as validity in explaining behavioral data. Earlier, we discussed research demonstrating a unique modulation of Sharp-Wave Ripples (SWR) events and in particular, reverse replays, by changing reward contingencies which have been established to be linked to value representation and Reward Prediction Error (RPE) [4]. While this models helps understand active navigation from an RL standpoint, it remains unclear if, and how replays fit into this learning framework.

Theoretical models have explored extensions of RL models to explain the phenomenon

of hippocampal replay. It is well known that replay trajectories aren't simple recapitulations of behavioral experiences [54], suggesting a mental simulation of spatial trajectories. Modeling work suggests that multiple RL models can explain hippocampal replay in the awake state [23], the simplest being a traversal of the value landscape. The Dyna family of algorithms [145] are well suited for explaining replays because of the use of Bellman backups for updating the policy Q in an offline state. [88] described a model incorporating the Dyna algorithm to provide a detailed model of replay by chaining Bellman backups. One of the key aspects of the model was selecting replay trajectories that would maximally improve the expected return resulting from the policy of the agent following Bellman backups, called prioritized replay. The authors decomposed the utility of accessing an experience, conditioned on the target state s_t , into a gain term which measured the expected improvement in return at the state s_t and a need term which measured the discounted expected number of visits to the state s_t . The latter was computed using the successor representation [31]. It was shown the using prioritized replays described above improved the rate of spatial learning. However, in the absence of prioritization, or even replays altogether, similar levels of performance could be reached by the agent, albeit at a slower rate. The interplay between the need and gain terms helps explain several of the experimental observations. For example, forward replays are prevalent at the start of a run, whereas, reverse replays are observed in larger numbers at the end of a run [34]. The large gain term at the end of the run near reward explains a prioritization of reverse replays, whereas, the dominance of the need term explains the prominence of forward replays at the start of the run. The over-representation of goal location (high gain) and local initiation bias of replays (high need) are similarly explained. The model, however, falls short in explaining the biological computation of gain. Moreover, how neural circuits within or outside the hippocampus drive these trajectories remains to be answered.

The interplay between oscillatory phenomena and neural activity is similarly not well understood. Recent research has also highlighted the advantages of oscillators as computational substrates [156, 155]. While oscillatory phenomenon have contributed vastly to a mechanical understanding of computation and inter-regional communication in the brain, it also remains unclear if and how brain employs oscillatory codes into computation. Finally, a comprehensive theoretical model of phenomena like oscillatory interactions and replays will allow us to design better experiments to probe their role in cognitive function.

Bibliography

- [1] Martin Abadi et al. “Tensorflow: A system for large-scale machine learning”. In: *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*. 2016, pp. 265–283.
- [2] Gautam Agarwal et al. “Spatially distributed local fields in the hippocampus encode rat position”. In: *Science* 344.6184 (2014), pp. 626–630.
- [3] David G Amaral and Menno P Witter. “The three-dimensional organization of the hippocampal formation: a review of anatomical data”. In: *Neuroscience* 31.3 (1989), pp. 571–591.
- [4] R Ellen Ambrose, Brad E Pfeiffer, and David J Foster. “Reverse replay of hippocampal place cells is uniquely modulated by changing reward”. In: *Neuron* 91.5 (2016), pp. 1124–1136.
- [5] Gheorghe Apostol and Otto D Creutzfeldt. “Crosscorrelation between the activity of septal units and hippocampal EEG during arousal”. In: *Brain research* 67.1 (1974), pp. 65–75.
- [6] DP Artemenko. “Role of hippocampal neurons in theta-wave generation”. In: *Neurophysiology* 4.5 (1972), pp. 409–415.
- [7] Nikolai Axmacher, Christian E Elger, and Juergen Fell. “Ripples in the medial temporal lobe are relevant for human memory consolidation”. In: *Brain* 131.7 (2008), pp. 1806–1817.
- [8] EH Baeg et al. “Dynamics of population code for working memory in the prefrontal cortex”. In: *Neuron* 40.1 (2003), pp. 177–188.
- [9] Mariano A Belluscio et al. “Cross-frequency phase–phase coupling between theta and gamma oscillations in the hippocampus”. In: *Journal of Neuroscience* 32.2 (2012), pp. 423–435.
- [10] Philipp Berens et al. “CircStat: a MATLAB toolbox for circular statistics”. In: *J Stat Softw* 31.10 (2009), pp. 1–21.
- [11] Elizabeth Bostock, Robert U Muller, and John L Kubie. “Experience-dependent modifications of hippocampal place cell firing”. In: *Hippocampus* 1.2 (1991), pp. 193–205.

- [12] Anatol Bragin et al. “Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat”. In: *Journal of neuroscience* 15.1 (1995), pp. 47–60.
- [13] Emery N Brown et al. “A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells”. In: *Journal of Neuroscience* 18.18 (1998), pp. 7411–7425.
- [14] Alessio P Buccino et al. “SpikeInterface, a unified framework for spike sorting”. In: *Elife* 9 (2020), e61834.
- [15] Eduardo R Butelman. “A novel NMDA antagonist, MK-801, impairs performance in a hippocampal-dependent spatial learning task”. In: *Pharmacology Biochemistry and Behavior* 34.1 (1989), pp. 13–16.
- [16] György Buzsáki. “Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning”. In: *Hippocampus* 25.10 (2015), pp. 1073–1188.
- [17] György Buzsáki. “Hippocampal sharp waves: their origin and significance”. In: *Brain research* 398.2 (1986), pp. 242–252.
- [18] György Buzsáki. “Theta oscillations in the hippocampus”. In: *Neuron* 33.3 (2002), pp. 325–340.
- [19] György Buzsáki, Costas A Anastassiou, and Christof Koch. “The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes”. In: *Nature reviews neuroscience* 13.6 (2012), pp. 407–420.
- [20] György Buzsáki, Cornelius H Vanderwolf, et al. “Cellular bases of hippocampal EEG in the behaving rat”. In: *Brain Research Reviews* 6.2 (1983), pp. 139–171.
- [21] Margaret F Carr, Shantanu P Jadhav, and Loren M Frank. “Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval”. In: *Nature neuroscience* 14.2 (2011), p. 147.
- [22] Margaret F Carr, Mattias P Karlsson, and Loren M Frank. “Transient slow gamma synchrony underlies hippocampal memory replay”. In: *Neuron* 75.4 (2012), pp. 700–713.
- [23] Romain Cazé et al. “Hippocampal replays under the scrutiny of reinforcement learning models”. In: *Journal of neurophysiology* 120.6 (2018), pp. 2877–2896.
- [24] Lee A Cenquizca and Larry W Swanson. “Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex”. In: *Brain research reviews* 56.1 (2007), pp. 1–26.
- [25] François Chollet et al. *keras*. 2015.
- [26] Jason E Chung et al. “A fully automated approach to spike sorting”. In: *Neuron* 95.6 (2017), pp. 1381–1394.
- [27] Laura Lee Colgin. “Rhythms of the hippocampal network”. In: *Nature Reviews Neuroscience* 17.4 (2016), pp. 239–249.

- [28] Laura Lee Colgin et al. “Frequency of gamma oscillations routes flow of information in the hippocampus”. In: *Nature* 462.7271 (2009), pp. 353–357.
- [29] Jozsef Csicsvari et al. “Mechanisms of gamma oscillations in the hippocampus of the behaving rat”. In: *Neuron* 37.2 (2003), pp. 311–322.
- [30] Heydar Davoudi and David J Foster. “Acute silencing of hippocampal CA3 reveals a dominant role in place field responses”. In: *Nature neuroscience* 22.3 (2019), pp. 337–342.
- [31] Peter Dayan. “Improving generalization for temporal difference learning: The successor representation”. In: *Neural Computation* 5.4 (1993), pp. 613–624.
- [32] Xinyi Deng et al. “Clusterless decoding of position from multiunit activity using a marked point process filter”. In: *Neural computation* 27.7 (2015), pp. 1438–1460.
- [33] Sachin S Deshmukh and James J Knierim. “Representation of non-spatial and spatial information in the lateral entorhinal cortex”. In: *Frontiers in behavioral neuroscience* 5 (2011), p. 69.
- [34] Kamran Diba and György Buzsáki. “Forward and reverse hippocampal place-cell sequences during ripples”. In: *Nature neuroscience* 10.10 (2007), p. 1241.
- [35] George Dragoi and Susumu Tonegawa. “Preplay of future place cell sequences by hippocampal cellular assemblies”. In: *Nature* 469.7330 (2011), p. 397.
- [36] George Dragoi et al. “Interactions between hippocampus and medial septum during sharp waves and theta oscillation in the behaving rat”. In: *Journal of Neuroscience* 19.14 (1999), pp. 6191–6199.
- [37] David Dupret et al. “The reorganization and reactivation of hippocampal maps predict spatial memory performance”. In: *Nature neuroscience* 13.8 (2010), p. 995.
- [38] Valérie Ego-Stengel and Matthew A Wilson. “Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat”. In: *Hippocampus* 20.1 (2010), pp. 1–10.
- [39] Valérie Ego-Stengel and Matthew A Wilson. “Spatial selectivity and theta phase precession in CA1 interneurons”. In: *Hippocampus* 17.2 (2007), pp. 161–174.
- [40] Guillaume Etter et al. “Optogenetic gamma stimulation rescues memory impairments in an Alzheimer’s disease mouse model”. In: *Nature communications* 10.1 (2019), pp. 1–11.
- [41] Ting Feng, Delia Silva, and David J Foster. “Dissociation between the experience-dependent development of hippocampal theta sequences and single-trial phase precession”. In: *Journal of Neuroscience* 35.12 (2015), pp. 4890–4902.
- [42] Antonio Fernández-Ruiz et al. “Entorhinal-CA3 dual-input control of spike timing in the hippocampus by theta-gamma coupling”. In: *Neuron* 93.5 (2017), pp. 1213–1226.

- [43] Antonio Fernández-Ruiz et al. “Gamma rhythm communication between entorhinal cortex and dentate gyrus neuronal assemblies”. In: *Science* 372.6537 (2021).
- [44] Antonio Fernández-Ruiz et al. “Long-duration hippocampal sharp wave ripples improve memory”. In: *Science* 364.6445 (2019), pp. 1082–1086.
- [45] David J Foster, Richard GM Morris, and Peter Dayan. “A model of hippocampally dependent navigation, using the temporal difference learning rule”. In: *Hippocampus* 10.1 (2000), pp. 1–16.
- [46] David J Foster and Matthew A Wilson. “Hippocampal theta sequences”. In: *Hippocampus* 17.11 (2007), pp. 1093–1099.
- [47] David J Foster and Matthew A Wilson. “Reverse replay of behavioural sequences in hippocampal place cells during the awake state”. In: *Nature* 440.7084 (2006), p. 680.
- [48] Mathias Franzius, Henning Sprekeler, and Laurenz Wiskott. “Slowness and sparseness lead to place, head-direction, and spatial-view cells”. In: *PLoS Comput Biol* 3.8 (2007), e166.
- [49] Gabrielle Girardeau, Anne Cei, and Michaël Zugaro. “Learning-induced plasticity regulates hippocampal sharp wave-ripple drive”. In: *Journal of Neuroscience* 34.15 (2014), pp. 5176–5183.
- [50] Gabrielle Girardeau et al. “Selective suppression of hippocampal ripples impairs spatial memory”. In: *Nature neuroscience* 12.10 (2009), p. 1222.
- [51] Joshua I Glaser et al. “Machine learning for neural decoding”. In: *Eneuro* 7.4 (2020).
- [52] Viviana Gradinaru et al. “Molecular and cellular approaches for diversifying and extending optogenetics”. In: *Cell* 141.1 (2010), pp. 154–165.
- [53] Igor Gridchyn et al. “Assembly-specific disruption of hippocampal replay leads to selective memory deficit”. In: *Neuron* 106.2 (2020), pp. 291–300.
- [54] Anoopum S Gupta et al. “Hippocampal replay is not a simple function of experience”. In: *Neuron* 65.5 (2010), pp. 695–705.
- [55] Anoopum S Gupta et al. “Segmentation of spatial experience by hippocampal theta sequences”. In: *Nature neuroscience* 15.7 (2012), p. 1032.
- [56] Torkel Hafting et al. “Microstructure of a spatial map in the entorhinal cortex”. In: *Nature* 436.7052 (2005), pp. 801–806.
- [57] Moritz Hardt and Benjamin Recht. “Patterns, predictions, and actions: A story about machine learning”. In: *arXiv preprint arXiv:2102.05242* (2021).
- [58] Andrew C Heusser et al. “Episodic sequence memory is supported by a theta–gamma phase code”. In: *Nature neuroscience* 19.10 (2016), pp. 1374–1380.
- [59] Daniel N Hill, Samar B Mehta, and David Kleinfeld. “Quality metrics to accompany spike sorting of extracellular signals”. In: *Journal of Neuroscience* 31.24 (2011), pp. 8699–8705.

- [60] Sepp Hochreiter and Jürgen Schmidhuber. “Long short-term memory”. In: *Neural computation* 9.8 (1997), pp. 1735–1780.
- [61] Jeffrey R Hollerman and Wolfram Schultz. “Dopamine neurons report an error in the temporal prediction of reward during learning”. In: *Nature neuroscience* 1.4 (1998), pp. 304–309.
- [62] David H Hubel and Torsten N Wiesel. “Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex”. In: *The Journal of physiology* 160.1 (1962), pp. 106–154.
- [63] Ahmed T Hussin, Timothy K Leonard, and Kari L Hoffman. “Sharp-wave ripple features in macaques depend on behavioral state and cell-type specific firing”. In: *Hippocampus* 30.1 (2020), pp. 50–59.
- [64] Shantanu P Jadhav et al. “Awake hippocampal sharp-wave ripples support spatial memory”. In: *Science* 336.6087 (2012), pp. 1454–1458.
- [65] Yangqing Jia et al. “Caffe: Convolutional architecture for fast feature embedding”. In: *Proceedings of the 22nd ACM international conference on Multimedia*. 2014, pp. 675–678.
- [66] Eric Jonas and Konrad Paul Kording. “Could a neuroscientist understand a microprocessor?” In: *PLoS computational biology* 13.1 (2017), e1005268.
- [67] Kenneth G Jordan. “Emergency EEG and continuous EEG monitoring in acute ischemic stroke”. In: *Journal of clinical neurophysiology* 21.5 (2004), pp. 341–352.
- [68] Norman P Jouppi et al. “In-datacenter performance analysis of a tensor processing unit”. In: *Proceedings of the 44th annual international symposium on computer architecture*. 2017, pp. 1–12.
- [69] MW Jung and BL McNaughton. “Spatial selectivity of unit activity in the hippocampal granular layer”. In: *Hippocampus* 3.2 (1993), pp. 165–182.
- [70] Clifford Kentros et al. “Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade”. In: *Science* 280.5372 (1998), pp. 2121–2126.
- [71] Diederik P Kingma and Jimmy Ba. “Adam: A method for stochastic optimization”. In: *arXiv preprint arXiv:1412.6980* (2014).
- [72] Kirsten Brun Kjelstrup et al. “Finite scale of spatial representation in the hippocampus”. In: *Science* 321.5885 (2008), pp. 140–143.
- [73] Eric B Knudsen and Joni D Wallis. “Closed-loop theta stimulation in the orbitofrontal cortex prevents reward-based learning”. In: *Neuron* 106.3 (2020), pp. 537–547.
- [74] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. “Imagenet classification with deep convolutional neural networks”. In: *Advances in neural information processing systems* 25 (2012), pp. 1097–1105.

- [75] Anders Krogh and John A Hertz. “A simple weight decay can improve generalization”. In: *Advances in neural information processing systems*. 1992, pp. 950–957.
- [76] Emilio Kropff et al. “Speed cells in the medial entorhinal cortex”. In: *Nature* 523.7561 (2015), pp. 419–424.
- [77] Bálint Lasztóczy and Thomas Klausberger. “Hippocampal place cells couple to three different gamma oscillations during place field traversal”. In: *Neuron* 91.1 (2016), pp. 34–40.
- [78] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. “Deep learning”. In: *nature* 521.7553 (2015), pp. 436–444.
- [79] Albert K Lee and Matthew A Wilson. “Memory of sequential experience in the hippocampus during slow wave sleep”. In: *Neuron* 36.6 (2002), pp. 1183–1194.
- [80] Inah Lee et al. “Comparison of population coherence of place cells in hippocampal subfields CA1 and CA3”. In: *Nature* 430.6998 (2004), pp. 456–459.
- [81] Timothy K Leonard et al. “Sharp wave ripples during visual exploration in the primate hippocampus”. In: *Journal of Neuroscience* 35.44 (2015), pp. 14771–14782.
- [82] Jill K Leutgeb et al. “Pattern separation in the dentate gyrus and CA3 of the hippocampus”. In: *science* 315.5814 (2007), pp. 961–966.
- [83] John Lisman. “The theta/gamma discrete phase code occurring during the hippocampal phase precession may be a more general brain coding scheme”. In: *Hippocampus* 15.7 (2005), pp. 913–922.
- [84] John E Lisman and Ole Jensen. “The theta-gamma neural code”. In: *Neuron* 77.6 (2013), pp. 1002–1016.
- [85] Evgeniy V Lubenov and Athanassios G Siapas. “Hippocampal theta oscillations are travelling waves”. In: *Nature* 459.7246 (2009), pp. 534–539.
- [86] Caitlin S Mallory et al. “Grid scale drives the scale and long-term stability of place maps”. In: *Nature neuroscience* 21.2 (2018), pp. 270–282.
- [87] David Marr and Tomaso Poggio. “From understanding computation to understanding neural circuitry”. In: (1976).
- [88] Marcelo G Mattar and Nathaniel D Daw. “Prioritized memory access explains planning and hippocampal replay”. In: *Nature neuroscience* 21.11 (2018), pp. 1609–1617.
- [89] Andrew P Maurer et al. “Organization of hippocampal cell assemblies based on theta phase precession”. In: *Hippocampus* 16.9 (2006), pp. 785–794.
- [90] R McDonald and N White. “A triple dissociation of memory systems: hippocampus, amygdala, and dorsal striatum.” In: (2013).
- [91] R McDonald et al. “NMDA-receptor blockade by CPP impairs post-training consolidation of a rapidly acquired spatial representation in rat hippocampus”. In: *European Journal of Neuroscience* 22.5 (2005), pp. 1201–1213.

- [92] Robert J McDonald and Norman M White. “Parallel information processing in the water maze: evidence for independent memory systems involving dorsal striatum and hippocampus”. In: *Behavioral and neural biology* 61.3 (1994), pp. 260–270.
- [93] Mayank R Mehta, Carol A Barnes, and Bruce L McNaughton. “Experience-dependent, asymmetric expansion of hippocampal place fields”. In: *Proceedings of the National Academy of Sciences* 94.16 (1997), pp. 8918–8921.
- [94] Mayank R Mehta, Michael C Quirk, and Matthew A Wilson. “Experience-dependent asymmetric shape of hippocampal receptive fields”. In: *Neuron* 25.3 (2000), pp. 707–715.
- [95] Frédéric Michon et al. “Post-learning hippocampal replay selectively reinforces spatial memory for highly rewarded locations”. In: *Current Biology* 29.9 (2019), pp. 1436–1444.
- [96] Steven J Middleton and Thomas J McHugh. “Silencing CA3 disrupts temporal coding in the CA1 ensemble”. In: *Nature neuroscience* 19.7 (2016), pp. 945–951.
- [97] Kenji Mizuseki et al. “Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop”. In: *Neuron* 64.2 (2009), pp. 267–280.
- [98] R Morris. “Synaptic plasticity and learning: selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-D-aspartate receptor antagonist AP5.” In: *The Journal of neuroscience : the official journal of the Society for Neuroscience* 9.9 (1989), pp. 3040–3057. ISSN: 0270-6474. DOI: 10.1523/JNEUROSCI.09-09-03040.1989.
- [99] R Morris et al. “Place navigation impaired in rats with hippocampal lesions”. In: *Nature* 297.5868 (1982), p. 681.
- [100] RG Morris. “Synaptic plasticity and learning: selective impairment of learning rats and blockade of long-term potentiation in vivo by the N-methyl-D-aspartate receptor antagonist AP5”. In: *Journal of Neuroscience* 9.9 (1989), pp. 3040–3057.
- [101] RGM Morris et al. “Selective impairment of learning and blockade of long-term potentiation by an N-methyl-D-aspartate receptor antagonist, AP5”. In: *Nature* 319.6056 (1986), pp. 774–776.
- [102] Richard Morris. “Developments of a water-maze procedure for studying spatial learning in the rat”. In: *Journal of Neuroscience Methods* 11.1 (1984), pp. 47–60. ISSN: 01650270. DOI: 10.1016/0165-0270(84)90007-4.
- [103] Richard GM Morris. “Spatial localization does not require the presence of local cues”. In: *Learning and motivation* 12.2 (1981), pp. 239–260.
- [104] Robert U Muller and John L Kubie. “The effects of changes in the environment on the spatial firing of hippocampal complex-spike cells”. In: *Journal of Neuroscience* 7.7 (1987), pp. 1951–1968.

- [105] Yitzhak Norman et al. “Hippocampal sharp-wave ripples linked to visual episodic recollection in humans”. In: *Science* 365.6454 (2019).
- [106] J O’Keefe. “Place units in the hippocampus of the freely moving rat”. In: *Experimental neurology* 51.1 (1976), pp. 78–109.
- [107] J O’Keefe and J Dostrovsky. “The hippocampus as a spatial map: preliminary evidence from unit activity in the freely-moving rat.” In: *Brain research* (1971).
- [108] John O’Keefe. “Hippocampus, theta, and spatial memory”. In: *Current opinion in neurobiology* 3.6 (1993), pp. 917–924.
- [109] John O’keefe and Lynn Nadel. *The hippocampus as a cognitive map*. Oxford: Clarendon Press, 1978.
- [110] John O’Keefe and Michael L Recce. “Phase relationship between hippocampal place units and the EEG theta rhythm”. In: *Hippocampus* 3.3 (1993), pp. 317–330.
- [111] John O’Keefe et al. “Fornix lesions selectively abolish place learning in the rat”. In: *Experimental neurology* 48.1 (1975), pp. 152–166.
- [112] Azahara Oliva et al. “Origin of gamma frequency power during hippocampal sharp-wave ripples”. In: *Cell reports* 25.7 (2018), pp. 1693–1700.
- [113] Bruno A Olshausen and David J Field. “Sparse coding of sensory inputs”. In: *Current opinion in neurobiology* 14.4 (2004), pp. 481–487.
- [114] Bruno A Olshausen and David J Field. “Sparse coding with an overcomplete basis set: A strategy employed by V1?” In: *Vision research* 37.23 (1997), pp. 3311–3325.
- [115] Marius Pachitariu et al. “Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels”. In: *BioRxiv* (2016), p. 061481.
- [116] M Packard and J McGaugh. “Double dissociation of fornix and caudate nucleus lesions on acquisition of two water maze tasks: further evidence for multiple memory systems.” In: *Behavioral neuroscience* 106.3 (1992), p. 439.
- [117] Adam Paszke et al. “Pytorch: An imperative style, high-performance deep learning library”. In: *arXiv preprint arXiv:1912.01703* (2019).
- [118] Jagdish Patel et al. “Local generation and propagation of ripples along the septotemporal axis of the hippocampus”. In: *Journal of Neuroscience* 33.43 (2013), pp. 17029–17041.
- [119] Jagdish Patel et al. “Traveling theta waves along the entire septotemporal axis of the hippocampus”. In: *Neuron* 75.3 (2012), pp. 410–417.
- [120] F. Pedregosa et al. “Scikit-learn: Machine Learning in Python”. In: *Journal of Machine Learning Research* 12 (2011), pp. 2825–2830.

- [121] H Petsche, Ch Stumpf, and G Gogolak. “The significance of the rabbit’s septum as a relay station between the midbrain and the hippocampus I. The control of hippocampus arousal activity by the septum cells”. In: *Electroencephalography and clinical neurophysiology* 14.2 (1962), pp. 202–211.
- [122] Brad E Pfeiffer and David J Foster. “Autoassociative dynamics in the generation of sequences of hippocampal place cells”. In: *Science* 349.6244 (2015), pp. 180–183.
- [123] Brad E Pfeiffer and David J Foster. “Hippocampal place-cell sequences depict future paths to remembered goals”. In: *Nature* 497.7447 (2013), p. 74.
- [124] Mark H Plitt and Lisa M Giocomo. “Experience-dependent contextual codes in the hippocampus”. In: *Nature Neuroscience* (2021), pp. 1–10.
- [125] Andressa Radiske et al. “Cross-frequency phase-amplitude coupling between hippocampal theta and gamma oscillations during recall destabilizes memory and renders it susceptible to reconsolidation disruption”. In: *Journal of Neuroscience* 40.33 (2020), pp. 6398–6408.
- [126] Blake A Richards et al. “A deep learning framework for neuroscience”. In: *Nature neuroscience* 22.11 (2019), pp. 1761–1770.
- [127] Adina L Roskies. “The binding problem”. In: *Neuron* 24.1 (1999), pp. 7–9.
- [128] Robert R Schaller. “Moore’s law: past, present and future”. In: *IEEE spectrum* 34.6 (1997), pp. 52–59.
- [129] Robson Scheffer-Teixeira and Adriano BL Tort. “On cross-frequency phase-phase coupling between theta and gamma oscillations in the hippocampus”. In: *Elife* 5 (2016), e20515.
- [130] Erik W Schomburg et al. “Theta phase segregation of input-specific gamma patterns in entorhinal-hippocampal networks”. In: *Neuron* 84.2 (2014), pp. 470–485.
- [131] William Beecher Scoville and Brenda Milner. “Loss of recent memory after bilateral hippocampal lesions”. In: *Journal of neurology, neurosurgery, and psychiatry* 20.1 (1957), p. 11.
- [132] Mijail D Serruya et al. “Instant neural control of a movement signal”. In: *Nature* 416.6877 (2002), pp. 141–142.
- [133] Liron Sheintuch et al. “Multiple maps of the same spatial context can stably coexist in the mouse hippocampus”. In: *Current Biology* 30.8 (2020), pp. 1467–1476.
- [134] Justin D Shin and Shantanu P Jadhav. “Multiple modes of hippocampal–prefrontal interactions in memory-guided behavior”. In: *Current opinion in neurobiology* 40 (2016), pp. 161–169.
- [135] Delia Silva, Ting Feng, and David J Foster. “Trajectory events across hippocampal place cells require previous experience”. In: *Nature neuroscience* 18.12 (2015), p. 1772.

- [136] Wolf Singer. “Synchronization of cortical activity and its putative role in information processing and learning”. In: *Annual review of physiology* 55.1 (1993), pp. 349–374.
- [137] William E Skaggs et al. “Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences”. In: *Hippocampus* 6.2 (1996), pp. 149–172.
- [138] Trygve Solstad, Edvard I Moser, and Gaute T Einevoll. “From grid cells to place cells: a mathematical model”. In: *Hippocampus* 16.12 (2006), pp. 1026–1031.
- [139] Trygve Solstad et al. “Representation of geometric borders in the entorhinal cortex”. In: *Science* 322.5909 (2008), pp. 1865–1868.
- [140] Marielena Sosa, Hannah R Joo, and Loren M Frank. “Dorsal and ventral hippocampal sharp-wave ripples activate distinct nucleus accumbens networks”. In: *Neuron* 105.4 (2020), pp. 725–741.
- [141] Nitish Srivastava et al. “Dropout: a simple way to prevent neural networks from overfitting”. In: *The journal of machine learning research* 15.1 (2014), pp. 1929–1958.
- [142] R Steele and R Morris. “Delay-dependent impairment of a matching-to-place task with chronic and intrahippocampal infusion of the NMDA-antagonist D-AP5”. In: *Hippocampus* 9.2 (1999), pp. 118–136.
- [143] Nicholas A Steinmetz et al. “Challenges and opportunities for large-scale electrophysiology with Neuropixels probes”. In: *Current opinion in neurobiology* 50 (2018), pp. 92–100.
- [144] Federico Stella et al. “Hippocampal reactivation of random trajectories resembling Brownian diffusion”. In: *Neuron* 102.2 (2019), pp. 450–461.
- [145] Richard S Sutton. “Integrated architectures for learning, planning, and reacting based on approximating dynamic programming”. In: *Machine learning proceedings 1990*. Elsevier, 1990, pp. 216–224.
- [146] Makoto Tamura et al. “Hippocampal-prefrontal theta-gamma coupling during performance of a spatial working memory task”. In: *Nature communications* 8.1 (2017), pp. 1–9.
- [147] Jeffrey S Taube. “Head direction cells and the neurophysiological basis for a sense of direction”. In: *Progress in neurobiology* 55.3 (1998), pp. 225–256.
- [148] Edward C Tolman. “Cognitive maps in rats and men.” In: *Psychological review* 55.4 (1948), p. 189.
- [149] Adriano BL Tort et al. “Theta-gamma coupling increases during the learning of item-context associations”. In: *Proceedings of the National Academy of Sciences* 106.49 (2009), pp. 20942–20947.
- [150] Alexandros T Tzallas, Markos G Tsipouras, and Dimitrios I Fotiadis. “Epileptic seizure detection in EEGs using time-frequency analysis”. In: *IEEE transactions on information technology in biomedicine* 13.5 (2009), pp. 703–710.

- [151] Daniel Valencia and Amirhossein Alimohammad. “A real-time spike sorting system using parallel OSort clustering”. In: *IEEE transactions on biomedical circuits and systems* 13.6 (2019), pp. 1700–1713.
- [152] Case H Vanderwolf. “Hippocampal electrical activity and voluntary movement in the rat”. In: *Electroencephalography and clinical neurophysiology* 26.4 (1969), pp. 407–418.
- [153] Alex P Vaz et al. “Replay of cortical spiking sequences during human memory retrieval”. In: *Science* 367.6482 (2020), pp. 1131–1134.
- [154] Mengni Wang, David J Foster, and Brad E Pfeiffer. “Alternating sequences of future and past behavior encoded within hippocampal theta oscillations”. In: *Science* 370.6513 (2020), pp. 247–250.
- [155] Tianshi Wang and Jaijeet Roychowdhury. “OIM: Oscillator-based Ising machines for solving combinatorial optimisation problems”. In: *International Conference on Unconventional Computation and Natural Computation*. Springer. 2019, pp. 232–256.
- [156] Tianshi Wang and Jaijeet Roychowdhury. “PHLOGON: Phase-based logic using oscillatory nano-systems”. In: *International Conference on Unconventional Computation and Natural Computation*. Springer. 2014, pp. 353–366.
- [157] Xiao-Jing Wang. “Neurophysiological and computational principles of cortical rhythms in cognition”. In: *Physiological reviews* 90.3 (2010), pp. 1195–1268.
- [158] John Broadus Watson and Karl Spencer Lashley. *Homing and related activities of birds*. Vol. 7. Carnegie institution of Washington, 1915.
- [159] John Widloski and Ila R Fiete. “A model of grid cell development through spatial exploration and spike time-dependent plasticity”. In: *Neuron* 83.2 (2014), pp. 481–495.
- [160] Andrew M Wikenheiser and A David Redish. “Hippocampal theta sequences reflect current goals”. In: *Nature neuroscience* 18.2 (2015), p. 289.
- [161] Matthew A Wilson and Bruce L McNaughton. “Reactivation of hippocampal ensemble memories during sleep”. In: *Science* 265.5172 (1994), pp. 676–679.
- [162] Laurenz Wiskott and Terrence J Sejnowski. “Slow feature analysis: Unsupervised learning of invariances”. In: *Neural computation* 14.4 (2002), pp. 715–770.
- [163] Xiaojing Wu and David J Foster. “Hippocampal replay captures the unique topological structure of a novel environment”. In: *Journal of Neuroscience* 34.19 (2014), pp. 6459–6469.
- [164] Haibing Xu et al. “Assembly responses of hippocampal CA1 place cells predict learned behavior in goal-directed spatial tasks on the radial eight-arm maze”. In: *Neuron* 101.1 (2019), pp. 119–132.

- [165] Michael M Yartsev, Menno P Witter, and Nachum Ulanovsky. “Grid cells without theta oscillations in the entorhinal cortex of bats”. In: *Nature* 479.7371 (2011), pp. 103–107.
- [166] Ofer Yizhar et al. “Optogenetics in neural systems”. In: *Neuron* 71.1 (2011), pp. 9–34.
- [167] Aarne Ylinen et al. “Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms”. In: *Journal of Neuroscience* 15.1 (1995), pp. 30–46.
- [168] Honghui Zhang and Joshua Jacobs. “Traveling theta waves in the human hippocampus”. In: *Journal of Neuroscience* 35.36 (2015), pp. 12477–12487.
- [169] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. “BIRCH: A new data clustering algorithm and its applications”. In: *Data Mining and Knowledge Discovery* 1.2 (1997), pp. 141–182.
- [170] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. “BIRCH: an efficient data clustering method for very large databases”. In: *ACM Sigmod Record* 25.2 (1996), pp. 103–114.
- [171] Chenguang Zheng et al. “Spatial sequence coding differs during slow and fast gamma rhythms in the hippocampus”. In: *Neuron* 89.2 (2016), pp. 398–408.
- [172] Mark C Zielinski, Justin D Shin, and Shantanu P Jadhav. “Coherent coding of spatial position mediated by theta oscillations in the hippocampus and prefrontal cortex”. In: *Journal of Neuroscience* 39.23 (2019), pp. 4550–4565.
- [173] Mark C Zielinski, Wenbo Tang, and Shantanu P Jadhav. “The role of replay and theta sequences in mediating hippocampal-prefrontal interactions for memory and cognition”. In: *Hippocampus* 30.1 (2020), pp. 60–72.
- [174] Stuart Zola-Morgan, Larry R Squire, and David G Amaral. “Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus”. In: *Journal of Neuroscience* 6.10 (1986), pp. 2950–2967.
- [175] Stuart M Zola-Morgan and Larry R Squire. “The primate hippocampal formation: evidence for a time-limited role in memory storage”. In: *Science* 250.4978 (1990), pp. 288–290.

Acronyms

KL	Kullback–Leible
c.i.	confidence interval
RL	Reinforcement Learning
TD	Temporal-Difference
AP	Anterior-Posterior
AP5	AnimoPhosponovaleric Acid
ML	Medial-Lateral
NAc	Nucleus Accumbens
PC	Principal Component
PCs	Principal Components
EC	Entorhinal Cortex
DG	Dentate Gyrus
DV	Dorsal-Ventral
NN	Neural Network
SR	Stratum Radiatum
WC	Weighted Correlation
CoM	Center of Mass
BMI	Brain Machine Interfaces
sos	Second-Order Sections
MAE	Mean Absolute Error

CPP	(6)-3-(2-Carboxypiperazin-4-yl) Propyl-1-Phosponic Acid
API	Application Processing Interface
REM	Rapid Eye Movements
RPE	Reward Prediction Error
GFP	Green Fluorescent Protein
EEG	Electroencephelogram
PFR	Phase-Frequency Response
CSD	Current Source Density
VHC	Ventral Hippocampal Commisure
LFP	Local Field Potential
MRV	Mean Resultant Vector
SWR	Sharp-Wave Ripples
SWS	Slow-Wave Sleep
ANN	Artificial Neural Network
ITI	Inter-Trial Interval
IEI	Inter-Event Interval
MEC	Medial Entorhinal Cortex
LEC	Lateral Entorhinal Cortex
LTP	Long-Term Potentiation
mPFC	medial PreFrontal Cortex
ReLU	Rectified Linear Unit
ACSF	Artificial CerebroSpinal Fluid
NMDA	N-Methyl-D-Aspartate
ANOVA	ANalysis Of VAriance
NMDAR	N-Methyl-D-Aspartate Receptor
BIRCH	Balanced Iterative Reducing and Clustering using Hierarchies
FIFO	First-In First-Out