

UNIVERSITY OF CALIFORNIA

Los Angeles

Transmission of Hippocampal Sharp Wave Ripples to Subcortical Brain Regions

A dissertation submitted in partial fulfillment of the requirements for the degree

Doctor of Philosophy in Neuroscience

by

Andrew Gregory Howe

2020

© Copyright

Andrew Gregory Howe

2020

ABSTRACT OF THE DISSERTATION

Transmission of Hippocampal Sharp Wave Ripples to Subcortical Brain Regions

By

Andrew Gregory Howe

Doctor of Philosophy in Neuroscience

University of California, Los Angeles, 2020

Professor Hugh T. Blair, Chair

The Sharp-Wave Ripple (SWR) is a signature electrophysiological phenomenon in the hippocampus, characterized by brief bursts of coordinated neural activity. SWR events are hypothesized to support learning, decision making and memory formation, but their exact role in these processes remains unclear. Since these higher level cognitive functions involve interactions between multiple brain regions, it is important to better understand how SWR events are relayed from the hippocampus to other brain structures.

Hippocampal outputs to cortical regions are routed through CA1 and the subiculum, while outputs to subcortical regions are routed through the lateral septum (LS), which in turn projects to the midbrain, striatum, and other subcortical regions involved in reward processing signaled by dopamine. In this dissertation, I report results of experiments investigating how hippocampal SWRs influence single cell activity in LS and striatum, and analyze what types of information are represented by single cells that are modulated by hippocampal SWRs in the LS and striatum (Chapter 2). I also report findings from experiments in which I attempted to perform simultaneous high-temporal resolution monitoring of dopamine flux, single unit activity and LFP

with a combination of fast-scan cyclic voltammetry(FSCV) and electrophysiology in the behaving rat (Chapter 3).

Hungry rats repeatedly performed an acquisition-and-reversal task for food rewards on a t-maze, while chronic implants recorded SWR events in the hippocampus and single-unit spike activity in LS and striatum. During periods of motor inactivity, SWRs triggered excitatory responses from 28% (64/226) and inhibitory responses from 14% (31/226) of septal neurons. By contrast, only 4% (14/378) of striatal neurons were excited and 6% (24/378) were inhibited during SWRs. In both structures, neurons which reduced firing during SWR exhibited greater spike coherence with hippocampal theta rhythm than neurons that did not respond to SWRs. In septum, neurons that were excited by SWRs fired at late phases of the theta cycle, whereas neurons that were inhibited by SWRs fired at early phases of the theta cycle. By contrast, SWR-responsive striatal neurons did not show consistent phase preferences during the theta cycle. A subset of SWR-responsive neurons in septum (55/95) and striatum (26/38) behaved as speed cells, with firing rates that were positively or negatively modulated by the rat's running speed. In both structures, firing rates of most SWR- excited speed cells were positively modulated by running speed, whereas firing rates of most SWR-inhibited speed cells were negatively modulated by running speed.

These findings are consistent with a growing body of evidence that SWRs might activate subcortical representations of motor actions in conjunction with hippocampal representations of places and states, which may be important for storing and retrieving values of state-action pairs during reinforcement learning and memory consolidation. Chapter 4 discusses implications of these results for our understanding of how SWR events contribute to learning, memory, and decision making. Further studies of SWR-evoked responses in subcortical reward circuits may build on the current findings to deepen our understanding of how SWR events contribute to cognitive functions.

The dissertation of Andrew Gregory Howe is approved,

Kate M. Wassum

Sotiris C. Masmanidis

Dean V. Buonomano

Hugh T. Blair, Committee Chair

University of California, Los Angeles

2020

TABLE OF CONTENTS

ABSTRACT OF THE DISSERTATION	ii
TABLE OF CONTENTS	v
LIST OF FIGURES	x
LIST OF TABLES	xi
LIST OF EQUATIONS	xi
LIST OF ACRONYMS & ABBREVIATIONS	xii
ACKNOWLEDGEMENTS	xiii
Dedication	xiii
Funding Sources	xiv
Appreciation	xv
BIOGRAPHICAL SKETCH	xviii
CHAPTER 1	1
1.1 INTRODUCTION	2
1.2 SPATIAL CODING BY HIPPOCAMPAL NEURONS	3
1.2.1 Positional Tuning of Place Cells	4
1.2.2 Anatomical Subregions of the Hippocampus	6
1.2.2.1 Dentate Gyrus: Sparse Coding and Pattern Separation	8
1.2.2.2 CA3 Region: Auto-Associativity and Pattern Completion	14
1.2.2.3 CA1 Region: Hippocampal Output to Cortex	18
1.3 PROCESSING STATES OF THE HIPPOCAMPUS	21
1.3.1 The Theta State	21
1.3.1.1 Phase Precession	22
1.3.1.2 Theta Sequences	25
1.3.2 Large, Irregular Activity (LIA) State	28
1.3.2.1 Sharp Waves	28
1.3.2.2 Ripples	29
1.3.2.3 Compressed Replay	30
1.4 FUNCTIONS OF SHARP-WAVE RIPPLES & REPLAYS	31
1.3.1 Memory Consolidation	33
1.3.2 Credit Assignment	35
1.3.3 Model-Based Decision-Making	38

CHAPTER 2	41
2.1 INTRODUCTION	42
2.2 RESULTS	44
2.2.1 Single Unit Responses during SWR Events	47
2.2.1.1 Excitatory and Inhibitory Responses	50
2.2.1.1.1 Hippocampal Units	52
2.2.1.1.2 Septum Units	53
2.2.1.1.3 Striatal Units	55
2.2.1.2 SWR Responses in the Bucket versus on the Maze	55
2.2.1.2.1 Hippocampal Units	56
2.2.1.2.2 Septum Units	57
2.2.1.2.3 Striatal Units	59
2.2.2 Predictors of Ripple Responsiveness	61
2.2.2.1 Coherence with Hippocampal Theta Rhythm	61
2.2.2.1.1 Hippocampus	63
2.2.2.1.2 Lateral Septum	68
2.2.2.1.3 Striatum	70
2.2.2.2 Speed Sensitivity	72
2.2.2.2.1 Hippocampus	72
2.2.2.2.2 Lateral Septum	73
2.2.2.2.3 Striatum	76
2.3. DISCUSSION	77
2.3.1 Activity of Septal Neurons during SWR Events	79
2.3.2 Activity of Striatal Neurons during SWR Events	81
2.3.3 Summary and Conclusions	83
2.4. METHODS	84
2.4.1 Subjects and Behavior	85
2.4.1.1 Subjects	85
2.4.1.2 Behavior Apparatus	85
2.4.1.3 T-Maze Task	86
2.4.1.4 Video Tracking	87
2.4.2 Surgery, Electrophysiology, and Histology	87
2.4.2.1 Surgery	87

2.4.2.2 Placement of LFP Electrodes	88
2.4.2.3 Recording Sessions and Tetrode Advancement	88
2.4.2.4 Histology	89
2.4.3 Data analysis	90
2.4.3.1 Spike Sorting	90
2.4.3.2 LFP Filtering and Analysis	90
2.4.3.3 Response Latency	91
2.4.3.4 Single Unit Responses During SWRs	91
2.4.3.5 Speed Analysis	92
CHAPTER 3	93
3.1 INTRODUCTION	94
3.2 DOPAMINE SIGNALING AND REWARD-MOTIVATED BEHAVIOR	96
3.2.1 The Midbrain Dopamine System	96
3.2.1.1 Dopamine and Hedonistic Pleasure	97
3.2.1.2 Dopamine and Incentive Motivation	98
3.2.1.3 Dopamine and Reward Prediction Error	99
3.2.2 Hippocampal Interactions with the Dopamine System	102
3.2.2.1 Hippocampal Outputs to the Dopamine System	103
3.2.2.2 Dopaminergic Inputs to the Hippocampus	105
3.2.2.3 Effects of Hippocampal Dopaminergic Manipulations on Behavior	107
3.3 HIPPOCAMPUS AND DECISION MAKING: A HYPOTHESIS	108
3.3.1 Theta State: M+ and M- Cells	109
3.3.1.1 Motor-Related Activity of LS Neurons	111
3.3.1.2 Motor-Related Activity of Striatal Neurons	112
3.3.2 LIA State: SWR-Evoked Responses	113
3.3.2.1 Model-Based Decision Making	114
3.3.2.2 Credit Assignment	115
3.3.2.3 Long Term Memory Consolidation	116
3.3.3 Striatal Involvement in Goal-Directed & Habitual Behaviors	117
3.4 SUMMARY AND CONCLUSIONS	119
3.4.1 Value-Based Decision-Making Outside the Lab	119
CHAPTER 4	122
4.1 INTRODUCTION & BACKGROUND	123

4.1.1 Motivation	124
4.1.2 Connecting Place to Reward	125
4.1.3 Experimental Plan	126
4.2 RESULTS	127
4.2.1 Histology	128
4.2.2 The FSCV Probe Occupied a Dopamine-Rich Region	129
4.2.3 Dopamine Recorded Simultaneously with Electrophysiology	130
4.2.4 Behavior Performance	132
4.2.5 Behavioral Correlates with Dopamine Flux	134
4.2.6 Electrophysiology Recorded Simultaneously with Fast-Scan Cyclic Voltammetry	138
4.2.7 Speed & Acceleration Correlations with Cell Firing	141
4.2.8 No Theta Synchronization	143
4.2.9 No Head Direction Preference	143
4.2.10 Combined Single-Unit Electrophysiology with FSCV	144
4.3 METHODS	145
4.3.1 Subjects	145
4.3.2 Targeting Considerations	146
4.3.2 Fast-Scan Cyclic Voltammetry Probes for the OvalDrive36 Headcap	147
4.3.2.1 Traditional FSCV Probes	148
4.3.2.2 OvalDrive FSCV Probes v1.0	149
4.3.2.3 OvalDrive FSCV Probes v2.0	150
4.3.2.4 Ag-AgCl Reference	153
4.3.3 Testing and Calibration	153
4.3.4 Additional FSCV Equipment & Supplies	154
4.3.5 Data Analysis for FSCV In Vivo	155
4.3.6 Electrophysiology	155
4.3.7 Resolving Electrical Noise	156
4.3.8 Behavior	158
4.3.8.1 The Figure 8 Maze	158
4.3.8.2 The Open Platform	161
4.3.9 System Synchronization	162
4.3.9 FSCV Voltage Sweep Artifact Removal	163
4.3.10 Speed & Acceleration Analysis	165

4.3.11 Theta Phase Analysis	166
4.3.12 Spatial Information	166
4.4 DISCUSSION	167
4.4.1 Comparison with Alternative Designs	168
4.4.2 Technical Recommendations for Future Combined FSCV Experiments	169
4.4.2.1 Experimental Design Considerations	169
4.4.2.2 Future Technical Developments	169
4.4.2.3 FSCV Voltage Pulse Cancellation	170
4.4.2.3 Alternatives to Electrophysiology	170
4.4.2.4 Alternatives to FSCV	170
4.4.2.5 Multichannel FSCV	171
4.4.2.6 Electrophysiology Data Collection Recommendations	171
WORKS CITED	172

LIST OF FIGURES

CHAPTER 1

Figure 1-1	: Example Place Cells	5
Figure 1-2	: Hippocampal subregions in the rodent	6
Figure 1-3	: Basic Hippocampal Circuit	7
Figure 1-4	: Sparse versus distributed coding	9
Figure 1-5	: Coding capacity as a function of percentage of active neurons	11
Figure 1-6	: Pattern separation schematic	12
Figure 1-7	: Autoassociative pattern completion	18
Figure 1-8	: Theta rhythm recorded in hippocampal CA1 region	22
Figure 1-9	: Schematic diagram of theta phase precession	23
Figure 1-10	: Theta phase precession examples	24
Figure 1-11	: Example of phase precession in a LS unit	27
Figure 1-12	: SWR electrophysiology schematic	29
Figure 1-13	: SWR Replay examples	32
Figure 1-14	: CA3-LS-VTA circuit schematic	37

CHAPTER 2

Figure 2-1	: Behavioral and neurophysiological data samples	46
Figure 2-2	: Example data from a single recording session	49
Figure 2-3	: Unit responses to SWR events	51
Figure 2-4	: Example cells recorded in septum	58
Figure 2-5	: Example cells recorded in striatum	60
Figure 2-6	: Theta coherence and SWR responsiveness	64
Figure 2-7	: Preferred theta phase and SWR responsiveness	65
Figure 2-8	: Speed modulation and SWR responsiveness	74

CHAPTER 3

Figure 3-1	: Dopamine and TD error	100
Figure 3-2	: Theta sequences, SWR responsive cells & speed tuning	110
Figure 3-3	: Real-world decision-making	120

CHAPTER 4

Figure 4-1	: Model-based decision making schematic	123
Figure 4-2	: Hippocampus CA3 -- LS -- VTA disinhibitory circuit	126
Figure 4-3	: Histology revealed recording sites in rat V4	128
Figure 4-4	: The FSCV Probe Occupied a Dopamine-Rich Region	129
Figure 4-5	: Representative dopamine data	131
Figure 4-6	: The rat is new to the platform task and expert at the Figure 8 task	132
Figure 4-7	: Average speed & acceleration by position during behavior	133
Figure 4-8	: Behavior event triggered dopamine	135
Figure 4-9	: Representative traces of dopamine and telemetry data	136
Figure 4-10	: Correlation of DA flux with spatial data	138
Figure 4-11	: The FSIEA artifact interferes with LFP, but can be cleanly removed	140
Figure 4-12	: The average spike waveform for the single unit in the LH	141
Figure 4-13	: The LH cell's firing rate correlates with the speed and velocity	142
Figure 4-14	: The channel 7 cell is not theta synchronized	143
Figure 4-15	: The LH cell is not tuned for head direction	144
Figure 4-16	: Spike triggered average of the dopamine trace	144
Figure 4-17	: Probes targeting schematic Nacc, CA1 and VTA	147
Figure 4-18	: FSCV through the OvalDrive required modifications to the classic	148
Figure 4-19	: FSCV probes modified for the OvalDrive36 function well in vitro	154
Figure 4-20	: The Faraday cage reduces noise	157
Figure 4-21	: Figure 8 Maze & Behavior	160
Figure 4-22	: Platter Maze & Behavior	162
Figure 4-23	: The FSIEA removal algorithm operating on pre-filtered data	165

LIST OF TABLES

Table 1-1	: SWR Activates Cortex	35
Table 3-1	: A partial list of the roles of LS	105

LIST OF EQUATIONS

Equation 1-1	: N-choose-K	10
Equation 2-1	: Speed Calculation	92
Equation 3-1	: The Temporal Difference Error Signal	100
Equation 4-1	: Spatial Information	166

LIST OF ACRONYMS & ABBREVIATIONS

ACC	-- anterior cingulate cortex
AC	-- auditory cortex
CA1	-- cornu ammonis 1
CA3	-- cornu ammonis 3
CS	-- conditional stimulus
DA	-- dopamine
DG	-- dentate gyrus
EEG	-- electroencephalogram
EC	-- entorhinal cortex
FSCV	-- fast scan cyclic voltammetry
LH	-- lateral hypothalamus
LIA	-- Large Irregular Activity
LFP	-- local field potential
LS	-- lateral septum
Nacc	-- nucleus accumbens (ventral striatum)
NR	-- non-responsive
PFC	-- prefrontal cortex
PPC	-- posterior parietal cortex
RMTg	-- rostromedial tegmental nucleus
RPE	-- reward prediction error
RSC	-- retrosplenial cortex
Str	-- striatum
SWR	-- sharp wave ripple
SSC	-- somatosensory cortex
TD	-- temporal difference
TT	-- tetrode
US	-- unconditional stimulus
VBRL	-- value based reinforcement learning
VC	-- visual cortex
VTA	-- ventral tegmental area

Abbreviations of standard anatomical names specify brain structures & their subregions, e.g. :

dIStr	-- dorsolateral striatum	dmStr	-- dorsomedial striatum
vStr	-- ventral striatum	NAcc	-- nucleus accumbens

Standard international scientific units and their abbreviations are employed, e.g. :

cm	-- centimeters	nM	-- nanomolar
mV	-- millivolts	Hz	-- Hertz (cycles per second)

ACKNOWLEDGEMENTS

Dedication

I dedicate this work to everyone who supported my interest and education in science and technology over the years.

Especially,

Ollie Douglas Howe-Estrada -- My son Ollie reminds me of the marvels of neurodevelopment and the awesome capabilities of the brain every day. Children seem to naturally be scientists!

Ana Maria Estrada Sanchez, PhD -- Ana is my love, spouse, partner, friend, collaborator, climbing companion, fellow adventurer and much more. Ana provided tremendous support, inspiration and advice throughout my graduate program and dissertation.

Allen Gregory Howe & Anne Bruce Howe – My parents provided a fun, educational childhood full of learning opportunities, supported my undergraduate education and continue to provide sage advice.

Grandmother Betty Howe – My Grandmother Betty was the first recipient of a doctorate in our family. She always reminded me to dream big, persevere, to remain voraciously curious and to enjoy fine dining.

Grandfather & Grandmother Talcott – My Grandparents Talcott always fostered scholarship, cultural exploration, and lively debates. They maintained an incredible library and generously supported my education throughout my life.

Funding Sources

NSF : NeuroNex Technology Hub (8091 . 7275)	2017 -- 2019
UCLA Depression Grand Challenge	2017 -- 2018
Neural Microcircuits Training Grant (T32-NS058280)	2016 – 2018
UCLA BRI Graduate Travel Award	2015, 2016
Blair Laboratory Graduate Student Researcher IX (MH079511)	2014 – 2016
UCLA BioSciences Council	2014 – 2016
UCLA NSIDP Graduate Fellowships	2013, 2014

Appreciation

Thank You for your help during my doctorate

Doctoral Committee

H. Tad Blair, PhD
Dean Buonomono, PhD
Sotiris Masmanidis, PhD
Kate Wassum, PhD

UCLA Professors

Avishek Adhikari, PhD
Robert M. Bilder, PhD
Aaron Blaisdel, PhD
Jason Cong, PhD
Michael Fanselow, PhD
Jack Feldman, PhD
Bill Grisham, PhD
David Hovda, PhD
Alicia Izquierdo, PhD
Pam Kennedy, PhD
Wentai Liu, PhD
Mayank R. Mehta, PhD
Stott D. Parker, PhD
Fred Sabb, PhD
Nicholas Wisniewski, PhD

Wassum Lab

Tara Aitken
Anne Collins, MA, PhD
Venuz Greenfield, MS
Nina Lichtenberg, MA, PhD
Ash Moore, PhD

Izquierdo Lab

Evan Hart, PhD

Kennedy Lab

Erik Harvey

Blair Lab

Sage Ben-David
Garret Blair, MA
Zhe Chen, PHD
Ryan Grgurich
Rose de Guzman, PhD
Austin Davis-Hunter
(Olufemi) Taiwo Dodo-Williams
Hanna Doh
Ryan King
Chris Liu
Amber Longo
Brian Odengardt
Shushanna Samvelian
Erik Soveriegn
Tammi Tse
Sandy Vu
Danielle Wenger

UCLA Classmates

Caitlin Aamodt
Konstantin Bakhurin, PhD
Hua Chai, MD, PhD
Stanislav Culaclii, PhD
Allison Furterer
Jill Haney
Cynthia He, MD, PhD
Lyle Kingsbury
Jannelle Liu, PhD
Jason Moore, PhD
Dan Nachun, PhD
Colin Terndrup, MS
Jennifer Tribble, PhD
Courtney Yeager, PhD

Fanselow Lab

Sarah Hersman, PhD
Katherine Fang
Zach Pennington, PhD

Evans Lab

Sadaf Mehrabani
Anna Taylor, PhD

Levine Lab

Mymy Huynh, MSN
Ana Parievsky, PhD

UVa

William B Levy, PhD

Rebec Lab (Indiana U)

Kendra Bunner, PhD

NSIDP Staff

Debra Kozal
Jennifer Lee
Mike Levine, PhD
Marrisa Ray
Felix Schweizer, PhD
Melissa Sherlock

Psychology Staff

Mahea Ayoso-Sadsad
Frank Falcon II
Areli Lucatero
Kevin Nygen
Cheyrl Polfus
Rae Anne Robinett
Tyler Tuionne
The business office staff
The finance office staff

Neural Microcircuits Training Grant

Cortney Wisneski

DLAM Staff

Numan Interiano
Psychology animal staff

Family

Ana M. Estrada Sanchez, PhD
Allen G. Howe
Anne B. Howe
Ollie Douglas Howe-Estrada

Special thanks to the following people :

My advisor, Professor Blair, for his enthusiastic support and involvement with all of my projects.

Professor Wassum for her invaluable contributions to the FSCV-electrophysiology project, for generously opening her lab, and for wonderful hand-me-down baby gear.

Professor Bilder, Professor Parker and Professor Sabb for their guidance and support of my research efforts prior to my doctorate with the Consortium for Neuropsychiatric Phenomics.

Professor Levy for teaching me a great deal about academic research, theoretical neuroscience and supporting my nascent research career for years.

Professor Kennedy and Professor Izquierdo for their advice, and lending me key equipment.

Dr. Anne Monks, the personnel in her cancer research lab and the Fort Detrick NCI community for introducing me to research science as a junior in high school.

My science instructors from middle and high school. Mr. Buckley was the first person to encourage me to seek a doctorate, as well as a true master of maintaining the attention of 8th graders with amazing demonstrations of scientific principles. Mr. Wilson's humble, yet seemingly effortless, ability to connect calculus and physics. Mrs. Ferrant, was a real-life Mrs. Frizzle, teaching us fun songs as mnemonic memory devices. Finally, my 7th grade science teacher, who taught us to sanitarily dissect edible, fresh, raw squid and then fried them for us to eat.

BIOGRAPHICAL SKETCH

Education

University of Virginia (UVA) May 2005
BA : Cognitive Science
Specialization & Certificate : Computer Science
Study Abroad : Pontificia U. Católica del Perú

Peer-Reviewed Journal Articles

Howe AG, Blair HT. Excitatory and inhibitory modulation of septal and striatal neurons during hippocampal sharp-wave ripple events. (submitted to *Hippocampus* June 2020)

Estrada-Sánchez AM, Blake CL, Barton SJ, **Howe AG**, Rebec GV. Lack of mutant huntingtin in cortical efferents improves behavioral inflexibility and corticostriatal dynamics in Huntington's disease mice. *J Neurophysiol.* 2019;122(6):2621-2629.

Chen Z, **Howe A**, Blair HT, Cong J. CLINK: Compact LSTM Inference Kernel for Energy Efficient Neurofeedback Devices. ISLPED '18 Proc Int Symp Low Pwr Elect and Desn. 2018 Jul; Article No. 2:1-6

Bilder RM, **Howe AG**, Sabb FW. Multilevel models from biology to psychology: mission impossible? *J Abnorm Psychol.* 2013;122(3):917-927.

Bilder RM, **Howe A**, Novak N, Sabb FW, Parker DS. The genetics of cognitive impairment in schizophrenia: a phenomic perspective. *Trends Cogn Sci.* 2011;15(9):428-435.

Levy WB, Chang KS & **Howe AG**. Progressively introducing quantified biological complexity into a hippocampal CA3 model that learns trace conditioning. *IJCNN '09 Int Joint Conf Neural Networks.* 2009; Jun:1777-1783.

Howe AG, Levy WB. A hippocampal model predicts a fluctuating phase transition when learning certain trace conditioning paradigms. *Cogn Neurodyn.* 2007;1(2):143-155.

Books

Levy WB, **Howe AG**. Supercomputing Simplified : The Bare Essentials for Parallel C Programming with MPI. Charlottesville, VA: Informed Simplifications; 2008.

Research Experience (selected)

Blair Lab, UCLA	Graduate Student Researcher	2014 -- 2020
Cons. Neuropsych. Phen., UCLA	Programmer/Analyst II	2009 -- 2011
Levy Lab, UVA	Programmer/Analyst I	2004 -- 2009

Professional Experience (selected)

Vetstreet.com	Sr. Software Engineer	2011 -- 2013
Informed Simplications, LLC	Co-Founder & Chief Research Officer	2005 -- 2009

Awards & Honors

Neural Microcircuits Training Grant	2016 -- 2018
Semel Inst. Neurosci. Grad. Travel Award	2015 -- 2016
High Performance Promotions, Vetstreet.com (3x)	2012 -- 2013

Presentations & Posters (Selected)

Blair G, **Howe AG**, Golshani P, Blair HT. Imaging long-term population dynamics of rat hippocampal place cells. Poster 604.05 : Soc Neurosci Ann Mtg; Nov 6, 2018; San Diego, CA.

Howe AG, Blair GJ, Blair HT. Hippocampal sharp-wave ripples precede high-effort movements. Poster 423.16 : Soc Neurosci Ann Mtg; Nov 5, 2018; San Diego, CA.

Howe AG, Blair HT. Sharp-wave ripples in a spatial, value-based decision-making task. Poster 303.6.13 : UCI Int Conf Learning & Memory; Apr 18 – 22, 2018; Huntington Beach, CA.

Howe AG. Neural correlates of spatial decision making. Oral presentation at: Learning & Behav Brown Bag; Feb 2, 2018; Los Angeles, CA, USA.

Howe AG. Sharp wave ripples and dopamine in a value-based spatial decision task. Oral presentation at: Neurosci Res Forum; Mar 2, 2017; Los Angeles, CA, USA.

Howe AG, De Guzman RM, Blair GJ, Blair HT. “Place cells in the septohippocampal nucleus of freely behaving rats.” Poster 356.26 : Soc Neurosci Ann Mtg; Nov 14, 2016; San Diego, CA.

Blair G, **Howe AG**, Aharoni D, Flores S, Shuman T, et al.,. Calcium imaging of hippocampal cell activity in behaving rats. Poster 560.04 : Soc Neurosci Ann Mtg; Nov 15, 2016; San Diego, CA.

Howe AG, Wassum KM, Blair HT. Simultaneous in-vivo single-unit recording and fast-scan cyclic voltammetry in the behaving rat.” Poster 267.16 : Soc Neurosci Ann Mtg; Nov 18, 2015; Chicago, IL.

Bilder RM, Parker DS, Sabb FW, **Howe A**. Evidence based science: Is working memory really impaired in schizophrenia? Poster. Abstract in : Biol Psych Suppl. 2011; May;69(9):278S.

Wang Y, Liu C, Chu WW, Sabb F, Parker DS & **Howe A**. PhenoMining : Text Mining Tools for CNP. Poster presented at: Consort Neuropsych Phenomics Retreat; 2010; Los Angeles, CA.

CHAPTER 1

A Brief Overview of Hippocampal Processing

1.1 INTRODUCTION

The hippocampus is among the most widely studied structures of the mammalian brain. As of July 2020, Pubmed lists over 160,000 publications for the search string “hippocampus,” dating as far back as 1888. And yet, much remains to be learned about this enigmatic structure.

Santiago Ramón y Cajal (1852-1934) -- widely regarded as the first modern neuroscientist -- took a special interest in the hippocampus in his anatomical drawings because of its highly organized circuitry. A major milestone in our understanding of the function of this circuitry arrived in 1957, when the patient H.M. underwent bilateral surgical resection of his medial temporal lobes to alleviate severe epileptic seizures (Scoville & Milner, 1957). In the process, both hippocampal lobes were ablated (in addition to other parts of the temporal lobe), leading to a range of symptoms that began to shed light on the role of the hippocampus in memory. Subsequent studies in human patients with hippocampal lesions (rev. Squire & Zola-Morgan, 2010) and experiments with animal models (Mizumori et al., 1989; Mishkin, 1978; rev. Wang & Morris, 2010; Tayler et al., 2013) yielded abundant evidence that the hippocampus is involved in storing and retrieving specific kinds of memories.

One type of memory that is now known to depend upon the hippocampus is *episodic memory* for events that occur at a specific time and place. In many ways, our episodic memories define us as individuals, since our sense of self is deeply rooted in memories of our life history. Episodic memories also play an important role in how we interact with the world, because some of the behavioral choices that we make are based upon episodic memories of the outcomes from similar choices that we have made in the past. Given that the hippocampus performs memory functions that are so central to who we are and how we behave, it is no surprise that the hippocampus remains a major focus of current research in modern neuroscience.

During certain behavioral states, the hippocampus generates distinctive bursts of neural activity known as *sharp-wave ripple (SWR) events*. A major objective of my thesis research has been to perform neurophysiological recording experiments in rats, to investigate how SWR events might contribute to the memory functions of the hippocampus. Before describing my research, this introductory chapter will review background information about hippocampal anatomy and neurophysiology that was key in motivating my studies.

1.2 SPATIAL CODING BY HIPPOCAMPAL NEURONS

Rodent models have played an outsized role in studies of hippocampal function. In 1971, while recording from neurons in the CA1 subregion of the rat hippocampus, O'Keefe and Dostrovsky observed "place cells" that reliably fired whenever the rat visited a specific location in a spatial environment. O'Keefe and Nadel (1976) subsequently hypothesized that these hippocampal place cells may form a "cognitive map" that stores geometrically accurate memory representations of spatial environments, of the kind previously predicted to exist by Tolman, (1948). A series of subsequent experiments demonstrated that disruption of the rodent hippocampus caused specific deficits in spatial memory (Morris et al., 1982), contextual fear conditioning (Kim, Rison & Fanselow, 1993), and a host of other memory tasks that have been widely used by researchers to probe the cellular and molecular basis of hippocampal dependent memories. Decades of research on rodent models have refined our understanding of how hippocampal neurons encode information about space and other variables. Indeed, hippocampal place cells have become one of the most important animal models in the field of neuroscience for studying not only the neural basis of spatial memory, but many other aspects of neural computation as well. John O'Keefe, Edvard Moser and May-Britt Moser shared the

2014 Nobel prize for physiology and medicine for their seminal contributions to the neuroscience of memory using rodent models of spatial memory coding.

1.2.1 Positional Tuning of Place Cells

As an animal navigates actively through its environment, a place cell is mostly silent unless the animal enters the cell's preferred region of space, called its "place field" (by analogy with the "receptive field" of a sensory neuron). Within its place field, the cell's firing rate has a maximum near the center of the field, which tapers off at the field's edges. After their initial discovery in rats (O'Keefe & Dostrovsky, 1971), place cells have also been described in mice {e.g., Wilson & Tonegawa, 1997; Ziv et al., 2013; Mou et al., 2018}, bats (Yartsev & Ulanovsky, 2013), and humans {e.g. Ekstrom et al., 2003}, suggesting that the hippocampus across mammals shares a similar neural architecture to encode and process spatial information.

Figure 1-1 shows some examples of rat place cells that I recorded on a T maze, which is the same apparatus employed for the experiments reported later in Chapter 2. Each place cell fires at its preferred location on the maze. As a group, the preferred locations of this relatively small set of cells nearly cover the entire maze. At the scale of the entire hippocampus, many place cells (each with its own place field) form a "population vector code" to represent the animal's location in space. Under such a coding scheme, the place cell population as a whole generates patterns of activity that are unique to specific locations in space. Each activity pattern




Figure 1-1 : Example Place Cells. Each graph depicts the position of the rat when a hippocampus CA1 unit fires (colored circles) on the maze (black background). Each unit appears to prefer a specific region on the maze. As a small ensemble, the group of units represent most of the space on the maze. The graph labeled “occupancy” depicts the relative occupancy time across spatial locations (intensity of black); occupancy is not uniform, and informs decisions about whether a unit is a place cell (see methods). Data are from 3 different tetrodes implanted in 1 rat from 1 day of recording and are speed filtered at > 10 cm/s. DA5, 2016-08-22, TT 2, 7, 8.

is a *firing rate vector*: a list of numbers that specifies the firing rate for every place cell in the population. The firing rate vector can be thought of as a “codeword” for labelling a specific location in space. Neural recording studies in animals actively navigating environments demonstrate that place cells in different subregions of the hippocampus have different coding properties, and may thus represent spatial information in different ways (Hargreaves et al., 2005; Mizuseki et al., 2012; Park et al., 2011).

1.2.2 Anatomical Subregions of the Hippocampus

The hippocampus is a relatively large brain structure with bilateral symmetry. In rats, it is located relatively posteriorly and ventral to the cortex, but dorsal to most subcortical structures (Figure 1-2). It resembles two bananas joined at the stalk. Hippocampal regions near the banana's stalk are called the "dorsal" or "septal" hippocampus, and regions near the tips are called the "ventral" or "septal" hippocampus. Slicing the "banana" perpendicular to the long axis reveals the stereotypical hippocampal circuit in cross section. Two curved layers of densely packed neurons, the granule cells of the dentate gyrus (DG) and pyramidal cells of the Ammon's horn (CA) region, form a pair of interlocking C's with dendrites radiating outward from the cell body layers. The CA region is further subdivided into four distinct subregions, numbered CA1-4, but the present discussion shall focus mainly upon the CA3 and CA1 subregions.




Figure 1-2 : *Hippocampal subregions in the rodent.* A) The rodent hippocampus resembles two bananas joined at their stalk. B) Photomicrograph of a hippocampal cross section. (unstained) C) Three main subregions of the hippocampus: DG, CA3, and CA1.

The hippocampus is separated by many layers of processing from both sensory input and motor output (Felleman & van Essen, 1991). It is reciprocally connected with much of the cortex, providing broad input including all major sensory areas. To a first approximation, information from the cortex flows through the hippocampus along a feed-forward processing

path (Figure 1-3). Layer II of the entorhinal cortex sends *perforant path* inputs which project to the dentate gyrus and CA3. In contrast, layer III of the entorhinal cortex sends *temporo-ammonal* projections to CA1 (Hartley et al., 2014; Rolls, 2010; Amaral & Lavenex, ch3, The Hippocampus Book, 2007). Dentate sends *mossy fiber* projections to CA3, and CA3 sends *Schaeffer collateral* projections to CA1. CA3 also sends *recurrent collaterals* back onto itself. This description of hippocampal processing includes only excitatory connections, and is highly simplified. As with

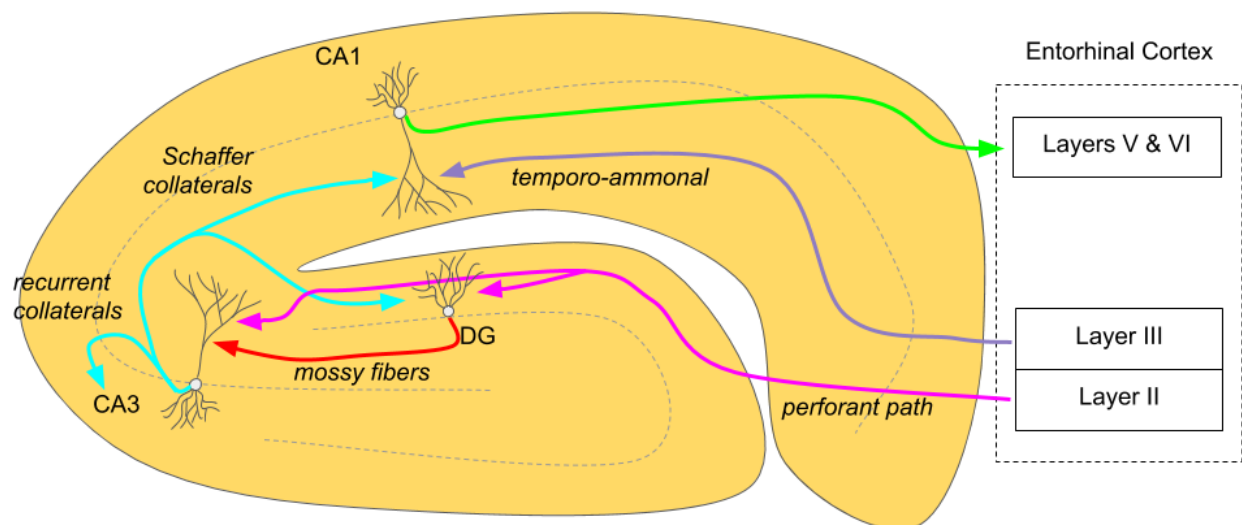


Figure 1-3 : Basic Hippocampal Circuit. Information flow in the hippocampus can be conceived of as a feedforward network. Input from the entorhinal cortex (EC) layer III synapses on the synaptic arbors of dentate gyrus (DG) granule cells and CA3 pyramidal cells. DG also projects to CA3. The CA3 provides feedback to itself, the DG and feeds forward onto the CA1 pyramidal cells. CA1 pyramidal cells synapse onto EC layers V and VI, completing the feedforward circuit. The pathway from DG >> CA3 >> CA1 along the perforant path, mossy fibers and Schaeffer collaterals is known as the trisynaptic circuit.

most of the brain, there are additional connections within and between structures, including complex connections among diverse networks of interneurons (Knierim 2015; van Strien et al., 2009).

1.2.2.1 Dentate Gyrus: Sparse Coding and Pattern Separation

The principal projection neurons of the DG are called *granule cells*. DG granule cells behave similarly to place cells, and are thus tuned to fire at specific locations in space in a given environment (Jung & McNaughton, 1993; Alme et al., 2010; GoodSmith et al., 2015; Senzai & Buzsáki, 2017; Diamantaki et al., 2016a/b; Danielson et al., 2016; van Dijk & Fenton, 2018; Leutgeb et al., 2007; Park et al., 2011). DG place cells are thought to store a very *sparse code*, with only 1-4% of granule cells being active in a given environment (the percent of active cells appears to correlate with novelty and environmental richness; Chawla et al., 2005; Tashiro et al., 2007; Alme et al., 2010; Chawla et al., 2018). Inactive cells are also largely silent outside of their preferred stimulus configuration, producing a low-noise code (Jung & McNaughton, 1993).




Figure 1-4 : *Sparse versus distributed coding.* An ideal sparse code will have no overlap, similar to the DG. An ideal distributed code can have the same exact neurons active;

The *coding capacity* of a neural population may be defined as the total number of uniquely discriminable activity patterns (or codewords) that the population can generate. The coding capacity of a neural population depends on the number of neurons permitted to be active at any given time. “Sparse coding” refers to an encoding scheme where only a very small number of neurons are active simultaneously, as observed in DG. “Distributed coding” refers to an encoding scheme where a larger number of neurons are active simultaneously. differentiation occurs due to the relative activation of each neuron.

In the extreme, the sparsest possible code would allow only a single neuron in the network to be active at any given time. Such a network would only have as many representational states (codewords) as there are neurons in the network (plus one if we count the state where no neurons are active). By contrast, in a distributed network where any number

of neurons is permitted to be active, the number of available codewords is vastly greater. For example, in a population of N neurons that can each assume M different activity states, the number of available codewords is M^N .

For simplicity, assume neural activity states are binary, restricted to firing or not firing. If only K neurons are permitted to be active at any given time, then the number of possible states is given by the N-choose-K formula:

$$f(K) = \frac{(N!)}{K!(N-K)!} \quad \text{Equation 1-1}$$

By this relation, the number of states is greatest when the ratio $K/N=0.5$ (corresponding to the case of a distributed code where about half of the neurons can be active at any given time). The number of states becomes smaller and smaller as K approaches either 1 or N (corresponding to the case of a sparse code, where only a single neuron is permitted to be in a different state from the others at a given time).

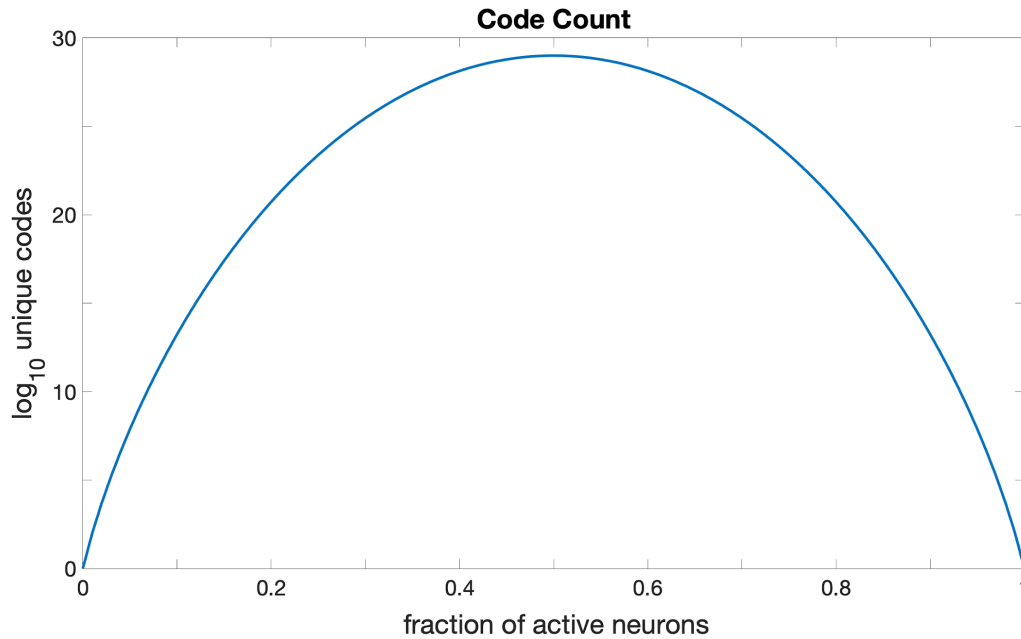


Figure 1-5 : Coding capacity as a function of percentage of active neurons. This graph plots the n-choose-k formula for a population of 100 neurons, with the y-axis on a log scale. Maximum coding capacity is achieved when 50% of neurons are permitted to be active simultaneously.

The n-choose-k formula implies that the sparse code of DG provides far fewer codewords than would be possible with more distributed representations. So apparently, DG only uses a tiny percentage of its potential storage capacity. Why should DG neurons use a sparse code and sacrifice so much of their potential storage capacity? Why not use a more distributed code, with a much greater coding capacity, so that more memories can be stored in DG networks? While sparse coding by DG granule cells does incur costs of low storage capacity, it also accrues some benefits as well. Sparse coding endows DG place cells with a useful property: each DG cell becomes almost like a “labelled line” or “grandmother cell” for a specific memory, such as a familiar location in a specific environment. Consequently, there is very little overlap between the set of DG cells that encode one memory versus another. This is thought to make DG good at performing *pattern separation*, which can be defined as the ability to discriminate the uniqueness of similar things that are different from one another.

In an optimal pattern separator, any change in the input will result in an orthogonal output code, clearly separating even very similar inputs. For example, within a given maze environment, the ideal pattern separator would produce completely different codes (with no overlap of active neurons) for each spatial location. Across maze environments, the ideal pattern separator would continue producing completely different codes for each spatial location. By contrast, in a distributed network where many neurons are active at both place A and B, there would necessarily be overlap between the neurons that are active at A versus B, making them harder to tell apart. The DG is believed to exhibit a higher degree of pattern separation than other hippocampal regions (Leutgeb et al., 2007; Yassa & Stark, 2011; Knierim & Nueneubel, 2016; although see Alme et al., 2010).

Pattern Separation




Figure 1-6 : *Pattern separation schematic*. Pattern separation allows an information processing system to amplify differences in 2 similar inputs to produce completely non-overlapping codes in an extreme case.

Evidence from behavior experiments suggests that DG is required for spatial pattern separation (Gilbert et al., 2001; McHugh et al., 2007; Keinath et al., 2020). When carefully controlled input codes partially overlap *in vitro*, DG can amplify differences at its output (Madar et al., 2019). In the behaving rat, when the environment is slowly morphed from one configuration to another, the DG code is more sensitive to the changes than CA3 (Leutgeb et al., 2007). Both granule and mossy cell types of the DG produce codes with low correlation to

separate a cue mismatch in the environment (Goodsmith et al., 2019), but mossy cells demonstrate stronger remapping than granule cells (Senzai et al., 2017).

This sparse coding property of DG granule cells has been exploited by researchers in “artificial memory” experiments, where a small subset of dentate gyrus cells that are active in one environment are tagged with an opsin during an experience (such as contextual fear conditioning), and then optically re-activated by light at a later time to produce artificial recall of a specific memory for the tagged experience (Liu et al., 2012; Ramirez et al., 2013; Redondo et al., 2014; Ramirez et al., 2015; Tonegawa et al., 2015). This kind of optical tagging experiment would be harder to accomplish in a neural population where place cells use a more distributed code (such as the CA3 and CA1 regions; see below), because there would be more overlap between the activity patterns that encode different environments and locations, making it harder to uniquely tag the representation of a specific location or environment.

As we have seen, DG pays a price for its ability to perform pattern separation. By using a sparse code, DG sacrifices a huge percentage of its potential storage capacity. However, the hippocampus itself is thought not to be a permanent site of memory storage. Rather, evidence suggests that the hippocampus is a temporary storage site where new memories are rapidly encoded when they are first acquired. After being encoded in the hippocampus, memories are thought to be “consolidated” to permanent storage sites in other brain regions, such as the cortex (Squire & Alvarez, 1995; Buzsáki, 1996; Eichenbaum, 2000; Lisman & Morris, 2001; Frankland et al., 2001; Franklin & Bontempi, 2005; Rothschild et al., 2017). SWR events have been hypothesized to play a key role in the consolidation of memories from hippocampus to other brain structures, as discussed further below.

What happens to old memories in the hippocampus after they have consolidated to other brain regions? Perhaps they are written over, or erased, to make room for new memories. If so,

then the hippocampus may not need a vast storage capacity; instead, it may only need enough capacity to store memories for a few weeks or months, until they have been consolidated to other areas. This might help to explain why DG can afford to sacrifice so much of its potential storage capacity by using a very sparse coding scheme: DG does not need to store memories that last throughout the lifespan, but only long enough to be consolidated. After they are consolidated, old memories can be erased or written over, to make room for new memories.

How are old memories in DG erased? Interestingly, DG granule cells are one of very few neural populations in the mammalian brain that undergo neurogenesis (birth of new neurons) throughout life (Aimone et al., 2014; Abbot & Nigussie, 2020). The overall number of DG cells remains relatively constant, despite the birth of new neurons, because older DG neurons die off at about the same rate that new neurons are born (Amaral et al., 2007; Lee et al., 2009; Alme et al., 2010). It is thus tempting to speculate that the death of old DG neurons might “erase” old memories that have been stored long enough to be consolidated. Meanwhile, ongoing neurogenesis of granule cells creates fresh sets of new codewords for representing newer incoming memories that need to be encoded (Aimone et al., 2011; Finnegan & Becker, 2015). Supporting this idea, manipulations of neurogenesis and adult born neurons have been shown to affect hippocampal circuit and memory function, especially in memory tasks that involve pattern separation (Lee et al., 2009; Sahay et al., 2011; Temprana et al., 2015; Danielson et al., 2016; Becker, 2017; Johnston et al., 2016).

1.2.2.2 CA3 Region: Auto-Associativity and Pattern Completion

The main output target of DG granule cells is CA3. Several studies have estimated that the total percentage of CA3 cells active in a given environment averages about 20-30% (Thompson & Best, 1989; Vazdarjanova & Guzowski, 2004; Alme et al., 2014; Marrone et al.,

2014; Skinner et al., 2014; Chawla et al., 2018), approximately an order of magnitude higher than the 1-4% of active cells in DG. In one study, after rats were allowed to explore 11 different environments, 39% of CA3 cells never fired (but these were active during sleep), and 30% formed only 1 place field in 1 room. 13% formed a place field in 2 rooms, and the remaining 18% expressed place fields in three or more rooms. (Alme et al., 2014). This shows that a single CA3 cell can be active in multiple locations and environments. Therefore, a single CA3 cell cannot really be regarded as a “grandmother cell” or “labelled line” that encodes a unique location. Instead, a CA3 cell might better be thought of a “feature detector” that represents a certain property that is possessed by some locations but not others. A specific location in space would thus be encoded in CA3 as a unique combination of properties, represented by a unique pattern of active CA3 cells. Since some locations have overlapping properties, there is also overlap in which CA3 cells are active at any two different locations. Similar to DG, CA3 cells are largely inactive outside their preferred stimulus configurations, but can fire quite vigorously when presented with a preferred stimulus configuration, producing a relatively low-noise, high-gain code (Barnes et al., 1990).

Since the neural code in CA3 is less sparse than in DG, the n-choose-k formula implies that CA3 might be able to utilize a much larger proportion of its potential storage capacity than DG. However, there is reason to believe that the storage capacity of the CA3 network may be reduced by a different phenomenon: *attractor dynamics*.

One of the most distinctive features of the CA3 anatomy is its recurrent connectivity. Unlike other hippocampal populations, CA3 pyramidal cells send direct monosynaptic excitatory connections to one another. CA3 neurons can also inhibit one another, via local inhibitory interneurons. This kind of network architecture (recurrent excitation and inhibition) tends to exhibit dynamics where neural activity settles into a limited set of activity patterns that occupy low-energy “attractor states” of the network (Knierim & Zhang, 2012). Most activity patterns are

not attractor states, so a recurrent network with attractor dynamics can only use a tiny fraction of its potential storage capacity (specifically, the fraction that belongs to the small subset of attractor states), even in cases where the network's code is not "sparse" in the classical sense of only permitting a few neurons to be active at one time (as in DG). The network's connections determine which specific activity patterns are attractor states, but in general, the vast majority of activity patterns (or codewords) in a recurrent network are not attractor states (Hedrick & Zhang, 2016). Thus, even though 20-30% of CA3 neurons can be simultaneously active, the CA3 network may nonetheless have a very limited storage capacity, because attractor dynamics render most activity patterns unstable, and therefore unusable for storing memories. But as discussed above, the hippocampus may only need to store memories long enough to be consolidated to other brain regions, and not throughout the entire lifespan. It was suggested above that this might help to explain why sparse coding in DG, which incurs large costs of reduced coding capacity, may be worth the benefits of enhanced pattern separation. Likewise, the storage capacity costs of attractor dynamics in CA3 may be worth the benefits they accrue.

But what are these benefits? Recurrent networks with attractor dynamics are thought to be ideal for supporting autoassociative memory storage and pattern completion using distributed "feature codes" (Rolls 1996, 2013; Treves & Rolls, 1994; McClelland & Goddard, 1996; O'Reilly & McClelland, 1994; Marr, 1971). A feature code assigns each neural unit to represent some specific property (or feature) of items stored in memory. For example, in a network designed to store memories of faces, individual neurons might encode features such as blue eyes, brown eyes, pointy nose, round nose, curly hair, straight hair, and so on. An individual person's face would then be represented as a specific combination of features (Chang & Tsao, 2017), which could be represented as a specific pattern of active neurons that encodes the conjunction of those features (brown eyes, dark skin, straight hair, pointed nose, high cheekbones, eyes far apart, etc.). Repeated activation of a specific pattern can cause the

neurons in the pattern to strengthen their recurrent excitatory connections to one another, as well as their recurrent inhibitory connections to inactive neurons. This is known as “autoassociative plasticity,” whereby an activity pattern becomes internally associated with itself, and thereby becomes “stamped in” as an attractor state of the network. If a partially complete or corrupted version of the pattern is then presented to the network, attractor dynamics will reconstitute the previous pattern, a process known as *pattern completion* (Fig. 1-7). For example, if a person who is familiar to us changes their appearance by putting on new clothes, or adjusting their hairstyle, we can still recognize that it is the same person, even though their appearance is somewhat changed. Likewise, if a place that is familiar to us changes its appearance (for example, if the walls of a familiar room are painted a different color, or a piece of furniture is moved), we can still recognize that it is the same location. At the level of neural computation, the ability to perform pattern completion is thought to be supported by autoassociative connections in feature-coding networks, and the recurrent connectivity of CA3 suggests that it may be one place in the brain where such computations are performed. Supporting this, lesions of CA3 disrupt a form of spatial pattern completion (Gold & Kesner, 2005) and disabling plasticity in the CA3 with a tissue-specific knockout mouse selectively impairs pattern completion while leaving other memory functions intact (Tsien et al., 1996).

Theoretical modeling demonstrates that CA3-like autoassociative architectures can perform a wide range of problem solving tasks (Levy 1996; Rolls, 1996; 2010), including spontaneous rebroadcast, one-trial learning, subsequence disambiguation, subsequence assembly, content-addressable memory recall, goal-finding, short-cut discovery, and compressed replay of spatial trajectories, which is thought to occur during SWR events that originate in CA3. As discussed further below, SWR events originating in CA3 (and the compressed replay that accompanies them) may be relayed to other brain structures via two main outputs from CA3: the Schaeffer collateral projection to CA1 (which then sends information

to cortex) and the fornical output to lateral septum (which then sends information to subcortical structures).

1.2.2.3 CA1 Region: Hippocampal Output to Cortex

The strategic location of the CA1 region—receiving input from CA3 and sending output to cortex via the subiculum—suggests that it may be a main route through which memories that are rapidly encoded and stored by the pattern separation and completion networks in DG and CA3, respectively, are slowly consolidated over time into longer term cortical memory stores.

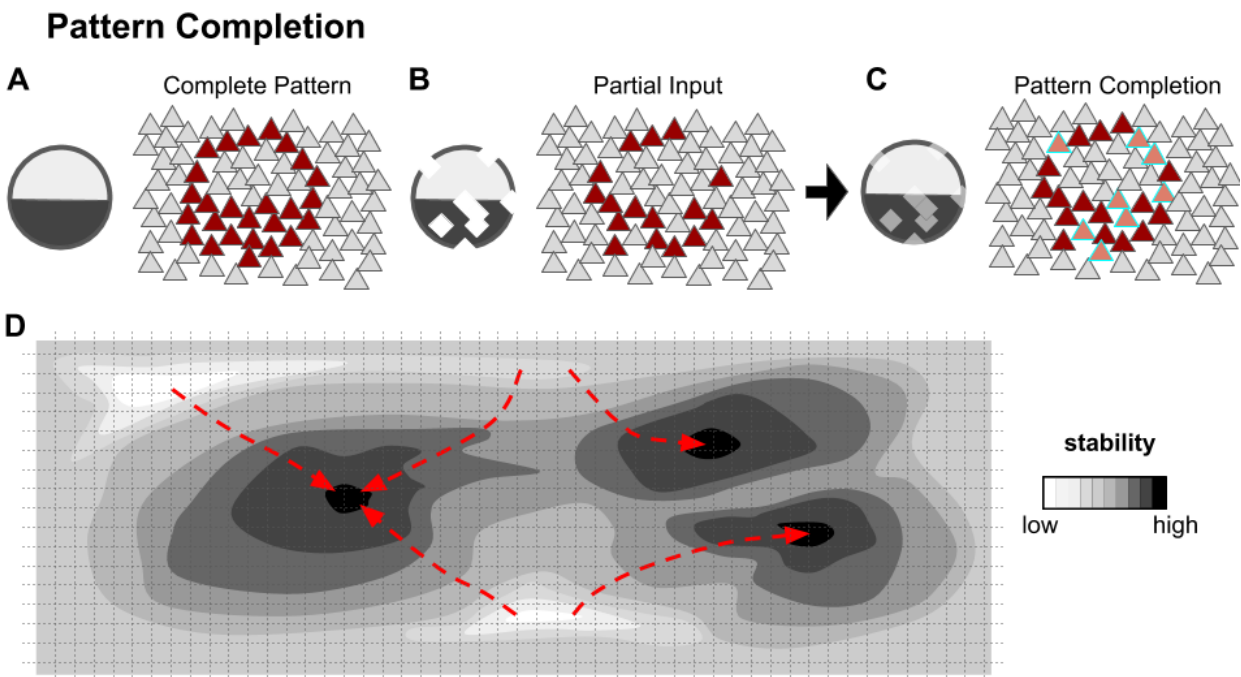


Figure 1-7 : *Autoassociative pattern completion.* Pattern completion allows an information processing system to overcome noisy or incomplete representations of a stimulus to recall a previously stored example, but it has the effect of reducing the effective number of codes available to the system. (A) The half-color circle represents the input and the triangles represent a population of neurons and their representation of this stimulus. This representation is stored in the network in a strongly recurrent network

as a “memory”. (B) When presented with a noisy version of the half-color circle, part of the network representing the half-color circle is activated and the network state quickly evolves to a completed version of the pattern (C). In (D), an 18 x 30 grid is overlaid onto a code space; each square represents a unique code that the representation system can express. Despite its capacity of 540 distinct squares, the space effectively only has 3 stable codes represented by the dark black blobs. All other possible states evolve to the nearest stable state. This arrangement is extremely useful given noisy input, but it substantially limits the number of codewords.

Of the estimated ~150,000 to 625,000 pyramidal cells in the rat CA1 (Miki et al., 2005; Miettinen et al., 2012; West et al., 1991), about 30-50% of them tend to be active in a typical experimental environment (Thompson & Best, 1989; Wilson & McNaughton, 1993; Guzowski et al., 1999; Vazdarjanova & Guzowski, 2004; Skinner et al., 2014; Marrone et al., 2014; Chawla et al., 2018; Muller, 1996; Rich et al., 2014; Leutgeb et al., 2004). A “typical” experimental environment would be a linear track spanning 1-2 meters, or an open field spanning 1-2 m². Most place cells have 1-6 place fields per environment, with increasing numbers of place fields in larger environments (O’Keefe 1976; Muller, 1996; Park et al., 2011; Rich et al., 2014; Mizuseki et al., 2012). An ensemble of as few as 10 CA1 place cells recorded in a behaving animal allows decoding of the position of the animal with an error of < 3cm (e.g. Wilson & McNaughton, 1993; Brown et al., 1998; Zhang et al., 1998), suggesting that even a very small fraction of all available hippocampal CA1 cells contain enough information for downstream structures to decode an animal's position in space with a precision significantly smaller than the size of the animal.

Of the three main hippocampal subregions, the CA1 code appears to be the least sparse (most distributed), and may thus use the largest proportion of its available capacity. If a limited number of cells are permitted to be active at any given time, then according to the n-choose-k formula, maximum capacity is achieved when the number of active cells is about 50%, which is

quite close to the percentage of cells observed to be active per environment in CA1 (Hargreaves et al., 2005; Vazdarjanova & Guzowski, 2004). One possible reason for this high coding capacity in CA1 could be that, as the main output from low-capacity hippocampal stores to high-capacity cortical stores, CA1 must be able to write a huge number of activity patterns (codewords) to multiple storage sites in cortical networks during memory consolidation.

Activity bursts that occur during SWR events have been hypothesized to provide a mechanism by which information could be transferred from hippocampus to cortex during memory consolidation (Buzsáki 1989; Wilson & McNaughton, 1994; Marr, 1971; Squire, 1992; McClelland, McNaughton O'Reilly 1995; Morris & McClelland, 1987; McClelland, 2013). SWR events are observed to occur during both the waking or sleeping states, and there is evidence that memory consolidation processes may occur preferentially during the sleeping state (Buzsáki 2015; Ego-Stengel & Wilson, 2010; Karlson & Frank, 2009; Karlson & Frank, 2008; O'Neill et al., 2008; Grosmark & Buzsáki, 2016; Roux et al., 2017; Girardeau et al., 2009; Gridchyn et al., 2020; Fernandez-Ruiz et al., 2019). Although SWR events are typically recorded in the CA1 layer, they are thought to originate as activity bursts in CA3 (Csicsvari et al., 1999, 2000; Nakashiba et al., 2009). Hence, detection of SWR events in CA1 may reflect the relay of information from low-capacity stores in DG and CA3 into higher capacity stores in cortical areas that CA1 projects via the subiculum.

Importantly, there are other output targets from CA3, besides its Schaeffer collateral outputs to CA1. CA3 sends dense projections to the lateral septum, which in turn projects to numerous other subcortical targets in the midbrain and hypothalamus. It is reasonable to assume that SWR events relayed from CA3 to CA1 may also be relayed from CA3 to subcortical regions via the lateral septum. Before discussing the possible functional importance of transmitting SWRs to subcortical regions, it is worthwhile to provide a more detailed review and background on the electrophysiological processing states of the hippocampal system.

1.3 PROCESSING STATES OF THE HIPPOCAMPUS

Early electrophysiological recordings from the hippocampus in rodents noted marked differences in the local field potential (LFP) that occurred when the animal was in different behavioral states (Vanderwolf 1969; Whishaw & Vanderwolf, 1973). When the animal was in motion – especially during navigation through the environment – the LFP was dominated by a strong theta oscillation in the 4-12 Hz band. Most spatially tuned neurons in the hippocampal system exhibit modulation of their spike train by theta rhythm during free locomotion. By contrast, when the animal is at rest, an LFP pattern known as Large Irregular Activity (LIA) is frequently observed. This consists of ongoing desynchronized activity punctuated by sharp bursts of activity. The characteristic sharp peak in the EEG during the burst events has given rise to the term “sharp wave.” As discussed below, sharp waves are often (but not always) accompanied by increased amplitude in high frequency bands of the EEG (150-300 Hz), known as “ripples.” For this reason, activity bursts during the LIA state are commonly called *sharp wave ripples*, or WR events.

1.3.1 The Theta State

Theta rhythm is a dominant 4-12 Hz rhythm found in the rat hippocampus during movement (Vanderwolf 1969), with a central frequency that is proportional to velocity (Vanderwolf 1969; Gupta et al., 2012; Richard et al., 2013; Hinman et al., 2011; Slawinska & Kasiki, 1998; Bender et al., 2015; Sheremet et al., 2016;). The specific frequency range corresponding to theta depends on species (Buzsáki 2006). Furthermore, the literature contains differences in the definition of the range of theta. Figure 1-8 shows theta filtered at 6-10 Hz from

one of my own experiments. Theta arises from a variety of synergistic mechanisms (rev. Buzsáki, 2002), including membrane potential oscillations (Kamondi et al., 1998), theta-rhythmic action potential production by individual neurons (e.g. Skaggs et al., 1996), local microcircuitry (e.g. Goutagny et al., 2009; Traub et al., 2004) and rhythmic pacemaking inputs from the medial septum (e.g. Zutshi, Brandon et al., 2018; Robinson et al., 2016) and habenula (Aizawa et al., 2013). Some cells exhibit theta phase locking (Ranck 1973; Feder & Ranck, 1973), but as a population, hippocampus cells are more likely to fire on specific phases of theta (e.g. O'Keefe & Recce, 1993; Skaggs et al., 1996; Klausberger et al., 2003; Fox, Wolfson, Ranck, 1986; Buzsáki, 2002).




Figure 1-8 : *Theta rhythm recorded in hippocampal CA1 region.* Graph shows raw LFP (gold), theta bandpass filtered LFP (6-10 Hz IIR filter, black), and running speed (green). When the animal is moving, the LFP contains strong, rhythmic theta. When the animal stops moving, LFP becomes desynchronized.

1.3.1.1 Phase Precession

A hippocampal place cell modulates its firing rate as a Gaussian-like function of position in 1D (e.g. Diba & Buzsáki, 2007; Kjelstrup et al., 2008), 2D (e.g. O'Keefe & Dostrovsky, 1971; Lever et al., 2002) and 3D environments (Yartsev & Ulanovsky, 2013) around its preferred point in space. This firing rate signal provides good information about the spatial location of the animal, but due to the symmetry of firing rates around the tuning curve peak, there are multiple locations with identical mean firing rates. The ambiguity can be resolved by a population code (extracting information from the firing rate of more than one cell), but it can also be resolved by

a phenomenon known as phase precession (O'Keefe & Recce, 1993; Skaggs et al., 1996), where the phase of theta at which a place cell fires advances as the animal moves along its trajectory through a place field (Figure 1-9). In effect, the burst frequency of the cell occurs at a slightly higher frequency than theta, resulting in phase precession. As the animal enters a cell's place field, the cell fires just prior to the valley of theta rhythm, and shifts toward earlier phases of theta on successive cycles. As the animal passes through the center of the place field, the cell spikes at the peak of theta. As the animal exits the place field, the cell spikes just after the valley of theta, completing a full 360-degree cycle of phase precession. Incorporating both firing rate and theta phase into decoding of spatial location from ensemble activity improves accuracy by >40% (Jensen & Lisman, 2000), demonstrating the potential benefits of this dual rate-phase coding scheme. Phase precession occurs every time the animal runs through the field (Fig 1-10).

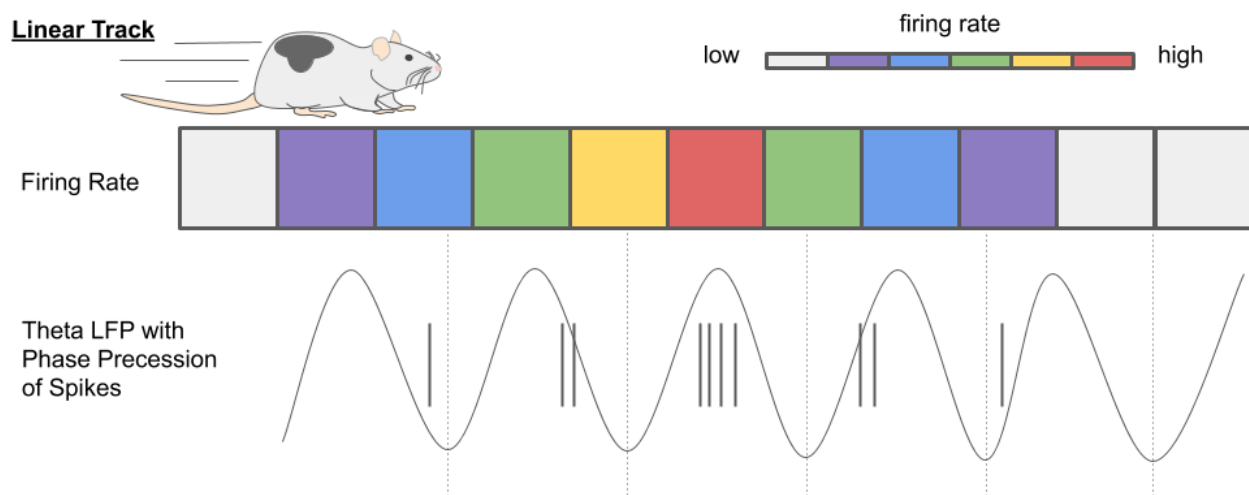


Figure 1-9 :. Schematic diagram of theta phase precession. As a rat moves left to right along a linear track (boxes), a place cell which prefers the center of the track fires with increasing and then decreasing vigor (colored boxes) in accordance with its place field. As the rat proceeds, it produces theta rhythm (grey signal) at approximately 8 Hz. The place cell fires bursts of spikes (grey marks) at a rate slightly faster than 8 Hz. The spikes thus shift in phase, firing on earlier phases of each subsequent theta cycle.

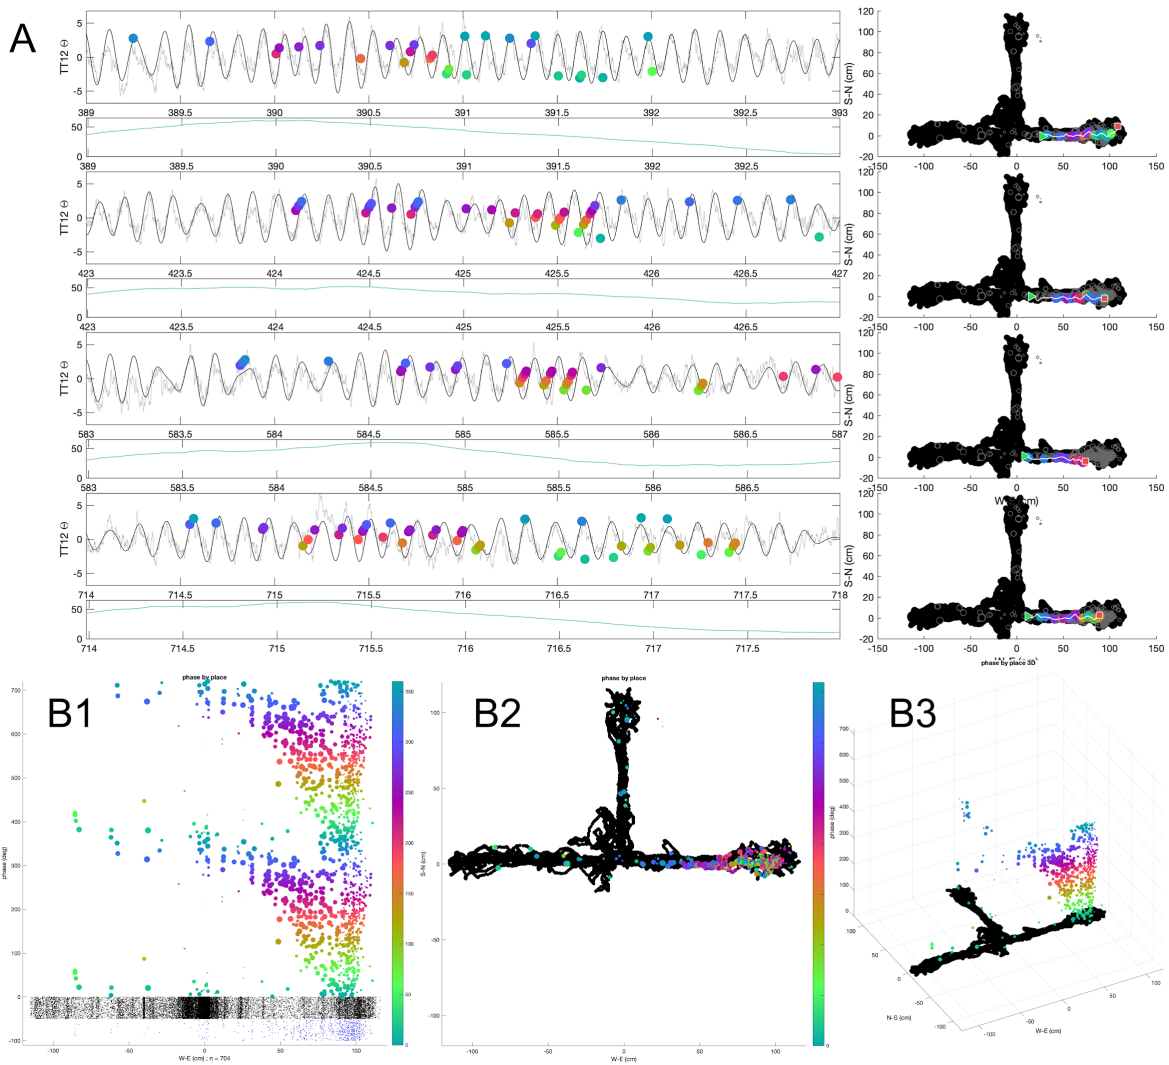


Figure 1-10 : (A) *Theta phase precession occurs on each passage through the place field. Four examples of phase precession by the same place cell. In each example, the rat runs from near the center towards the East most position on the maze. Each example consists of an unfiltered theta LFP (grey line), a filtered Theta LFP (black line), and phase at the time of firing (colored circles; color represents phase, as does the y-axis position; radians). Speed (green line) appears below each LFP-unit plot. The maze occupancy map (black circles), all cell firings (grey rings), phase + firing position (colored circles), trajectory (white line), start position (green arrow with white border) and end position (red square with white border) appear next to each pair of LFP-unit and speed plots. (B) Phase coding of space in a hippocampal CA1 unit. The example unit is a place cell that preferentially fires on the East arm of the*

maze. (A) A depiction of phase precession in one spatial dimension (East-West; x-axis; cm) and theta phase (y-axis; degrees). The cell shifts in phase to represent space as the rat proceeds East, as seen by the ovoid-shaped blob of colorful dots. Black stippling represents all spatial points occupied by the rat on the maze in the East-West axis. Blue stippling represents the position of the unit when it fires, filtered for speed. The position on the y-axis as well as the color of the point represent phase in degrees. In order to clearly demonstrate the cyclic nature of the phase code, the phase positions are doubled. (B) A depiction of theta phase and position at the time of firing projected onto a 2D representation of the rat's occupancy on the maze. Black circles represent the position of the maze. Colorful circles represent the phase at the time of firing and simultaneously depict the position. (C) An isometric view of a 3D rendering of the position and phase data. Data from DA5, 2016-08-22 : TT7, unit 1; phase measured against theta on LFP from channel 44; units are speed filtered > 10 cm/s.

1.3.1.2 Theta Sequences

As phase precession occurs simultaneously in many individual place cells, it gives rise to temporal structuring of spike activity at the population level, resulting in the generation of *theta sequences*. Briefly, phase precession in individual place cells produces an temporally ordered sequence of firing among multiple place cells which lie adjacent to one another in space (Foster & Wilson, 2007). On a linear track, the temporal sequence of spikes fired by place cells within each theta cycle closely matches the spatial sequence of their place fields along the track, in the direction that the animal is currently traveling (Johnson & Redish, 2007; Regier et al., 2015; Wikenheiser & Redish, 2015a/b; Kay et al., 2020). Within each theta cycle, the first place cells to fire are those that fire maximally behind the animal (thus, the animal is currently leaving their place fields), the next cells to fire are those that fire maximally at the animal's current position, and the last to fire are those that fire maximally at locations ahead of the animal (thus, the

animal is just starting to enter their place fields). Hence, spike sequences within each theta cycle generate a time-compressed replay of the rat's past, current, and future positions.

These time-compressed theta sequences have been proposed to assist with decision making, as rats can be observed to actively sample different possible future trajectories at intersections in a maze by moving their heads (Johnson & Redish, 2007), or even abruptly switch direction after seeming to commit to a trajectory choice (Schmidt et al., 2013; Redish, 2016; Papale et al., 2016). Recent decoding analyses suggest that online alternation between theta sequences encoding different future paths occurs on alternating theta cycles before the rat commits to a path (Kay et al., 2020), suggesting that the animal may be deliberating over which path to take by generating different theta sequences. When commitment to a chosen path occurs, the rat stops encoding both choices on alternating theta cycles, and encodes only the chosen path (Kay et al., 2020).

If it is true that theta sequences are involved in decision making, then temporally ordered place cell spikes should be relayed out of the hippocampus to other brain structures that are involved in decision making. Hippocampal outputs from CA1 to cortical regions could be responsible for transmitting theta sequences to frontal cortical areas that are involved in decisions. Outputs from CA3 to the lateral septum, which in turn projects to the midbrain and hypothalamus, could be involved in relaying theta sequences to subcortical structures that are involved in decision making and motivated behaviors. Supporting this latter possibility, phase precession by lateral septal neurons has been observed in our lab and others (Figure 1-11).




Figure 1-11 : *Example of phase precession in a LS unit.* Phase precession occurs outside the hippocampus in the LS, and it occurs in 2D environments. Phase precession in 2D environments is more difficult to identify and interpret. The left panel depicts T-maze occupancy (grey); note the square open field in the center. Selecting paths from the North arm to the West arm (black) and plotting the rat's location and theta phase (colored circles) at the time of unit firing reveals a gradual shift in phase across the entire path. Linearizing the path by collapsing it onto the arc of a circle allows a 1D projection of place in the right panel. A clear ovoid shape with a downward slope indicates the phase precession phenomenon. The dotted grey line depicts a linear regression fit to the linearized position and the phase. The linearized position appears on the x-axis, and the phase appears as both y-axis position and color. Point size is proportional to velocity at the time of firing. Phase is doubled to the range 0-720 degrees to depict an entire cycle. The thin middle panel depicts a histogram of the phases, revealing a slight phase bias for firing. The upper middle histogram depicts an occupancy histogram of the linearized projection. The upper right panel depicts a histogram of speed-filtered occupancy (blue boxes) and also firing rate in Hz (red line; right axis). The firing rate modulation suggests that this LS cell operates like a place cell from the hippocampus with a large firing field.

Along 1D pathways, phase precession is clearly visualized on a position vs spike phase plot (Fig *Phase-position-plots-example; phase-precession-examples* ; Kempter et al., 2012).

Along a trajectory through a 2D space, phase precession is more complex to visualize, requiring linearization relative to the trajectory through the place field (Huxter et al., 2008; Jeewajee et al., 2014). Phase precession has also been described in 3D space in bats in the absence of a strong rhythmic theta with a steady frequency (Eliav et al., 2018).

1.3.2 Large, Irregular Activity (LIA) State

When a rodent stops navigating, the hippocampal EEG transitions out of the theta state and into the LIA state, during which SWR events occur in the local field potential (LFP) of the CA1 region (Vanderwolf 1969; O'Keefe 1976; Buzsáki, 1986; Suzuki & Smith, 1988 a/b/c; Csicsvari et al., 1999; Foster & Wilson, 2006; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Gupta et al., 2010; Pfeiffer & Foster, 2013). SWR is an acronym of "sharp wave-ripple," a compound term referring to two distinct electrophysiological phenomena that occur in separate layers of the hippocampus: the sharp wave, and the ripple.

1.3.2.1 Sharp Waves

Some have argued that sharp wave (SW) production is mutually exclusive with the theta state (Buzsáki 1986; Suzuki & Smith, 1988 a/b/c). Optogenetic activation of the medial septum to induce hippocampal theta suppresses SWR production (Vandecasteele et al., 2014). These observations impl that the rat will not be in motion when a SW is produced. The SW is easily detected in LFP signals that are bandpassed between 1-50 Hz, and manifest as a large negative deflection in the LFP in the stratum radiatum of CA1 (Buzsáki, 2015). The SW results from a burst of input to CA1 arising from CA3, and arriving via Schaffer collaterals to the CA1 stratum radiatum (Csicsvari et al., 1999, 2000; Nakashiba et al., 2009). The burst of input

synchronously depolarizes dendrites in the radiatum, producing the characteristic “sharp” deflection in extracellular potential (see Fig SWR anatomy).

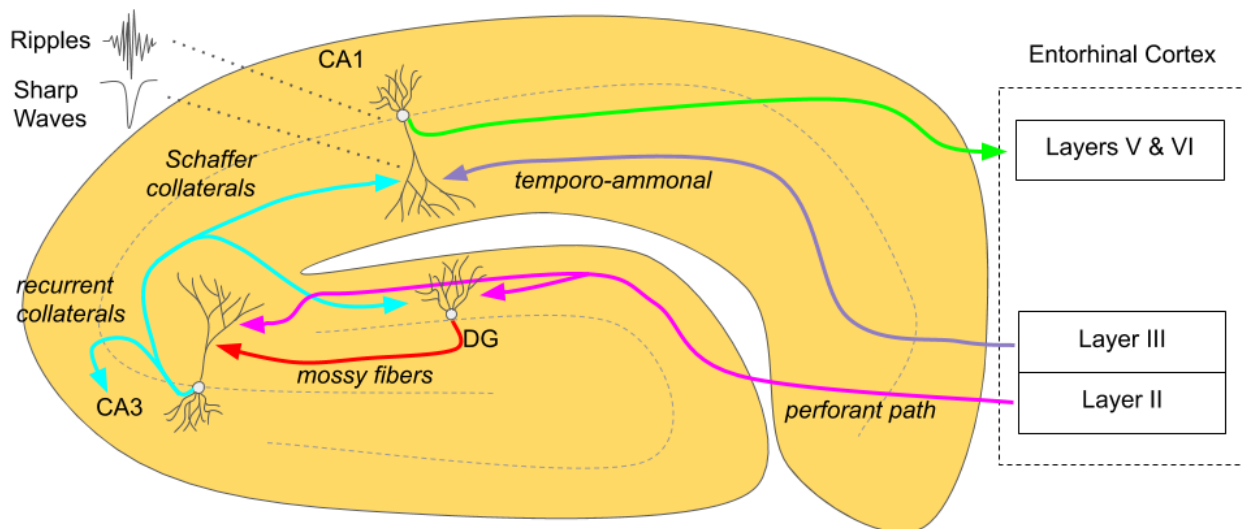


Figure 1-12 : SWR electrophysiology schematic. Coordinated bursts of multi-unit activity in CA3 produce depolarization in CA1 dendrites, generating a 5-40 Hz “sharp wave” in the LFP at the dendritic layer. Ripples reflect coordinated bursts of multi-unit activity in the CA1, generating a 150-250 Hz “ripple” in the LFP at the somatic layer. When both LFP events coincide, it is known as a sharp-wave ripple (SWR).

1.3.2.2 Ripples

The ripple arises from a population spike burst in the hippocampal CA1 cell layer, which produces a large, transient power increase in high frequency bands of the LFP. The specific band for measuring ripples varies across the literature, but filtering between 125-250 Hz is typical (O’Keefe & Nadel, 1978, p150; Suzuki & Smith, 1988 a/b/c; Buzsáki, 2015; O’Neill et al., 2006). Johnson & Redish (2007) argue in favor of constraining the band to 180-250 Hz to avoid possible high gamma contamination. Informal experimentation with band limits suggests that

180-250 Hz performs at least as well as 125 or 150 to 250, so this work describes ripples detected in the 180-250 Hz band. SW and ripples can occur independently: a burst of input from CA3 that fails to trigger a CA1 population spike will yield a SW without a ripple, whereas a population spike in CA1 that is triggered by input other than a coordinated CA3 volley will yield a ripple without a SW. However, SW and ripples are often observed together while the animal is at rest in the non-theta state. Events detected in the ripple band are sometimes called High Frequency Events (HFEs) to distinguish the high-frequency (ripple) from the lower frequency (SW) component of SWR events, especially in cases where recording methods do not permit recording both phenomena simultaneously.

1.3.2.3 Compressed Replay

SWR events occur when populations of hippocampal neurons burst together, and these population bursts often contain “compressed replay” of hippocampal place cells embedded within them (Lee & Wilson, 2002; Foster & Wilson, 2006; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Gupta et al., 2010; Pfiefer & Foster, 2013). “Compressed replay” refers to ordered sequences of hippocampal place cell spikes that occur at an accelerated timescale. Much like theta sequences, compressed replay events can be decoded as time-compressed traversals along trajectories through space (Skaggs & McNaughton, 1996; Nadasdy et al., 1999). But unlike theta sequences--which occur during navigation and encode forward trajectories along the animal’s current motion path--compressed replay events tend to occur during stillness and seem to encode a wider range of trajectories than just forward motion trajectories.

Replay events may occur during waking or sleeping states (Buzsáki 2015; Ciscsvari et al., 2007; Ego-Stengel & Wilson, 2010; Karlsson & Frank, 2009; O'Neill et al., 2008; O'Neill et al., 2006; Grosmark & Buzsáki, 2016). Replay events often traverse trajectories within the context that the animal currently occupies. For example, while navigating a maze, trajectories through the maze may be replayed during a pause in motion (Foster & Wilson, 2006; Diba & Buzsáki, 2007). Alternatively, replay events may traverse remote trajectories from another environment, such as replaying trajectories from environment B while the animal is in environment A (Karlsson & Frank, 2009). During sleep, replay events tend to represent traversal of trajectories through environments that the animal has recently visited before sleep, rather than the environment in which the animal is sleeping (Lee & Wilson, 2002). Given the large number of possible ordered sequences that the hippocampus could produce, and the precisely sequenced nature of cell firings, the most parsimonious explanation for replay events is that they represent a mental deliberation over a previous or planned experience.

1.4 FUNCTIONS OF SHARP-WAVE RIPPLES & REPLAYS

There are three commonly accepted hypotheses in the literature concerning the possible function of SWR events and the accompanying compressed replay activity : memory consolidation, credit assignment, and model-based planning. All three hypotheses can co-exist without difficulty (they are not mutually exclusive). Each of these three functions might be subserved by different hippocampal output pathways to other brain regions.

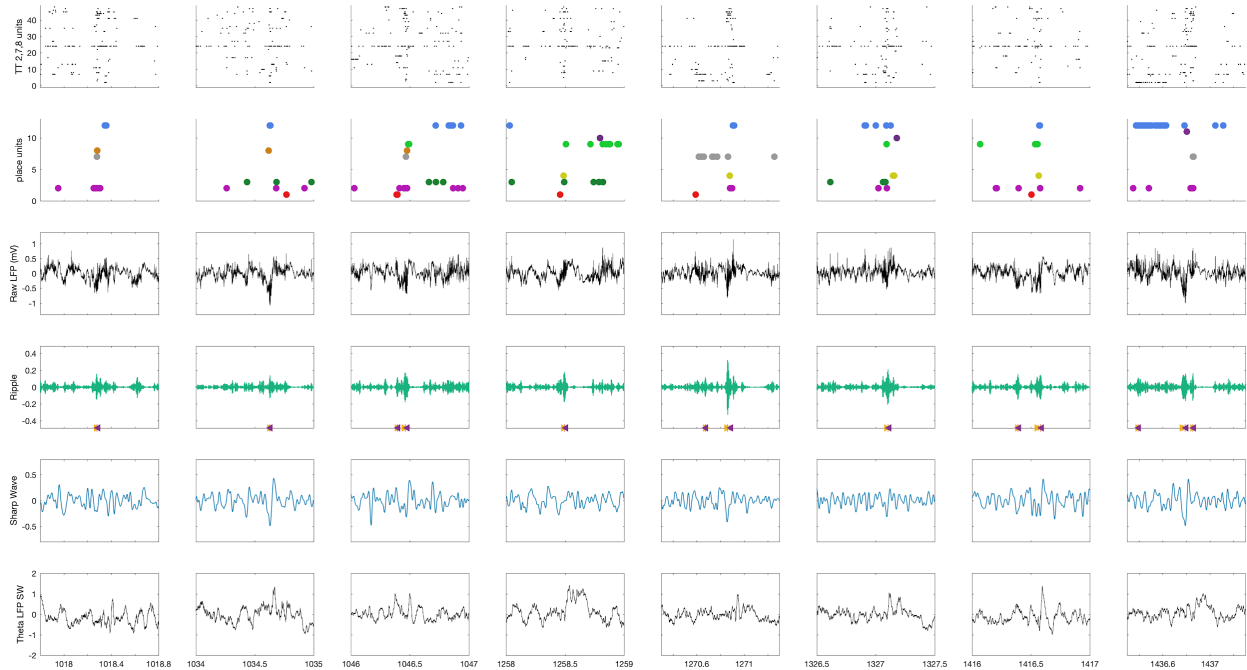


Figure 1-13 : Several examples of ripple events that coincide with temporally ordered sequences of place cell firing. Each column contains one example. The top row contains rastergrams of all units for 3 different tetrodes; population firing bursts are evident during events. The second row contains rastergrams for 11 place-like units ordered from bottom to top such that a continuous trajectory from the North arm to the East arm would appear as a line with a large positive slope. Cells that prefer the West arm are at the very top. Colors correspond to the figure depicting place-like units. Referencing the place unit figure reveals several sequences that appear to represent paths on the maze. It is not possible to determine if these replay sequences are forward or reverse due to the rat engaging in exploratory behavior, engaging in bi-directional behavior. The third row contains the unfiltered SWR LFP from one of the tetrodes with units. Events coincide with a “drop” in the signal, which is the sharp wave, and a higher frequency component, which is the ripple. The fourth row contains the filtered ripple band (180-250 Hz); signal power in this band correlates the multi-unit firing rate seen in the top row. The fifth row contains the filtered “sharp wave” band (5-40 Hz); strong sharp waves are not always obvious, although the ripple events coincide with a trough in the signal in every case. The sharp wave corresponds to coordinated input from CA3 depolarizing the dendrites of the CA1 pyramidal cells (see text on “Sharp Waves”). The sixth row contains

the unfiltered signal from the theta LFP from a tetrode about 800 microns rostral to the unit tetrode.
(unpublished data)

1.3.1 Memory Consolidation

The ability to recall episodic events from another time and place is a powerful mental mechanism. This out-of-context recall can manifest itself during waking sleeping states. For example, in the waking state, I might recall details of a mid-day conversation in the lab while riding home on the bus later in the evening. In the sleeping state, I may dream of surfing in the ocean at Venice beach the night after a trip to the beach. In both cases, my brain mentally reviews an earlier activity, and perhaps this helps the activity to be consolidated into long-term memory, rather than forgotten. Presumably, non-human animals may also possess similar waking and sleeping episodic memory recall capabilities. It is believed that hippocampal replay during SWR events might be a neural biomarker for this sort of “mental time travel” to different spatial contexts and times.

During sleep, the LFP signature for SWR becomes more prominent (Buzsáki 2015, Fig 4), and ripple production rate is several-fold higher during non-REM sleep (Ego-Stengel & Wilson, 2010. Fig 2b). Awake replays of a remote environment are more robust than either replay in the present environment or replays during quiescence (Karlsson & Frank, 2009). Restful SWRs have been proposed to play a role in binding hippocampal ensembles together after novel experience and learning (Karlsson & Frank, 2008; O’Neill et al., 2008; Grosmark & Buzsáki, 2016; Tang et al., 2017). Disruption of SWR during a rest period after memory training temporarily disrupts task performance and can reduce the rate of learning across days, suggesting a role for SWRs in memory consolidation (Girardeau et al., 2009; Ego-Stengel & Wilson, 2010; Roux et al., 2017). Recent technical advances allowed disruption of only SWR

events accompanied by replays that encode trajectories in the learning context, sparing SWR events that encode trajectories in other contexts; this precision approach also disrupts task performance in the learning context (Gridchyn et al., 2020), reinforcing the hypothesis that the content of the SWR replay event is critical for consolidation. Complementing these results, artificial prolongation of SWR events -- which appears to extend the physical distance encoded by the concurrent replay trajectories -- enhances performance in spatial memory tasks (Fernandez-Ruiz et al., 2019). Hence, memory performance can be bidirectionally modulated by disrupting or enhancing SWR events, providing strong evidence that SWR replay events play an important role in memory consolidation.

Decades of theory predict that interplay between the hippocampus and cortex gradually establishes strong cortical representations of episodic memories, largely removing dependency on the hippocampus (Marr 1971; Buzsáki, 1989; Buzsáki, 1996; Thierry et al., 2000; Frankland & Bontempi, 2005; Willshaw et al., 2014; Kitamura et al., 2017). The primary output of the hippocampus is through the subiculum and the entorhinal cortex (Fig 1.3; Witter & Groenewegen, 1990; O'Mara 2005; Amaral & Lavenex, 2007; Cappaert et al., 2015), although CA1 also directly innervates subregions of prefrontal cortex and the medial nucleus accumbens (Jay & Witter, 1991; Thierry et al., 2000; Floresco et al., 2001; Cenuiza & Swanson, 2007; Cappaert et al., 2015). The entorhinal cortex is heavily interconnected with the cortex, providing a two-way interface for hippocampal-cortical dialogue (Felleman & van Essen, 1991; Cappaert et al., 2015). A substantial body of evidence demonstrates that hippocampal-cortical synchronization occurs during SWRs (see Table 1). This coordination presumably provides a mechanism by which initially hippocampal-dependent memories gradually become effectively independent from the hippocampus.

Table 1-1: SWR Activates Cortex

Cortex	multi-region		Battaglia et al., 2004; Isomura et al., 2006; Logothetis et al., 2012; Logothetis 2015 Abadchi et al., 2020
	Medial Entorhinal	mEC	Chrobak & Buzsáki, 1996; Olafsdottir et al., 2016; O'Neil et al., 2017
	Prefrontal	PFC	Jadhav et al., 2016; Tang et al., 2017; Khodagholy et al., 2017
	Anterior Cingulate	ACC	Wang & Ikemoto, 2016
	Retrosplenial	RSC	Alexander et al., 2018; Nitzan et al., 2020
	Posterior Parietal	PPC	Khodagholy et al., 2017; Wilber et al., 2017
	Somatosensory	SSC	Sirota et al., 2003
	Visual	VC	Ji & Wilson, 2007
	Auditory	AC	Rothschild et al., 2017

1.3.2 Credit Assignment

Value-based reinforcement learning (VBRL) is the process of learning to perform whatever available actions are predicted to yield the greatest amount of future reward (Montague et al., 1996; Sutton & Barto, 1998; Colombo, 2014). In both natural and laboratory environments, animals readily learn to perform actions that lead to rewarding outcomes. But in complex VBRL tasks, such as learning to perform a long sequence of left and right turns to navigate from start to goal in a large maze, the reward is not delivered until long after certain actions have been performed. The first turn that a rat makes out of the start box might be critical

for finding the reward, but by the time the rat reaches the goal, many seconds or even minutes may have elapsed since that first turn was made. How can a reward signal experienced at the goal location reinforce actions that were taken much earlier at other locations in the maze? The attribution of an outcome to behaviors performed in the remote past is a conundrum known in learning theory as the *credit assignment problem* (Sutton & Barto, 1998).

Replay of spatial trajectories during SWR events is hypothesized to help animals to solve the credit assignment problem (Foster & Wilson, 2006; Lee et al., 2012). One type of awake SWR replay event encodes trajectories that extend backwards from the animal's present location on a maze. That is, place cells fire sequentially in reverse order to replay the path most recently taken by the animal in the backwards direction (Foster & Wilson, 2006; Diba & Buzsáki, 2007; Gupta et al., 2010). These reverse replay events occur at reward locations, and thus may provide a mechanism whereby the animal can reflect back upon all of the navigational choices that it has made to arrive at a goal, thereby helping to solve the credit assignment problem.

If SWR events are involved in assigning credit (or blame) for past choices that have led to rewarding (or non-rewarding) outcomes, then SWR signals should be delivered to brain structures that modify stored associations that guide future behavior. Different types of associations are thought to guide different types of behaviors (Sutton & Barto, 1998): associations between actions and outcomes are thought to guide “model free” instrumental choice behaviors, associations between states and their values are thought to guide “model-based” planning and decision making, and associations between discriminative stimuli and actions are thought to guide “habitual” behaviors that become automatized through extensive repetition. Some of these associations may be stored in frontal cortical regions, and as we have already seen, CA1 provides an output pathway via which SWR signals can be relayed from the hippocampus to cortical areas. Hence, SWR events initiated by CA3 place cells may be relayed

to CA1, which in turn sends them to cortical regions where neural associations are modified by learning (Csicsvari et al., 1999, 2000; Nakashiba et al., 2009).

Other associations that guide behavior are likely to be stored in subcortical regions, such as the striatum and midbrain. The same CA3 place cells that project to CA1 also project down into the LS, which in turn projects down to the RMTg of the VTA (Raisman 1966; Siegel et al., 1975; Risold & Swanson, 1997; Lou et al., 2011; Gomperts et al., 2015; Tingley & Buzsáki, 2018; Wirtshafter & Wilson, 2019; Tingley & Buzsáki, 2020). The VTA is well associated with signaling rewarding stimuli and reward prediction error (Schultz 1986; Schultz, 1998; Schultz et al., 1993; Schultz et al., 1997; Pan et al., 2005; Cohen et al., 2012), and sends dense projections to the dorsal and ventral striatum. SWR replay events initiated by CA3 cells could thus influence credit assignment and neural plasticity via outputs to these structures that are routed through LS. This pathway is summarized in Fig. 11-14.




Figure 11-14 : *CA3-LS-VTA circuit schematic*. Hippocampal CA3 cells project to LS which project down to inhibitory cells in the tail of the VTA (the rMTG). These connections constitute a feedforward disinhibitory circuit whereby CA3 MUA associated with a CA1 SWR replay can enable coordinated firing of dopaminergic projection neurons of the VTA.

1.3.3 Model-Based Decision-Making

During decision-making tasks, plans for future action are often based upon predictions derived from past outcomes of similar decisions. Hence, decision making relies upon the ability to recall the past to predict the outcome of future choices.

Very simple or highly overtrained decision tasks can be acquired as stimulus-response associations, whereby the animal simply learns that in the presence of particular discriminative stimuli, a specific action leads to a rewarding outcome. The rewarding action is reinforced over trials, and becomes automatic (insensitive to devaluation) after many trials. This kind of “model-free” reinforcement learning is thought to be controlled largely by habit learning circuits in the dorsal striatum (see Chapter 3 for further discussion).

More complex decision tasks may require cognitive deliberation over various possible action choices, and ongoing comparisons between the expected reward for each possible action. Such deliberation may not only require making one-step predictions about what reward will immediately follow the currently selected action, but also making “chained” predictions about what new choices will be made available as a consequence of the currently selected action, and what rewards and/or choices those newly available actions might lead to. In other words, complex decision tasks may require animals to mentally navigate along trajectories through a “decision space” that is represented by an internal model of the causal structure of the world.

Navigation through decision space may depend upon the ability of rats to “pre-play” valid sequences of place cells through paths on a maze which have never yet been experienced by the rat (Gupta et al., 2010; Dragoi & Tonegawa, 2011; 2013). The capability to produce viable behavior sequences that the rat has never experienced is strong evidence that the hippocampus contains a world-model and is not simply limited to memory recording. During learning, a rat must actively engage in decision-making, accessing a mental model of the decision space and also updating that model prior to the development of habitual behavior.

In support of this idea, Singer et al., (2013, Fig 5b) report increased SWR preceding successful trials while the rat is still improving at the task. When rodents are engaged in decision tasks, such as spatial navigation through a maze where a choice must be made between two or more possible arms, SWRs accompanied by replay events are observed at decision points (Karlsson & Frank, 2009) or before initiation of a trial (Diba & Buzsáki, 2007; Davidson et al., 2009; Pfeiffer & Foster, 2013). In a task involving a choice between a near or far food well, forward replays project to either of the two wells (Wikenheiser & Redish, 2013). Hippocampal place fields sometimes cluster around goal zones (Hok et al., 2007; Durpet et al., 2010; Olafsdottir et al., 2015), as do decoded SWR replays (Durpet et al., 2010). SWR preplay can predict the immediate future behavioral trajectory (Pfeiffer & Foster, 2013). There even appears to be a very small population of cells (~1%) dedicated to representing reward/goal location (Gauthier & Tank, 2018). Together, these findings paint a compelling picture that perhaps the hippocampus can directly participate in model-based planning of actions that lead to goals and rewards.

However, place cells do not always cluster around goal locations (Poucet & Hok, 2017). Even more problematic, when the rat must dynamically choose between competing goal locations, it is difficult (at best) to decipher the rat's goal from place cell theta sequences or SWR replay content (Johnson & Reddish, 2007; Singer et al., 2013; Papale et al., 2016; Kay et al., 2020). If the hippocampus is not principally responsible for encoding goal and reward information, then how can the content of SWR replay sequences be used to guide reward-motivated decisions? SWR signals may have to be relayed to other brain structures outside of the hippocampus.

Projections from CA1 and subiculum to either PFC (mPFC & OFC) or to Nacc may help to connect hippocampal spatio-temporal processing with goals and rewards (Lisman & Grace, 2005; Cenquizca et al., 2007; Poucet & Hok, 2017). Indeed mPFC cells display joint place-goal activity in a task with a goal, but not in a random foraging task, in contrast to the more evenly distributed hippocampal place field map (Hok et al., 2005; Poucet & Hok, 2017). But model based planning may also depend upon projections from CA3 to LS, which in turn projects to subcortical areas involved in reward processing, prediction error, and value estimation.

CHAPTER 2

Bidirectional firing rate modulation of lateral septal and dorsomedial striatal neurons during hippocampal sharp-wave ripple events

(This chapter is a version of a research report submitted for publication.)

2.1 INTRODUCTION

The lateral septum (LS) is a major subcortical output target of hippocampal projection neurons (Raisman 1966). LS sends descending projections to midbrain regions such as the lateral hypothalamic area, substantia nigra, and ventral tegmental area (Risold & Swanson, 1997), which in turn send diffuse projections to the ventral and dorsal striatum. It has been proposed that this septal output pathway may be an important route via which the hippocampus exerts influence over behaviors that are regulated by the midbrain dopamine system, including motor actions, reward-seeking, attention, arousal, and decision making (Luo et al., 2011; Gomperts et al., 2015; Tingley & Buzsáki, 2018, 2020; Wirtshafter & Wilson, 2019). To further investigate how hippocampal output influences the activity of septal and striatal neurons, the present study analyzed how hippocampal EEG states were correlated with single-unit spikes recorded in hippocampus, septum, and striatum.

The rodent hippocampus exhibits distinct patterns of local field potential (LFP) activity during different behavioral states (Vanderwolf 1969). While an animal is actively navigating through its environment, the LFP is synchronized by theta oscillations in the 4-12 Hz band. By contrast, when the animal is at rest, the LFP enters a state of desynchronization punctuated by phasic bursts, a pattern known as large irregular activity (LIA). During LIA, transient synchronization events produce peaks in lower frequency bands (1-50 Hz) of the LFP, known as *sharp waves*. Sharp waves often co-occur with bursts of power in higher bands (125-300 Hz) known as *ripples*. Sharp waves and ripples can occur independently, but they are often observed together in the low and high frequency bands of the LFP (Buzsáki, 2015). Hence, they are commonly referred to together as *sharp-wave ripple* (SWR) events.

Theta and SWR states of the LFP are accompanied by distinct firing patterns of hippocampal neurons. During the theta state, as the animal navigates through its environment,

hippocampal *place cells* fire selectively at preferred spatial locations (O'Keefe & Dostrovsky, 1973). Place cells are hypothesized to encode cognitive maps of familiar spatial environments (O'Keefe & Nadel, 1978; Redish, 1999), and supporting this, an animal's momentary position can be accurately decoded from population vectors of place cell activity as it navigates through space (Wilson & McNaughton, 1993). When an animal is at rest, the LFP switches from the theta state to the LIA state, during which SWRs are accompanied by brief population bursts of place cell activity—referred to *compressed replay* events—that can be decoded as “imagined” spatial trajectories through an environment (Skaggs & McNaughton, 1996; Lee & Wilson, 2002; Foster & Wilson, 2006; Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009). While an animal is running on a maze, replay events occur during pauses in motor activity and tend to encode trajectories that start or end at the animal's current location (Jackson et al., 2006; Johnson and Redish, 2007; Diba & Buzsáki, 2007; Karlsson & Frank, 2009; Pfeiffer & Foster, 2013; Wu et al., 2017; Kay et al., 2020). These on-maze replay events have been hypothesized to aid in deliberative decision making about where to travel next (Yu & Frank, 2015), and also in assessing outcomes of prior navigational choices (Foster & Wilson, 2006). After an animal is removed from a maze, replay events that occur during subsequent periods of rest often encode trajectories from the recently visited maze environment, and these post-behavioral replay events have been hypothesized to aid in the consolidation of recent experiences to long-term memory (Wilson & McNaughton, 1994; Buzsáki, 1998; Ego-Stengel & Wilson, 2010; Girardeau & Zugaro, 2011).

Here, we analyzed responses of neurons in septum and striatum during SWRs, while freely behaving rats ran trials on a T maze and rested in a bucket between trials. We found that a subset of septal neurons were either excited or inhibited during SWRs; these SWR-responsive septal neurons often fired coherently with hippocampal theta rhythm, and some were modulated by the animal's running speed, in agreement with other recent reports (Wirtshafter & Wilson,

2019). A small percentage of neurons in striatum were also excited or inhibited during SWRs, and as in septum, these neurons spiked coherently with theta rhythm and were commonly modulated by running speed. In both structures, neurons that were excited during SWRs tended to be positively modulated by running speed, whereas neurons that were inhibited during SWRs tended to be negatively modulated by running speed. In septum (but not striatum), SWR-excited neurons fired during late phases of the theta cycle, whereas SWR-inhibited neurons fired during early phases of the theta cycle. After describing these findings in detail, we discuss their possible implications for understanding how the hippocampus modulates subcortical circuits for behavioral learning and decision making.

2.2 RESULTS

Rats ($n=3$) were trained to run repeated acquisition and reversal trials on a T-maze (Fig. 2-1). At the start of each session, recording cables were connected and the rat was placed for 5 minutes in a white plastic bucket located next to the maze (during behavior, the bucket and experimenter always remained stationary and in the same arrangement behind the start arm such that the bucket and experimenter provided no cues about reward location) for a period of baseline recording. The rat was then placed by the experimenter at the start location on a T-maze apparatus consisting of 4 arms extending 90 cm at right angles from a 30x30 cm central platform (Fig. 2-1A). Throughout each block of trials, a barrier was placed at the entrance to one of the four arms, while the three remaining arms served as the start, baited, and unbaited arms for the T-maze task (see Methods). After each trial on the maze, the experimenter returned the rat to the bucket for 2-5 minutes while the maze was cleaned and baited for the next trial.

Over 6-8 days of initial training, rats learned to find food on one arm of the T-maze. During this initial training period, hippocampal tetrodes were advanced until robust SWRs and theta rhythm were detected on two different tetrodes in the same hemisphere (see Methods).

These two tetrodes were assigned as the ripple and theta recording electrodes, respectively, and neither was advanced further during the remainder experiment. Starting with the next session, the goal and/or start arm was changed each time the rat achieved a criterion of 7/8 correct responses (see Methods). Rats spent a median of 3.7 minutes on the maze and 19.9 minutes in the bucket during each session (Fig. 2-1B). In the bucket and on the maze, SWR events were only measured during periods of stillness when the rat's running speed remained <2 cm/s for 3 s or more (Fig. 2-1E). During these periods of stillness, the mean rate of SWR generation was significantly higher (paired $t_{48}=7.97$, $p=2.4 \times 10^{-10}$) on the maze (0.43 Hz) than in the bucket (0.28 Hz; Fig. 2-1C), whereas the mean peak amplitude of SWR events was significantly higher (paired $t_{48}=16.7$, $p=1.2 \times 10^{-21}$) in the bucket (84 mV) than on the maze (63 mV; Fig. 2-D).




Figure 2-1 : Behavioral and neurophysiological data samples. A) Maze apparatus and holding bucket. **B)** Distributions of total time spent sitting still on the maze versus in the bucket during each recording session (N=53). **C)** Distributions of mean SWR rates during stillness on the maze versus in the bucket during each recording session. **D)** Distributions of peak ripple amplitudes during stillness on the maze

versus in the bucket during each recording session. **E)** Cumulative spatial distributions across all recording sessions of locations visited (black), locations of sitting still (blue), and locations where SWR events occurred (red) for each rat. **F)** Number of cells recorded in each brain area (top graph) and total number of recording sessions (bottom) for each of the 3 rats in the study. **G)** Septal and striatal recording sites for each rat; symbols indicate recording sites for cells that were excited (stars) versus inhibited (triangles) by SWR events.

2.2.1 Single Unit Responses during SWR Events

Concurrent with SWR event detection, single units were recorded from the hippocampus (n=216 units from 3 rats), septum (n=226 units from 2 rats), and striatum (n=378 units from 3 rats) during each maze session (Fig. 2-1F). Hippocampal single units were only analyzed in the hemisphere contralateral from the SWR detection site, to prevent confounds in the analysis that might arise from anatomical proximity between the SWR detection and single-unit recording sites. Hence, in all three brain regions from which single units were recorded, the tetrodes were positioned several mm away from the SWR detection site. To maximize the number of unique cells that were recorded throughout the experiment, tetrodes in septum and striatum were advanced by 150 μm after each behavior session, so that different units would be recorded from these tetrodes in every session. By contrast, hippocampal tetrodes were advanced by at most 83 μm per day (and usually not at all), so that these tetrodes would remain within the hippocampal region throughout the entire experiment. Consequently, most hippocampal units in the dataset were recorded more than once over multiple sessions, whereas each septal and striatal unit was recorded exactly once, during a single session, before the tetrode was advanced to find new cells. Since many hippocampal units were recorded over multiple sessions, the analyses below include only from the first session during which a hippocampal

unit was recorded. Thus, in all three structures (septum, striatum, and hippocampus), single-unit responses were always analyzed using a single session's worth of data for each cell.

Fig. 2-2 shows example data from Rat 1, obtained from a session during which SWR events were recorded in the right hemisphere of CA1, while hippocampal units were recorded contralaterally in the left hemisphere of CA1 (Fig. 2-2A). Example recordings are also shown for units recorded in left septum (Fig. 2-2B) and right striatum (Fig. 2-2C). Most septal units included in the study were recorded in the lateral septum, but some were recorded in the septofimbrial region; few if any were recorded in medial septum (Fig. 2-1G). A small number of striatal units were recorded from the ventral striatum (nucleus accumbens), but most were recorded in dorsal striatum (Fig. 2-1G). Hippocampal unit data came from the CA1 region near the pyramidal layer (see examples in Fig. 2-2A). Single-unit responses to SWR events were measured only during periods of stillness (running speed < 2cm/s), both in the bucket (panels D1-F1) and on the maze (panels D2-F2). By contrast, coherence of single unit spike trains with hippocampal theta rhythm was measured only during periods of active behavior on the maze (running speed >10 cm/s, circular spike phase distributions in panels D3-F3). We adopt the convention that the valley and peak of theta rhythm occur at phases of 0° and 180° , respectively. For the example neurons shown in Fig. 2-2, SWR events excited CA1 and septum units, but inhibited the unit in striatum.

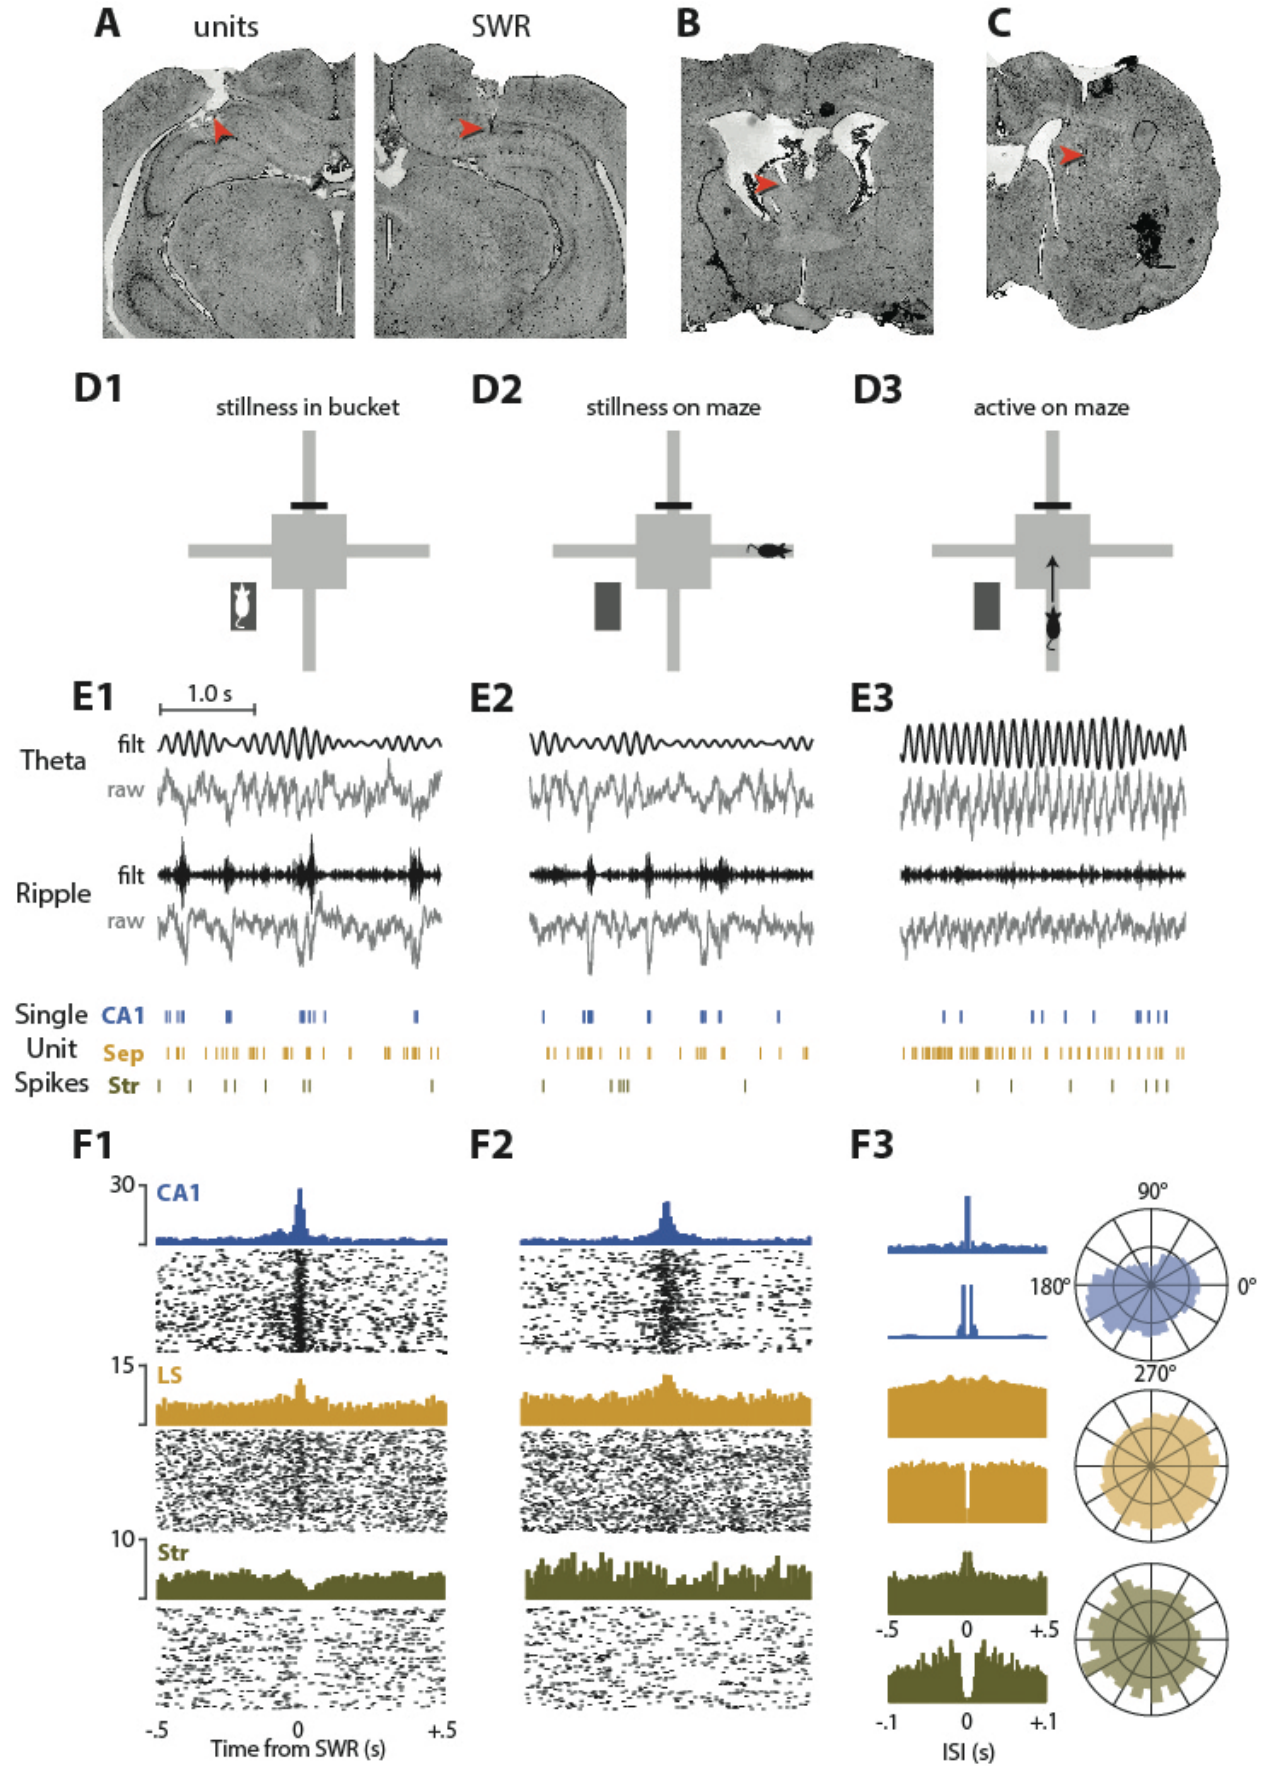


Figure 2-2. Example data from a single recording session. **A)** Red arrows indicate the example unit recording site in the left CA1 (left panel) and the SWR detection site in right CA1 (right panel). **B)** Example unit recording site in left septum. **C)** Example unit recording site in right striatum. **D)** Schematic diagrams for three different behavior conditions: stillness in bucket (D1), stillness on maze (D2), and running on maze (D3). **E)** Traces show 3 s of raw and filtered LFP data from the theta (top row) and ripple (middle row) channels, aligned with examples of single unit spike rasters (bottom row) from a CA1, septum (Sep), and striatum (Str); sample data is shown for stillness in bucket (E1), stillness on maze (E2), and running on maze (E3). **F)** Example cell PETHs and rastergrams aligned to SWR events that occurred in the bucket (F1) or on the maze (F2). Autocorrelograms of interspike intervals (ISIs) are shown (F3, left) between -0.5 and $+0.5$ s (top graph in each row) to illustrate theta rhythmicity, and between -0.1 and $+0.1$ s (bottom graph in each row) to illustrate spike refractory periods. Polar plots (F3, right) show distributions of each example cell's spike phase relative to hippocampal theta.

2.2.1.1 Excitatory and Inhibitory Responses

To test how neurons responded to SWR events, a signed rank test was performed to compare each neuron's spike rate in a time window spanning ± 50 ms from the SWR peak against the baseline firing rate when SWRs were not occurring (see Methods). Two separate tests were performed to analyze whether a neuron responded to SWRs in the bucket versus on the maze, and a neuron was classified as excited or inhibited during SWR events if either of the two tests yielded $p < 0.01$, indicating that the SWR response rate was greater or less than the baseline rate, respectively (Fig. 2-3). We never observed units with opposing SWR responses for the two behavior conditions (that is, no cells were excited on the maze but inhibited in the bucket, or vice versa). However, some units did exhibit significant SWR excitation or inhibition only in the bucket but not on the maze, or vice versa (see below).

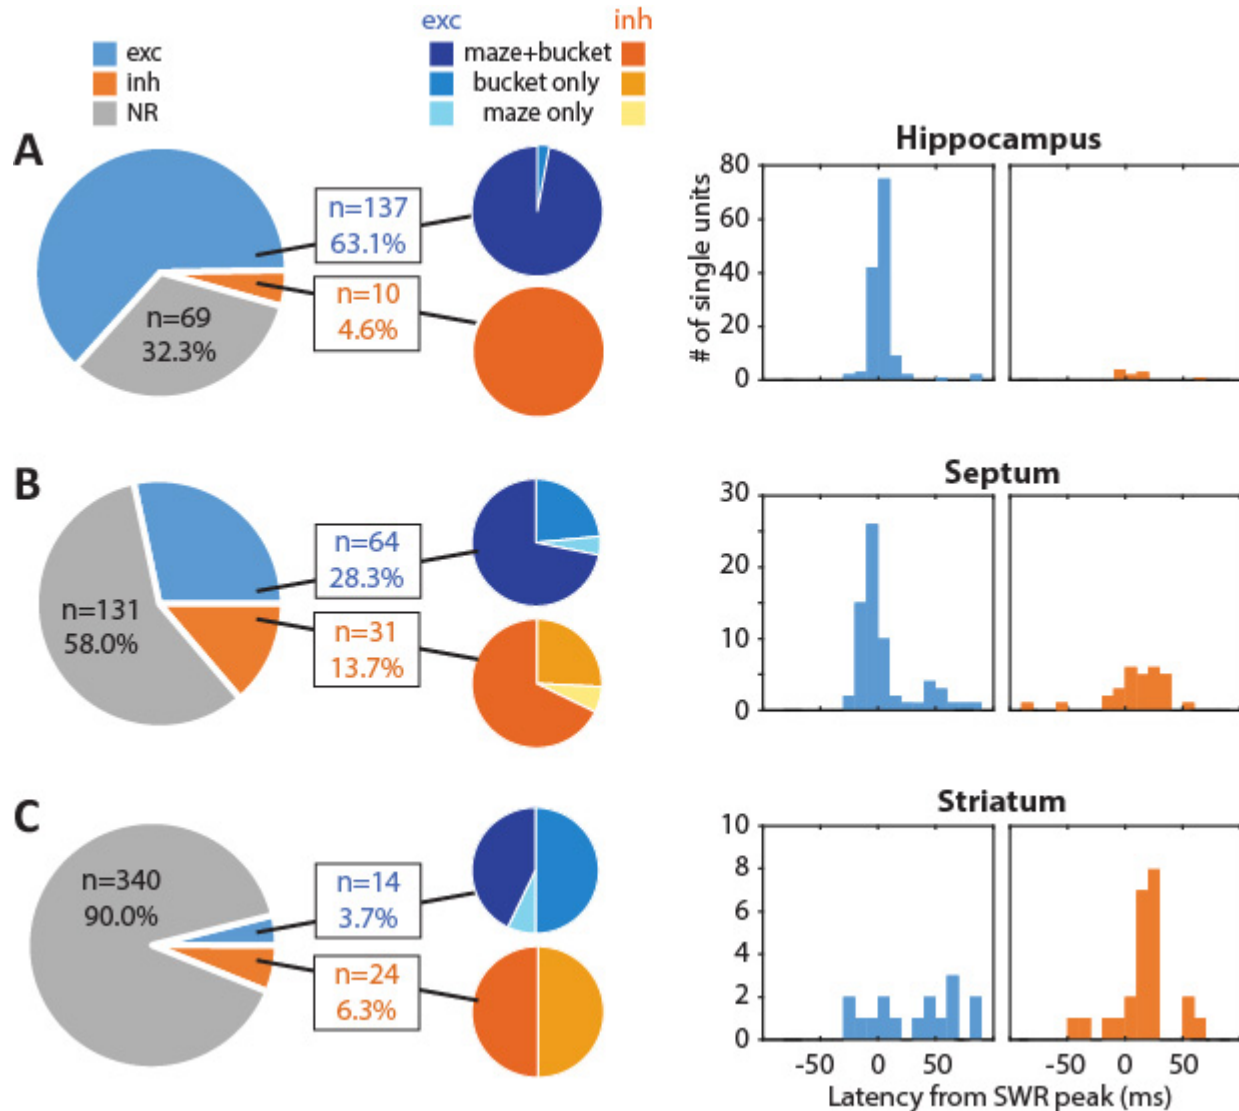


Figure 2-3. *Unit responses to SWR events.* Large pie charts in the left column show proportions of neurons in hippocampus (A), septum (B), and striatum (C) that were excited (exc), inhibited (inh), or non-responsive (NR) to SWR events. Small pie charts in the middle column show proportions of SWR-excited (blue) and inhibited (orange) cells that responded in the bucket, on the maze, or both. Histograms at right show distribution of response latencies for cells that were excited or inhibited by SWRs.

2.2.1.1.1 Hippocampal Units

Hippocampal neurons were recorded from all three rats in the study. Excitatory responses to SWR events were observed in 137/216 (63.4%) of hippocampal neurons, while 69/216 (31.9%) of hippocampal neurons were not significantly responsive during SWR events (Fig. 2-3A). Only 10/216 (4.7%) hippocampal neurons were inhibited during SWR events, but at least one SWR-inhibited neuron was observed in the hippocampus of each rat. The percentage of hippocampal neurons that were excited during SWRs was similar in all three rats (67.5%, 54.2%, 66.7%). Hippocampal neurons that were excited during SWRs usually exhibited their peak spike response within ± 5 ms of the SWR onset (Fig. 2-3D), with a mean response latency of $+4.3 \pm 1.4$ ms. For this analysis, SWR onset time was measured at the moment when the ripple envelope crossed a standard deviation threshold on the SWR electrode (see Methods). The average latency for inhibitory SWR responses was slightly longer ($+10.0 \pm 7.1$ s), but did not differ significantly from the latency of excitatory responses ($t_{145}=1.04$, $p=.3$; the statistical power of this comparison was limited by the small sample of SWR-inhibited neurons).

The excitation of hippocampal neurons that we observed during SWRs is consistent with findings from prior studies (Wilson and McNaughton, 1994; Skaggs & McNaughton, 1996; Kudrimoti et al., 1999; Foster & Wilson, 2006; Davidson et al., 2009). Here, we did not decode replay trajectories from hippocampal unit bursts, since this phenomenon has been well studied in prior experiments. Instead, our main objective was to analyze neural activity in septum and striatum during SWRs (see below), and confirming the occurrence of hippocampal unit bursts was necessary to validate accurate isolation of SWR events for these analyses. Note that we only analyze hippocampal unit responses in the hemisphere opposite from where SWRs were recorded. Evidence suggests that SWR events typically originate in CA3, from which they are monosynaptically transmitted to CA1 via Schaffer collaterals (Csicsvari et al., 2000; Nakashiba

et al., 2009; Buzsáki, 2015). If SWR events occur simultaneously in both hemispheres of CA3, then they should arrive nearly simultaneously in both hemispheres of CA1, in which case it would be expected that the peak of the LFP ripple response in CA1 should occur nearly simultaneously with SWR-evoked unit responses in the opposite hemisphere, as we observed.

2.2.1.1.2 Septum Units

Septum neurons were recorded from two of the three rats in the study (n=121 units from rat 1, n=105 units from rat 2). A majority of septum neurons (131/226, or 57.9%) were non-responsive during SWRs, but 64/226 (28.3%) of septal neurons were excited during SWR events (Fig. 2-3B), which was a significantly smaller proportion of SWR-excited neurons than we observed in the hippocampus, (1,N=442)=54.9, $p=1.27 \times 10^{-13}$, and 31/226 (13.7%) of septal neurons were inhibited during SWR events, which was a significantly higher proportion of inhibited neurons than we observed in the hippocampus, (1,N=442)=10.8, $p=9.9 \times 10^{-4}$. The overall percentage of SWR-responsive neurons was higher in Rat 1 (39% excited, 30% inhibited) than in Rat 2 (16.2% excited, 18.9% inhibited), possibly owing to the different paths that tetrodes followed through septum in the two rats (see Fig. 2-1G). The mean latency for excitatory SWR responses in septum was -4.6 ± 3.1 s, which did not differ significantly from the mean latency in hippocampus ($t_{202}=.07$, $p=.94$). Even though many septum neurons had small negative response latencies, this does not mean that they fired prior to the initiation of hippocampal sharp waves. A small negative response latency simply suggests that the sharp wave signal (which presumably originates from CA3) can be detected a bit earlier by unit spikes in septum than by LFP ripples in CA1. Large negative response latencies (<20 ms) would be more likely to indicate cells that actually fire prior to the initiation of sharp waves, but these were rarely observed. Interestingly, the latency distribution for excitatory septum responses appeared to be bimodal (Fig. 2-1E): most units responded at short latencies (<10 ms) as would be

expected if SWRs originating in CA3 were relayed to septum neurons via a monosynaptic pathway, but some units responded at longer latencies (>35 ms) as might be expected from a polysynaptic pathway (or alternatively, from a slower time constant for integrating excitatory inputs in some septum units). When neurons with long response latencies (>35 ms) were excluded from both the septum and hippocampal datasets, SWR-excited neurons were observed to have significantly shorter response latencies in septum than hippocampus ($t_{66}=3.65$, $p=5.2 \times 10^{-4}$). One possible explanation for this could be that on average, septum neurons might have shorter membrane time constants than hippocampal neurons; if this were the case, then even if SWR events originating in CA3 arrived simultaneously in CA1 and septum, the septum responses could appear slightly earlier than CA1 responses. The mean latency for inhibitory SWR responses in septum was $+10.0 \pm 5.1$ s, which did not differ significantly from the latency of excitatory responses when all excitatory cells were included in the comparison ($t_{95}=.97$, $p=.34$), but did differ significantly when latencies >35 ms were omitted from both datasets ($t_{62}=3.16$, $p=.002$). This suggests that excitatory input from hippocampus may monosynaptically drive most of the excitatory SWR-evoked responses in septum, followed at a short delay by polysynaptic feedforward inhibitory responses, and a small number of polysynaptic excitatory responses.

2.2.1.1.3 Striatal Units

Striatal neurons were recorded from all three rats in the study. A majority (340/378, or 90%) of striatal neurons were non-responsive to SWRs (Fig. 2-3C). Only 14/378 (3.7%) of striatal neurons were excited during SWR events (but at least two SWR-excited neurons were observed in each of the three rats), whereas 24/378 (6.3%) of striatal neurons were inhibited during SWR events (with roughly similar percentages in all three rats: 6.0%, 7.7%, 0%). The latency for excitatory SWR responses in striatum was more variable and on average longer ($+36.8 \pm 10.8$ ms) than excitatory responses in hippocampus ($t_{152}=6.05$, $p=1.1 \times 10^{-8}$) or septum ($t_{82}=4.04$, $p=1.18 \times 10^{-4}$; all SWR-excited septum neurons were included in this comparison). The mean latency for inhibitory SWR responses in striatum was $+15.0 \pm 8.4$ ms, which did not differ significantly from the latency of inhibitory responses in hippocampus ($t_{38}=0.0$, $p=1$) or septum ($t_{54}=0.053$, $p=.6$) and also did not differ from the latency of excitatory responses in striatum ($t_{41}=1.62$, $p=.11$).

2.2.1.2 SWR Responses in the Bucket versus on the Maze

SWR responses were analyzed separately in the bucket versus on the maze. Across all brain regions, $n=52$ neurons responded selectively to one behavioral condition but not the other (bucket but not maze, or vice versa), with the majority of these neurons (48/52, or 92.3%) located outside of the hippocampus. Of these behavior-selective neurons, 88.5% (46/52) responded only in the bucket but not on the maze, whereas 11.5% (6/52) responded on the maze but not in the bucket. Hence, neurons that responded to SWRs only in the bucket were much more prevalent than neurons that responded only on the maze ($p<.0001$ for binomial test against 50% probability for each preference). One possible explanation for this bias could be that rats spent more time sitting still in the bucket than on the maze during each session (Fig. 2-

1B), and therefore, the statistical power for detecting SWR responses was generally higher in the bucket than on the maze. To control for such bias in statistical power, a downsample-and-shuffle approach was used to equalize the statistical power for detecting SWR responses on the maze and in the bucket (see Methods). After controlling for sampling bias in this way, the total number of behavior-selective cells in all brain regions fell from 52 to 32; among these, the proportion of cells that responded to SWRs only in the bucket but not on the maze was 87.5% (28/32), and the proportion of cells that responded to SWRs only on the maze but not in the bucket was 12.5% (4/32). Hence, controlling for SWR sampling bias did not change the proportion of behavior-selective cells that were responsive in the bucket versus on the maze, $(1, N=84)=.018$, $p=.89$. It should also be noted that the amplitude of SWR responses was larger in the bucket than on the maze (Fig. 2-1D), which could have influenced detection accuracy.

Across all SWR-excited and SWR-inhibited neurons combined, the proportion of neurons that responded selectively to SWRs under only one behavioral condition (either on the maze but not in the bucket, or vice versa) differed significantly among brain regions, $(1, N=280)=61.0$, $p<.00001$. A more detailed analysis of responding in the maze versus bucket for each brain region is presented below.

2.2.1.2.1 Hippocampal Units

Fig. 2-3A shows that in the hippocampus, 97% (133/137) of SWR-excited and 100% (10/10) of SWR-inhibited neurons responded to SWRs both on the maze and in the bucket (see examples in Fig. 2-2F). For excited and inhibited neurons combined, 97.3% (143/147) of hippocampal units responded to SWRs both on the maze and in the bucket. Hence, hippocampal neurons showed very little selectivity of their SWR responses for one behavioral condition over the other; almost all hippocampal neurons that responded to SWRs were responsive both in the bucket and on the maze.

2.2.1.2.2 Septum Units

Fig. 2-3B shows that in septum, 71.9% (46/64) of SWR-excited neurons (examples shown in Fig. 2-4A) and 67.7% (21/31) of SWR-inhibited neurons (examples shown in Fig. 2-4D) responded to SWRs in both behavior conditions, whereas 23.4% (15/64) of SWR-excited neurons (examples shown in Fig. 2-4B) and 25.8% (8/31) of SWR-inhibited neurons (see example, Fig. 2-4E) responded to SWRs only in the bucket but not on the maze, and 4.7% (3/64) of SWR-excited neurons (see example, Fig. 2-4C) and 6.5% (2/31) of SWR-inhibited neurons responded to SWRs only on the maze but not in the bucket. The proportion of septum units exhibiting behavior-selective SWR responses did not differ significantly for SWR-excited versus SWR-inhibited neurons, ($1, N=123$) = .09, $p = .76$. Hence, excitatory and inhibitory responses to SWRs in septum were similarly dependent upon the rat's behavioral state.

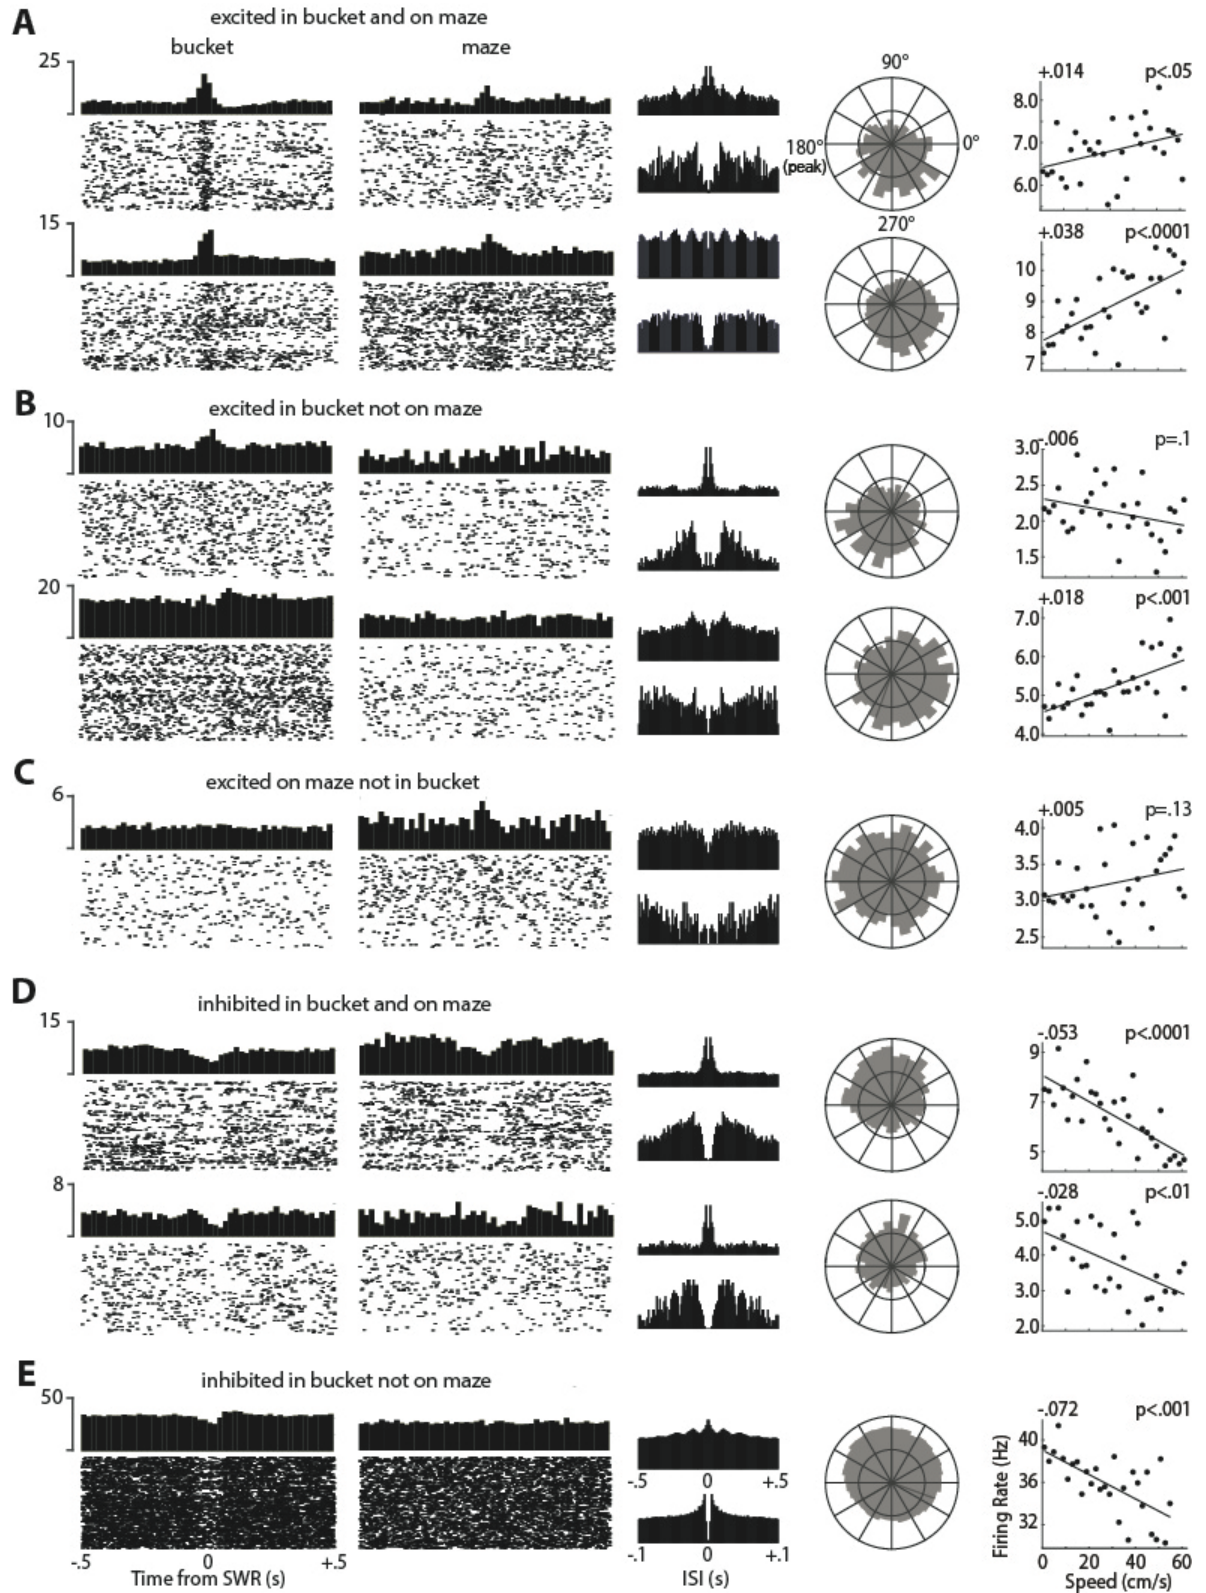


Figure 2-4. *Example cells recorded in septum.* Two left columns show PETHs and rastergrams aligned to SWR events that occurred in the bucket (first column) or on the maze (second column). Third column shows autocorrelograms of interspike intervals (ISIs) between -.5 and +5 s (top graph in each row) to illustrate theta rhythmicity, and between -.1 and +.1 s (bottom graph in each row) to illustrate spike refractory periods. Fourth column shows circular distribution of spike phases relative to theta rhythm in the hippocampal LFP. Fifth column shows a scatterplot in which each point is the cell's mean firing rate (y-axis) at a given running speed (x-axis); regression line shows linear fit to the scatter points, with slope of speed modulation (in Hz/cm/s) at upper left and p-value of Pearson correlation at upper right. Examples are shown for septal neurons that were excited by SWRs in both the bucket and the maze (**A**), excited by SWRs in the bucket but not the maze (**B**), excited by SWRs on the maze but not in the bucket (**C**), inhibited by SWRs in both the bucket and the maze (**D**), and inhibited by SWRs in the bucket but not on the maze (**E**).

2.2.1.2.3 Striatal Units

Fig. 2-3C shows that in striatum, 42.9% (6/14) of SWR-excited neurons (examples shown in Fig. 2-5A) and 50% (12/24) of SWR-inhibited neurons (examples shown in Fig. 2-5C) responded to SWRs in both behavior conditions, whereas 50% (7/14) of SWR-excited neurons (examples shown in Fig. 2-5B) and 50% (12/24) of SWR-inhibited neurons (see example, Fig. 2-5D) responded to SWRs only in the bucket but not on the maze. The proportion of striatal units exhibiting behavior-selective SWR responses did not differ significantly for SWR-excited versus SWR-inhibited neurons, ($1, N=58$) = .06, $p=.81$. Hence, excitatory and inhibitory responses to SWRs in striatum were similarly dependent upon the rat's behavioral state.

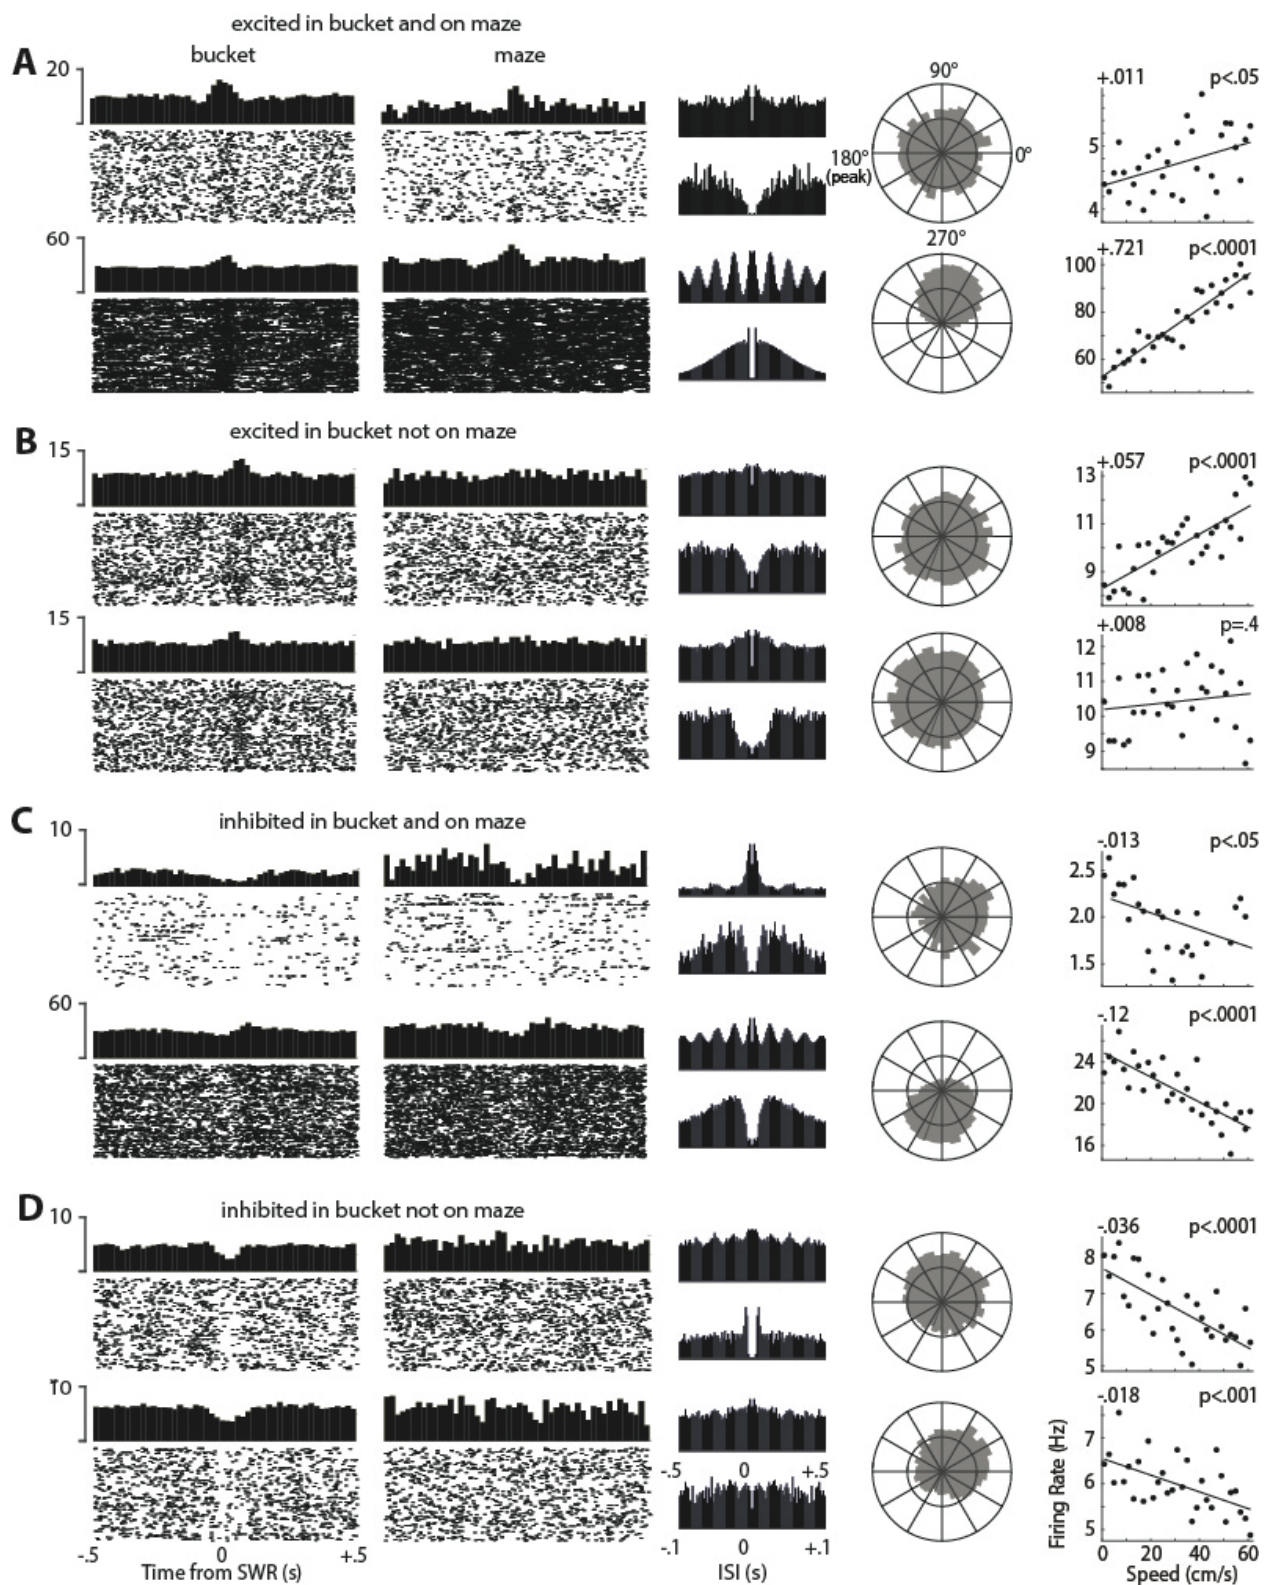


Figure 2-5. *Example cells recorded in striatum.* Graphs are the same as described in Figure 2-4. Examples are shown for striatal neurons that were excited by SWRs in both the bucket and the maze **(A)**, excited by SWRs in the bucket but not the maze **(B)**, inhibited by SWRs in both the bucket and the maze **(C)**, and inhibited by SWRs in the bucket but not on the maze **(D)**.

2.2.2 Predictors of Ripple Responsiveness

2.2.2.1 Coherence with Hippocampal Theta Rhythm

To quantify coherence of each neuron's spike train with hippocampal theta rhythm, we measured the phase of theta rhythm at which individual spikes occurred, and then plotted the circular distribution of phases over all spikes generated by the neuron during active movement on the maze (running speed >10 cm/s). As in prior analyses (see above), we retain the convention that the valley and peak of theta rhythm occur at phases of 0° and 180°, respectively. For population analyses, the length of the phase distribution's resultant vector (which ranged between 0 and 1) was taken as a measure of each cell's spike coherence with hippocampal theta rhythm, where values near 0 indicate low coherence, and values near 1 indicate high coherence. It should be noted that coherence does not measure theta rhythmicity in a cell's spike train, but instead measures the tendency of the cell to spike at a specific phase of the hippocampal LFP. That is, coherence measures periodicity in the cross-correlation between the spike train and the LFP, not in the auto-correlation of the spike train with itself, and it is possible to have one without the other (Zeitler et al., 2006). In some cases, neurons that exhibited strong spike coherence with hippocampal theta rhythm also exhibited strong theta rhythmicity of their own spike trains (for example, see the striatal unit in the bottom panel of Fig. 2-5A). But in other cases, neurons exhibited strong coherence with hippocampal theta rhythm

while exhibiting very little evidence of theta rhythmicity in their own spike train (for example, see the septal unit in the top panel of Fig. 2-4D).

To compare theta coherence of different neural populations, parametric statistical tests were performed upon the \log_{10} of the resultant lengths (rather than raw resultant lengths), because \log_{10} resultants were normally distributed whereas raw resultants were not. A 3-way independent ANOVA revealed that \log_{10} resultant lengths differed significantly for neurons recorded in hippocampus, septum, and striatum ($F_{2,819}=71.47$, $p=2.5e-29$). Post-hoc comparisons revealed that \log_{10} resultants were not very different for hippocampal versus septal neurons ($t_{441}=2.16$, $p=.03$; not significant at the .05 level after Bonferroni correction), reflecting the fact that both hippocampal and septal neurons were strongly coherent with hippocampal theta. However, striatal neurons exhibited much less theta coherence than either hippocampal neurons ($t_{591}=10.81$, $p=5.4e-5$) or septal neurons ($t_{602}=8.75$, $p=2.1e-17$).

For analyses of individual cells, a neuron was classified as “theta coherent” if a Rayleigh test for circular non-uniformity of its phase distribution yielded $p<.01$. The preferred firing phase for cells meeting this theta coherence criterion was then measured as the circular mean of the phase distribution (see below). This criterion classified 87% (188/216) of all hippocampal neurons as theta coherent, whereas 65.9% (149/226) of septal neurons and 25.1% (95/378) of striatal cells were classified as theta coherent. To further investigate whether any relationship existed between theta coherence and SWR responsiveness, additional analyses were carried out on cell populations from each area.

2.2.2.1.1 Hippocampus

A 3-way independent ANOVA (Fig. 2- 6, left) revealed that \log_{10} resultant lengths differed for hippocampal neurons that were excited by, inhibited by, or non-responsive to SWR events ($F_{2,215}=3.08$, $p=.0478$). Post-hoc comparisons indicated that SWR-excited cells exhibited significantly larger \log_{10} resultants (and thus, more theta coherence) than non-SWR responsive cells ($t_{204}=2.4$, $p=.0174$). By contrast, SWR-inhibited cells did not exhibit significantly different theta coherence from non-SWR responsive cells ($t_{77}=1.22$, $p=.22$) or from SWR-excited cells ($t_{145}=.54$, $p=.58$). However, the sample size of SWR-inhibited neurons was quite small ($n=10$), so statistical comparisons of these neurons against other populations may have been underpowered. Consistent with these population analysis results, Fig. 2- 7A (left column) shows that the proportion of individual SWR-excited cells that were theta coherent was 92.7% (127/137), and the proportion of SWR-inhibited cells that were theta coherent was similarly high at 90% (9/10). By contrast, only 75.3% (52/69) of non-SWR responsive cells were theta-coherent. A 2x2 chi-square test (coherent vs non-coherent, SWR-responsive vs non-responsive) indicated that for septum cells, a significantly greater proportion of SWR-responsive than non-responsive cells were theta coherent, ($1, N=216$)= 12.24 , $p<.001$.

We next analyzed the phase of theta rhythm at which theta-coherent neurons fired within each individual rat. Rayleigh tests revealed that the preferred phases of theta coherent hippocampal neurons that were excited by SWRs (Fig. 2-7A, top row) were non-uniformly distributed in all three rats (Rat 1: $Z_{73}=6.3$, $p=.0017$; Rat 2: $Z_{28}=16.6$, $p=2.8e-9$; Rat 3: $Z_{26}=5.2$, $p=.0047$). Moreover, the mean preferred phases of SWR-excited cells concentrated near the peak of theta at 180° (mean phase for Rat 1: 176.4° , Rat 2: 220.4° , Rat 3: 226.1°). Circular V-tests




Figure 2-6. *Theta coherence and SWR responsiveness.* Graphs show mean \log_{10} resultant lengths of preferred spike phases for cells that were excited (exc), inhibited (inh), or non-responsive (NR) to SWR events in hippocampus (left), septum (middle), and striatum (right). Symbols: † $p < .1$, * $p < .05$, ** $p < .01$, **** $p < .0001$.

revealed that the mean preferred phase of SWR-excited hippocampal neurons was not distinguishable from 180° for any of the rats (Rat 1: $V_{73}=21.4$, $p=2.0e-4$; Rat 2: $V_{28}=16.4$, $p=5.7e-6$, Rat 3: $V_{26}=8.0$, $p=.0128$). Hence, SWR-excited hippocampal cells tended to fire near the peak of theta rhythm in all three rats. Theta coherent hippocampal cells that were inhibited by SWRs (Fig. 2-7A, middle row) were recorded from 2 of the 3 rats, and also exhibited mean preferred phases near the peak of theta rhythm (Rat 1: 154.8° , Rat 2: 147.7°). Sample sizes of SWR-inhibited hippocampal cells were too small in each rat to perform adequately powered within-animal Rayleigh or V tests (but see below for pooled analysis across all rats). Theta coherent hippocampal neurons that did not respond to SWRs were recorded from all three rats, and again exhibited mean preferred phases near the peak of theta rhythm (Rat 1: 151.2° , Rat 2: 181.5° , Rat 3: 195.7°). Circular statistics were only performed on non-responsive cell data from

Rats 1 and 2, since only 2 non-responsive hippocampal cells were recorded from Rat 3.

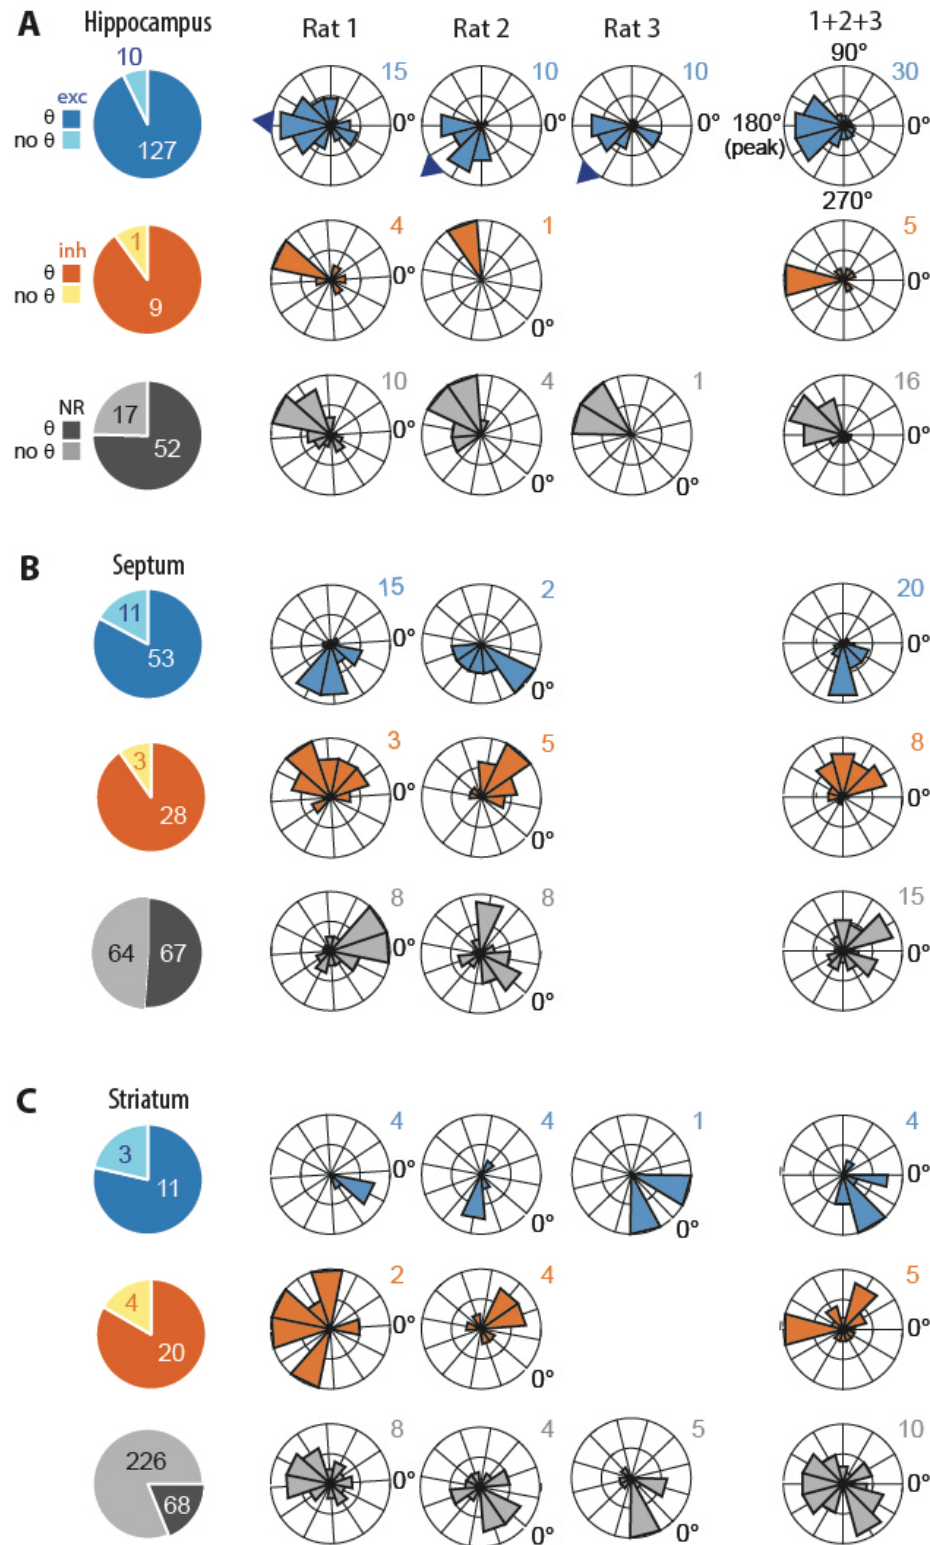


Figure 2-7. Preferred theta phase and SWR responsiveness. A) Pie charts at left show proportions of hippocampal neurons that were excited (exc), inhibited (inh), or non-responsive (NR) to SWR events that spiked coherently (c) or non-coherently (nc) with LFP theta rhythm. Polar plots at right show circular distributions of preferred firing phase for c-coherent cells in each individual rat (middle columns), and pooled across rats (right column). Arrows at rims of polar plots for individual rats in the top row show mean theta phase of SWR-excited hippocampal cells in each rat. These means were used as a correction angle to align spike phases prior to averaging across rats (see main text). Below the top row, polar plots for individual rats in each column are rotated by the correction angle for that rat, to show how spike phases were aligned with one another prior to being pooled in the rightmost column. **B,C)** Same as 'A' for neurons in septum (**B**) and striatum (**C**).

Preferred phases of non-responsive neurons were non-uniformly distributed in Rat 1 ($Z_{42}=8.42$, $p=1.5e-4$) and Rat 2 ($Z_{16}=6.81$, $p=5.8e-4$), and the mean cluster phase was not distinguishable from 180° in Rat 1 ($V_{37}=15.4$, $p=1.7e-4$) or Rat 2 ($V_{13}=10.9$, $p=9.6e-6$).

The above results from within-rat analyses indicated that most hippocampal neurons were theta coherent, and most tended to fire near the peak of theta rhythm (regardless of whether they responded to SWRs). There was some variability in the precise phase at which cells preferred to fire in different rats, and since the phase of the theta LFP is known to vary at different recording locations within the hippocampus (Lubenov & Siapas, 2009; Patel et al., 2012), some of the between-rat variability in mean preferred firing phases may have resulted from differences in the precise locations at which theta LFP reference electrodes were positioned in each rat. To compensate for this, phase preference data was corrected to a new reference frame before being pooled across rats. The preferred phase of each neuron from a given rat was corrected by an angle equal to the mean phase at which SWR-excited hippocampal cells fired in that rat, plus 180° (to retain the convention that the peak of theta falls at 180°). In this corrected reference frame, cells that tended to fire in phase with SWR-excited

hippocampal cells had a phase preference of 180° , and cells that tended to fire in antiphase with SWR-excited hippocampal cells had a phase preference of 0° . When pooled across rats, SWR-excited hippocampal neurons exhibited significant clustering of their corrected phases near 180° ($V_{127}=54.6$, $p=3.6e-12$), but of course this is entirely expected, since the mean phase of SWR-excited cells was used to correct the phases in each rat. SWR-inhibited neurons also showed a trend to cluster near 180° ($V_9=3.1$, $p=.07$), even though the phase of SWR-excited cells was used for correction. However, this analysis was underpowered because of the small sample size for SWR-inhibited hippocampal neurons ($n=9$). Theta coherent hippocampal neurons that were non-responsive to SWRs also exhibited significant clustering of their corrected phases near 180° ($V_{52}=26.1$, $p=1.5e-7$), so these cells tended to fire near the same phase as SWR-excited cells.

In summary, the firing phases of hippocampal neurons tended to cluster near the peak of the LFP theta rhythm, regardless of how the cells responded to SWR events. Since more than half of hippocampal neurons in each rat were excited by SWRs, and since the preferred firing phase of SWR-excited neurons tended to cluster near a common LFP phase in each rat, we used the mean phase of hippocampal SWR-excited cells in each rat to correct the LFP phases of cells recorded in other areas (septum and striatum) into a common angular reference frame. In analyses presented below, we shall use the term “corrected phase” to refer to firing phases that have been shifted by a rat’s mean phase of SWR-excited hippocampal neurons, whereas “uncorrected phase” shall refer to the raw spike phase measured against each rat’s unshifted LFP.

2.2.2.1.2 Lateral Septum

A 3-way independent ANOVA (Fig. 2-6, middle) revealed that \log_{10} resultant lengths differed significantly for septal neurons that were excited by, inhibited by, or non-responsive to SWR events ($F_{2,225}=12.53$, $p=6.9e-6$). Post-hoc comparisons revealed that non-SWR responsive cells exhibited significantly smaller \log_{10} resultants (and thus, less theta coherence) than either SWR-excited cells ($t_{194}=4.31$, $p=2.6e-5$) or SWR-inhibited cells ($t_{161}=3.49$, $p=6.2e-4$), but SWR-excited cells did not exhibit significantly different resultant lengths from SWR-inhibited cells ($t_{93}=.02$, $p=.98$). Hence, on average, SWR-excited and SWR-inhibited septal neurons both exhibited significantly greater theta coherence than non-SWR responsive septal neurons. Consistent with these population analysis results, Fig. 2-7B (left column) shows that the proportion of individual SWR-excited septal cells that were theta coherent was 82.8% (53/64), and the proportion of SWR-inhibited cells that were theta coherent was similarly high at 90.3% (28/31). By contrast, only 51.2% (67/131) of non-SWR responsive cells were theta-coherent. A 2x2 chi-square test on these proportions indicated that theta coherence of septum neurons was significantly contingent upon whether or not the cell was responsive to (that is, either excited or inhibited by) SWRs, ($1, N=226$)= 35.9 , $p<.00001$.

The uncorrected mean preferred phases of SWR-excited septal neurons (Fig. 2-7B, top row) in Rat 1 and Rat 2 were 273° and 316° , respectively. Rayleigh tests revealed that the preferred phases of theta coherent SWR-excited neurons were non-uniformly distributed in Rat 1 ($Z_{47}=20.9$, $p=5.4e-11$) and Rat 2 ($Z_6=2.6$, $p=.064$). When phase values from Rat 1 and Rat 2 were corrected by the mean phase of SWR-excited hippocampal cells in each rat (see Methods) and then pooled across rats, the mean phase at which SWR-excited septal neurons fired was 277° , which is $+97^\circ$ from the 180° peak of hippocampal theta. Circular V-tests revealed that in both individual rats, the mean preferred phase of SWR-excited septal neurons was not distinguishable from being shifted by $+90^\circ$ relative to the peak of hippocampal SWR-excited

cells (Rat 1: $V_{47}=31.3$, $p=5.3e-11$; Rat 2: $V_6=3.9$, $p=.0122$), and this remained true when data was pooled across both rats ($V_{53}=35.1$, $p=4.8e-12$).

The preferred phases of SWR-inhibited neurons (Fig. 2-7B, middle row) were also non-uniformly distributed in both rats (Rat 1: $Z_{13}=4.8$, $p=.006$; Rat 2: $Z_{15}=-8.2$, $p=8.8e-5$). The mean uncorrected phases of SWR-inhibited septal neurons in Rat 1 and Rat 2 were 92.8° and 95.3° , respectively. When phase data from Rat 1 and Rat 2 were corrected to the hippocampal reference frame and pooled together, the mean phase at which SWR-inhibited septal neurons fired was 70° , or -110° from the 180° peak of hippocampal theta. Circular V-tests revealed that in both individual rats, the mean preferred phase of SWR-inhibited septal neurons was not distinguishable from being shifted by -90° relative to the peak of hippocampal theta (Rat 1: $V_{13}=7.9$, $p=9.6e-4$; Rat 2: $V_{15}=11.1$, $p=2.7e-5$), and this remained true when data was pooled across both rats ($V_{28}=14.34$, $p=4.8e-5$).

The preferred phases of neurons that were non-responsive to SWRs (Fig. 2-7B, bottom row) were non-uniformly distributed in Rat 1 ($Z_{35}=7.0$, $p=6.9e-4$), and trended toward non-uniformity in Rat 2 ($Z_{33}=-2.4$, $p=.09$). The mean uncorrected phases of non-SWR responsive septal neurons in Rat 1 and Rat 2 were 345.7° and 356.4° , respectively. When phase data from Rat 1 and Rat 2 were corrected into the hippocampal reference frame and pooled together, the mean phase at which non-responsive neurons fired was 355.9° , or $+176^\circ$ from the peak of hippocampal theta. Circular V-tests revealed that in both individual rats, the mean preferred phase of SWR-inhibited septal neurons was not distinguishable from being in antiphase with hippocampal theta (Rat 1: $V_{35}=15.1$, $p=.0025$; Rat 2: $V_{33}=14.9$, $p=.0073$), and the same was true of the pooled data ($V_{68}=24.5$, $p=1.3e-5$).

In summary, a majority of all septal neurons were coherent with hippocampal theta rhythm. Similar to the hippocampus, about 80-90% of SWR-responsive cells were theta coherent, whereas only half of non-SWR responsive cells were theta coherent. But unlike in the

hippocampus, SWR-excited versus inhibited neurons in septum tended to fire at distinct phases of hippocampal theta. SWR-excited cells clustered at +90° and SWR-inhibited cells clustered at -90° from the peak of hippocampal theta. Theta coherent cells that did not respond to SWRs clustered near the 0° valley of hippocampal theta.

2.2.2.1.3 Striatum

A 3-way ANOVA (Fig. 2-6, right) revealed that \log_{10} resultant lengths differed significantly for striatal neurons that were excited by, inhibited by, or non-responsive to SWR events ($F_{2,376}=11.23$, $p=1.8e-5$). Post-hoc comparisons revealed that non-SWR responsive striatal cells exhibited significantly smaller \log_{10} resultants (and thus, less theta coherence) than either SWR-excited cells ($t_{351}=3.07$, $p=.0023$) or SWR-inhibited cells ($t_{361}=3.79$, $p=1.7e-4$), but SWR-excited cells did not exhibit significantly different resultant lengths from SWR-inhibited cells ($t_{36}=.04$, $p=.96$). Hence, on average, SWR-excited and SWR-inhibited striatal neurons both exhibited significantly greater theta coherence than non-SWR responsive striatal neurons. Consistent with these population analysis results, we found that the proportion of individual SWR-excited cells (Fig. 2-7C, left column) that were theta coherent was 78.6% (11/14), and the proportion of SWR-inhibited cells that were theta coherent was similarly high at 83.3% (20/24). By contrast, only 18.9% (64/340) of non-SWR responsive cells were theta-coherent. A 2x2 chi-square test on these proportions indicated that theta coherence of striatal cells was significantly contingent upon whether or not the cell was responsive to SWRs, ($1, N=348$)= 63.3 , $p<.00001$.

Fig. 2-7C (top row) shows that the mean uncorrected phases of SWR-excited striatal neurons were somewhat similar in all three rats: 311.0° (Rat 1), 313.4° (Rat 2), and 4.0° (Rat 3). The sample size of SWR-excited neurons in the striatum each rat was too small to perform within-rat analysis of preferred phase distributions, but when phase data from all three rats was

corrected into the hippocampal reference frame and pooled together, SWR-excited striatal neurons exhibited significant non-uniformity ($Z_{11}=6.0$, $p=.001$) with a mean phase of 304.9° . This suggests that these neurons may have preferred to fire just prior to the valley of hippocampal theta. SWR-inhibited striatal neurons (Fig. 2-7C, middle row) that were theta coherent were recorded in only two of the three rats. In Rat 1, the preferred phases of these neurons showed only a trend for non-uniformity of their circular distribution ($Z_{10}=2.5$, $p=.08$), with a mean preferred phase of 150.8° . In Rat 2, the preferred phases beat significance for non-uniformity on the circle ($Z_{10}=3.2$, $p=.04$), with a mean preferred phase of 85.6° . When phase data from both rats was shifted into the hippocampal reference frame and pooled together, SWR-inhibited septal neurons did not beat significance for non-uniformity of their preferred phases ($Z_{20}=1.3$, $p=.28$). The mean uncorrected phases of theta coherent septal neurons that were not responsive to SWRs (Fig. 2-7C, bottom row) were 127.8° (Rat 1), 301.6° (Rat 2), and 291.7° (Rat 3). Preferred phases of non-responsive neurons were non-uniformly distributed in Rat 3 ($Z_{12}=4.0$, $p=.0154$), but not in Rat 1 ($Z_{37}=2.3$, $p=.1$) or Rat 2 ($Z_{15}=.82$, $p=.44$). When phase data from all three rats was corrected into the hippocampal reference frame and pooled together, the preferred phases of non-SWR responsive striatal neurons did not beat significance for non-uniformity of their preferred phases ($Z_{64}=.60$, $p=.55$).

In summary, a majority (about 80%) SWR-responsive striatal neurons spiked coherently with hippocampal theta rhythm, even though a minority (about one fourth) of all striatal neurons were theta coherent. Hence, theta coherence was significantly more common among SWR-responsive than non-responsive striatal neurons. SWR-excited striatal neurons showed a tendency to fire prior to the valley of hippocampal theta, but SWR-inhibited and non-SWR responsive striatal neurons did not exhibit significant tendencies to fire at a specific preferred phase of theta.

2.2.2.2 Speed Sensitivity

To analyze modulation of neural firing rates by running speed, we performed a linear regression analysis upon plots of firing rate versus running speed for each recorded cell (see Methods). The slope of the regression line (in units of Hz/cm/s) was taken as an estimate for the slope of speed modulation, and the intercept of the regression line (in units of Hz) was taken as an estimate for the cell's mean firing rate during stillness.

2.2.2.2.1 Hippocampus

Less than half of the neurons recorded in the hippocampus (85/216, or 39.3%) met criterion for inclusion in the analysis of speed modulation. This was because many hippocampal neurons were spatially tuned (data not shown), and cells that fired selectively at a specific location often did not fire across a wide enough range of running speeds to meet criterion for inclusion in the speed analysis. We found that 65/85 (76.5%) of eligible hippocampal neurons exhibited a significant linear correlation ($p < .05$) of their firing rates with running speed, and of these, 62/65 (95%) were positively and 3/65 (5%) were negatively correlated with speed (Fig. 2-8A, left). When the sign of the SWR response was ignored, a 2x2 chi-square test found no contingency between sharp wave responsiveness (responsive vs non-responsive) and speed modulation (modulated vs unmodulated), $(1, N=85) = .47$, $p = .49$. Hence, hippocampal cells that responded to SWRs were no more or less likely to be speed modulated than cells that did not respond to SWRs. Of the cells that were positively correlated with running speed, 39/65 (60%) were excited by SWRs and 1/65 (1.5%, a single cell) were inhibited by SWRs. For the subset of hippocampal cells that were both speed modulated and SWR-responsive (Fig. 2-8A, right), the mean slope of speed modulation was positive ($.078 \pm .014$ Hz/cm/s) and significantly greater than zero ($Z_{39} = 5.52$, $p = 3.3e-8$). In summary, a considerable majority of the hippocampal cells that

were eligible for speed analysis had firing rates that were positively correlated with running speed, regardless of whether they were responsive to SWRs.

2.2.2.2.2 Lateral Septum

Of the 226 neurons recorded in septum, 198 (87.6%) met criterion for inclusion in the analysis of speed modulation. We found that 95/198 (48%) of these neurons exhibited a significant linear correlation ($p < .05$) of their firing rates with running speed. Of the septal neurons that were speed modulated, 46/95 (48%) were positively and 49/95 (52%) were negatively correlated with running speed (Fig. 2-8B, left). Hence, speed-modulated neurons in septum were almost perfectly split in half between positively and negatively modulated cells, and this was true in both individual rats (rat #1: 26 positive, 25 negative; rat #2: 20 positive, 24 negative). When the sign of the SWR response was ignored, a 2x2 chi-square test found a significant contingency between sharp wave responsiveness (responsive vs non-responsive) and speed modulation (modulated vs unmodulated), ($1, N=198$)=10.4, $p = .0013$. Hence, septal cells that responded to SWRs were more likely to be speed modulated than cells that did not respond to SWRs, in accordance with prior results (Wirtshafter & Wilson, 2019). A 3x2 chi-square test found a strong contingency between sharp wave responsiveness (excited, inhibited, or non-response to SWR) and speed modulation (positive vs negative), $z(2, N=95) = 29.2$, $p < .00001$. When the sign of the SWR response was ignored, a 2x2 chi-square test found no contingency at all between sharp wave responsiveness (responsive vs non-responsive) and sign of speed modulation (positive vs negative), ($1, N=95$)=.02, $p = .88$; hence, the sign of speed modulation did not depend upon whether or not a cell was responsive to SWRs. However, a chi-square test on just the SWR-responsive cells found that the sign of the SWR response was highly contingent upon the sign of speed modulation, ($1, N=55$)=29.1, $p < .00001$.

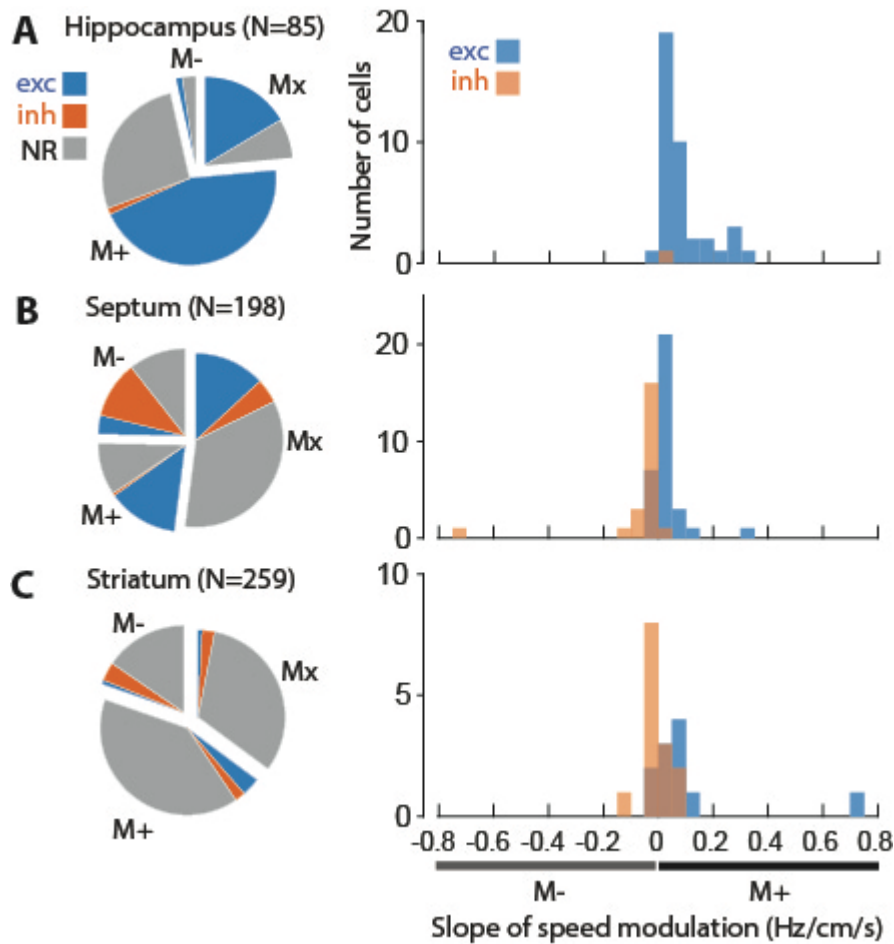


Figure 2-8. Speed modulation and SWR responsiveness. **A)** Pie chart shows proportions of hippocampal neurons that were eligible for speed analysis (N=85) that were positively (M+), negatively (M-), or not significantly (Mx) modulated by running speed. Within each speed classification, shading of wedges indicates proportions of cells that were excited (exc), inhibited (inh), or non-responsive (NR) to SWR events. Histogram at right shows the distribution of speed slopes for SWR-excited and SWR-inhibited cells that were significantly modulated by running speed. **B,C)** Same as 'A' for neurons in septum (**B**) and striatum (**C**).

Indeed, only 1/27 (3.7%) of SWR-responsive septal cells with positive speed modulation slopes were inhibited by SWRs (the rest were excited by SWRs), and only 7/28 (25%) of SWR-responsive septal cells with negative speed modulation slopes were excited by SWRs (the rest were inhibited by SWRs). Consistent with this chi-square analysis, among cells that were both SWR responsive and speed modulated, the average slope of speed modulation for SWR-excited cells was significantly greater than zero (mean $.028 \pm .002$ Hz/cm/s; $Z_{32}=2.41$, $p=.016$), whereas the average slope for SWR-inhibited cells was significantly less than zero (mean $-.06 \pm .007$ Hz/cm/s; $Z_{21}=-1.88$, $p=.06$). Moreover, an independent t-test revealed that slopes of speed modulation for SWR-excited septal neurons were significantly more positive than slopes for SWR-inhibited septal neurons ($t_{53}=2.98$, $p=.0044$). To make sure that this significant result did not arise solely from two outlying slopes with large values (see Fig. 2-8B, right), the t-test was re-run without these two outlying values, and despite a smaller difference between means, the result became even more statistically significant ($t_{51}=5.18$, $p=3.8e-6$) because variance was reduced by eliminating the outliers.

These results show that among the subset of septal neurons that were both SWR responsive and speed modulated, cells with positive speed slopes were almost always excited by SWRs, and a considerable majority (about 75%) of cells with negative speed slopes were inhibited by SWRs. One possible confound is that this result might arise from a statistical power artifact, because SWRs occurred during stillness, and therefore, cells with negative speed slopes might tend to fire at a higher rate during stillness than cells with positive speed slopes. If so, then this may confer greater statistical power to detect SWR-induced inhibition of cells with negative speed slopes against their higher background firing rates during stillness, and greater statistical power to detect SWR-induced excitation of cells with positive speed slopes against their lower background firing rates during stillness. However, the median firing rate during stillness (estimated as the y-intercept of the speed slope line) did not differ (rank sum test

Z=0.682, $p=.5$) for cells with positive (median 4.4 Hz) versus negative (median 3.8 Hz) speed slopes, nor did it differ (rank sum test $Z=1.33$, $p=.18$) for SWR-excited cells (median 4.6 Hz) versus SWR-inhibited cells (median 3.3 Hz). Hence, the correlation between the sign of the speed slope and the sign of the SWR response was unlikely to be a statistical power artifact.

2.2.2.2.3 Striatum

Of the 378 neurons recorded in striatum, 259 (68.5%) met criterion for inclusion in the analysis of speed modulation. We found that 168/259 (64.8%) of these neurons exhibited a significant linear correlation ($p<.05$) of their firing rates with running speed, and of these, 117/168 (69.6%) had a positive and 51/168 (30.4%) had a negative slope of speed modulation (Fig 8C, left). Hence, in striatum, positively modulated cells were about twice as common as negatively modulated cells, although percentages varied somewhat across the three individual rats (rat #1: 32 positive, 33 negative; rat #2: 66 positive, 18 negative; rat #2: 19 positive, 0 negative). A 3x2 chi-square test found a significant contingency between sharp wave responsiveness and the sign of speed modulation for striatal neurons, $(2, N=168)=9.32$, $p=.0094$. When the sign of the SWR response was ignored, a 2x2 chi-square test found only a weak trend for contingency between sharp wave responsiveness and sign of speed modulation, $(1, N=168)=2.31$, $p=.13$; this trend indicated a modest tendency for negatively modulated speed cells to be more common among SWR-responsive than non-responsive cells in striatum. A chi-square test on just the SWR-responsive cells found that the sign of the SWR response was contingent upon the sign of speed modulation, $(1, N=26)=6.0$, $p=.014$. Indeed, only 2/11 (18.2%) of SWR-responsive striatal cells with negative speed modulation slopes were excited by SWRs (the rest were inhibited by SWRs), and 5/15 (33.3%) of SWR-responsive cells with positive speed modulation slopes were inhibited by SWRs (the rest were excited by SWRs). Among cells that were both SWR responsive and speed modulated, the mean slope of speed

modulation was positive for SWR-excited cells ($.101 \pm .019$ Hz/cm/s), but not significantly greater than zero ($Z_{32}=1.59$, $p=.11$). The mean slope of speed modulation was negative for SWR-inhibited cells ($-.009 \pm .003$ Hz/cm/s), not significantly less than zero ($Z_{32}=-0.74$, $p=.45$). However, an independent t-test revealed a trend for SWR-excited striatal neurons to have more positive speed slopes than SWR-inhibited neurons ($t_{23}=1.91$, $p=.069$), and this effect reached significance ($t_{22}=2.63$, $p=.015$) when variance was reduced by eliminating one extreme outlying slope which, despite being an outlier, did fit the overall trend for SWR-excited cells to have positive speed slopes (see Fig. 2-8C, right).

Taken together, these results show that among striatal neurons that were both SWR responsive and speed modulated, positively modulated speed cells tended to be excited by SWRs, and negatively modulated speed cells tended to be inhibited by SWRs. This correlation between speed slope and SWR responsiveness was less pronounced for striatal neurons than for septal neurons, but was nonetheless evident in both structures.

2.3. DISCUSSION

A growing body of evidence suggests that hippocampal projections to the septum may be an important route via which the hippocampus relays information to the midbrain and other subcortical regions to exert influence over behaviors such as reward-seeking, motor actions, reinforcement learning, and decision making (Luo et al., 2011; Gomperts et al., 2015; Tingley & Buzsáki 2018; Wirtshafter & Wilson 2019). Prior studies have demonstrated that septal neurons can encode an animal's position in their firing rates (Takamura et al., 2006) as well as their spike phases (Tingley & Buzsáki 2018). Septal projections to the midbrain may thus relay position information from the hippocampus to dopaminergic and hypothalamic circuits that attach motivational value to specific spatial locations and environmental states (Luo et al., 2011;

Gomperts et al., 2015; Wirtshafter & Wilson 2019; Tingley & Buzsáki 2020). Value and prediction error signals computed in the midbrain might then be relayed to cortical and striatal regions that govern learning, memory, and decision making, thereby allowing reinforcement learning processes to be influenced by hippocampally encoded locations and states.

Compressed replay events that occur during SWRs have been hypothesized to play three distinct but related roles in reinforcement learning. First, it has been proposed that during navigation, forward replay of alternative future trajectories supports deliberation over the best path for the animal to take from its current location (Johnson and Redish, 2007; Pfeiffer & Foster, 2013; Yu & Frank, 2015; Wu et al., 2017; Kay et al., 2020). Second, it has been proposed that when reward outcomes are obtained, compressed replay of prior trajectories that have been traversed in the recent past may help to solve the “credit assignment” problem in reinforcement learning, which is the problem of assigning credit or blame for reward outcomes to decisions that were made in the remote past, before the outcome was obtained (Foster & Wilson, 2006). Third, it has been proposed that during sleep, compressed replay during SWRs may be necessary for consolidating short-term memories of recent experiences to long-term storage (Wilson & McNaughton, 1944; Buzsáki, 1996; Ego-Stengel & Wilson, 2010; Girardeau & Zugaro, 2011).

The septal output pathway from the hippocampus could play an important role in all three of these hypothesized functions for SWR events. Previous reports have demonstrated SWR-evoked responses in subpopulations of septal neurons (Wirtshafter & Wilson 2019; Tingley & Buzsáki 2020). It has also been shown that SWRs are sometimes accompanied by activation of midbrain neurons that respond to reward (Gomperts et al., 2015), as might be expected if dopamine circuits are computing value or prediction errors signals derived from replay of navigational trajectories during SWRs. Here, we recorded single units in the hippocampus, septum, and striatum while freely behaving rats ran trials in a T-maze task and

rested in a holding bucket between trials. A large proportion of hippocampal neurons were excited during SWRs, as reported in prior studies (Wilson and McNaughton, 1994; Skaggs & McNaughton, 1996; Kudrimoti et al., 1999; Foster & Wilson, 2006; Davidson et al., 2009). We also identified several novel properties of SWR-responsive neurons in septum and striatum.

2.3.1 Activity of Septal Neurons during SWR Events

In agreement with prior findings (Wirtshafter & Wilson 2019), we observed that SWR-responsive septal neurons tended to fire coherently with hippocampal theta rhythm during periods of locomotion. However, we also found that spikes of SWR-responsive septal neurons were segregated in time across the theta cycle, in such a way that SWR-excited neurons fired late in the cycle, whereas SWR-inhibited neurons fired early in the cycle. It has previously been shown that some septal neurons exhibit spatial phase precession against the hippocampal LFP as a rat runs on a maze (Tingley & Buzsáki 2018), but to our knowledge, it has not been previously reported that the valence (excitation versus inhibition) of a septal neuron's SWR response during stillness is predictive of its preferred firing phase during theta rhythm when the animal is moving, as we found here.

In further agreement with prior findings (Wirtshafter & Wilson 2019), we observed that some SWR-responsive septal neurons behaved as speed cells, since their firing rates were positively or negatively modulated by the rat's running speed. But interestingly, we also found (for the first time, as far as we know) that SWR-excited septal neurons tended to show positively sloped modulation of their firing rates by running speed, whereas SWR-inhibited septal neurons showed negatively sloped modulation of their firing rates by running speed. Taken together, our findings suggest there may be two distinct types of SWR-responsive neurons in septum: SWR-excited cells, which fire late in the hippocampal theta cycle and are biased to show positive

modulation of their firing rates by running speed, and SWR-inhibited cells, which fire early in the theta cycle and are biased to show negative modulation of their firing rates by running speed.

It is interesting to consider how the spiking of SWR-excited versus inhibited septal neurons may align with the timing of place cell spikes in the hippocampus. As an animal passes through a place cell's preferred firing location (or place field), the place cell bursts rhythmically at a slightly higher frequency than the LFP theta frequency, causing spikes to exhibit *phase precession* against the LFP (O'Keefe & Recce, 1993). At the population level, phase precession segregates place cell spikes in time, so that cells with place fields that lie ahead of the animal's current location fire at late phases of LFP theta, whereas cells with place fields behind the animal's current location fire early phases of LFP theta (Skaggs et al., 1996; Dragoi & Buzsáki, 2006; Wikenheiser & Redish, 2013). Phase coding of spatial locations occurs in the lateral septum as well as the hippocampus (Tingley & Buzsáki, 2018; Monaco et al., 2019). Here, we found that SWR-excited septal neurons (which tend to be positively correlated with running speed) fired late in the theta cycle, so they presumably fired together with place cells that encoded locations ahead of the animal along its current motor trajectory. Conversely, we found that SWR-inhibited septal neurons (which tend to be negatively correlated with running speed) fired early in the theta cycle, so they presumably fired together with place cells that encoded locations behind of the animal along its current motor trajectory. Speed cells in the septum thus appear to have firing rates that are positively correlated with whichever behavior (running versus stopping) is most appropriate for reaching spatial locations that are encoded by place cells that co-fire with the speed cell during the theta cycle. This suggests that there may be a phase code for motor actions in septum that complements phase coding for position: speed cells that are positively correlated with movement (and excited during SWRs) may fire in phase with place cells whose preferred locations are reachable via continued movement, whereas speed cells

that are negatively correlated with movement (and inhibited during SWRs) may fire in phase with place cells whose preferred locations are reachable via cessation of movement.

2.3.2 Activity of Striatal Neurons during SWR Events

In the present study, neurons were recorded from both ventral and dorsal striatum. A prior study has shown that ventral striatal neurons exhibit phasic responses during dorsal hippocampal SWRs (Sosa et al., 2020), and we observed similar SWR responses in ventral striatal neurons. We also observed SWR responses in dorsal striatum, which to our knowledge has not been reported before. Previous single-unit recording studies have reported that firing rates of striatal neurons in rodents are correlated with the animal's running speed (Ruede-Orozco & Robbe, 2015), and we similarly observed that a subset of SWR-responsive striatal neurons were modulated by running speed. Interestingly, SWR-excited striatal neurons tended to show positively sloped modulation of their firing rates by running speed, whereas SWR-inhibited striatal neurons showed negatively sloped modulation of their firing rates by running speed, similar to the results we obtained in septum.

Projection cells from the striatum are GABAergic medium spiny neurons (MSNs), which can be broadly subdivided into two main classes expressing D1 versus D2-type dopamine receptors. Classical models of the basal ganglia posit that D1 MSNs are the origin of a “direct” striatonigral motor output pathway which excites motor behavior, whereas D2 MSNs are the origin of an “indirect” striatopallidal motor output pathway which inhibits motor behavior. It would be worthwhile in future studies to investigate whether a significant proportion of SWR-responsive striatal neurons are MSNs, and if so, how excitatory versus inhibitory responses during SWRs are distributed among D1 versus D2 subtypes of MSNs. Neural recording and imaging studies have consistently failed to find evidence that D1 and D2 MSNs behave simply

as motor-on and motor-off cells, as classical models would predict. Instead, both types of MSNs seem to fire together during initiation and execution of voluntary motor behaviors (Cui et al., 2013; Isomura et al., 2013), and combined with other evidence, these findings have led to speculation that D1 MSNs may help to drive the execution of selected actions, while D2 MSNs may simultaneously inhibit the execution of competing non-selected actions (Tecuapetla et al., 2016). It has been hypothesized that compressed replay by place cells during SWRs might provide a mechanism for animals to “deliberate” over decisions about which actions to select, and which actions to suppress (Yu & Frank, 2015). If so, then this could be regarded as tantamount to sorting out which actions should be excited by the D1 population and which should be suppressed by the D2 population during an impending motor decision. Consistent with this idea, it has been reported that reward-responsive midbrain dopamine neurons tend to fire synchronously with SWRs during wakeful stillness on a maze (but not during sleep), as might be expected if the animal were assessing the values of potential action plans during SWRs that occur on the maze (Gomperts et al., 2015).

Here, we recorded SWR events during stillness while overt motor actions were not being performed. Hence, one possibility to consider is that the SWR-responsive striatal cells we observed might be MSNs that are involved in inhibiting motor behavior, and thereby preventing actual motor actions from being performed during “virtual” navigation. Another possibility arises from prior evidence that acquisition of maze learning tasks is impaired by disruption of SWRs during both waking and sleep states (Girardeau et al., 2009; Jadhav et al., 2012), suggesting that SWRs may be directly involved in programming specific patterns of action selection that are required to achieve correct performance in such tasks. Since the striatum plays a key role in action selection and behavioral decision making, it could be that the SWR-evoked responses we observed in a small percentage of striatal neurons are reflective of a process by which MSNs become “programmed” to either excite or inhibit specific actions in the future, based upon value

estimates for those actions that are generated during SWRs and compressed replay. This possibility could be further investigated in the future by experiments in which striatal unit activity is selectively disrupted during SWRs.

2.3.3 Summary and Conclusions

SWRs are frequently accompanied by compressed replay of spatial trajectories within hippocampal place cell populations (Skaggs & McNaughton, 1996; Lee & Wilson, 2002; Foster & Wilson, 2006; Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009). Findings presented here support the view that, in addition to being accompanied by hippocampal replay of spatial trajectories, SWR events might also be accompanied by activation of subcortical motor representations (Wirtshafter & Wilson, 2019). Hence, when mental representations of a particular location become active within hippocampal place cell populations—either during an SWR event or during a “theta sequence” driven by phase precession—a corresponding representation of the motor action necessary to reach that location may become concurrently activated within subcortical regions, including septum and striatum. Concurrent activation of hippocampal state representations and subcortical action representations might support neural computations that are essential for reinforcement learning and value-based decision making.

Reinforcement learning theory (Sutton & Barto, 1998) suggests that value-based decision policies can be optimized by attaching values not just to particular states (such as residing at a specific spatial location) or particular actions (such as performing a specific motor behavior), but rather to *state-action pairs* (such as performing a specific action at a specific location). SWRs might therefore support reinforcement learning and decision making by activating representations of spatial trajectories and motor actions at the same time. For example, deliberation over alternative future trajectories during SWRs might not only involve

activating hippocampal representations of spatial locations that lie along those trajectories (Johnson and Redish, 2007; Pfeiffer & Foster, 2013; Yu & Frank, 2015; Kay et al., 2020), but could additionally require activating representations of motor actions that must be performed at each location to adhere to a given trajectory. Similarly, when assigning credit for outcomes to recent behavioral choices (Foster & Wilson, 2006), re-activation of recently traversed trajectories during SWRs may require concurrent re-activation of the motor actions that were performed at each location along the trajectory. Finally, memory consolidation processes that require re-activation of recent experience during sleep (Wilson & McNaughton, 1944; Buzsáki, 1996; Ego-Stengel & Wilson, 2010; Girardeau & Zugaro, 2011) might necessitate concurrent reactivation of recently navigated spatial trajectories as well as motor actions performed along those trajectories, so that memories of decision policies can be consolidated by attaching values not just to states or to actions, but to state-action pairs that have previously yielded positive outcomes during waking experience.

2.4. METHODS

All animal research protocols were reviewed and approved in advance by the UCLA Animal Research Committee, and conducted in accordance with United States federal guidelines. The data that support the findings of this study are openly available at DOI: [10.17632/rg3xjbgyjx.2](https://doi.org/10.17632/rg3xjbgyjx.2).

2.4.1 Subjects and Behavior

2.4.1.1 Subjects

Long-Evans rats (Charles River Laboratories, Hollister, CA, USA) were housed in a temperature and humidity controlled vivarium with a 12-12 reverse light-dark cycle, and fed ad lib until they attained a weight of ~550 grams, after which they were reduced to 85% of their ad lib weight by limited daily feeding. The 3 rats used in the study were selected from a larger cohort of 6 rats that were all trained to perform a Figure 8 maze task prior to surgery. The three rats that were selected for surgery were the first three rats to reach a performance criterion (1 reward per minute over 20 minutes) on the Figure 8 maze.

2.4.1.2 Behavior Apparatus

After recovery from surgery, rats were trained on a T-maze task. The three-arm T-maze was formed by blocking one of the arms on a four-arm plus maze apparatus. Throughout each block of trials, a barrier was placed at the entrance to one of the four arms, while the three remaining arms were assigned as the start, baited, and unbaited arms for the T-maze task. The maze was 218 cm wide with a 30 cm square platform in the center (see Fig. 2-1), located in a 3x3 m room with matte black walls and ceiling. Four 70 cm high posters with distinctive high-contrast black-and-white designs hung on the wall at the end of each arm to provide orienting landmark cues. The room was dimly lit by a 15 W light bulb aimed at the ceiling of the room. The reward was a ~1 g piece of fresh banana. To make sure the rat was not guided by the strong odor of the banana, a dish containing a small amount of banana was always placed underneath the non-baited arm, inaccessible to the rat. Rats spent intertrial intervals in a holding bucket, from which they were not able to observe experimenters placing reward for the next trial.

2.4.1.3 T-Maze Task

Rats were trained to run repeated acquisition and reversal trials on a T-maze (Fig. 2-1). At the start of each session, recording cables were connected and the rat was placed for 5 m in a white plastic bucket located next to the maze (the bucket always remained stationary in the same location, even as the start and goal arms were switched during different trial blocks) for a period of baseline recording. The rat was then placed by the experimenter at the designated start location for the current trial block, where it could immediately begin exploring the maze. The rat was free to run on the maze until it reached the end of either the baited or unbaited arm, at which point the experimenter placed the bucket behind the animal so that it the only available exit from the arm was to walk into the bucket. The rats usually climbed into and out of the bucket voluntarily, minimizing handling stress. The experimenter then placed the bucket in its assigned location on the floor beside the maze for a period of 1-3 m while the maze was cleaned and baited for the next trial. We cleaned the maze after each trial with 70% ethanol, and baited the reward arm for the next trial. When the rat completed 7/8 correct choice trials in a row, the baited and unbaited arms were swapped, and the rat began a reversal learning phase from the same start position. When the reversal criterion of 7/8 correct choice trials was reached, the barrier on the plus maze was moved to a different arm, so that the start, baited, and unbaited arms of the T-maze were reassigned. Another round of acquisition and reversal trials then began with the new maze configuration. This sequence of acquisition, reversal, and maze reconfiguration blocks continued throughout the entire duration of the recording experiment.

2.4.1.4 Video Tracking

The rat's position was sampled at 30 Hz and tracked at a resolution of 2.2 pixels/cm by an overhead color video camera (JVC TK-C1480) outfitted with Tamron 2.8-12mm cctv CS aspherical lenses. The video signal was relayed to a position tracking system built into the electrophysiological data acquisition system (Neuralynx, Bozeman, MT). A custom offline algorithm compensated for lens distortion prior to analyzing the 2D position data.

2.4.2 Surgery, Electrophysiology, and Histology

2.4.2.1 Surgery

Under deep isoflurane anesthesia, each rat was surgically implanted with a skull-mounted microdrive containing an array of 36 independently moveable probes. The 36 probes were grouped into 4 clusters, each consisting of 9 probes (8 tetrodes plus one reference) arranged in a diamond-shaped pattern where individual probes were spaced 400 μ m from their nearest neighbors. Hence, the entire microdrive contained a total of 32 tetrodes and 4 reference wires. Of the 32 tetrodes, 16 were targeted at the dorsal hippocampus (8 per hemisphere), 6 were targeted at the lateral septum (3 per hemisphere), and 10 were targeted at the striatum (5 per hemisphere). In Rats 2 and 3, bilateral skull holes (each \sim 2 mm in diameter) were centered at $-/+3.2$ ML and AP +1.1 (right) for dmStr/LS probe clusters, and at $-/+3.4$ ML and AP -4.5 for CA1 probe clusters. In Rat 1, skull holes were centered at $-/+2.6$ ML and AP +1.3 for dmStr/Nacc probe clusters, and at $-/+2.0$ ML and AP -2.8 for CA1 probe clusters. All coordinates are relative to Bregma. Rats recovered from surgery for at least 10 days before experiments began.

2.4.2.2 Placement of LFP Electrodes

After recovery from surgery, recordings were obtained while rats ran on the T-maze. On the first and second recording day, rats freely explored the maze for 15 minutes with no food rewards to acclimate to the environment. On the third day, rats began the initial acquisition phase of learning on the dual choice T-maze. During this initial training period, hippocampal tetrodes were advanced slowly into the CA1 layer of the hippocampus, until robust SWRs were detectable in the LFP on some of the tetrode wires, and robust 6-8 Hz theta rhythm was detectable on other tetrode wires. Data was not recorded from septal or striatal tetrodes during this initial training period (nor were the tetrodes advanced in these regions). When a hippocampal tetrode with robust SWRs was identified, and another tetrode wire with robust theta was found in the same hemisphere, these two tetrodes were chosen as the ripple and theta recording electrodes, respectively. Neither of these two tetrodes were advanced further during the remainder of the experiment. Starting with the next session, the goal and/or start arm was changed each time the rat achieved a criterion of 7/8 correct responses.

2.4.2.3 Recording Sessions and Tetrode Advancement

Throughout each maze session, a 128 channel DigitalLynx SX data acquisition system (Neuralynx, Bozeman, MT) was used to record LFP signals and single units at a sampling rate of 32 KHz per channel. LFP channels were high pass filtered above 1 Hz, and single-unit channels were bandpass filtered between 600-6000 Hz. Recording sessions varied in duration from 1 to 2 hours. At the end of a recording session, the rat occupied the bucket for 5 minutes before disconnection from the recording system. Tetrodes in septum and striatum were advanced by 165 μm after each session, so that different units would be recorded from these

tetrodes in every session (Henze et al., 2000; Shoham et al.,). By contrast, hippocampal tetrodes were advanced by at most 83 μm per day (and usually not at all), so that these tetrodes would remain within the hippocampal region throughout the entire experiment. Rats remained in the experiment area and rested in their home cages for at least 1 h before being returned to the vivarium for weighing and feeding.

2.4.2.4 Histology

One day prior to euthanasia, the rat was deeply anesthetized with isoflurane and marking lesions were made on one tetrode wire per probe cluster by passing a 50 μA current through a lesion maker (Grass Instruments, West Warwick, RI) for 10 seconds at each polarity. 24 h after marking lesions were made, the rat was perfused transcardially with formalin and the fixed brain was carefully separated from the tetrode bundles, which were still positioned at their final advancement locations (we measured each probe's linear excursion from the guide cannula to corroborate the advancement logs kept during the experiment). Brains were fixed in a solution of 30% sucrose formalin, sectioned at a thickness of 40 μm , and mounted on slides for imaging on a semi-automated digital light microscope (Keyence, Osaka, Japan). Slice images were referenced by overlaying them onto plates from the rat atlas of Paxinos & Watson (2004). Based upon marking lesions and track positions, the trajectory of each probe through the tissue was reconstructed by serial examination of all slices. The position of each tetrode on each recording day was estimated from the reconstructed trajectories.

2.4.3 Data analysis

2.4.3.1 Spike Sorting

Manual spike sorting was performed offline using SpikeSort 3D (Neuralynx, Bozeman, MT). Cluster cutting was primarily performed based on the peak and valley amplitudes of spikes across all tetrode channels. In some cases, spike energy and PCA components 1, 2 and 3 were analyzed to achieve better separation. Clusters containing interspike intervals <1 ms were removed from analysis for lack of a refractory period.

2.4.3.2 LFP Filtering and Analysis

On the assigned SWR probe for each animal, SWR events were detected as threshold crossings of the ripple band envelope which occurred when the rat was sitting still (movement speed < 2 cm/s). An 8th order IIR filter was applied to extract signals in the 180-250 Hz band from LFP channel data sampled at 32 KHz. The envelope of the ripple band was taken as the absolute value of the Hilbert transform of the bandpass filtered signal. SWR events were detected as upward crossings of the ripple envelope amplitude past a threshold equal to 4 standard deviations above the mean envelope amplitude. The mean, standard deviation, and SWR threshold were calculated separately for data collected on the maze versus in the bucket, because SWR amplitudes differed for these two conditions (see Results). A lockout period of 100 ms was imposed after each SWR event, so that the next SWR event could not be detected until the lockout period had expired.

On the assigned theta LFP probe for each animal, a bidirectional 8th order IIR filter was applied (using MATLAB's 'filtfilt' command) to extract signals in the 4-12 Hz theta band from LFP channel data sampled at 32 KHz. Theta phase was derived using MATLAB's 'angle' command from the Hilbert transform of the bandpass filtered signal.

2.4.3.3 Response Latency

To measure the latency between SWR events and a neuron's spike responses, we computed a peristimulus histogram of spike responses (10 ms bins, spanning ± 0.5 s) triggered at the peak of each SWR event's ripple band LFP envelope. Two iterations of smoothing with a 50 ms (5 bin) boxcar window were performed, and the peak of the smoothed histogram was taken as the time of the peak unit response. As explained in Results, septal neurons and striatal neurons were never recorded during more than one session, but hippocampal neurons were often recorded over multiple sessions. In cases where a hippocampal neuron was recorded during more than one session, we identified the session for which the cell exhibited the most statistically significant (lowest p-value) spike response to SWRs, and used that session to measure the response latency.

2.4.3.4 Single Unit Responses During SWRs

To quantify a neuron's response to SWR events, we counted the number of spikes that the neuron fired within a ± 50 ms window surrounding each SWR, and divided by the width of the time window (100 ms) to compute the unit's response rate (in Hz) during each SWR event. To measure the neuron's baseline firing rate, we summed the number of spikes fired within two baseline windows on either side of the SWR event (-500 to -300 ms, and +300 to +500 ms), and divided by the summed width of both windows (400 ms) to compute a baseline firing rate (in Hz) for each SWR event. A neuron was considered to be responsive during SWRs if the response rate during SWRs was significantly different from the baseline response rate across all SWRs during which the neuron was recorded (see Results).

2.4.3.5 Speed Analysis

To analyze modulation of neural firing rates by running speed, position data from the video tracker (sampled at 30 Hz) was smoothed by convolution with a boxcar window 7 samples wide. Speed at each sample time t was then estimated at seven different lag times: $L=33, 66, 99, 122, 155, 168,$ and 201 ms. The median of these seven estimates was then taken as the measure of speed at time t . The following formula estimated speed at sample time t using lag L :

$$St = \frac{R}{P} \sqrt{(X_{t+L/2} - X_{t-L/2})^2 + (Y_{t+L/2} - Y_{t-L/2})^2} \quad \text{Equation 2-1}$$

where St is the estimated speed, $R = 1000/L$ is the lag frequency in Hz, $P = 2.2$ cm/pixel is the tracking resolution, and (X_t, Y_t) is the interpolated position in pixels at time t . Linear interpolation of the speed time series (sampled at 30 Hz) was used to estimate the rat's running speed at each spike time (sampled at 32 KHz). A cell's firing rate at each running speed was computed by binning spike-triggered speed measurements in the range 0 to 60 cm/s using bins 2 cm/s wide, and then dividing the total number of spikes in each speed bin by the total time spent running at that speed. Linear regression then calculated the slope and intercept of the best linear fit to points on the speed curve. Bins containing <10 spikes or <2 s of occupancy time were omitted, and at least 4 valid bins were required for inclusion in the regression analysis.

CHAPTER 3

Hippocampal Coordination with Neural Populations Outside the Hippocampus

3.1 INTRODUCTION

Decision-making involves interactions between multiple cognitive processes, each of which is supported by a distinct subset of brain regions. Animals learn the associations between states, actions and outcomes from experience; they attach value to those states-action pairs according to the desirability of the outcomes they predict (Sutton & Barto, 1998). For example, a human toddler that is thirsty will learn to seek and lift a cup to drink because this action slakes thirst. An alternative course of action for the thirsty baby might be flailing in frustration due to thirst, which results in spilled liquid. Flailing acquires a negative outcome association due to several undesirable consequences (e.g. continued thirst, wetness, upset parents). Once the actor learns state and action values, some process of deliberation is necessary to compare the value of the current state against the predicted values of other states that can be reached by performing currently available actions. The “best” action to perform is then chosen based on the likelihoods and desirability of the outcomes they portend (Mowrer 1947; Rangel et al., 2008; van der Meer et al., 2012). The learned associations between states, actions and outcomes constitute a learned internal model that the animal can use to evaluate plans of action prior to engaging in those actions.

As discussed in more detail in Chapters 1, hippocampal sharp-wave ripple (SWR) events have been hypothesized to play a role in deliberation over choices during value-based decision making (Diba & Buzsáki, 2007; Davidson et al., 2009; Karlsson & Frank, 2009; Wikenheiser & Redish, 2013; Pfeiffer & Foster, 2013). Specifically, compressed replay of navigational trajectories by place cells during SWR events is hypothesized to provide a neural mechanism by which the outcomes of different actions are predicted and compared (Karlsson & Frank, 2009). For example, if a rat navigating through a maze arrives at a T intersection, then it must decide whether to travel right or left. To choose between the leftward versus rightward paths, the animal may preplay (“imagine”) each path, generating predictions about where it will end up after following one path versus the other. This sort of predictive ability depends upon a

predictive mental model of the environment. The hippocampus has long been thought to store a type of mental representation known as a “cognitive map” (O’Keefe and Nadel, 1978; Redish, 1999). A cognitive map is an internal allocentric representation of a familiar spatial environment, which is exactly the type of mental model needed for envisioning future states that might result from following a navigational path. Hence, when SWR events occur at a T intersection in a maze, the hippocampus may draw upon the information stored in its cognitive map to generate compressed replay sequences of the rightward versus leftward trajectories, helping the rat to predict which of the two trajectories will likely lead to a better outcome, and thereby allow the rat to decide whether to go right versus left.

While the hippocampus is thought to be a site where cognitive maps are stored, it is not widely regarded as a site where information about the *values* of states and actions are stored. Evidence suggests that information about state and action values is instead stored in other brain regions that contain reward circuitry, including the midbrain dopamine system, the dorsal and ventral striatum, the amygdala, and frontal cortical areas. Hence, while model-based predictions about future states may be derived from hippocampal replay sequences during SWR events, the *expected values* of those predicted states may be encoded in other brain regions. If so, then hippocampal predictions of future states would presumably need to be relayed to other brain regions for an assessment of value. Transmission of SWR signals out of the hippocampus to targets in the brain’s reward systems may be necessary for this to occur.

Chapter 2 presented results from experiments in which neural activity was recorded in the lateral septum (LS) and striatum during hippocampal SWR events. Some LS and striatal neurons exhibited robust excitatory or inhibitory responses during SWRs. One of the most intriguing findings from this study was that LS appeared to contain two distinct populations of neurons. The first population contained neurons that were excited during SWRs that occurred when the rat was immobile, fired preferentially on the downslope of hippocampal theta rhythm when the rat was running on a maze, and exhibited a positive correlation of their firing rates with

the rat's running speed. Conversely, neurons in the second population were inhibited during SWRs that occurred when the rat was immobile, fired preferentially on the upslope of hippocampal theta rhythm when the rat was running on a maze, and exhibited a negative correlation of their firing rates with the rat's running speed. In this chapter, I shall propose a model for how these two distinct populations of LS neurons might be involved in relaying information from the hippocampus to the brain's reward circuitry, and thereby participate in model-based decision making during spatial navigation.

3.2 DOPAMINE SIGNALING AND REWARD-MOTIVATED BEHAVIOR

The details of which brain regions outside the hippocampus evaluate the options presented by the hippocampus are murky, and may depend on specific details (evaluating actions after a painful electric shock will involve the amygdala, while maximizing reward in a safe environment may not.) In the context of exploiting a situation for reward, one possibility is a disinhibitory pathway connecting place cells of the hippocampus CA3 to dopaminergic neurons of the VTA via the Lateral Septum (Risold & Swanson, 1997; Lou et al., 2011). Midbrain DA neurons of the VTA are involved in reward-motivated learning and decision making (Berridge & Robinson, 1998; Salamone & Correa, 2012). These neurons project to the Nacc and support Reward Prediction Error signaling.

3.2.1 The Midbrain Dopamine System

Dopamine (DA) is a modulatory neurotransmitter that is known to play a major role in reward-motivated learning and behavior. There are two major DA projections in the mammalian brain: the mesostriatal DA pathway originates in substantia nigra and projects mainly to dorsal striatal areas, and the mesolimbic DA pathway originates in the ventral tegmental area (VTA) and projects to ventral striatal areas such as the nucleus accumbens (NAcc), as well as to

prefrontal areas and a minor projection to hippocampus. The role of dopamine in learning and behavior is still a very active area of research. Much of this research has been motivated by three influential theories of dopamine function: the hedonistic pleasure (or “liking”) theory, the incentive motivation (or “wanting”) theory, and the prediction error theory.

3.2.1.1 Dopamine and Hedonistic Pleasure

Early theories of dopamine function tended to equate dopamine release with “feeling good.” Human subjects anecdotally report feelings of dysphoria after taking drugs that block dopamine function, such neuroleptics, and conversely report feelings of pleasure and euphoria after taking drugs that stimulate dopamine function, such as psychostimulants (Healy, 1989; Hollister et al., 1960; Bellmaker and Wald, 1977). Drugs or brain lesions that impair dopamine function also tend to impair reward-motivated learning and behavior in rats, consistent with the idea that reward signals and hedonic pleasure may be blunted by these manipulations (see Berridge & Robinson, 1998).

If dopamine release is a basic neurobiological mechanism for feeling pleasure, then it may be responsible for our “liking” of things that we enjoy. Thus, if a person likes chocolate or warm baths or kissing, then it could partly be because those activities trigger the release of dopamine in their brain. And if dopamine release is equatable with pleasure, it follows that any activity or stimulus that triggers dopamine release should be positively reinforcing. Therefore, behaviors that produce dopamine release as an outcome will become more likely to be repeated in the future, as a consequence of instrumental learning.

Although vastly oversimplified, there may be some validity to the idea that dopamine release can produce feelings of pleasure. Valid or not, this idea has become somewhat ensconced in popular psychology; it is not uncommon for people to say that enjoyable activities

give them a “burst of dopamine.” But equating dopamine release with hedonistic pleasure is a gross oversimplification of the role of dopamine in reward-motivated learning and behavior.

3.2.1.2 Dopamine and Incentive Motivation

A competing theory, known as the incentive motivation theory, proposes that dopamine is involved in “wanting” rather than “liking” things (Berridge & Robinson, 1998). In this theory, rather than producing feelings of hedonic pleasure during the consummation of a rewarding experience, dopamine is responsible for our feelings of desire for those consummatory experiences when we are not actively engaged in them. This desire for pleasurable experiences is thought to motivate many of our actions and behaviors, which is why this view is called the “incentive motivation” theory of dopamine function. The basic idea is that our willingness to perform an action depends upon how much we want or desire that action’s expected outcome. If we desire the outcome of an action very much (such as the sweet taste that results from eating chocolate, or the pleasant relaxation that results from drinking a well-crafted beer), then dopamine is released when we contemplate that action, and this dopaminergic boost motivates and incentivizes us to perform the actions required to receive the desired outcome.

Evidence for the incentive motivation theory comes from studies showing that impairment of dopamine signaling reduces an animal’s willingness to exert effort to obtain reward, and thus appears to make the animal “want” the reward less (for review, see Berridge & Robinson, 1998; Salamone & Correa, 2012). Conversely, drugs or other manipulations that enhance dopamine signaling can invigorate motivated effort to obtain rewards.

Addiction may be viewed as a pathologically strong motivation to seek and consume a particular reward, such as drugs. Indeed, there is strong evidence that dopamine plays a role in the pathology of drug addiction. Drugs of abuse are known to alter dopaminergic signaling pathways (Saal et al., 2003; Everitt & Robbins, 2005; Edwards et al., 2007; Chen et al., 2008;

Schmitt & Reith, 2010). Dopaminergic manipulations can also lead to pathological changes in incentives to perform other behaviors, such as gambling, in patients undergoing dopaminergic drug therapies for disorders like Parkinson's or Restless Leg Syndrome (Pinder 2007; Pirritano et al., 2014; Shapiro et al., 2007; Tippmann-Peikert et al., 2007; Voon et al., 2007).

In summary, the incentive motivation theory proposes that dopamine release is related to the wanting or desiring of a rewarding outcome, which in turn invigorates motivated effort invested in behaviors that can help to attain such outcomes.

3.2.1.3 Dopamine and Reward Prediction Error

A third theory of dopamine function, which combines some elements from the previous two, is the Reward Prediction Error (RPE) theory. The basic tenets of this theory are illustrated by a classic experimental design from the lab of Wolfram Schulz (Fig. 3-1).

This experiment involves a monkey trained in a Pavlovian conditioning paradigm, where a visual stimulus (the conditioned stimulus, or CS) is followed a few seconds later by a squirt of juice into the monkey's mouth (the unconditioned reward stimulus, or US). During early trials, before the monkey has learned that the juice US follows the visual CS, each juice reward is an unexpected pleasurable outcome. During this phase, dopamine neurons of the VTA exhibit a marked, transient increase in their firing rate at the moment of juice delivery. This is very much what might be expected from the "hedonistic pleasure" theory described above. The juice is pleasurable, and dopamine neurons fire when this pleasure is being experienced.




Figure 3-1: Dopamine and TD error. (A) A monkey sees a light (visual CS) which reliably predicts delivery of juice from a tube (reward US). (B) Prior to learning, or in the absence of the CS, VTA neurons increase their firing in response to reward receipt. (C) After learning, VTA neurons respond to the CS, but not the reward US. (D) After learning, if juice delivery does not occur as expected after the CS presentation, VTA neurons fire vigorously to the CS, but then drop briefly below their tonic firing level when the reward is not delivered. The drop reflects a negative change (decrease) in expected future reward. Solid lines represent cell firing in Hz; dotted lines represent zero Hz. Adapted from Shultz, 1997.

But during later trials, after the monkey has learned to expect the rewarding US to follow the visual CS, the dopamine neurons no longer respond to the reward. Instead, they respond to the CS. More generally, the excitement of dopamine neurons tends to migrate backward in time to the earliest reliable predictor of the pleasurable outcome. The lack of response to the desirable juice reward contradicts the simple explanation that dopamine purely signals hedonistic pleasure. To explain why dopamine neurons respond to a CS that predicts the reward, it has been hypothesized that they encode a specific type of RPE signal, called a temporal difference (TD) error signal (Montague et al., 1996).

The TD error signal can be written mathematically as follows (Sutton & Barto, 1998):

$$\lambda = R_t + (V_t - V_{t-\Delta t}), \quad \text{Equation 3-1}$$

where R_t is the actual reward being experienced at the current time t , V_t is the total amount of reward that the animal expects to accumulate in the future starting from the current time t , and

$V_{t-\Delta t}$ is the total amount of future reward that the animal was expecting to accumulate at a moment in the very recent past, $t-\Delta t$. The term in parenthesis, $(V_t - V_{t-\Delta t})$, is the time derivative of the animal's expected future reward. If the animal expects the same amount of reward now as it did a moment ago, then regardless of whether this expected reward is big or small, it is always the case that $(V_t - V_{t-\Delta t}) = 0$.

At the beginning of the Pavlovian conditioning experiment (Fig. 3-1), the animal has no reason to expect any reward, so V_t , $V_{t-\Delta t}$, and $(V_t - V_{t-\Delta t})$ are constantly zero. During delivery of an unexpected reward that follows the CS, R_t becomes positive, so the TD error λ becomes positive as well. The occurrence of the positive prediction error trains predictive networks in the brain to associate the CS with the US. Over trials, the animal learns to expect that the rewarding US will occur after the CS. After the monkey has been trained, at the moment when the CS comes on, the animal's expectation of future reward increases (even though no reward has been delivered yet, and thus R_t is still zero). This causes $(V_t - V_{t-\Delta t})$ to briefly become positive, and thus by Equation 3-1, the prediction error λ becomes positive as well. But in the next moment after that, the animal continues to have the same elevated level of expected reward, so $(V_t - V_{t-\Delta t})$ falls back to zero for the remainder of the interval between the CS and US.

What function is served by this learned response to the CS? The TD error signal was first proposed to explain the phenomenon of second order conditioning, whereby a CS that has become associated with a US can itself behave as a reinforcer (Sutton & Barto, 1998). If dopamine drives predictive learning, then by acquiring the ability to activate dopamine, a CS would also acquire the ability to drive predictive learning (Montague et al., 1996). However, another possible function of the CS-evoked dopamine response may be postulated as a corollary to the incentive motivation theory of dopamine function: the CS-evoked dopamine burst may also prime the brain to perform motivated actions that are required to attain the desirable outcome. In a simple Pavlovian task like that shown in Fig 3-1, the monkey is not

required to perform any instrumental action to obtain reward. But in most real life situations, rewards are not squirted directly into an animal's mouth. Some kind of motor action (e.g. reaching for the reward, searching for the reward, competing with other animals to obtain the reward) are almost always necessary. Given the involvement of dopamine in incentive motivation (see above) and motor action (the motor deficits of Parkinson's disease results from a failure of dopaminergic signaling in the striatum), it is reasonable to conclude that CS-evoked dopamine responses may not only serve as second order reinforcement signals, but also as incentive motivation signals that invigorate motor behavior when there is an increase in an animal's expectation of future reward.

Why do dopamine neurons stop responding to a reward after it becomes expected? At the precise moment when the animal learns to expect the US, the expectation of future reward falls back to zero, because the animal transitions from expecting reward to no longer expecting reward. In that moment, $(V_t - V_{t-\Delta t})$ briefly becomes negative. If the expected reward is not delivered in that moment, then R_t remains zero while $(V_t - V_{t-\Delta t})$ becomes negative, so that the prediction error λ becomes negative as well. VTA dopamine neurons actually reduce their firing rate if an expected reward fails to occur (Schultz et al., 1997; Columbo, 2014; Watanabe-Uchida et al., 2017), exactly as would be expected if their firing encodes a TD prediction error. But if the expected reward is delivered at the expected moment, then R_t becomes positive at exactly the same moment when $(V_t - V_{t-\Delta t})$ becomes negative. They cancel out, and there is no response; this is thought to explain why dopamine neurons stop responding to expected rewards.

3.2.2 Hippocampal Interactions with the Dopamine System

Hippocampal influences upon reward-motivated learning and behavior may depend upon bidirectional interactions between the hippocampus and the dopaminergic system.

Hippocampal outputs to the dopamine system may be necessary for the hippocampus to exert influence over behaviors including motor actions, reward-seeking, attention, arousal, and decision making (Luo et al., 2011; Gomperts et al., 2015; Tingley & Buzsáki, 2018, 2020; Wirtshafter & Wilson, 2019). Dopaminergic inputs to the hippocampus may be important for hippocampal function as well, although the presence and source of hippocampal DA is sometimes contested (Gasbarri et al., 1997; Kempadoo et al., 2016).

3.2.2.1 Hippocampal Outputs to the Dopamine System

Direct projections from the hippocampus to midbrain dopamine neurons are scarce or non-existent, but lateral septum (LS) is a major subcortical output target of hippocampal projection neurons (Raisman 1966; Swanson & Cowan, 1975; Swanson et al., 1981; Tingley & Buzsáki, 2018). A growing body of evidence suggests that hippocampal projections to the lateral septum may be an important route via which the hippocampus relays information to the midbrain and other subcortical regions to exert influence over behaviors such as reward-seeking, motor actions, reinforcement learning, and decision making (Luo et al., 2011; Gomperts et al., 2015; Tingley & Buzsáki, 2018; Wirtshafter & Wilson, 2019).

The lateral septum (LS) is a gateway to subcortical regions such as the Lateral PreOptic area (LPO), Lateral Hypothalamic Area (LHA), Substantia Nigra (SN) and ventral tegmental area (VTA) (Risold & Swanson, 1997; Luo et al., 2011). The connections of the LS presumably allow it to influence many cognitive functions including reward sensation, orientation, locomotion, attention, arousal, and behavior selection (Risold & Swanson, 1997). Manipulation of LS suggests that it contributes to a variety of behavioral and homeostatic responses (Table 3-1). Each of these functions is important for online evaluation of behavioral plans during active behaviors.

Compressed replay by place cells during SWRs that occur during pauses in active behavior is hypothesized to provide a mechanism for animals to “deliberate” over decisions about which actions to perform (Yu & Frank, 2015). If so, then when a compressed replay sequence is generated to represent an “imagined” spatial trajectory, a TD prediction error encoded by dopamine might occur at the moment when the replay occurs, to report how much the animal’s expectation of future reward would increase if the animal were to navigate along the imagined trajectory. Consistent with this idea, reward-responsive midbrain dopamine neurons tend to fire synchronously with SWRs during wakeful stillness on a maze, as might be expected if the animal were assessing the values of potential action plans during SWRs that occur on the maze (Gomperts et al., 2015). Interestingly, this study reported that SWR-evoked dopamine responses appear to be less pronounced during sleep, perhaps because SWRs perform a memory consolidation function rather than a decision making function during sleep.

Table 3-1 A Partial List of the Roles of LS

behavioral	locomotion	Bender et al., 2015, Fig 4
	spatial	Rawlins & Olton, 1982 M'Harzi & Jarrard, 1992 Leutgeb & Mizumori, 1999 Zhou et al., 1999 Leutgeb & Mizumori, 2002 Takamura et al., 2006 Tingley & Buzsáki, 2018 Monaco et al., 2019 Wirtshafter & Wilson, 2020
mixed	reward	Cazala et al., 1988
	feeding	Sweeney & Yang, 2015
homeostatic	stress	Menard & Treit, 1995 Chee & Menard, 2011 Singewald et al., 2011
	fear	Feldon et al., 1982 Endres & Fendt, 2008 Reis et al., 2010 Parfitt et al., 2017
	aggression	Clarke & File, 1982 Wong et al., 2016

3.2.2.2 Dopaminergic Inputs to the Hippocampus

Several reports suggest that rat VTA projects to hippocampus and that some portion of these projections are dopaminergic (TH+) cells in (Wyss et al., 1979; Scatton et al., 1980; Swanson, 1982; Verney et al., 1985; Gasbarri et al., 1994a/b; Gasbarri et al., 1997; Cenquizca & Swanson, 2006; Rosen et al., 2015) and in mice (McNamara et al., 2014; Yang et al., 2017; Broussard et al., 2016.) Substantia Nigra may also contribute some of the DA in the hippocampus with a gradient that increases along the dorsal-ventral axis (Scatton et al., 1980;

Gasbarri et al., 1994). Microdialysis demonstrates that dopamine is present in the hippocampus. DA concentrations are $\sim 1/25$ hippocampal norepinephrine (Laatikainen et al., 2012), and $\sim 1/90$ the concentration of DA in the striatum (Scatton et al., 1980). Intriguingly, DA concentrations in the hippocampus nearly double during bouts of running (Goekint et al., 2012), corresponding with online spatial processing and strong theta rhythm.

In at least mice, DA is co-released with NE in the hippocampus from innervation originating in the LC (Kempadoo et al., 2016; Takeuchi et al., 2016). In mice, hippocampal inputs from the VTA may be less important than those from the LC for at least some behaviors (Kempadoo et al., 2016; Takeuchi et al., 2016). The presence of a transmitter in a brain region does not establish that it has any function in the structure; to help establish a potential role, it is especially useful to check that appropriate receptors are present. Hippocampus contains both RNA and receptor proteins for dopamine receptors. The rodent hippocampus expresses RNA for DRD1 (Lazarov et al., 1998; Gingrich et al., 1992; Tiberi et al., 1991) and DRD2 (Brouwer et al., 1992). In one report, DRD5 RNA expression dominates over DRD1 and DRD2 (Laurier et al., 1994). Mouse hippocampus expresses GFP constructs under the control of DRD1 & DRD2 receptor promoters (Gangarossa et al., 2012; Azeveda et al., 2019) and rat hippocampal neurons express proteins for DRD4 (but not much DRD2 or DRD3, Khan et al., 2000; Goldsmith & Joyce, 1994). Specific details of which receptors and subtypes exist in which regions in which species could be further clarified and the field would benefit from further confirmation of these results.

Application of dopamine to slices modifies the electrophysiological behavior of CA1 pyramidal cells (Bernardo & Prince, 1982a/b; Gribkoff & Ashe, 1984a/b; Stanzione et al., 1984) and impairs LTP in hippocampal slices (Frey et al., 1990; substantially through DRD1 receptors, Frey et al., 1991). Dopamine produces hyperpolarization (Bernardo & Prince, 1982a/b; Stanzione et al., 1984), producing a short-term inhibitory effect (on the order of 10's of seconds to minutes).

However, over a longer time frame (hours), dopamine appears to disinhibit the hippocampal cell population, producing increases in activity (Gribkoff & Ashe, 1984a/b).

3.2.2.3 Effects of Hippocampal Dopaminergic Manipulations on Behavior

Hippocampal-dependent spatial navigation task performance and spatial representations in the hippocampus are affected when dopamine transmission is altered. Dopamine receptor antagonist infusion into the hippocampus impairs behavior performance in spatial tasks (Kempadoo et al., 2016; Retailleau & Morris, 2018) and alters place cells (Retailleau & Morris, 2018). 6-OHDA lesions in the hippocampus that spare NE fibers produce deficits in a Morris water maze task (Gasbarri et al., 1996). VTA disruption impairs CA1/2 spatial representation (Martig & Mizumori, 2011) in rats. (This result suggests that VTA projections to the hippocampus in rat and mouse may differ.) VTA fiber stimulation in the dorsal hippocampus improves some aspects of task performance and supports reinstatement of recent spatial map ensembles (McNamara et al., 2014.)

The short term inhibitory effect of dopamine on hippocampal neurons could sharpen the hippocampal map by ensuring that only the most powerfully driven cells are active; such a mechanism would maximize ensemble binding and reduce noise. The tonic increase in hippocampal DA during a bout of running suggested by the doubling of DA concentration during the theta state (Goekint et al., 2012) should globally inhibit firing, producing a sparse ensemble consisting of highly-driven neurons, and improving spatial representation precision. Phasic DA release from the VTA during receipt of an unexpected reward could produce an acute sharpening of spatial representation, emphasizing the current location if the rat is actively exploring. Alternatively, a phasic DA rise during non-theta consummatory behavior that

corresponds with a SWR replay may transiently restore the high DA state associated with active theta rhythm, helping to sharpen the spatial representation.

Intriguingly, the long term increase in population activity observed in slices after short term inhibition of cells likely suggests an exciting mechanism explaining why new learning correlates with an increase in SWR production -- dopamine released from the RPE system primes the hippocampus for later memory consolidation. Experiments designed to observe the electrophysiological effects of DA inhibition in the hippocampus during active behavior would help address this possibility.

3.3 HIPPOCAMPUS AND DECISION MAKING: A HYPOTHESIS

Chapter 2 reported results from experiments which analyzed how neurons in LS and striatum respond during SWR events, and how they are synchronized with hippocampal theta rhythm. In this section, I integrate the findings of these experiments with prior findings from the literature to propose a hypothesis about how hippocampal outputs to subcortical brain regions may be involved in regulating reward-motivated behavior and model-based decision making.

Chapter 2 reports that LS appears to contain two distinct populations of neurons. The first population--henceforth referred to as M+ cells--consists of neurons that are excited during SWRs, fire preferentially on the downslope of hippocampal theta rhythm, and exhibit positive correlations of their firing rates with the rat's running speed. Conversely, the second population--henceforth referred to as M- cells--consists of neurons that are inhibited during SWRs, fire preferentially on the upslope of hippocampal theta rhythm, and exhibited a negative correlation of their firing rates with the rat's running speed. Striatum was also found to contain populations of M+ and M- cells, with firing rates that were positively and negatively correlated with running speed, respectively. As in LS, striatal M+ and M- tended to be excited and inhibited, respectively, during SWRs (although striatal neurons were not observed to fire at preferred

phases of theta as LS neurons did). This suggests that across subcortical brain regions (LS and striatum), neurons exhibit a correlation between the valence of their SWR-evoked responses and the valence of their firing rate modulation by running speed. How might these subcortical responses to SWRs be involved in learning, memory, and decision making?

3.3.1 Theta State: M+ and M- Cells

As reviewed in Chapter 1, hippocampal place cells exhibit a phenomenon called “phase precession” during active locomotion. As a rat runs through a place cell’s firing field, the cell spikes late in the theta cycle (during the downslope) as the rat enters the field, in the center of the cycle (near the peak) as the rat passes through the middle of the field, and early in the theta cycle (during the upslope) as the rat leaves the field. At the population level, phase precession by individual neurons gives rise to a phenomenon of “theta sequences” while rats are running through an environment (Skaggs et al., 1996). Within a given theta cycle, place cells that encode locations behind the rat fire first (on the upslope of the cycle), cells that encode the animal’s current location fire next (at the peak of the cycle), and cells that encode locations ahead of the rat fire last (on the downslope of the cycle). Consequently, as a rat locomotes through space, hippocampal place cells are constantly alternating back and forth between representing locations ahead of the animal versus behind the animal, at the theta frequency of 6-8 Hz (Figure 3-2).

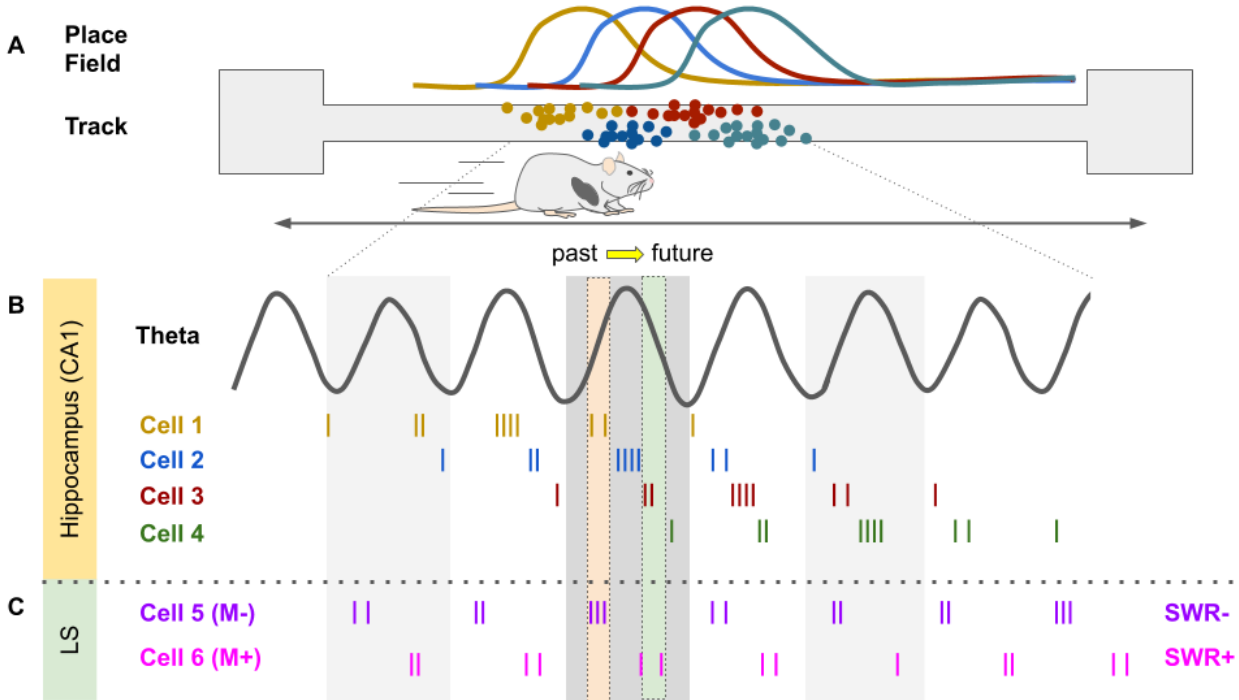


Figure 3-2 : *Theta sequences reflect the immediate past, present and future during behavior.* (A) Rat running a linear track. Four place cell tuning curves are displayed above the track (place fields). The position of the rat's head reflects the theta cycle with the darker background. (B) Idealized spike rastergrams of place cells on successive cycles of theta (black line) as the rat progresses down the track. During the cycle shaded in gray, the rat is in the center of Cell 2's place field. The temporal ordering of the cells' spikes during the cycle are 1, 2, 3, 4. The rat has already passed cell 1's preferred location, is presently in the preferred location for cell 2 and will soon occupy Cell 3's field center before moving on to cell 4's field center. Thus, the field of cell 1 is behind the rat, cell 2 is at the rat's current position, and cells 3 and 4 are ahead of the rat. (C) Speed modulated LS cells prefer theta phases that correspond to the action associated with reaching place cells associated with those phases. M- cells fire on early phases of theta, concurrently with place cells that encode locations behind the animal's current position. M+ cells fire on late phases of theta, concurrently with place cells that encode locations ahead of the animal. Furthermore, M+ cells tend to fire more during SWR events, while M- cells pause during SWR events. Diagram adapted from Diba & Buzsáki, 2007 and Skaggs et al., 1996.

3.3.1.1 Motor-Related Activity of LS Neurons

If this alternating spatial representation signal were relayed to the midbrain dopamine system, then it could be used to generate a continuous series of TD prediction errors that may help the animal to constantly evaluate whether it should continue to run along its current trajectory, versus stop or change its trajectory. By Equation 3-1, dopamine neurons might compare the expected future reward during the current phase of the theta cycle, V_t , against the future reward that was previously expected during an earlier phase the theta cycle, $V_{t-\Delta t}$. Assuming that Δt is equal to one half of a theta cycle, then the time derivative of expected future reward ($V_t - V_{t-\Delta t}$) would compare the expected future reward on the current theta phase against the reward expected half a cycle earlier. Suppose the current theta phase is the downslope (when place cells encoding locations ahead of the animal are spiking), so that half a cycle earlier was the upslope (when place cells encoding locations behind the animal were spiking). In this situation, V_t would encode the “value” attached to locations ahead of the animal (that is, the reward expected to accrue from continuing to run forward), whereas $V_{t-\Delta t}$ would encode the value attached to locations behind the animal (that is, the reward expected to accrue from stopping or going backward). Exactly the opposite would be true when the current theta phase is the upslope (when place cells encoding locations behind the animal are spiking): V_t would encode the value attached to locations behind the animal (that is, the reward expected to accrue from stopping), whereas $V_{t-\Delta t}$ would encode the value attached to locations ahead of the animal (that is, the reward expected to accrue from continuing to run forward).

It follows from these observations that the prediction error signal, λ , would alternate back and forth at the theta frequency between encoding the value of continuing the current trajectory (which would coincide with the upslope of theta) versus the value of stopping or changing trajectory (which would coincide with the downslope of theta). Each cycle of theta rhythm might thus be thought of as a “frame” during which a specific action plan (running versus stopping) is

evaluated. The upslope would be the frame for computing the expected value of running, and the downslope would be the frame for computing the expected value of stopping.

Intriguingly, Chapter 2 reports that M+ neurons in LS showed a positive correlation of their firing rates with running speed (and with SWR events), and fired on the downslope of theta, which is the “frame” during which the expected value of running would be computed by a TD error signal. Conversely, M- neurons in LS neurons showed a negative correlation of their firing rates with running speed (and with SWR events), and fired on the upslope of theta, which is the “frame” during which the expected value of stopping would be computed by a TD error signal. One possible interpretation for this finding is that M+ neurons are part of a network for behavioral activation, whereas M- neurons are part of a network for behavioral inhibition. If so, and if the expected value of behavioral activation versus inhibition are computed on the downslope versus upslope of the theta cycle, respectively, then it would stand to reason that M+ cells should fire on the downslope and M- cells on the upslope of the theta cycle, as reported Chapter 2.

3.3.1.2 Motor-Related Activity of Striatal Neurons

Previous single-unit recording studies have reported that firing rates of striatal neurons in rodents are correlated with the animal’s running speed (Ruede-Orozco & Robbe, 2015). Consistent with this finding, Chapter 2 reports that a subset of SWR-responsive striatal neurons are also modulated by running speed. Unlike LS neurons, striatal neurons did not show a preference to fire at any particular phase of theta rhythm.

Projection cells from the striatum are GABAergic medium spiny neurons (MSNs). MSNs are often categorized according to whether they express either D1 family or D2 family dopamine receptors. Classical models of the basal ganglia posit that D1 MSNs are the origin of a “direct”

striatonigral motor output pathway which releases motor behavior, whereas D2 MSNs are the origin of an “indirect” striatopallidal motor output pathway which inhibits motor behavior. It is tempting to speculate that the striatal M+ cells reported in Chapter 2 might be D1 cells that excite motor activity, and the striatal M- cells might be D2 cells that inhibit motor activity. However, prior neural recording and imaging studies have consistently failed to find evidence that D1 and D2 MSNs behave simply as motor-on and motor-off cells, as classical models would predict. Instead, both types of MSNs seem to fire together during initiation and execution of voluntary motor behaviors (Cui et al., 2013; Isomura et al., 2013). Combined with other evidence, these findings suggest that perhaps D1 MSNs may help to drive the execution of selected actions, while D2 MSNs may simultaneously inhibit the execution of competing non-selected actions (Tecuapetla et al., 2016). Hence, the M- cells reported in Chapter 2 do not behave similarly to previously studied D1 or D2 neurons, and might not be medium spiny neurons at all. Instead, they could belong to one of several interneuron populations known to exist in striatum (Gonzales & Smith, 2015; Dudman & Gerfen, 2015; Benhamou et al., 2014). Further study would be necessary to clarify this.

3.3.2 LIA State: SWR-Evoked Responses

Data presented in Chapter 2, and two recent publications from other labs (Wirtshafter & Wilson, 2019; Tingley & Buzsáki, 2020) show that LS neurons respond during SWR events. It was shown in Chapter 2 that M+ cells are excited during SWR events, whereas M- cells are inhibited. A prior study has also shown that ventral striatal neurons exhibit phasic responses during dorsal hippocampal SWRs (Sosa et al., 2020), and Chapter 2 reports similar SWR responses in ventral striatal neurons. Chapter 2 also reports SWR-evoked responses in dorsal striatum, which has not been reported before.

What might be the functional significance of these subcortical responses during SWRs? To address this question, it is helpful to separately consider the three main functions that have been proposed for SWR events: model-based decision making, credit assignment during reinforcement learning, and long term consolidation of memories from hippocampus to cortex.

3.3.2.1 Model-Based Decision Making

One hypothesized role of SWR replay is to provide a mechanism for animals to “deliberate” over decisions about which actions to select (Yu & Frank, 2015). If so, then SWR signals may be relayed to the brain’s reward system, so that the animal can estimate how much reward to expect in the future if it chooses to navigate along a replayed trajectory. Consistent with this idea, it has been reported that reward-responsive midbrain dopamine neurons of the VTA fire synchronously with SWRs during wakeful stillness on a maze (but not during sleep), as might be expected if the animal were assessing the values of potential action plans during SWRs that occur on the maze (Gomperts et al., 2015); these cells also synchronize with theta.

In reinforcement learning theory, a distinction is drawn between two types of value signals: *state* values and *action* values (Sutton & Barto, 1998). Compressed replay of place cell trajectories has been hypothesized as a mechanism for generating model-based predictions of what future location would be reached by following a replayed trajectory, and thus estimating the state value of that future location. This is an example of a state value signal, because being at a particular future location is tantamount to being in a particular future state. But it stands to reason that when a rat is “predicting a future state” during compressed replay of a spatial trajectory, a neural representation of the actions required to reach that state might be concurrently activated. This could explain why neurons that are positively correlated with

locomotor activity (and thus may represent motor actions) excited during SWRs, as reported in Chapter 2 and elsewhere (Wirtshafter & Wilson, 2019).

There are reasons to believe that representations of actions (possibly encoded by subcortical M+ cells) as well as states (encoded by compressed replay of place cells) should become activated while an animal is deliberating over a choice of future action. The state value would carry information about how much reward to expect in the future, while the action value might carry additional information that is not encoded by the state value alone, such as how much effort is required to follow the trajectory that leads to the future state. Further study is warranted to investigate how representations of motor actions might be activated during SWRs.

3.3.2.2 Credit Assignment

When an animal reaches a goal location and experiences an unexpected positive outcome, the dopamine system generates a reward prediction error (RPE) signal. As discussed above, this signal moves backwards to become associated with the first reliable predictor of reward. How can this happen? The brain must associate stimuli in the past with the present outcome to accomplish this feat.

Generalized, this phenomenon is known as the Credit Assignment problem, whereby the immediate desirable (and perhaps unexpected) outcome must become associated with a potentially long chain of prior states, stimuli and actions so that these groups of states-stimuli-actions may become associated with later desirable outcomes. Solving this problem allows an agent to shape its behavior to achieve future desirable outcomes.

The reverse replays of place cell sequences during SWRs that arise during consummatory behavior provide a compelling mechanism for animals to reflect back upon the sequence of decisions that led to the outcome in question, and thus modify value representations that influence the likelihoods of repeating those decisions again in the future (Foster & Wilson, 2006). In this way, SWR events may participate in solving the credit assignment problem during reinforcement learning (Sutton & Barto, 1998).

Just as it may be beneficial to read out action as well as state values during model-based decision making (see prior section), it may likewise be beneficial to modify action as well as state values during reinforcement learning. Hence, the credit assignment problem may be viewed as a two-part problem: updating the values of states visited along the way to the current outcome, and updating the values of actions performed along the way to the current outcome. If the credit assignment problem is viewed in this way, then it stands to reason that neural representations of prior actions should be activated in conjunction with reverse replay of prior trajectories during credit assignment. This might help to explain why subcortical M+ cells are activated during SWR events. More broadly, a wide range of cortical and subcortical regions respond during SWR bolstering the argument that these events may help bind state, stimuli and action representations across the brain (see Table 1-1).

3.3.2.3 Long Term Memory Consolidation

Human patients with hippocampal damage suffer from partial retrograde, and complete anterograde semantic and episodic memory loss, and yet they can recall events from the distant past (years prior to the damage.) To rephrase this, *memories appear to become independent of the hippocampus with time.* SWR replays may facilitate this shift in brain structure dependence.

After an animal is removed from a maze, replay events that occur during subsequent periods of rest often encode trajectories from the recently visited maze environment. Several authors hypothesize that these post-behavioral replay events may contribute to the consolidation of recent experiences to long-term memories (Wilson & McNaughton, 1994; Buzsáki, 1998; Ego-Stengel & Wilson, 2010; Girardeau & Zugaro, 2011). Supporting this idea, acquisition of maze learning tasks is impaired by disruption of SWRs after behavior is complete, during sleep (Girardeau et al., 2009), suggesting that SWRs may be involved in consolidating hippocampal memories.

It is possible that activation of M+ and deactivation of M- neurons during SWRs could help to consolidate instrumentally learned motor actions to spatial trajectories and outcomes, forming stable long-term memories and strengthening specific patterns of action selection that are required to achieve desirable outcomes in a stable environment.

After substantial training in a stable environment, rats begin to express behavior termed “habitual”, in that the expression of this behavior appears to depend less on the value of the reward to the animal. A large body of evidence suggests that “procedural” memories for motor skills and actions do not normally depend upon the hippocampus (e.g. Zola-Morgan et al., 1986; Teuber et al., 1968; Packard & McGaugh, 1996; Purves et al., 2001; Packard & Goodman, 2016).

3.3.3 Striatal Involvement in Goal-Directed & Habitual Behaviors

The dorsal striatum supports normal behavioral performance in a rewarded navigation task (Yin & Knowlton, 2004). The medial and lateral divisions of the dorsal striatum (dmStr and dlStr, respectively) support different behavior strategies for solving behavior tasks (Packard & McGaugh, 1996). The dmStr (Castañé et al., 2010; Yin & Knowlton, 2004; Yin et al., 2004,

2005; Ragozzino et al., 2002; Devan et al., 1999; Devan & White, 1999; Castañé et al., 2010; rev. Smith & Graybiel, 2016; rev. Balleine et al., 2007) enables spatially-aware, goal-directed behavior strategies. The dlStr supports a stimulus-response, habit-driven pattern of behavior (Yin et al., 2004; Yin et al., 2006; Smith & Graybiel, 2013; rev. Smith & Graybiel, 2016; Balleine et al., 2007).

Behavioral evidence suggests that goal-directed strategies predominate early in training, while habit-driven behaviors predominate late in training (Ritchie et al., 1950; Hicks, 1964; Packard & McGaugh, 1996). The dlStr develops cells which fire to mark transition points during the behavior required to complete a task trial, such as start, decision point and end (Kubota et al., 2009; Atallah et al., 2014). In spatial navigation mazes, these cells may appear similar to place cells, but this is due to the confound of space with specific task phases. Certain spaces on the maze, such as the start arm, decision point and reward area are all associated with specific behavioral action patterns, so these cells are believed to mark these distinct phases of behavior rather than expressing spatial tuning. Multi-unit activity in the dlStr rises to mark the start, decision & end of trials, while activity in the dmStr ramps up over the course of the trial until the decision is made (Thorn et al., 2010, Fig 2). These action patterns develop prior to the expression of devaluation-insensitive, habit-driven behaviors (Kubota et al., 2009 ; Atallah et al., 2014), supporting the hypothesis that rats attend to all available cues during behavior (Restle 1957; Goodroe et al., 2018) and that both “habit” and “goal” systems are online and available simultaneously. Similarly, the results of Packard & McGaugh (1996) illustrate that disabling the “habit” system simply returns the rat to “goal-directed” performance strategies.

The hippocampus is also required for spatial performance on navigation tasks (Packard & McGaugh, 1996; Packard, 1999; Stringer et al., 2005), but intriguingly, only the ventral striatum (Nacc) receives significant input from the hippocampus (Cappaert et al., 2015).

3.4 SUMMARY AND CONCLUSIONS

Given the key role of the striatum plays in action selection and behavioral decision making, it could be that the SWR-evoked responses we observed in a small percentage of striatal neurons reflect a process by which MSNs become “programmed” to either excite or inhibit specific actions in the future, based upon value estimates for those actions that are generated during SWRs and compressed replay. This possibility could be further investigated in the future by experiments in which striatal unit activity is selectively disrupted during SWRs.

3.4.1 Value-Based Decision-Making Outside the Lab

The discovery of a population of cells in two brain regions outside the hippocampus which respond to SWR and appear to also encode for aspects of behavior suggests pathways through which hippocampal activity, and especially SWR replays, might influence value-based decision making. This connection is particularly intriguing given the hypothesis that SWR replays may contribute to planning, which precedes execution of an action plan. During a pause in behavior, an animal will evaluate both its internal state and the state of the environment to settle on a goal for the subsequent series of behaviors. It must then evaluate its options, which will require specific patterns of behavior to achieve. The animal must evaluate the desirability of state-goal-behavior tuples before choosing a course of observable action.

Outside of a highly controlled laboratory environment, there are many tradeoffs to consider in the course of making a decision. For the purposes of illustration, imagine a rat in a field with a stream and two potential sources of food. A rat may be both hungry and tired; here, the rat prioritizes feeding over rest. The two sources of food are different. One source is relatively nearby and consists of apples hanging in a tree. The second source is further away

and consists of blueberries hanging on a bush. Acquiring apples requires a potentially dangerous climb, while acquiring the blueberries requires a potentially dangerous swim. If the rat considers these two foodstuffs equally valuable, he must then base his decision on an evaluation the actions required to achieve his goal of eating.

In such a situation, I hypothesize that SWRs generated in the hippocampus would both represent the navigation required through space in the form of replay, and also would activate ensembles in other brain regions that encode information concerning the actions required to acquire either food. The data presented in Chapter 2 represents an early piece of evidence of such a link.




Figure 3-3 : *Real-world decision making requires evaluating state, outcomes and actions. The rat evaluates its state and determines that above all else, it is hungry, so eating is the most desirable outcome. It evaluates the available options in the environment and identifies two paths (yellow and blue lines), that lead to different food rewards -- apples or blueberries. Acquiring*

either reward entails engaging in different actions. The decision to pursue apples or blueberries involves balancing many factors ranging from food preference, to perceived danger, to the value of performing the actions. Perhaps the rat loves blueberries, but hates swimming, or anticipates being cold due to the ambient temperature, driving it to choose the apples.

CHAPTER 4

A Technique for Simultaneous Fast-Scan Cyclic Voltammetry & Extracellular Electrophysiology

Some of the results contained in this chapter were presented in the form of a poster at Society for Neuroscience in Chicago, IL, USA in 2015.

4.1 INTRODUCTION & BACKGROUND

It is clear that organisms evaluate internal & external states to prioritize an action sequence which is then executed, often in the pursuit of some desirable outcome. This process presumably accesses an internal model to make predictions about the world prior to action selection. Once the action is complete, and the animal will evaluate the outcome of the actions and learn. Learning from the outcome updates the internal model (visualized in Figure 4-1).

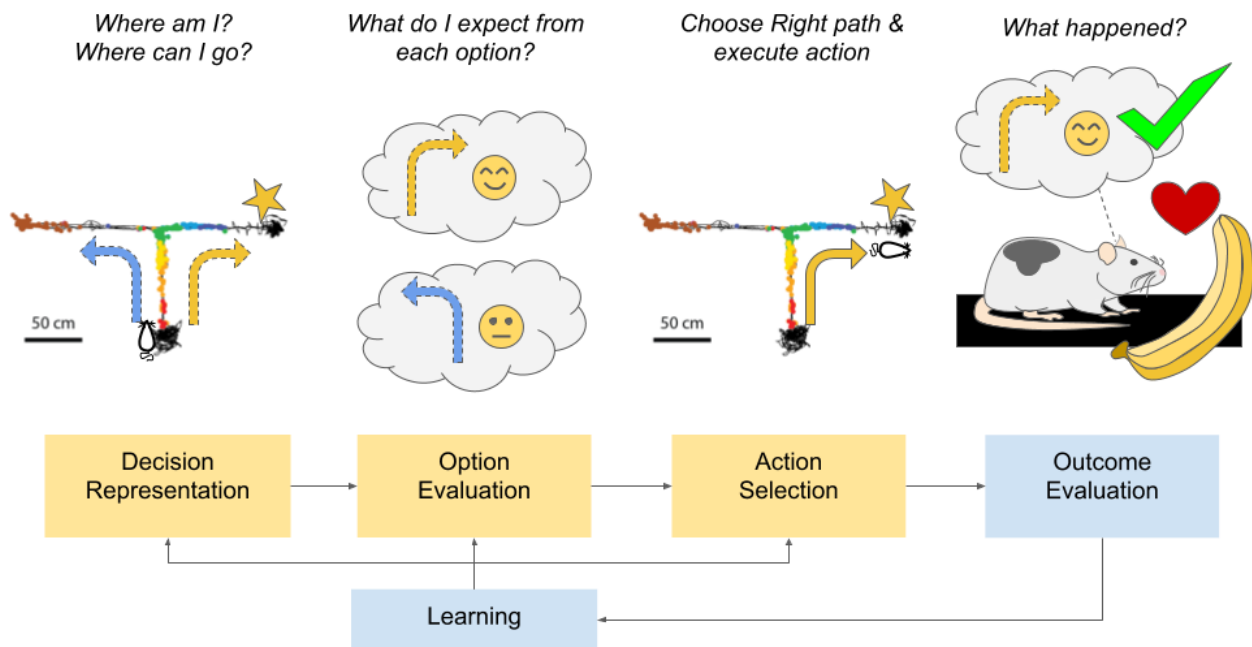


Figure 4-1 : Model-based decision making may be abstracted into the 5 step process illustrated along the bottom. A concrete example of a rat navigating in a maze in search of a reward is depicted for each step.

Assume that the rat evaluated its internal state and prioritized obtaining food. Place cells of the hippocampus support decision representation. Sharp-Wave Ripple Forward Replay of the available direction options supports the option evaluation and action selection steps. Sharp-Wave Ripple Backward Replay (not illustrated) supports outcome evaluation and learning; in this diagram, the rat's mental model and decision were accurate and yielded a delicious reward. The hippocampus does not exclusively support any of these steps; other regions participate. Of particular interest is how the VTA

may interact with SWR replay during these stages of model based decision making. Partially based on Rangel et al., 2008.

4.1.1 Motivation

How does a brain implement the feedback loop of prediction, action and learning?

Here, we restrict this question to the investigation of an animal navigating through a simple environment with limited behavioral choices in an effort to obtain a desirable outcome. It is well established that the hippocampus provides the brain with a cognitive map and memory functions, both of which contribute components of state evaluation and learning (see Chapters 1, 2 & 3). Of particular interest is the potential for dopaminergic signaling arising from the VTA to contribute to this process.

DA neurons in VTA fire phasically when an unexpected positive event occurs, but less so when an expected positive event occurs (Schultz et al., 1993). Further, VTA DA neurons also respond to predictors of reward, such as the “bell” in a Pavlovian conditioning paradigm (Schultz et al., 1997; Tobler et al., 2005). These findings led to the reward prediction error (RPE) hypothesis, which postulates that phasic mesolimbic DA encodes mismatches between expected reward and actual outcomes (Montague et al., 1996; Day et al., 2007; Colombo, 2014). But unexpected rewards and predictors of reward may also activate incentive salience signals that motivate animals to exert effort to pursue rewards (Berridge & Robinson, 1998; Salamone & Correa, 2012).

Some FSCV studies from behaving animals report that phasic DA release in NAcc resembles a prediction error signal (Hart et al., 2014), while other studies have reported that NAcc DA responses are more consistent with a signal for incentive motivation (Howe et al., 2013; Collins et al., 2016). We hypothesize that simultaneous monitoring of DA release and

multichannel electrophysiology in the behaving animal will yield further insight into the role of DA.

4.1.2 Connecting Place to Reward

Hippocampal subregions CA1 and CA3, both project to lateral septum (LS) (Swanson & Cowen, 1979; Luo et al., 2013; Tingley & Buzsáki, 2019), providing a relatively short pathway for spatial information to flow down into the VTA. The circuitry is disinhibitory, with glutamatergic inputs from the CA3 exciting GABAergic LS cells. These GABAergic LS cells inhibit GABAergic cells in the tail of the VTA (the rMTG), disinhibiting dopaminergic projection neurons (Lou et al., 2013).

Our lab (Howe et al., 2016), and other labs (Zhou et al., 1999; Leutgeb & Mizumori, 2002; Wirtshafter & Wilson, 2019; 2020) report place cell-like activity in the LS. The LS projects to the VTA (Risold & Swanson, 1997; Luo et al., 2013). The VTA contains some cells that may be place-like (Glykos & Fujisawa, 2016), although this response may be interpreted as an association with the earliest predictor of reward, as the 'place' occurs immediately after a cue for a spatial decision. DA in the NAcc, presumably arising from the VTA, has been reported to rise at a sub-phasic rate ("ramping") with proximity to reward in a spatial navigation task (Howe et al., 2013), and also during a sequence learning task (Collins et al., 2016b); these DA "ramps" might reflect increasing motivational drive to complete the behavior as the rat approaches the goal, or, alternatively, as a value or reward proximity signal (these alternative interpretations are not necessarily incompatible). Figure 4-2 summarizes the circuit of interest.




Figure 4-2 : *Hippocampus CA3 participates in a disinhibitory circuit hypothesized to boost VTA dopaminergic output. References appear in the text.*

It is also important to note that a subpopulation of NAcc neurons ramp up their firing rates as proximity to rewarding sites increases, consistent with ramp up of DA in

NAcc, and that some of these neurons phase precess more strongly against hippocampal theta than against NAcc theta (Van der Meer & Redish, 2011), supporting the view that a functional circuit coordinates hippocampal activity and DA release. Finally, forward SWR replay events during quiet wakefulness are associated with the activation of reward-responsive VTA neurons (Gomperts et al., 2015), suggesting that replay sequences may sometimes activate representations of reward. Manipulation of the posterior VTA (rMTG) can produce conditioned place preference by increasing the level of DA in the NAcc (Hipolito et al., 2011). Taken together, these findings suggest possible roles for coordination between hippocampal SWR replay and phasic DA in learning, action-planning and decision-making.

4.1.3 Experimental Plan

Hippocampus CA1 produces SWR replay events, which represent trajectories through space. These replays are hypothesized to support action planning prior to behavior, decision making during pauses in behavior, solving the credit assignment problem at the time of a rewarding outcome and memory consolidation after behavior is complete. (See chapter 1 for further details.) The CA1 SWR are initiated by volleys of input from CA3, which we hypothesize

are transmitted into the LS and then down to the VTA. This circuit may yield bursts of dopamine in the NAcc in response to SWR, as it is known that dopaminergic cells of the VTA respond to SWR by increasing their firing rates. We developed a system intended to observe SWR replays while simultaneously observing NAcc dopamine flux by combining extracellular tetrode electrophysiology with fast-scan cyclic voltammetry (FSCV) in a rat engaging in a spatial navigation task.

4.2 RESULTS

Two animals received surgeries yielding the data presented below, providing data from the first successful simultaneous FSCV and electrophysiology recordings.

One animal (v4) yielded single unit recordings from the lateral hypothalamus (Figure 4-3) alongside dopamine flux data from the nucleus accumbens (discussed in detail below). This animal also yielded LFP and some additional single units in the hippocampus (not shown).

The second animal (DA5) yielded abundant electrophysiological data, but the FSCV probes did not yield dopamine flux data due to positioning outside the Nacc. Data from this animal demonstrates recovery of multi-frequency components of LFP.

4.2.1 Histology

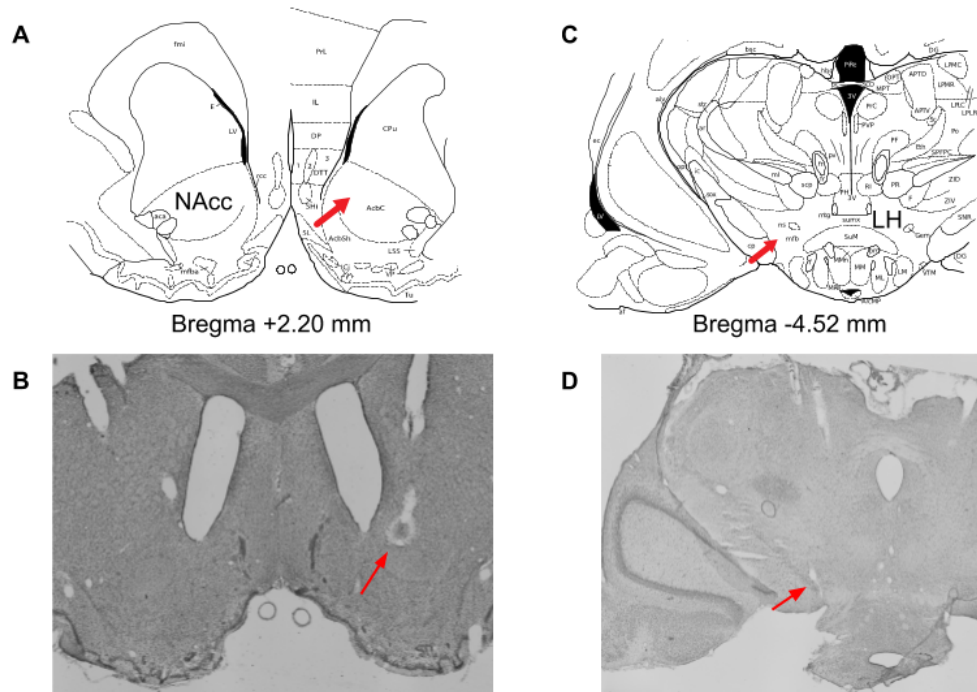


Figure 4-3 : Histology revealed recording sites in rat V4. Panels A & B depict the position of the FSCV probe in the tissue (B) at the nucleus accumbens, according to the corresponding atlas figure (A; adapted from Paxinos & Watson, 2007). The halo in B results from a marking lesion. Panels C & D depict the location of a tetrode in the lateral hypothalamus (LH) that yielded a prominent single unit. No halo accompanies the LH tetrode; sequential slices allowed tracking of probes from entry to endpoint. Red arrows indicate the location of the probes.

The headcap design bilaterally directed probes to Nacc, hippocampus and an area spanning the VTA. For the data that demonstrates simultaneous recording of dopamine and a single unit, the single animal (V4) providing this data received a headcap that successfully guided probes to the nucleus accumbens (Nacc) and lateral hypothalamus (LH). Probes also recorded from the hippocampus, but these data from this animal are not discussed further.

The multiband LFP demonstration (Figure 4-11a) comes from a different animal (codename da10); its histology is not depicted in this chapter. Electrophysiological

characteristics and histological examination both confirm electrophysiological recording in the hippocampus. The FSCV probe did not reach the Nacc in this animal, preventing simultaneous recording of a strong dopamine flux signal.

4.2.2 The FSCV Probe Occupied a Dopamine-Rich Region

Nacc contains relatively richer and poorer regions of dopamine (Wightman et al., 2007). Tonic dopamine release in the Nacc results in a constant source of dopamine for oxidation and reduction via FSCV. In regions with relatively high tonic dopamine, a characteristic distortion appears in the FSCV background current trace. We drove the probe down until we located such a prominent dopamine distortion. Figure 4-4 displays the same probe *in vitro* in PBS and *in vivo* in a dopamine rich region.

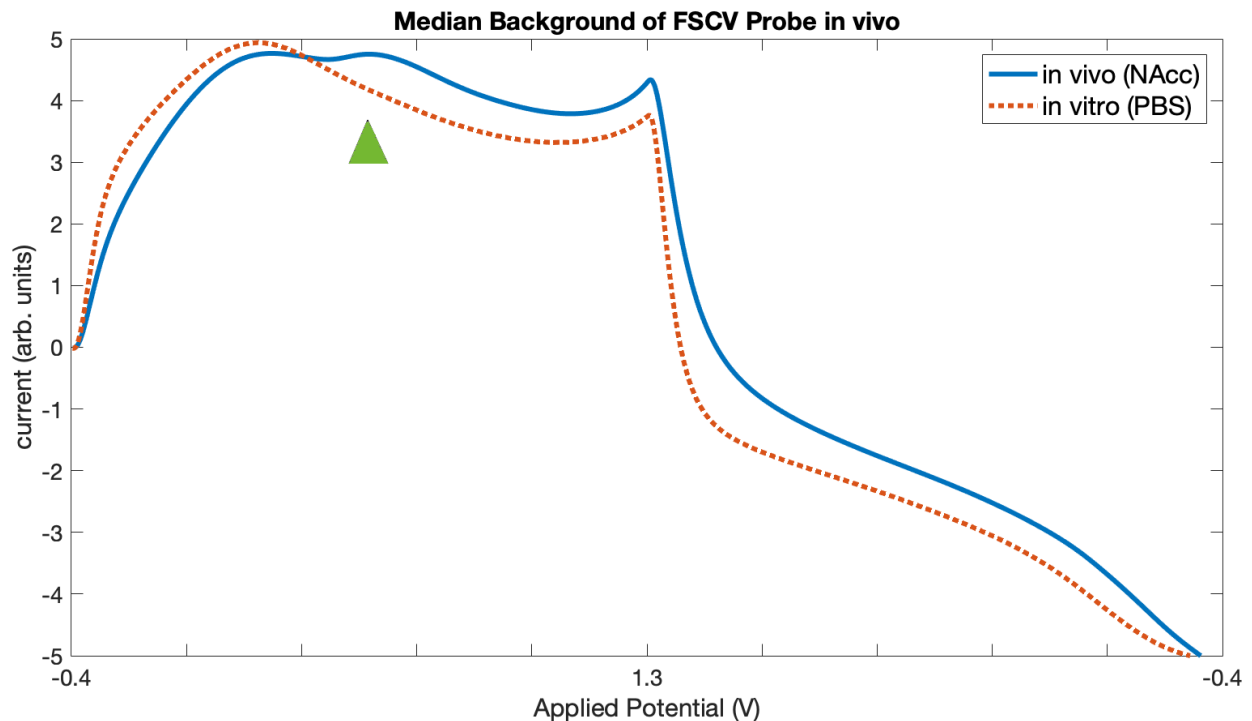


Figure 4-4 : *In vivo*, the total trace contains a prominent distortion in the waveform indicated by the green arrow. This distortion occurs at the expected potential for dopamine oxidation, suggesting that there is a relatively high tonic level of dopamine in the tissue near the probe. Each background is an average of 600 samples collected continuously during recording.

4.2.3 Dopamine Recorded Simultaneously with Electrophysiology

We simultaneously recorded dopamine in the NAcc and electrophysiology in the LH. FSCV is very sensitive to electrical noise, and the simultaneous technique increases the noise in the FSCV signal. Figure 4-5 depicts representative FSCV data from a simultaneous recording experiment. Although there is clear noise in the signal arising from recording from a behaving animal outside of a Faraday cage and connected to a separate ephys system, it is nonetheless possible to find DA voltammograms in the data. The two example voltammograms correspond to the reward triggers, which produce both sound and vibration on the maze. The animal had ample experience with the behavior, and learned to predict the reward receipt based on these conditional stimuli, as seen in electrophysiology from the VTA (Schultz et al., 1993; Schultz et al., 1997; Tobler et al., 2005) and in DA signaling from the Nacc (Hart et al., 2014). The angled lines in the color plot and cyclical noise in the DA signal trace most likely represent interference from cyclic electromagnetic interference such as the ~60 Hz wall noise, and perhaps other sources including fluorescent lights, wireless internet transmitters and ground loops with the ephys system.

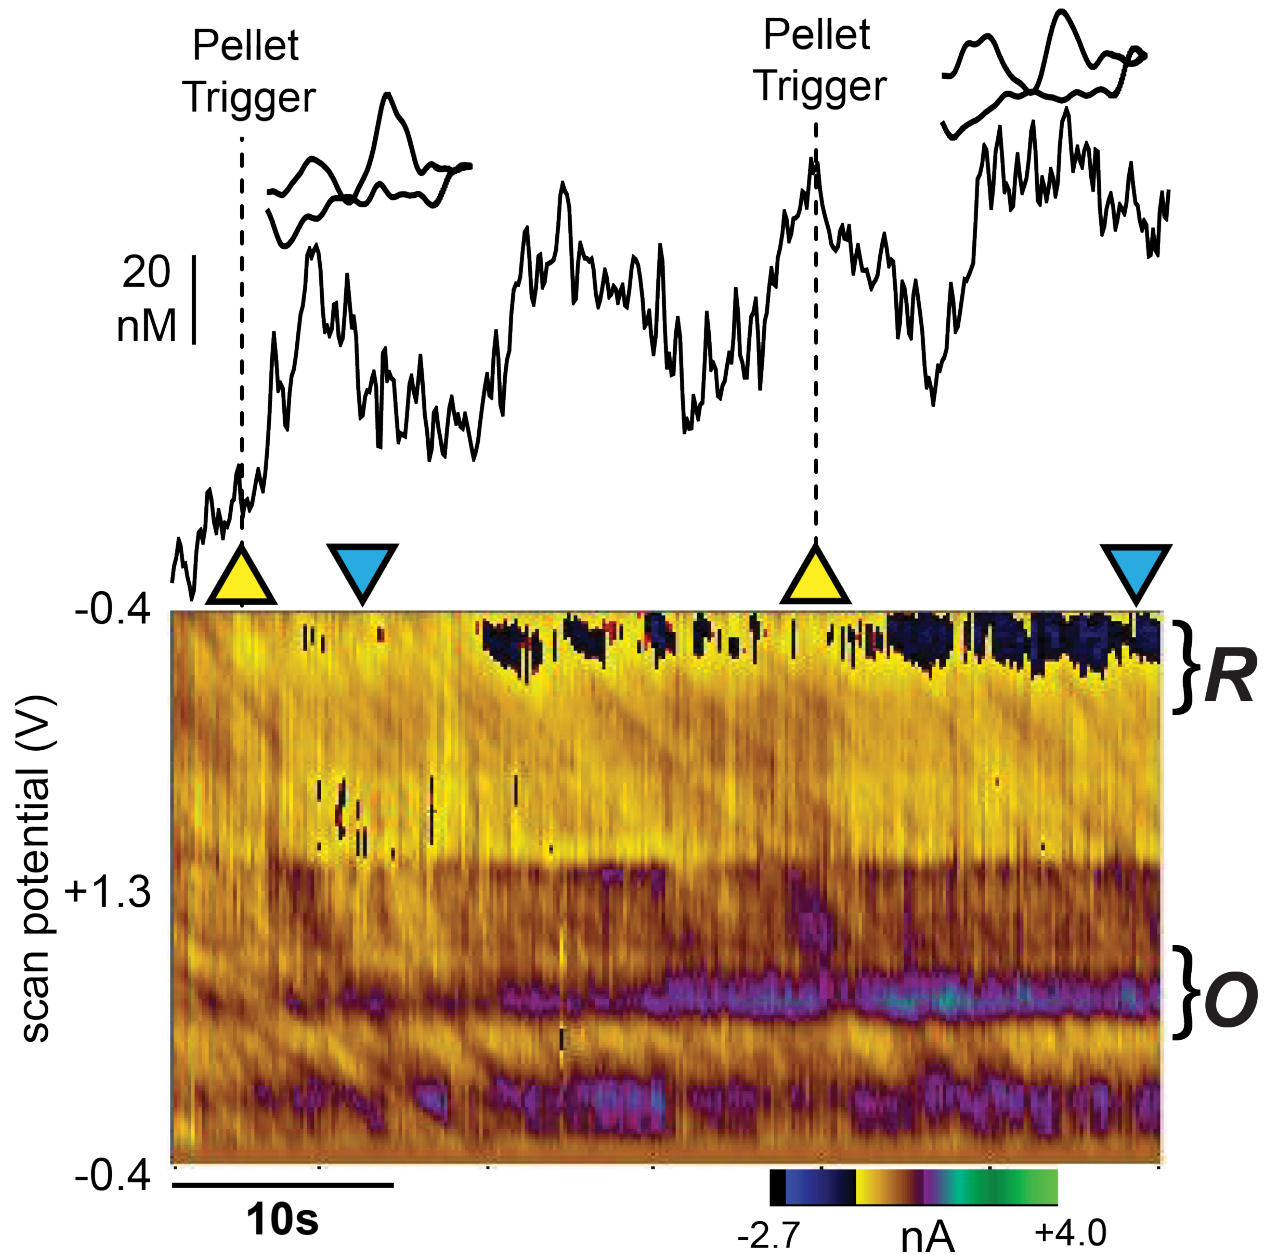


Figure 4-5 : *Representative FSCV data.* The upper half of the figure depicts the variation in the peak of the DA signal as a black line. The lower half of the figure depicts the voltage applied on the y-axis, the elapsed time on the x-axis and the background-subtracted signal as the color (z-axis). The rat produces recognizable DA voltammogram traces at the time when the reward is triggered. Yellow pyramids indicate reward delivery (“click” sound); inverted blue pyramids indicate reward collection. The behavior was not performed in a Faraday cage. DA reduction and oxidation are indicated by R & O respectively.

4.2.4 Behavior Performance

In an effort to overcome any reductions in DA signaling that might arise due to complete mastery of a task, the rat performed 2 tasks. The first task -- the platform task -- was brand new, while the second task -- the Figure 8 task -- was overtrained. The Platform task involved running to an unmarked zone on a large circular platform placed immediately overtop of the familiar Figure 8 task apparatus. Upon entry to the unmarked zone for the first day (the data shown here), dispensers automatically delivered rewards onto the maze. In the Figure 8 task, the rat navigated a Figure 8 shaped maze in an alternating pattern for rewards. (see Methods for further details.)

The rat was highly trained on the Figure 8 maze, as evidenced by its stereotyped behavior, and bias towards faster movements (Figures 4-6 & 4-7). The rat displays no such stereotyped behavior on the platform, and instead explores the entire maze, with a slight bias towards the reward delivery zones (Figures 4-6 & 4-7).

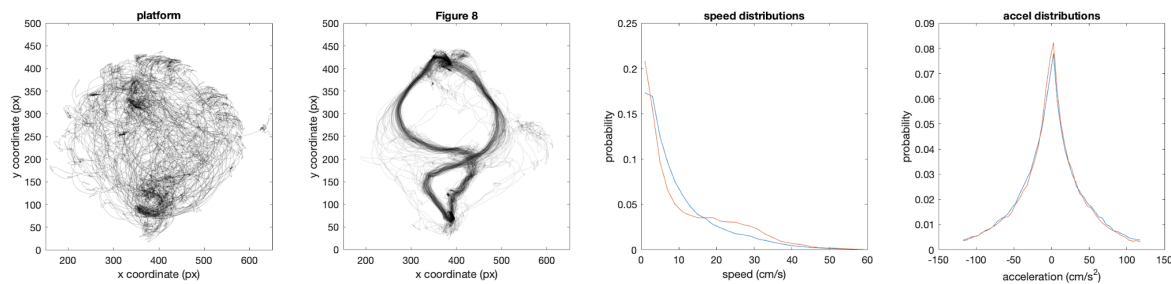


Figure 4-6 : *The rat is new to the platform task and expert at the Figure 8 task.* The 2 panels depict the behavioral traces for the rat behaving on the maze. The platter is leftmost; it is clear that the rat explores the platter significantly. The two darker zones near the center line and towards the edges of the maze are the reward delivery areas, so the rat spent more time exploring these regions. The middle left panel contains the behavioral trace for the Figure 8 maze; the North and South points are the reward delivery points. The rat's trajectory on the Figure 8 maze is very well practiced and is seemingly optimized. The

asymmetry in the 8 trajectory resulted from a behavior where the rat jumped over part of the maze. Video observation suggests that the rat adopted precise and elaborate body positioning strategies to produce the observed trajectory. The lower half of the maze is significantly brighter than the upper half, which perhaps allowed the rat to accurately judge a jump across the lower part that allowed it to shave seconds off the laps. Comparisons of velocity and acceleration appear in the right two panels. Acceleration (rightmost) does not noticeably differ. Velocity (middle right) shows a difference. The Figure 8 maze velocity probability curve reveals that the rat spends relatively more time running between 20-35 cm/s, consistent with the more stereotyped, expert-level behavior suggested by the Figure 8 position trace. Blue corresponds to the platform and red to the Figure 8 maze.

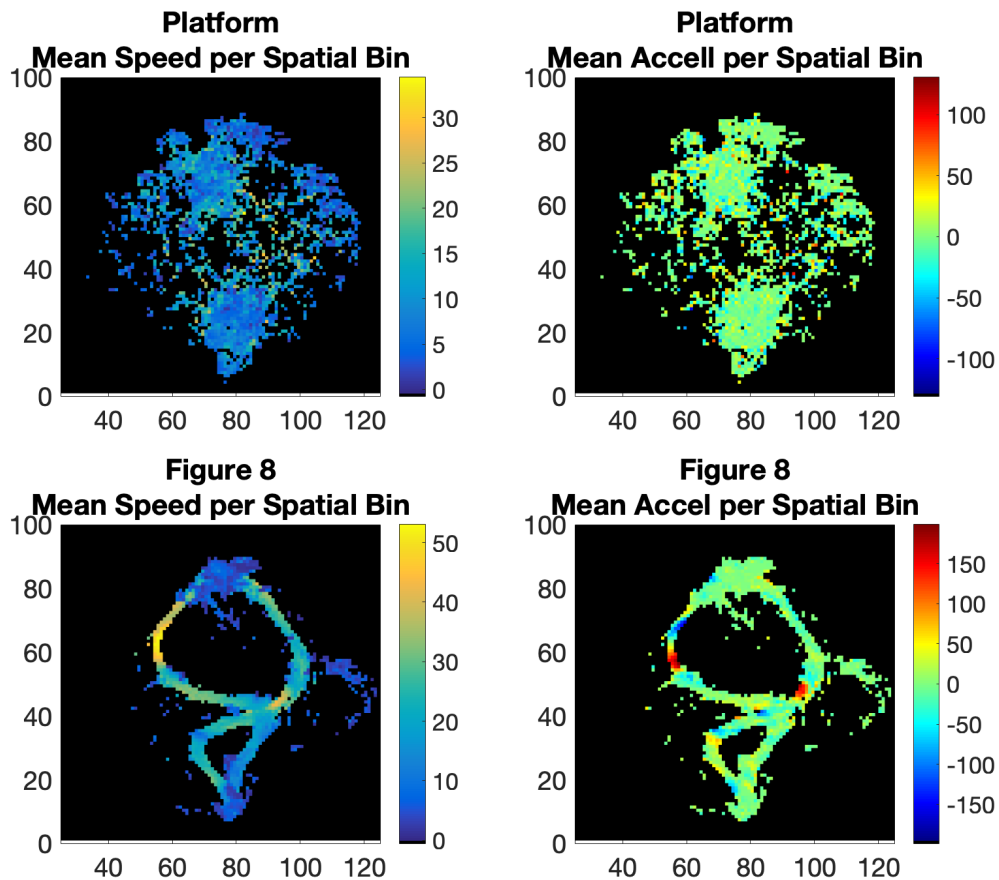


Figure 4-7 : Average speed and acceleration maps per spatial bin show no pattern on the platform, but show very clear stereotyped patterns in the Figure 8 Maze. The clear pattern on the Figure 8 maze suggests that the rat is an expert at that task.

4.2.5 Behavioral Correlates with Dopamine Flux

We observe a gradual ramp up in dopamine concentration during each “trial” (Figure 4-8). The behavior is continuous and self-paced, so this analysis treats the repeated series of behavioral steps as trials. Our observations roughly match those of other reports employing trials with readily definable trials (Howe et al., 2013; Collins et al., 2016).

An important difference arises in the timing of the dopamine ramp. In other reports, dopamine rises throughout the trial as if to motivate completion of the behavioral sequence (Salamone & Correa, 2012; Niv, 2013; Howe et al., 2013; Collins et al., 2016). This rat needed to run through space, so velocity helps mark trials and assists in the interpretation of the data. The rat’s velocity is highest prior to the onset of the rise, suggesting that most of the rise in dopamine occurs after the rat has already completed the bulk of the spatial navigation required for a trial and has begun decelerating near the reward. It is more difficult to attribute this rise in DA to a motivational component in these data.

Behavior-Triggered Dopamine

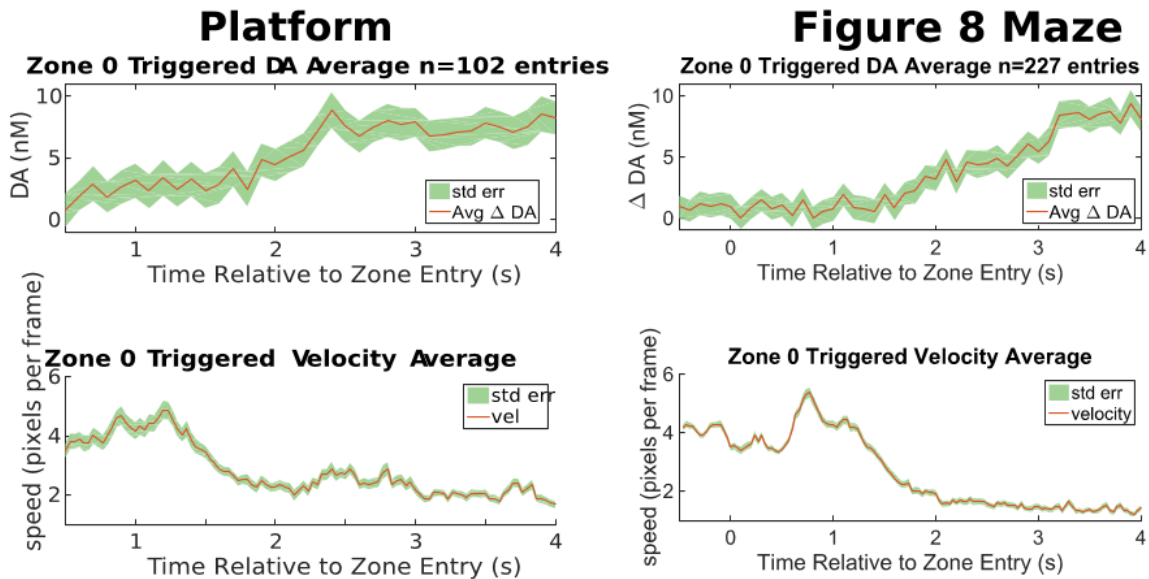


Figure 4-8 : The rat's dopamine ramps up after he enters into the "start zone" (Zone 0) of the sequence he must run. Red lines represent the mean; green bands represent standard error around the mean.

The rat is an expert in the performance of the figure 8 task, but the platform task is novel (Figure 4-9). The extremely regular, clocklike pattern of distance data from the two reward zones demonstrates that the rat is expert at the Figure 8 maze. After 16 minutes in the new task, the rat's behavior remains somewhat irregular, showing hesitations.

The proximity metric is highly correlated with the velocity in the sense that velocity drops to nearly zero as the rat dwells at a feeding site to collect a reward.

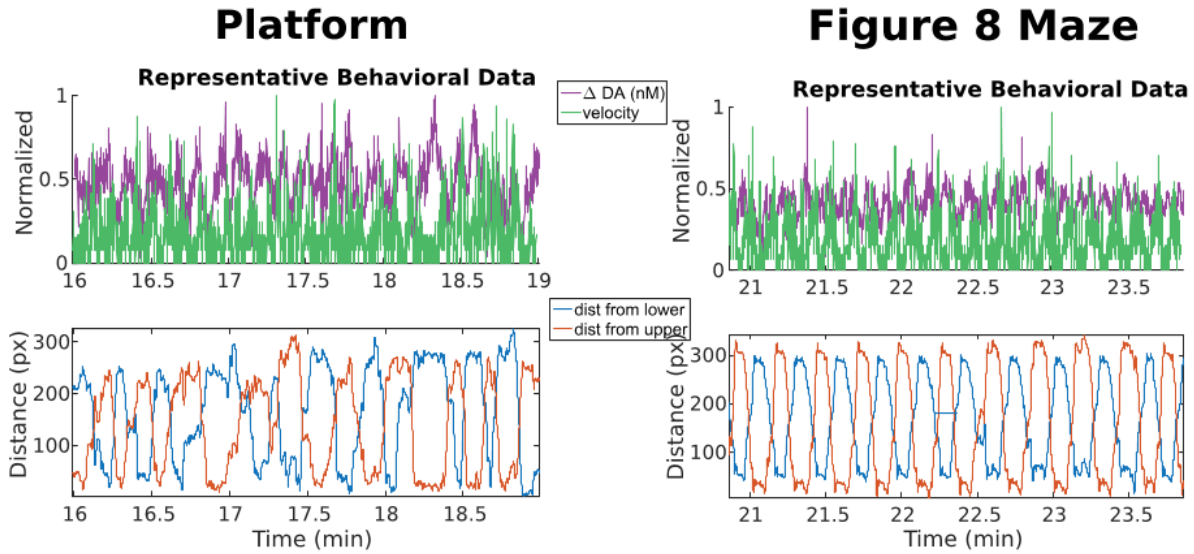


Figure 4-9 : The change in dopamine is maximal when both the velocity is minimal and the animal is closest to one of the two reward dispensers, suggesting that the dopamine peaks at each reward consumption. Comparing the plots of distance from the reward sites between the platform and the Figure 8 maze illustrates that although the rat is fairly good at exploiting the new platform task, his trajectory is nowhere near as regular as in the well-trained Figure 8 Maze.

Visually examining the traces of DA flux and velocity in the upper row of Fig 4-9 suggest that anti-correlation may exist between them. This is intuitively appealing, as the rat must at least briefly stop moving to collect its sugar pellet rewards, which one expects may result in a dopamine burst, at least in the Platform task due to the lack of internal model available to predict the reward. Figure 4-10 depicts the results of this correlation.

Similarly, it is intuitive that being close to a reward site should be associated with high dopamine. However, for parallelism with velocity, the distance from the rewards is calculated instead of the proximity such that this metric is also anticorrelated. Both mazes are bilaterally symmetric, which means that when the rat is in the middle he is equally near either reward site. Velocity is anti-correlated at a temporal offset of about +1 s, and highly correlated with about a 5 s delay, perhaps reflecting the cyclical nature of the task, and the trial times of about 7-8 s. The

rat spends about 2 s in stillness, consuming the reward, so a 4-5 s delay to peak velocity relative to DA makes sense if the DA is highest at the reward consumption site -- the rat would have a peak velocity every ~10 seconds as he transitions between reward sites.

On the platform, the correlation of the distance from the upper or lower reward site relative to the dopamine trace appears to be phase-shifted. That is, the distance traces are fairly regular, but opposite due to physical constraints. The peak correlation is about ~1 s before or after, suggesting that the dopamine trace peaks are slightly offset in time from the physical reward zones themselves. This may reflect that rat's arrival, and subsequent ~1 s search for the pellets, where DA is maximal at consumption.

On the Figure 8 maze, the correlation of the distance traces with the dopamine trace do not appear to be phase shifted versions of one another. The rat traverses less distance and performs trials faster on the lower half of the maze relative to the upper half, creating an asymmetry in the behavior. This asymmetry may contribute to the asymmetry observed here. Furthermore, the lower half of the maze is more brightly lit, and the rat engages in the riskier behavior of jumping over the gap in the maze on the lower part. The rat is engaged in this while furthest away from the upper reward site. In Fig 4-10c, DA and distance from the upper site are maximally correlated at a delay of about 2 s, and the overall correlation for the upper distance site is higher. It is possible that this upward shift represents higher motivation as the rat executes the lower half more rapidly, and with a more intense behavioral strategy (the jump). The lower trace appears to be largely anti-correlated with the behavior. Given that this is an individual rat on a specific day, it is also possible that the upper reward dispenser partially or completely jammed or ran out of rewards, leading to a situation where the rat would only experience an expected reward on the lower half of the maze, despite having to engage in the complete alternation behavior to experience any reward at all.

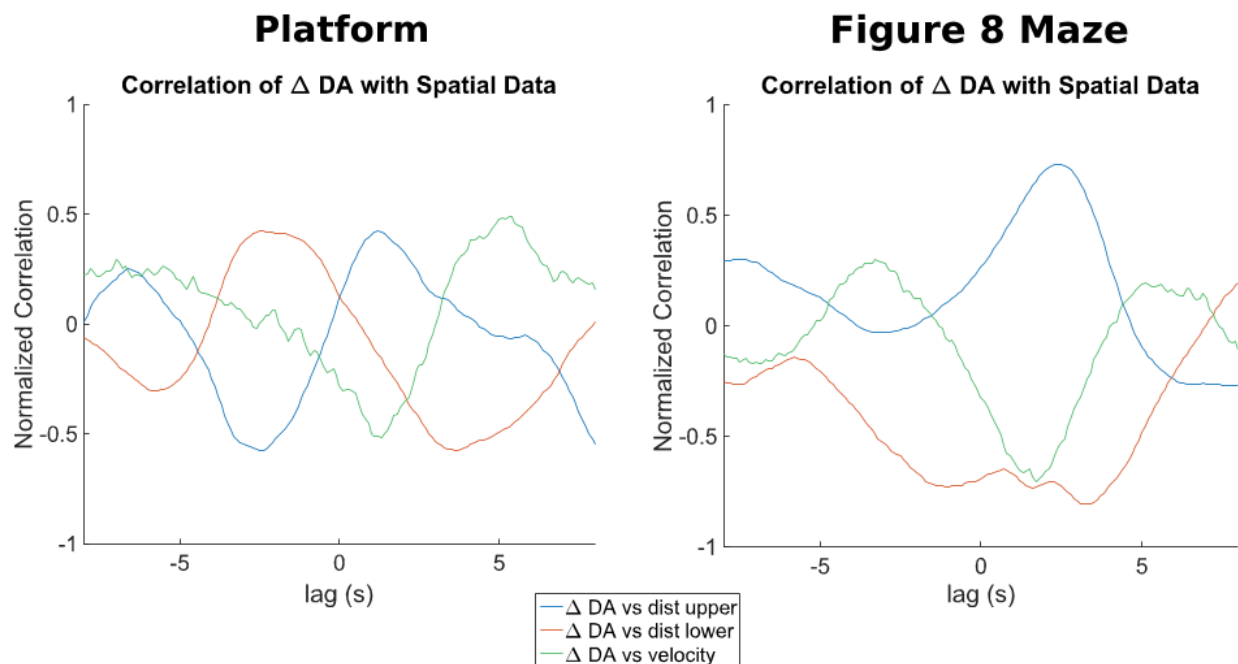


Figure 4-10 : Consistent with the prior figure, the correlation at lag 0 is zero, as the data are near antiphase. See text for a detailed interpretation.

4.2.6 Electrophysiology Recorded Simultaneously with Fast-Scan Cyclic

Voltammetry

Electrophysiological recording during FSCV scans contains regular, prominent artifacts that reflect the voltage sweep introduced by the FSCG pulse. While these pulses are significantly larger than the LFP recorded, they are significantly smaller than the capabilities of the Neuralynx recording system. These FSCV Scan-Induced Electrophysiology Artifacts (FSIEA) measured 1-3 mV in size, depending on the animal and preparation. The AtoD for the Neuralynx system provides a range of +/- 132 mV at a resolution of 24 bits, which is >100 times larger than the FSIEA observed in our recordings. No amplifiers were harmed in the collection of this data.

The FSIEA is much smaller than the full 1.7 V range of the FSCV voltage sweep due to relative referencing. The electrophysiological recording system provides great flexibility for selecting an analog reference. For general electrophysiology, we find that references in contact with the CSF spanning several millimeters of brain surface provides the best performance. This preferred reference spans 4-6 different broad sites over the cortex, and consists of all the stainless steel guide cannula, electrically unified. Analog relative referencing reduces the size of the FSCV voltage sweep to $< 1/1000$ of its true size (from 1.7 V to $\sim 1.5 \times 10^{-3}$ V).

The FSIEA itself is extremely regular, allowing recovery of LFP and single unit spiking data (Figures 4-8.a) with the algorithm described in Methods. As Figure 4-8.a depicts, signal recovery is possible for both low (Theta band, 6-10 Hz) and high (ripple band, 180 - 250 Hz) LFP. Spikes can also be cleanly recovered (Figures 4-11a/b). The algorithm works both on unfiltered raw data (Figure 4-11a) and pre-filtered data (Figure 4-11b).




Figure 4-11 : *The FSIEA artifact interferes with LFP, but can be cleanly removed. (A) Raw 24 bit LFP. The large >1 mV spikes are the FSCV artifact. (B) The unfiltered LFP (blue) and Theta band after artifact correction. (C) The unfiltered LFP with an without the artifact (grey, blue) and the SWR band (orange). (D) The SWR band alone, demonstrating a SWR event. (SWR events are significantly smaller than Theta.)*

(E) The unfiltered LFP with an without the artifact (grey, blue) and the single unit band (orange). (F) High temporal resolution depictions of spikes recovered from the de-artifacted LFP. De-artifacting does not distort the spike waveforms, as 2 distinctly different spikes appear on this channel. All data comes from a single channel in a single rat. Panels A & B, C & D and E & F share common timescales.

A single cell appeared on the LH tetrode. It was well isolated, as shown in Figure 4-12, where its average waveform and ISI are displayed.

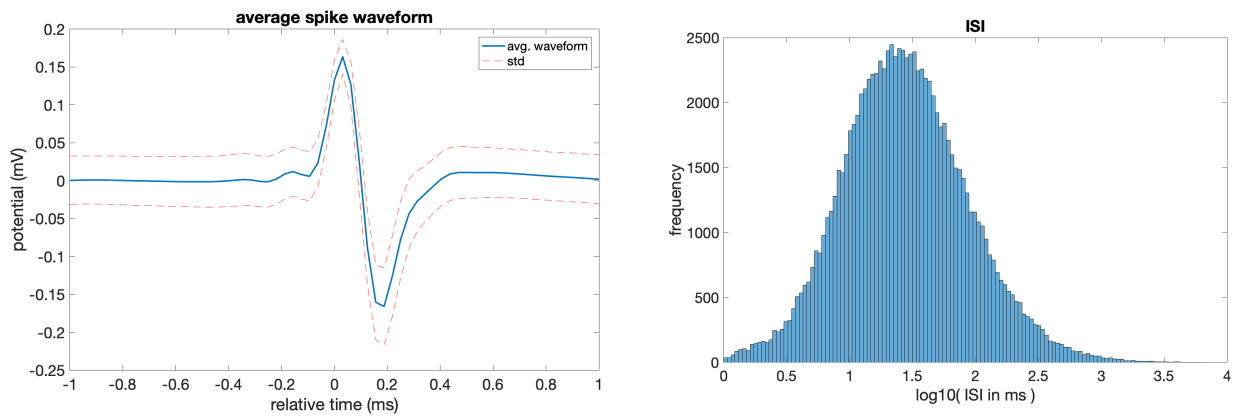


Figure 4-12 : *The average spike waveform for the single unit in the LH.* The left figure depicts the average waveform and its standard deviation over a 2 ms long window. The right figure depicts the interspike interval for the isolated cell.

4.2.7 Speed & Acceleration Correlations with Cell Firing

In order to determine what the cell might encode, spatial and behavioral correlations were examined. The cell appears to be tuned for both speed and acceleration, but it does not appear to be tuned for space (Figure 4-13). Given the confound for space with reward receipt, it is a

little surprising that the cell does not appear to obviously prefer spatial locations associated with reward.

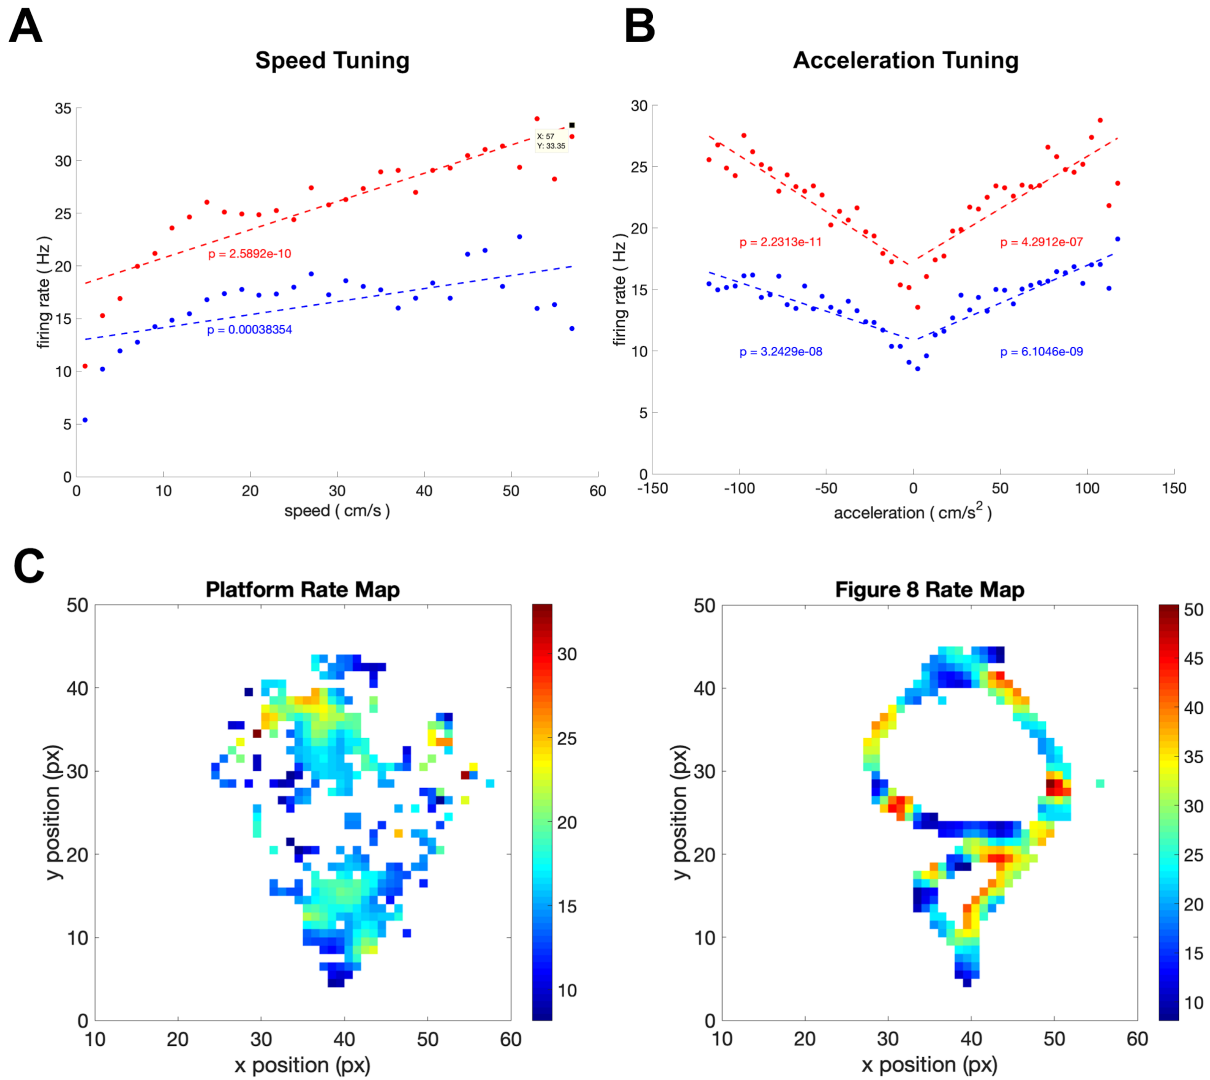


Figure 4-13 : *The LH cell's firing rate correlates with the speed and velocity of the rat in both behaviors, but it does not show strong spatial preferences. (A)* Red data are from the Figure 8 maze and blue data are from the platform. Both are $p < 0.001$. *(B)* Red data are from the Figure 8 maze and blue data are from the platform. All linear regressions are $p < 0.001$. *(C)* Rate maps for the cell in both behaviors. The left panel shows the platform, where the cell shows mildly biased firing clustered at the spatial locations where reward is delivered; despite this clustering, the cell has a dismal spatial information score of 0.0561

bits/spike (using the method from Skaggs et al., 1993). The right panel shows the rate map for the Figure 8 maze; the cell has no clear place fields in this environment, although it may have some periodicity. It has a poor spatial information score of 0.1636 here. The Figure 8 maze high firing rate zones are confounded with acceleration and preferred velocity (see above 4-7).

4.2.8 No Theta Synchronization

The LH cell does not synchronize with theta. It has no theta phase preference nor does the autocorrelegram show any periodicity (Fig. 4-14).





Figure 4-14 : *The channel 7 cell is not theta synchronized.* A theta phase histogram is essentially uniform for speeds >10 cm/s, Autocorrelegrams shown on the right hand side do not show any periodicity. Neither environment nor running speed affect this characterization.

4.2.9 No Head Direction Preference

The LH cell displays no preference for head direction, as shown in the corresponding figure (4-15).

Figure 4-15 : *The LH cell is not tuned for head direction in either environment nor at any relevant behavioral speed.*



4.2.10 Combined Single-Unit Electrophysiology with FSCV

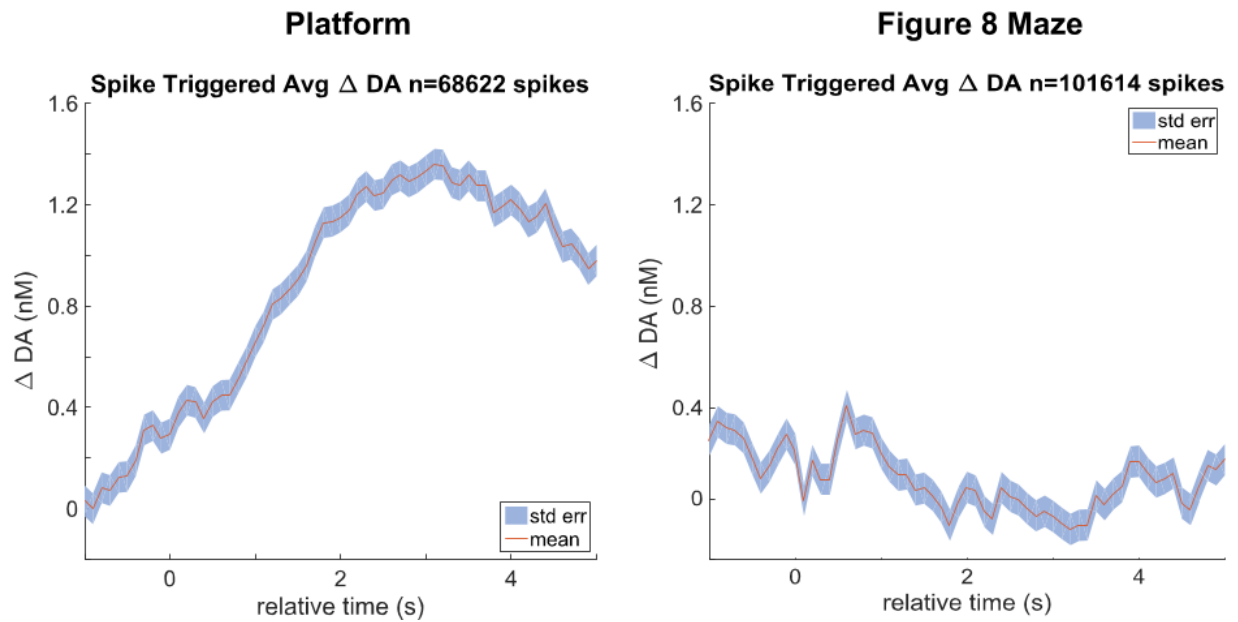


Figure 4-16 : A spike triggered average of the dopamine trace shows that levels of dopamine rise after the cell fires on the platform task, but not after the cell fires on the Figure 8 task. The rat is an expert in

the Figure 8 task, so we do not expect much DA flux in this environment because his prediction of rewards is perfect. However, the rat has no prior experience with the platform task, expressing rapid learning as the task progresses. The correlation of the cell firing with dopamine flux most likely reflects a circuit from LH to VTA to Nacc. The LH cell may respond to feeding behavior, helping to predict unexpected reward as the animal learns the task. Red lines represent the mean value of the dopamine flux. The blue surround represents standard error.

The availability of spiking neurons allows for interesting new analysis of the correlates of dopaminergic flux with cell activity. In Figure 4-16, the activity of a LH cell correlates with a rising trend of DA release in the Nacc during a novel task (platform), but not during a well-learned task (Figure 8). LH projects to VTA, and seems to disinhibit DA release in the Nacc (Stuber & Wise, 2016). LH cells are sensitive to feeding (Tang et al., 2016). Presumably, LH input concerning food consumption contributes to VTA signaling of unexpected reward receipt (Schultz et al., 1993; Schultz et al., 1997; Tobler et al., 2005). This spike-dopamine flux correlation likely reflects this signaling; it is visible in the platform because the animal is not an expert and expresses dopamine flux in response to unexpected reward, whereas there are no unexpected rewards in the Figure 8 maze, so there is no such interaction visible.

4.3 METHODS

4.3.1 Subjects

Long-Evans rats (Charles River Laboratories, Hollister, CA, USA) were housed in a temperature and humidity controlled vivarium with a 12-12 reverse light-dark cycle, and fed ad lib until they attained a weight of ~550 grams, after which they were reduced to 85% of their ad

lib weight by limited daily feeding. The rats used in the study were selected from a larger cohort of 8 rats that were all trained to perform a Figure 8 maze task prior to surgery. The three rats that were selected for surgery were the first three rats to reach a performance criterion (1 reward per minute over 20 minutes) on the Figure 8 maze.

4.3.2 Targeting Considerations

The brain regions targeted for simultaneous recording in this work are the hippocampus (CA1), the VTA and the nucleus accumbens (Nacc; summarized in Figure 4-17). We collected electrophysiology recordings in all three brain regions, and FSCV in the Nacc in hopes of addressing the following experimental questions :

- 1) VTA electrophysiology and Nacc FSCV would provide data that can help resolve some of the ambiguity surrounding to what extent control of dopamine release in the Nacc resides with the VTA vs with local (synapto-synaptic) signaling in the Nacc (Cragg et al., 2006; Cragg, 2006; Threlfell & Cragg, 2011; Threlfell et al., 2012; Cachepe et al., 2014).
- 2) Hippocampal electrophysiology and Nacc FSCV would provide data that could help clarify the results of Howe et al., (2013), and deepen understanding of how navigational task performance, the “cognitive map” of the hippocampus and dopaminergic reward signaling shape behavior learning and performance.
- 3) Nacc electrophysiology and Nacc FSCV would provide data to address the role of endogenous dopaminergic signaling on the neurons of the Nacc in an intact, behaving animal.




Figure 4-17 : Electrophysiology probes targeted Nacc, CA1 and VTA. FSCV probes targeted Nacc. Lines do not represent the number of probes.

4.3.2 Fast-Scan Cyclic Voltammetry Probes for the OvalDrive36 Headcap

Modifying the implants to accommodate both recording technologies involved several design decisions. The decisions weighed the following :

- conventional electrophysiology tetrodes are mounted on mobile shuttles
- conventional FSCV probes are fixed in place
- mobile tetrodes allow higher yields of neurons than fixed probes
- fixed probes severely limit the number of recording sites
- Striatum (including Nacc) contains microdomains with varying DA richness (Wightman et al., 2007; Schwerdt et al., 2017, 2018)
- FSCV probes are enveloped by gliosis, compromising recording quality over time (Scwerdt et al., 2017, 2018)
- for single unit studies, n is the number of cells, such that hundreds of relatively independent observations can be made with a 3-4 rats by moving probes to new sites throughout recording

- for FSCV, n is the number of animals because bulk extracellular dopamine measurements of FSCV is less independent than single units

Modifying conventional FSCV probes to fit into the Blair Lab's OvalDrive36 headcap offered several advantages. No modification for tetrode electrophysiology probes is required, reducing the amount of uncertainty in a complex experiment. Mounting FSCV probes on mobile shuttles provided a mechanism to break through gliosis and to seek dopamine rich microdomains in striatum.

To the author's knowledge, FSCV probes mounted on drivable shuttles in a behaving rat is unprecedented. The following several pages detail two major revisions of the design for these probes.




Figure 4-18 : FSCV through the OvalDrive required modifications to the classic (A) design to accommodate the omnetics connector and long cannula (B,C). Both redesigns (B,C) yield viable probes. The OvalDrive v2.0 (C) has a sensor tip which is fabricated in an identical way to the Classic design (A), but is potentially more difficult to fabricate.

4.3.2.1 Traditional FSCV Probes

Traditional FSCV probes are about 10 mm long total and consist of a carbon fiber, an epoxy cap, a borosilicate tube, silver conducting epoxy and a gold connecting pin (Fig FSCV Probes, "Classic"). These probes are implanted during stereotaxic surgery and are immobile.

4.3.2.2 OvalDrive FSCV Probes v1.0

OvalDrive v1.0 FSCV probes consisted of borosilicate capillary tubing (*Borosilicate capillary tubing, Molex PolyMicro Tech, #1068150018, ID 74 μm , OD 152 μm*) approximately 50 mm long. Under 2-propanol, I threaded carbon fiber of 7 microns in diameter and at least 60 mm long through the borosilicate tube. Tungsten wire (*Tungsten 99.95%, CFW-211-002-HML 0.002" insulated diameter, California Fine Wire Company, Grover Beach, CA*) with about 2 mm of insulation stripped off with a scalpel blade was swiped through freshly mixed 2-part silver epoxy (*Pure Silver Epoxy, MG Chemicals, 8331-14g, Toronto, Ontario, Canada*) before insertion (~2 mm) into the borosilicate tube containing the carbon fiber. The joint is cleaned by rolling a cotton-tipped swab upward, removing excess silver epoxy, leaving the probe narrow and smooth. The carbon fiber at the other end was trimmed to protrude 70-125 microns from the end of the borosilicate tube. After 24 hours of curing, the carbon fiber probe was carefully passed through the OvalDrive36 cannula backwards (tungsten fiber first entering through the side to be implanted). The carbon fiber probe was then sealed with careful application of molten wax, taking care to form a dome top and maintain the carbon fiber free of wax. (This step involved a soldering iron on the lowest heat setting, with a wire wrapped around the heating end. The tip of the wire was shaved into a sharp tip. Wax application worked best when applied to the sharp edge of the borosilicate tube, such that capillary action allowed the wax to flow into the borosilicate tube, and not up the carbon fiber.) OvalDrive v1.0 FSCV probes had a high attrition rate. I produced around 100 probes in batches, found 7 of sufficient quality to retain in the headcap and of these 7, only 2 produced signals of moderately acceptable quality in the rat.

OvalDrive36 probes are custom sized for each individual cannula. A set of informal measurements of about 36 cannula in one drive suggested that the median is approximately 48 mm +/- 2 mm. It is necessary to significantly lengthen the FSCV probes to accommodate the OvalDrive36 design. Furthermore, FSCV probes do not survive forward passage through the OvalDrive36 cannula due to the bends in the cannula and the fragility of the carbon fiber probe.

The gold connecting pin was not compatible with navigation behavior or with the OvalDrive36 requirements. It was therefore necessary to significantly modify multiple aspects of the FSCV probe design and develop a technique for mounting the probes into the drive. (Figure FSCV Probes, "Classic")

4.3.2.3 OvalDrive FSCV Probes v2.0

The final design employs the standard narrow gauge, dome shaped epoxy tip for the carbon fiber, and the tungsten connector.

Sections of borosilicate tubing (*Molex PolyMicro Tech, #1068150381, ID 20 μm , OD 89 μm*) are cut to ~30 mm with a very sharp, fresh and undamaged scalpel blade. A large (≥ 6 " diameter, low walled) petri dish (Pyrex) was filled with 2-propanol and several long carbon fibers. Filling borosilicate tubes with 2-propanol by inserting the tubes into a syringe, and expelling 2-propanol helps ensure that no air bubbles remain. This dish was placed under a dissection microscope on a white surface with an bright IKEA LED reading spotlight for illumination. Tubes were inspected at high magnification (50x) under a dissection microscope to ensure clean cuts; any cracks or breaks were discarded. The investigator passes a long carbon fiber through a narrow gauge borosilicate tube (~30 mm long) under 2-propanol. Borosilicate tubes must be completely filled with 2-propanol; loading several into a needle attached to a syringe and forcing 2-propanol through them as they are prevented from exiting the needle by abutting the ends against the corner of the petri dish ensures that they are filled with 2-propanol. After the carbon fiber is passed through the borosilicate tube, the assembly is allowed to dry.

Next, 2-part 5-minute epoxy (Devcon 5 minute epoxy #14250; ITW Polymers Adhesives, Danvers, MA) is freshly prepared and allowed to stand briefly (~30 s) before application to the carbon fiber with a tiny spherical blob on an insect pin. (The fiber is inserted into the blob and the blob is run along 5 mm of the fiber before being pulled off. The fiber is then retracted,

accumulating droplets of epoxy into a dome shaped cap at the end of the probe. This assembly is allowed to dry for at least 12 hours before the next step.

The carbon fiber is trimmed to ~70-125 microns with a sharp scalpel. Next, the other end of the probe assembly now consists of a long, free carbon fiber and the tubing. This end must be mated to a second, larger-diameter borosilicate tube (*Borosilicate capillary tubing, Molex PolyMicro Tech, #1068150020, ID 98 μm , OD 163 μm*). For clarity, the narrow borosilicate tube shall be referred to as the “carbon fiber housing”, and the larger diameter borosilicate tube shall be known as the “tungsten wire housing.” The carbon fiber housing assembly is aligned with the tungsten wire housing. To do so, the carbon fiber is placed on an extra section of narrow gauge tubing to lift it above the copy-paper work surface (the carbon fiber floats a few dozen microns above the surface). The tungsten wire housing is carefully advanced over the carbon fiber, and then over the narrow gauge tube until the free carbon fiber is protruding from the end and the total probe length is approximately 48 mm. When it is time to pass the carbon fiber housing into the tungsten wire housing, care must be taken to ensure smooth entry -- cleanly cut borosilicate tubes can easily cut the carbon fiber, making it impossible to glue the tungsten wire to the carbon fiber.

The joint between the two borosilicate tubes is sealed with molten bee’s wax; the wax is allowed to flow as far up the joint as possible through capillary action. Next, approximately 2 mm of insulation are stripped off a ~50 mm long wire (*Tungsten 99.95%, CFW-211-002-HML 0.002” insulated diameter, California Fine Wire Company, Grover Beach, CA*) with a scalpel. Tungsten wire length varies with routing considerations in the headcap. The stripped end is swiped through freshly prepared silver epoxy (*Pure Silver Epoxy, MG Chemicals, 8331-14g, Toronto, Ontario, Canada*) and then inserted into the wide end of the borosilicate tube to 2 mm depth. The carbon fiber should become embedded in the silver epoxy -- producing a conductive bond - - and the epoxy should bind tightly to the walls of the glass tubing. A strong mechanical joint is

essential for loading the probe into the drive. Excess silver epoxy must be cleared off the joint, as the joint will pass through the cannula tubing.

This assembly must dry for 24 hrs. The tungsten wire pulls out if the glue is improperly mixed, or is not allowed sufficient drying time. However, long curing times can leave the silver epoxy brittle. Prepare the drive by installing the shuttle and ensuring that the corresponding cannula is free of dust and debris. Once the probe is cured just right, feed the tungsten connector backwards up the appropriate cannula. This is easiest holding the probe in the dominant hand with latex gloves (because latex offers better grip) under a dissecting microscope with good lighting. Feed the probe as far as possible pushing from the bottom without endangering the carbon fiber probe tip. Switch to gently pulling on the tungsten connecting wire from the top to fully retract the probe until it is just inside the cannula.

An OvalDrive36 cannula experiences stress and bending during assembly, which tends to restrict the cannula slightly where the cannula exits the top of the guide cone floor. Care in planning cannula assignments, building the OvalDrive and installing the probes minimizes but cannot eliminate the possibility of breaking probes. The tight tolerance of the cannula is the main reason that one must clean all excess silver epoxy off the probe exterior, lest it become bound up in the cannula bend.

Once in place, a glob of beeswax holds the probe on the shuttle for initial testing. Once initial testing is passed, the probe is affixed to the shuttle with 2 part epoxy such that both the top exit part and also a section of the tungsten wire are captured in the glue cap to protect the fragile top joint.

The tungsten wire is threaded through the Omnetics connector and then pinned in place with a gold pin, as is done with the tetrodes.

4.3.2.4 Ag-AgCl Reference

All 3 probe types require a reference. Pure silver wire (*A-M Systems Silver Wire, PFA, 0.003" bare, 0.0055" coated, #785500, Sequim, WA*) served as a reference for both probe designs. References were prepared the day before surgery, allowing for overnight curing in bleach (Clorox). I stripped 1-2 mm of insulation off both ends of a ~50 mm length of silver wire. I carefully threaded the wire through the OvalDrive assembly and pinned the top into the Omnetics connector board. The tetrode shuttle drove the silver wire out of the drive, and into a container of fresh bleach to produce the requisite Ag-AgCl reference.

4.3.3 Testing and Calibration

Testing occurred in phases. In the first phase, the probe must pass a basic background signal shape test in a saline solution contained in a beaker. Probes that stabilized their signal and passed a basic form and size test moved on to an *in vitro* flow cell. A peristaltic pump fed 0.1 M (1x) PBS into a custom made hydraulic oscillation damper. The oscillation damper consisted of a sealed chamber filled entirely with PBS and a flexible bladder inflated with air. PBS flowed into the cell on one side and out another side. The flexibility of the bladder acts as a low pass filter, removing all evidence of the peristaltic action from the flow, minimizing noise in the flow cell from pressure waves. The smooth flow enters a mixer and distributor, allowing for the injection of boluses of dopamine-laden PBS into the stream. Finally, the solution enters the base of a custom build flow cell that mounts into ear bars on a stereotaxic device, allowing precise lowering of the headcap rig into the flow.

Each probe measured dopamine concentrations of freshly prepared 2, 4 and 8 micromolar dopamine solution 3 to 5 times. The peak detection values from each sample provided data to generate a calibration curve for each probe.

Any unstable probes or poor calibration characteristics were replaced. Figure 4-19 shows the Version 2 OvalDrive probes sensing a bolus of dopamine *in vitro*. Despite the modifications described, these probes provide high-quality signals *in vitro*.




Figure 4-19 : FSCV probes modified for the OvalDrive36 function well *in vitro*. In the center (B) is a voltammogram where a bolus of dopamine laden saline is injected into the flow cell apparatus. The prominent green ovoid shape represents the oxidation of dopamine. The left panel (A) depicts the trace over time of the peak oxidation voltage for dopamine; the dramatic rise shows the saturation and washout of the dopamine bolus. The right panel (C) depicts a single cyclic voltammogram near the peak saturation. The trace depicts a high quality signature profile for dopamine.

4.3.4 Additional FSCV Equipment & Supplies

Scott Ng-Evans custom-built the power supply, potentiostat and headstages employed in this work. UNC developed the Tarheel CV and Chemometric software for recording and analyzing FSCV data. Dell supplied the computer hardware, and National Instruments supplied two A2D/D2A boards for controlling and recording the data.

Several Matlab scripts were developed for processing FSCV data; interested parties may locate them on the internet, or can request copies directly from the author.

4.3.5 Data Analysis for FSCV In Vivo

FSCV Analysis: FSCV data was analyzed with prepackaged software written in LabVIEW (National Instruments), which extracts DA and pH changes via chemometric analysis (Heien et al., 2004, Heien et al., 2005; Keithley & Wightman, 2011). The software derives a new baseline for normalization from each 1 minute data block and the FSCV signal drifts over minutes, resulting in discontinuities between each data block. I developed custom software to remove the discontinuities and detrend the slow FSCV drift.

4.3.6 Electrophysiology

Data collection for electrophysiology employed the same methods as detailed elsewhere in this document (Chapter 2; see also [open-ephys.net OvalDrive36](#)). Briefly, custom-built tetrodes populated 16-24 of the cannula in the OvalDrive. Tetrode cannula bilaterally targeted dorsal hippocampus CA1, the nucleus accumbens and the ventral tegmental area. Targeting resulted from the entry position and angle applied to each bundle of cannula. Virtual planning and design in SolidWorks (France) allowed visualization of probe trajectories through an imported Paxinos and Watson atlas (2007). Prior to surgery, we made adjustments to the positioning of the headcap relative to Bregma if probe trajectories differed in the realized headcap from the virtual plan. Prior to sacrifice, probes remained in place and measurement data and examination helped confirm turning logs. Finally, histological sections confirmed the trajectories and final destinations of the probes.

A Neuralynx DigitalLynx SX (Bozeman, MT, USA) system acquired electrophysiology data during experiments at 32,000 Hz.

4.3.7 Resolving Electrical Noise


Electrophysiology and FSCV are not naturally compatible recording techniques. FSCV is extremely sensitive to electrical noise from both radio frequency interference and ground loop phenomenon. FSCV necessarily injects a relatively large (1.7 V sweep vs $\sim 100\text{-}250 \times 10^{-6}$ V action potential) voltage sweep into the brain to perform detection. This voltage sweep produces significant interference with the electrophysiological recordings.

To resolve some electrical noise issues likely caused by ground loops, we electrically isolated the Neuralynx system by removing it from the building's electrical grid. A 1000 CC Amp, 12V marine battery provided power to the DigitalLynx. The DigitalLynx is electrically isolated from its companion computer via fiber optic network cabling. All TTL signals traversed opto-isolator circuits. These steps completely separated the electrophysiological system from the building electrical system, removing ground loop effects.

To remove RF noise, FSCV is frequently performed in a Faraday cage. Many laboratories employ copper wire mesh for shielding, which is more effective against high frequency RF noise (Evans 1997). Testing with the neuralynx system disconnected from a subject revealed a noticeable peak around the 60 Hz, and several harmonics, suggesting that the AC frequency of the USA electrical system was the strongest source of interference (see Figure 4-20 A). Lower frequency interference appeared in the first successful recordings with FSCV and electrophysiology, suggesting contamination from the 60 Hz wall noise affected the FSCV system (Figure 4-20 B). A custom built 120x180x216 cm Faraday cage almost entirely eliminates the electrical noise according to test data from a disconnected electrophysiology amplifier. The Faraday cage consists of 1/16" steel sheet metal that forms a 6-wall, fully electrically connected box the size of a small walk in closet. A custom connector grounds the

cage to the wall or to the recording system(s), depending on which provides the lowest recording noise.

A Faraday Cage electrophysiology noise reduction



B Faraday cage FSCV noise reduction

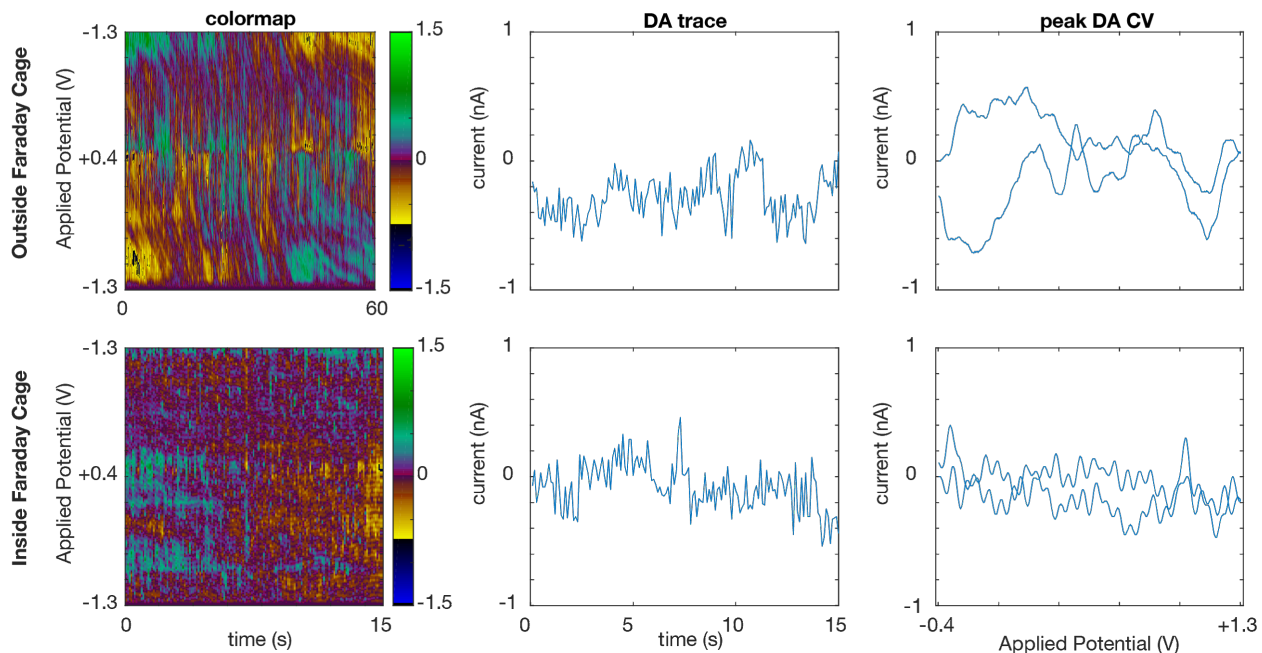


Figure 4-20 : *The Faraday cage reduces noise.* (A) Fourier transform on 223 s of data collected from a disconnected headstage in, and out, of the Faraday cage. A strong peak and many harmonics can be seen out of the cage (orange); the only peak in the cage is at 60 Hz (blue). The upper panel depicts the whole Fourier transform, and the lower panel is a zoom of the lower frequencies. (B) The upper series of plots depicts data collected from a rat at rest outside the cage, and the lower series of plots depicts data

collected from a rat at rest inside the cage. Colormaps are as in Figures 4-5 & 4-19. The noise is evident in the rope-like pattern snaking through the upper plot; the lower plot lacks this pattern. The trace across the dopamine peak voltage (middle panels) look relatively similar. The cyclic voltammogram at the peak DA trace value are noticeably different -- the plot out of the cage is much noisier than the plot in the cage.

4.3.8 Behavior

4.3.8.1 The Figure 8 Maze

The Figure 8 Maze apparatus is custom built and consists of a rigid, square, flat surface with 2 identical slots cut out to form a shape frame like shape with one long (corner-to-corner) axis connected. The maze measures 70 cm x 70 cm. The middle path connects two opposite corners. The maze is mounted on 4 detachable metal legs with leveling adjusters, raising the surface 1 m above the floor. Pathways are ~10 cm wide. The maze has ~2 cm beveled edges to discourage the rat from falling off. The surface has been painted with a waterproof black paint and is not slippery.

The maze is positioned under a tent in the center of a room with black walls. Monochrome cue posters positioned on the 4 walls provide distal cues. The two corners that are not connected by the long central span serve as reward delivery locations, and are outfitted with Med Associates pellet dispensers. The maze is entirely symmetrical. Pellets roll through a 30 cm vinyl tube, landing in a shallow glass dish glued to the maze. Illumination is provided through a LED reading light (IKEA) directed at the wall, and producing strong differential from one side to another.

Rats trained on this maze to progressively acquire an alternation behavior that consisted of unidirectional transit through the center followed by transit along one of the two “L” shaped

arms (see Figure 4-21). The task is continuous and self-paced; rats do not have to wait to initiate behavior, and high performing rats execute loops as fast as possible.

Shaping was necessary, as rats were reluctant to enter the central arm on their own. On the first and second days of training, rats explored the maze and collected randomly scattered food pellets to encourage exploration. On the third and fourth days, food pellets were exclusively scattered in the central arm, forming a Hansel and Gretel style breadcrumb trail for them to follow. No order was imposed on the rats, but shaping revealed a preference for directional passage along the central arm; this directional preference became the unidirectional requirement depending on the rat. Once rats began to enter the central arm regularly, shaping shifted to a sequence based requirement.

Sequences consist of restricted behavioral zones on the maze. Five zones defined the maze, one for each corner, and a fifth zone for the center of the central arm. In the first sequence, rats must enter the “start corner” (an unrewarded corner that determines the unilateral direction on the maze) and then the central area to trigger a reward. In the second sequence, rats needed to start in the start corner, move through the central zone and enter the corner at the far end of the central arm to trigger reward. In the third sequence, rats needed to occupy a sequence from start, center, opposite and then the opposite reward area from the previous trial. For any sequence, if the sequence is broken at any time, the rat must restart the entire sequence before a reward is dispensed. Reward is dispensed on opposite sides of the maze to encourage alternation behavior. Rats can hear and feel reward dispensing when it occurs through a subtle vibration and a click noise. Two small chocolate reward pellets dispensed for each successful round of behavior. A custom Matlab program, Neuralynx “Zoned Video” program, Neuralynx behavior tracking program and an AMPI Master 8 monitored the rat’s behavior and controlled the environment. Each rat wore a small, lightweight bicycle LED light powered by a small watch battery for automated tracking. Each rat wore a backpack-like rig

to hold the LED light in place for each session. The backpack consisted of a piece of parachute cord with several central strands removed to flatten the profile of the backpack. One side of the backpack was tied in a “Figure 8” knot (familiar to rock climbers), The other side was an open length of material that could be wrapped under the rat’s front arm such that 2 straps were formed to hold the backpack in place. Rats adjusted quickly to the backpack, and the pre-session backpack fitting helped to pre-train rats for the post-surgical process of connecting cabling during recording sessions.

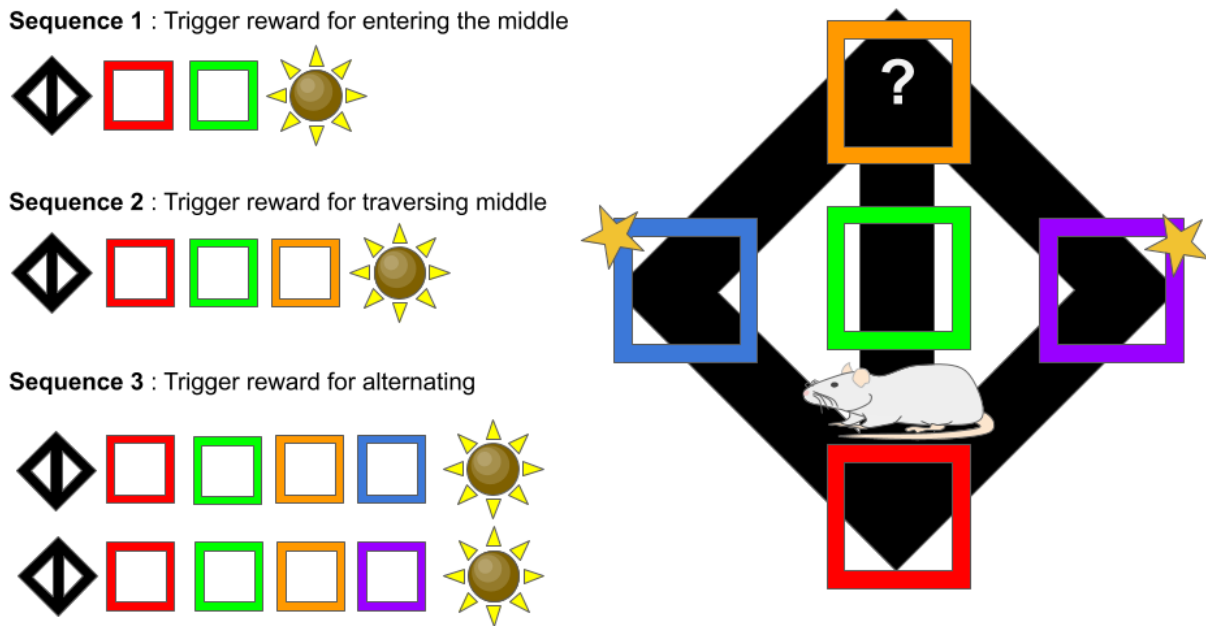


Figure 4-21 : Training sequences for the rat behavior (left) and a representation of the maze (right) with the spatial zones drawn over the maze design. Each colored box represents a spatial zone. The gold stars represent reward dispenser placements. The question mark indicates the spatial region where the rat must decide whether it will turn left or right towards a reward zone (or return down the central stem).

Dopamine responses diminish as rats become increasingly familiar with tasks (Clark et al., 2013; Collins et al., 2016). In order to ensure strong dopaminergic signaling, we introduced a simple automated task to rats after attaining expertise in the prior figure 8 maze during a few FSCV+ephys sessions.

4.3.8.2 The Open Platform

This behavior took place on a 2 m open circular platform placed directly on top of the previously described Figure 8 maze. Reward tubing output relocated to identical glass dishes on the platform.

The maze behavior similarly involved shaping. Expert rats expressed high motivation to find reward on any maze apparatus, so no initial random foraging took place (see Figure 4-22). In sequence 1, the rat simply needed to find the “start” zone to dispense a reward to both rewarded areas. After collecting these rewards, the maze reset. In the second stage, rats completed short sequences from start to reward to trigger a reward dispenser. Unlike the previous Figure 8 behavior, reward dispensing requires 2 repetitions to one side before alternation. Rats learned this behavior very quickly.

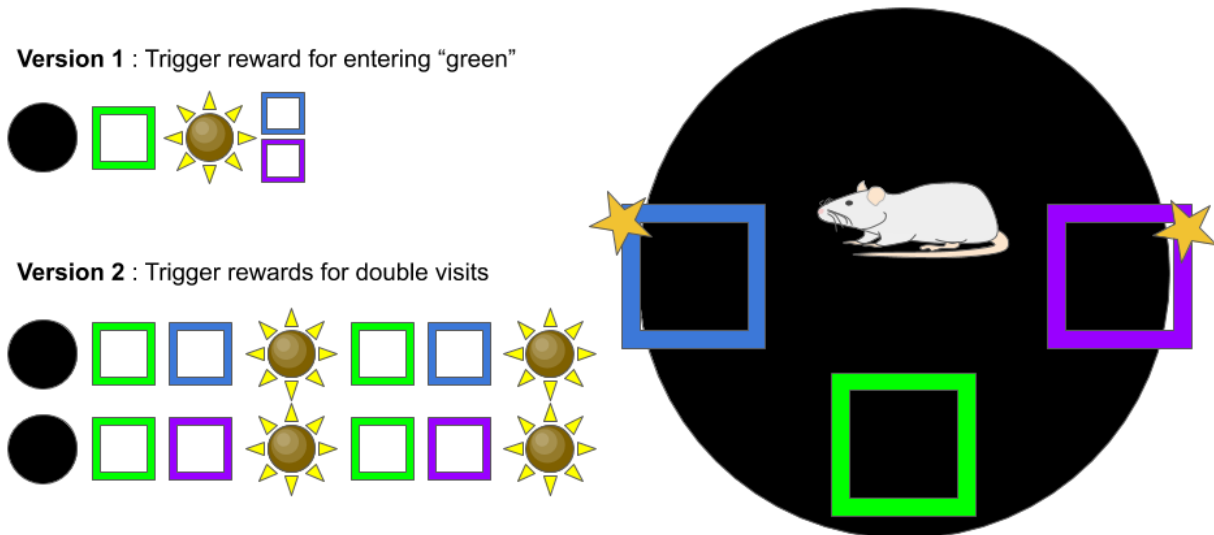


Figure 4-22 : Training sequences for the rat behavior (left) and a representation of the maze (right) with the spatial zones drawn over the maze design. Each colored box represents a spatial zone. The gold stars represent reward dispenser placements. In the first version of the task, the rat enters the green zone and rewards are dispensed at both locations. After the rat understands where rewards are dispensed (at least 1 training session), the rat must engage in a double visit alternation task to receive rewards. That is, the rat completes the sequence depicted in Version 2, where he visits blue 2x and then purple 2x in a repeating pattern, always resetting the maze by returning to the unmarked green zone between each reward visit. After attaining expertise on the Figure 8 maze, rats very rapidly acquired this new behavior.

4.3.9 System Synchronization

The TarheelCV FSCV system, the Neuralynx electrophysiology system and the Matlab environment controller are not natively synchronized. Neuralynx automatically provides unified timestamp data to synchronize the video behavior data to the electrophysiology data. The Matlab environment controller connects to the Neuralynx computer over a network. It logs all the timestamps locally, and also sends event messages back to the Neuralynx recording system,

rendering behavior and electrophysiology synchronized to the same timestamp clock. There is no direct way to link the Neuralynx timestamps to the TarheelCV FSCV system through networking. Instead, the Matlab environment controller generates a pulse train with random duration on/off periods, which is recorded as changes in the TTL signals. Furthermore, reward TTLs are mirrored to the TarheelCV system, providing multiple channels of data to synchronize the systems offline.

The pulse train durations are relatively low in temporal resolution. To obtain precise alignment, the pulse trains from Neuralynx and TarheelCV are cross correlated. The maximal correlation is taken as the proper coarse offset. A custom algorithm obtains the fine alignment by identifying the voltage pulses in the electrophysiology data, which correspond to each FSCV sample event. An acute disconnection of the FSCV headstage during recording, but after completion of behavior provided unequivocal evidence that the alignment method works properly -- a brief disconnection abruptly zeros out the FSCV sampling as well as the pulse in the electrophysiology data.

Data from FSCV and electrophysiological recordings is synchronized by aligning a Poisson train of TTL pulses (mean rate 1 Hz) that is recorded simultaneously by each system. For the present study, we will compute similar triggered averages of DA responses, but using SWR and place cell replay events as the trigger stimulus.

4.3.9 FSCV Voltage Sweep Artifact Removal

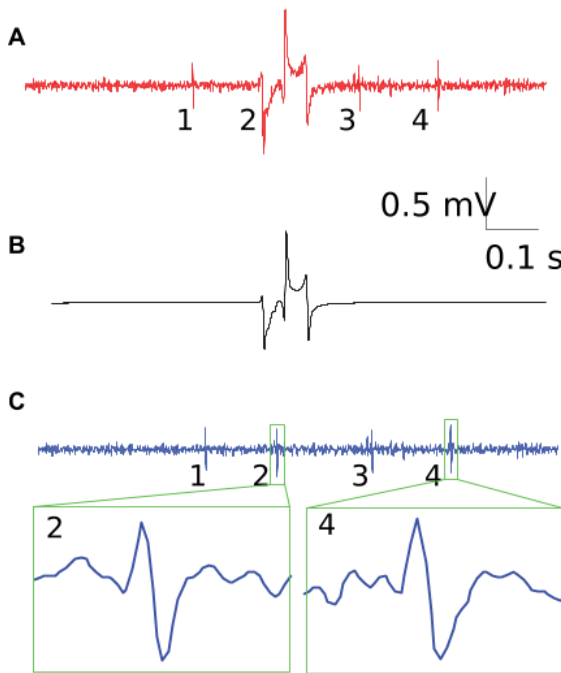
The FSCV sampling technique necessarily injects a large voltage sweep signal into the brain. Relative referencing eliminates the great majority of this signal (see Results), but not all of it. Figures 4-8 and 4-23 depict representative examples of the FSCV artifact as it appears in the electrophysiology data. This artifact signal is highly regular in time, but as it is superimposed on

the brain's electrical signaling, producing a situation where individual examples are subtly different, but the underlying signal can be recovered by averaging many examples of the artifact signal together.

The algorithm proceeds as follows :

1. Examine a 100 ms window of signal
2. Locate the maximum
3. Align the window on the max value
4. Add this window to a buffer
5. Jump forward 100 ms and repeat steps 1-4 until the end of the signal
6. Divide the buffer by the number of observed windows to obtain the average signal
7. Return to the start of the signal
8. For every 100 ms window, align and scale the average signal to the window
9. Subtract the average signal to obtain the underlying signal without distortion

Figure 4-23 : The FSIEA removal algorithm operating on pre-filtered data. Panel A depicts a 100 ms



stretch of filtered LFP trace bandpassed for the 600 - 6000 Hz range. The FSIEA appears prominently in the center. Panel B depicts the expected FSIEA shape, found by averaging every FSIEA together. Subtracting the expected FSIEA shape aligned to the raw waveform results in a clean corrected LFP (Panel C) and reveals a spike previously hidden within the FSIEA. Comparing the waveforms in the boxes suggests that Spike 2 is not simply a noise artifact born of the FSIEA correction procedure, but rather that spike waveforms can “ride” on the large FSIEA as a small distortion. The spike waveforms

can therefore be recovered when the FSIEA is removed.

4.3.10 Speed & Acceleration Analysis

Speed is estimated by calculating the distance separating the points occupied by the rat at intervals of 66, 134 and 200 ms, centered around the current time point. The median of these estimates is the speed, providing a speed estimate at approximately half the rate of positional sampling (~15 Hz). The resulting speed trace is relatively smooth while retaining relatively high temporal resolution. Acceleration is the discrete differentiation of the speed estimate over a 66.7 ms window. Speed and acceleration at each spike time is estimated by linear interpolation from the speed or acceleration trace at the time of spike production. Histograms with the bins of 2 cm/s and 5 cm/s² are constructed for all speed and acceleration states (respectively), and for spikes. The spike count is divided by the total behavioral time spent in each bin to produce an

estimate of the firing rate per bin. Linear correlation estimated whether the cell was speed or acceleration tuned. The author recognizes that the tuning curve may not be linear. This approach follows conventions used in a recent article (Wirtshafter & Wilson, 2019).

A 3D histogram mapped speed and acceleration for every time point recorded onto 5x5 pixel spatial bins. The median of the list of speeds or accelerations for a given spatial bin estimates the central tendency for the distribution of speed or acceleration at that spatial location.

4.3.11 Theta Phase Analysis

A digital filter isolated theta from the de-artifacted LFP with the strongest theta signal. The filter is an IIR of order 10 with half frequencies at 6 and 10 Hz and ran over the 32 kHz data. The angle between the filtered signal and its Hilbert transform estimated theta phase. At each spike time, linear interpolation estimated theta phase. Circular histograms show approximately uniform distributions, and Rayleigh tests suggest no significant deviation from uniformity.

4.3.12 Spatial Information

Spatial information per spike was estimated with the following equation :

$$information = \sum P_i (R_i / R) \log_2(R_i / R) \quad \text{Equation 4-1}$$

P_i is the probability of occupying a spatial bin (Skaggs et al.,; Markus et al., 1994)

R_i is the firing rate of the cell for each spatial bin

R is the average firing rate of the cell

The output is in units of bits of spatial information per spike; higher numbers correspond to greater spatial sensitivity.

4.4 DISCUSSION

This chapter describes and demonstrates a novel combined experimental technique that allows simultaneous, fully contiguous recording of extracellular electrophysiology and fast-scan cyclic voltammetry (FSCV) in the behaving rat. Combining the techniques is technically challenging; several recommendations for future practitioners appear at the end of this section.

The main technical challenges arise from two factors :

(1) FSCV is extremely sensitive to electrical noise, and connecting two separate recording systems to a single rat simultaneously can introduce noise.

(2) FSCV necessarily introduces a large voltage pulse into the brain, which directly interferes with electrophysiological recording of voltage.

The issue of noise for FSCV was resolved here by building a large Faraday cage to remove 60 Hz interference from building electricity and electrically isolating the electrophysiology recording system by providing DC power from a battery to prevent ground loop issues.

Here, we found that FSCV directly interferes with electrophysiology, but that relative referencing removes ~99.9% of the voltage deviation from baseline. The remaining deviation of about 1-3 mV easily fits within the ranges of modern electrophysiology recording equipment. Furthermore, we found that the remaining distortion introduced by the FSCV system can be removed from electrophysiology recordings with a simple technique that averages all FSCV pulses together to produce a baseline average distortion. This average artifact can be linearly

subtracted from each event to recover an undistorted electrophysiological recording trace. The extracellular electrophysiology is superimposed on the much larger FSCV voltage pulse.

4.4.1 Comparison with Alternative Designs

The system demonstrated in the following results has interest in the field. The Wightman lab developed a system which combines FSCV and electrophysiology (Takmakov et al., 2011; Cheer et al., 2007; Cheer et al., 2005). The technique employs a single carbon fiber electrode as the sensor for both electrophysiology and FSCV. Sharing an electrode allows for investigators to monitor the local DA flux in the exact spot where electrophysiological recordings occur, and provides a clever *in vivo* technique to the effect of changes in DA at endogenous levels on neurons. However, the systems switch off control over the electrode, introducing 20 ms gaps in the electrophysiological recording for every FSCV sample obtained. The single channel complicates spike sorting, and it is not clear whether obtaining LFP is possible. It is also not possible to collect large-scale electrophysiology with the implementations described by the Wightman lab.

The system implemented for this study has no data gaps for either technique. It is also possible to acquire electrophysiology data anywhere in the brain from up to 34 tetrodes (although we only attempted a maximum of 24 tetrodes, to allow for multiple carbon fibers.) Although it is possible to record from the same structure, it would be challenging to record from a tetrode and a carbon fiber in exactly the same location without further technical developments using this technique. Overall, the two approaches are complementary as currently implemented, and future investigators selecting between the two approaches should consider which approach best addresses their experimental questions.

4.4.2 Technical Recommendations for Future Combined FSCV Experiments

4.4.2.1 Experimental Design Considerations

FSCV is a technique that greatly benefits from data from at least 4-8 animals because n depends on the number of animals. Single-unit electrophysiology can be successful with only 3 animals, provided that single unit yield is high, because n depends on the number of recorded neurons, not the number of animals. Custom-made single-unit headcap designs such as the OvalDrive36 are costly both in terms of construction time and financial expenditures, but they are worth it due to the yield of single units possible. However, combining both technologies together presents a tradeoff. Optimizing for units dramatically increases the cost of every successful FSCV n .

Future progress with this combined recording technique may require finding opportunities to make tetrode recordings more efficient to allow for higher throughput. One possible solution might be chronic silicon probe arrays, allowing for simpler headcap designs that allow stereotaxic placement, rather than relying on shuttle-driven placements which are more difficult to guide for FSCV probes.

4.4.2.2 Future Technical Developments

No unified commutator exists that can pass signals from both electrophysiology and FSCV simultaneously. The lack of commutator requires a behavior that mitigates the risk of the rat twisting the cabling, such as a continuous alternation task, or a trial based task allowing for untangling to occur between each trial. Development of a commutator would open up additional behavior possibilities and help improve cable management during experiments.

Wireless electrophysiology systems could be combined with wireless FSCV systems (Roham et al., 2008; Bledsoe et al., 2009) to produce compact, unified implants. Such a system

would likely be best integrated as a System-on-a-Chip solution. Battery power for both systems would ensure low interference noise between the two systems. A wire-free system should reduce signal noise and help the rat engage in more natural behavior. Recording to a large microSD card would provide a system that could be fully isolated inside a faraday cage to minimize electrical noise without compromising data quality.

4.4.2.3 FSCV Voltage Pulse Cancellation

It may be possible to further diminish or entirely eliminate the FSCV artifact imposed on electrophysiology signals with recent technological advances in circuit design (Culaclii et al., 2018). It is also possible that recording electrophysiology with a silicon probe array and estimating the current source density to detect spikes

4.4.2.3 Alternatives to Electrophysiology

The rise of calcium imaging techniques such as fiber photometry (Adelsberger et al., 2005; Hamid et al., 2016; Mohebi et al., 2019; Mohebi & Berke, 2020) and *in vivo* micro-endoscope recordings (Ziv et al., 2013; Cai et al., 2016; Hart et al., 2019) open compelling alternatives to electrophysiology. Calcium imaging technology relies on photon transmission to monitor cell activity, not electrical potential, providing a technology that can provide insight into single cell activity without any bi-directional interference. Fiber-photometry offers higher temporal resolution, although it cannot offer not single cell resolution, and could be employed as an alternative to LFP. Furthermore, calcium imaging techniques rely on viral transfection or genetically-engineered animals, enabling the interrogation of genetically-specified cell types.

4.4.2.4 Alternatives to FSCV

Two recent genetically-encoded, light-emitting molecular monitoring tools for dopamine recently became available (Patriarchi et al., 2018; Sun et al., 2018). These tools provide an

alternative mechanism to FSCV for monitoring dopamine flux, and rely on light transmission, so they do not conflict with electrophysiology. Furthermore, these technologies enable monitoring of dopamine over a 2D surface, rather than at a single point, dramatically increasing the data dimensionality for dopamine.

4.4.2.5 Multichannel FSCV

The opportunity to record from multiple sites in striatum simultaneously is physically possible with the implant design described here. We found that our recording system was not able to record full resolution off multiple probes. Development of a system capable of multichannel recording could be interesting to study subregion differences in striatum. The design of Schwerdt et al., (2017) might be adaptable for this purpose.

4.4.2.6 Electrophysiology Data Collection Recommendations

It is possible that this combined technique worked due to the wide AtoD range of the Neuralynx system. We recommend utilizing Neuralynx's raw recording format to capture the full bit depth of the AtoD output, which enables optimal voltage resolution to capture spikes.

WORKS CITED

- Abadchi, J. K., Nazari-Ahangarkolaee, M., Gattas, S., Bermudez-Contreras, E., Luczak, A., McNaughton, B. L., & Mohajerani, M. H. (2020). Spatiotemporal patterns of neocortical activity around hippocampal sharp-wave ripples. *ELife*. <https://doi.org/10.7554/eLife.51972>
- Abbott, L. C., & Nigussie, F. (2020). Adult neurogenesis in the mammalian dentate gyrus. In *Journal of Veterinary Medicine Series C: Anatomia Histologia Embryologia*. <https://doi.org/10.1111/ahe.12496>
- Adelsberger, H., Garaschuk, O., & Konnerth, A. (2005). Cortical calcium waves in resting newborn mice. *Nature Neuroscience*. <https://doi.org/10.1038/nn1502>
- Aimone, J. B., Li, Y., Lee, S. W., Clemenson, G. D., Deng, W., & Gage, F. H. (2014). Regulation and function of adult neurogenesis: from genes to cognition. In *Physiological reviews*. <https://doi.org/10.1152/physrev.00004.2014>
- Aizawa, H., Yanagihara, S., Kobayashi, M., Niisato, K., Takekawa, T., Harukuni, R., McHugh, T. J., Fukai, T., Isomura, Y., & Okamoto, H. (2013). The synchronous activity of lateral habenular neurons is essential for regulating hippocampal theta oscillation. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.4369-12.2013>
- Alexander, A. S., Rangel, L. M., Tingley, D., & Nitz, D. A. (2018). Neurophysiological signatures of temporal coordination between retrosplenial cortex and the hippocampal formation. *Behavioral Neuroscience*. <https://doi.org/10.1037/bne0000254>
- Alme, C. B., Buzzetti, R. A., Marrone, D. F., Leutgeb, J. K., Chawla, M. K., Schaner, M. J., Bohanick, J. D., Khoboko, T., Leutgeb, S., Moser, E. I., Moser, M. B., McNaughton, B. L., & Barnes, C. A. (2010). Hippocampal granule cells opt for early retirement. *Hippocampus*. <https://doi.org/10.1002/hipo.20810>
- Alme, C. B., Miao, C., Jezek, K., Treves, A., Moser, E. I., & Moser, M. B. (2014). Place cells in the hippocampus: Eleven maps for eleven rooms. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1421056111>
- Amaral, D., & Lavenex, P. (2007). Hippocampal Neuroanatomy. In *The Hippocampus Book*. <https://doi.org/10.1017/CBO9781107415324.004>
- Atallah, H. E., McCool, A. D., Howe, M. W., & Graybiel, A. M. (2014). Neurons in the ventral striatum exhibit cell-type-specific representations of outcome during learning. *Neuron*. <https://doi.org/10.1016/j.neuron.2014.04.021>

- Azevedo, E., Pomeranz, L., Cheng, J., Schneeberger, M., Vaughan, R., Stern, S., Tan, B., Doerig, K., Greengard, P., & Friedman, J. (2019). A Role of Drd2 Hippocampal Neurons in Context-Dependent Food Intake. *Neuron*, *102*(4), 873–886. <https://doi.org/10.1016/j.neuron.2019.03.011>
- Balleine, B. W., Delgado, M. R., & Hikosaka, O. (2007). The role of the dorsal striatum in reward and decision-making. In *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.1554-07.2007>
- Battaglia, F. P., Sutherland, G. R., & McNaughton, B. L. (2004). Hippocampal sharp wave bursts coincide with neocortical “up-state” transitions. *Learning and Memory*. <https://doi.org/10.1101/lm.73504>
- Bellmaker, R., Wald, D., Belmaker, R. H., & Wald, D. (1977). Haloperidol in normals. *Brazilian Journal of Psychiatry*, *131*, 222–223. <https://doi.org/10.1192/bjp.131.2.222b>
- Benardo, L. S., & Prince, D. A. (1982). Dopamine action on hippocampal pyramidal cells. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.02-04-00415.1982>
- Benardo, L. S., & Prince, D. A. (1982). Dopamine modulates a Ca²⁺-activated potassium conductance in mammalian hippocampal pyramidal cells. *Nature*. <https://doi.org/10.1038/297076a0>
- Bender, F., Gorbati, M., Cadavieco, M. C., Denisova, N., Gao, X., Holman, C., Korotkova, T., & Ponomarenko, A. (2015). Theta oscillations regulate the speed of locomotion via a hippocampus to lateral septum pathway. *Nature Communications*. <https://doi.org/10.1038/ncomms9521>
- Benhamou, L., Kehat, O., & Cohen, D. (2014). Firing pattern characteristics of tonically active neurons in rat striatum: Context dependent or species divergent? *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.1798-13.2014>
- Berridge, K. C., & Robinson, T. E. (1998). What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? In *Brain Research Reviews*. [https://doi.org/10.1016/S0165-0173\(98\)00019-8](https://doi.org/10.1016/S0165-0173(98)00019-8)
- Blair, H. (Tad), DeGuzman, R. M., Kim, E., Howe, A., & Blair, G. (n.d.). *OvalDrive36 / OvalDrive18-ES*. <https://Open-Ephys.Org/>. <https://open-phys.atlassian.net/wiki/spaces/OEW/pages/27197443/OvalDrive36+OvalDrive18-ES>
- Bledsoe, J. M., Kimble, C. J., Covey, D. P., Blaha, C. D., Agnesi, F., Mohseni, P., Whitlock, S., Johnson, D. M., Horne, A., Bennet, K. E., Lee, K. H., & Garris, P. A. (2009). Development of the Wireless Instantaneous Neurotransmitter Concentration System for intraoperative neurochemical monitoring using fast-scan cyclic voltammetry: Technical note. *Journal of Neurosurgery*. <https://doi.org/10.3171/2009.3.JNS081348>

- Broussard, J. I., Yang, K., Levine, A. T., Tsetsenis, T., Jenson, D., Cao, F., Garcia, I., Arenkiel, B. R., Zhou, F. M., De Biasi, M., & Dani, J. A. (2016). Dopamine Regulates Aversive Contextual Learning and Associated In Vivo Synaptic Plasticity in the Hippocampus. *Cell Reports*. <https://doi.org/10.1016/j.celrep.2016.01.070>
- Brouwer, N., Van Dijken, H., Ruiters, M. H. J., Van Willigen, J. D., & Ter Horst, G. J. (1992). Localization of dopamine D2 receptor mRNA with non-radioactive in situ hybridization histochemistry. *Neuroscience Letters*. [https://doi.org/10.1016/0304-3940\(92\)90378-K](https://doi.org/10.1016/0304-3940(92)90378-K)
- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., & Wilson, M. A. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.18-18-07411.1998>
- Bubic, A., Yves von Cramon, D., & Schubotz, R. I. (2010). Prediction, cognition and the brain. In *Frontiers in Human Neuroscience*. <https://doi.org/10.3389/fnhum.2010.00025>
- Buzsáki, G. (2006). Rhythms of the Brain. In *Rhythms of the Brain*. <https://doi.org/10.1093/acprof:oso/9780195301069.001.0001>
- Buzsáki, G. (2002). Theta oscillations in the hippocampus. In *Neuron*. [https://doi.org/10.1016/S0896-6273\(02\)00586-X](https://doi.org/10.1016/S0896-6273(02)00586-X)
- Buzsáki, G. (2015). Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*, 25(10), 1073–1188. <https://doi.org/10.1002/hipo.22488>
- Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/6.2.81>
- Buzsáki, G., & Buzsáki, G. (1998). Memory consolidation during sleep: a neurophysiological perspective. *Journal of Sleep Research*, 7(S1), 17–23. <https://doi.org/10.1046/j.1365-2869.7.s1.3.x>
- CA, B., BL, M., SJ, M., BW, L., & LH, L. (1990). - Comparison of spatial and temporal characteristics of neuronal activity in. - *Prog Brain Res* 1990;83:287-300.
- Cachope, R., & Cheer, J. F. (2014). Local control of striatal dopamine release. *Frontiers in Behavioral Neuroscience*. <https://doi.org/10.3389/fnbeh.2014.00188>

- Cai, D. J., Aharoni, D., Shuman, T., Shobe, J., Biane, J., Song, W., Wei, B., Veshkini, M., La-Vu, M., Lou, J., Flores, S. E., Kim, I., Sano, Y., Zhou, M., Baumgaertel, K., Lavi, A., Kamata, M., Tuszynski, M., Mayford, M., ... Silva, A. J. (2016). A shared neural ensemble links distinct contextual memories encoded close in time Recent studies suggest that a shared neural ensemble may link distinct memories encoded close in time. *Nature*. <https://doi.org/10.1038/nature17955>
- Cappaert, N. L. M., Van Strien, N. M., & Witter, M. P. (2015). Hippocampal Formation. In *The Rat Nervous System: Fourth Edition*. <https://doi.org/10.1016/B978-0-12-374245-2.00020-6>
- Castañé Anna, A., Theobald, D. E. H., & Robbins, T. W. (2010). Selective lesions of the dorsomedial striatum impair serial spatial reversal learning in rats. *Behavioural Brain Research*. <https://doi.org/10.1016/j.bbr.2010.02.017>
- Cazala, P., Galey, D., & Durkin, T. (1988). Electrical self-stimulation in the medial and lateral septum as compared to the lateral hypothalamus: Differential intervention of reward and learning processes? *Physiology and Behavior*. [https://doi.org/10.1016/0031-9384\(88\)90345-9](https://doi.org/10.1016/0031-9384(88)90345-9)
- Cenquizca, L. A., & Swanson, L. W. (2007). Spatial organization of direct hippocampal field CA1 axonal projections to the rest of the cerebral cortex. *Brain Research Reviews*, 56(1), 1–26. <https://doi.org/10.1016/j.brainresrev.2007.05.002>
- Chang, L., & Tsao, D. Y. (2017). The Code for Facial Identity in the Primate Brain. *Cell*. <https://doi.org/10.1016/j.cell.2017.05.011>
- Chawla, M. K., Guzowski, J. F., Ramirez-Amaya, V., Lipa, P., Hoffman, K. L., Marriott, L. K., Worley, P. F., McNaughton, B. L., & Barnes, C. A. (2005). Sparse, environmentally selective expression of Arc RNA in the upper blade of the rodent fascia dentata by brief spatial experience. *Hippocampus*. <https://doi.org/10.1002/hipo.20091>
- Chawla, M. K., Sutherland, V. L., Olson, K., McNaughton, B. L., & Barnes, C. A. (2018). Behavior-driven arc expression is reduced in all ventral hippocampal subfields compared to CA1, CA3, and dentate gyrus in rat dorsal hippocampus. *Hippocampus*. <https://doi.org/10.1002/hipo.22820>
- Chee, S. S. A., & Menard, J. L. (2011). Lesions of the dorsal lateral septum do not affect neophagia in the novelty induced suppression of feeding paradigm but reduce defensive behaviours in the elevated plus maze and shock probe burying tests. *Behavioural Brain Research*. <https://doi.org/10.1016/j.bbr.2011.02.027>
- Cheer, J. F., Aragona, B. J., Heien, M. L. A. V., Seipel, A. T., Carelli, R. M., & Wightman, R. M. (2007). Coordinated Accumbal Dopamine Release and Neural Activity Drive Goal-Directed Behavior. *Neuron*. <https://doi.org/10.1016/j.neuron.2007.03.021>

- Cheer, J. F., Heien, M. L. A. V., Garris, P. A., Carelli, R. M., & Wightman, R. M. (2005). Simultaneous dopamine and single-unit recordings reveal accumbens GABAergic responses: Implications for intracranial self-stimulation. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.0509607102>
- Chen, B. T., Bowers, M. S., Martin, M., Hopf, F. W., Guillory, A. M., Carelli, R. M., Chou, J. K., & Bonci, A. (2008). Cocaine but Not Natural Reward Self-Administration nor Passive Cocaine Infusion Produces Persistent LTP in the VTA. *Neuron*. <https://doi.org/10.1016/j.neuron.2008.05.024>
- Clark, J. J., Collins, A. L., Sanford, C. A., & Phillips, P. E. M. (2013). Dopamine encoding of pavlovian incentive stimuli diminishes with extended training. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.5119-12.2013>
- Clarke, A., & File, S. E. (1982). Selective neurotoxin lesions of the lateral septum: Changes in social and aggressive behaviours. *Pharmacology, Biochemistry and Behavior*. [https://doi.org/10.1016/0091-3057\(82\)90334-3](https://doi.org/10.1016/0091-3057(82)90334-3)
- Cohen, J. Y., Haesler, S., Vong, L., Lowell, B. B., & Uchida, N. (2012). Neuron-type-specific signals for reward and punishment in the ventral tegmental area. In *Nature*. <https://doi.org/10.1038/nature10754>
- Collins, A. L., Aitken, T. J., Greenfield, V. Y., Ostlund, S. B., & Wassum, K. M. (2016). Nucleus accumbens acetylcholine receptors modulate dopamine and motivation. *Neuropsychopharmacology*. <https://doi.org/10.1038/npp.2016.81>
- Collins, A. L., Greenfield, V. Y., Bye, J. K., Linker, K. E., Wang, A. S., & Wassum, K. M. (2016). Dynamic mesolimbic dopamine signaling during action sequence learning and expectation violation. *Scientific Reports*. <https://doi.org/10.1038/srep20231>
- Colombo, M. (2014). Deep and beautiful. The reward prediction error hypothesis of dopamine. *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences*. <https://doi.org/10.1016/j.shpsc.2013.10.006>
- Cragg, S. J., Exley, R., & Clements, M. A. (2006). Striatal Acetylcholine Control of Reward-Related Dopamine Signalling. In *The Basal Ganglia VIII*. https://doi.org/10.1007/0-387-28066-9_9
- Csicsvari, J., Hirase, H., Mamiya, A., & Buzsáki, G. (2000). Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron*, 28(2), 585–594. [https://doi.org/10.1016/S0896-6273\(00\)00135-5](https://doi.org/10.1016/S0896-6273(00)00135-5)

- Csicsvari, J., O'Neill, J., Allen, K., & Senior, T. (2007). Place-selective firing contributes to the reverse-order reactivation of CA1 pyramidal cells during sharp waves in open-field exploration. *European Journal of Neuroscience*. <https://doi.org/10.1111/j.1460-9568.2007.05684.x>
- Cui, G., Jun, S. B., Jin, X., Pham, M. D., Vogel, S. S., Lovinger, D. M., & Costa, R. M. (2013). Concurrent activation of striatal direct and indirect pathways during action initiation. *Nature*, *494*(7436), 238–242. <https://doi.org/10.1038/nature11846>
- Culaclii, S., Kim, B., Lo, Y. K., Li, L., & Liu, W. (2018). Online Artifact Cancellation in Same-Electrode Neural Stimulation and Recording Using a Combined Hardware and Software Architecture. *IEEE Transactions on Biomedical Circuits and Systems*. <https://doi.org/10.1109/TBCAS.2018.2816464>
- Danielson, N. B. B., Kaifosh, P., Zaremba, J. D. D., Lovett-Barron, M., Tsai, J., Denny, C. A. A., Balough, E. M. M., Goldberg, A. R. R., Drew, L. J. J., Hen, R., Losonczy, A., & Kheirbek, M. A. A. (2016). Distinct Contribution of Adult-Born Hippocampal Granule Cells to Context Encoding. *Neuron*. <https://doi.org/10.1016/j.neuron.2016.02.019>
- Davidson, T. J., Kloosterman, F., & Wilson, M. A. (2009). Hippocampal Replay of Extended Experience. *Neuron*, *63*(4), 497–507. <https://doi.org/10.1016/j.neuron.2009.07.027>
- Day, J. J., Roitman, M. F., Wightman, R. M., & Carelli, R. M. (2007). Associative learning mediates dynamic shifts in dopamine signaling in the nucleus accumbens. *Nature Neuroscience*. <https://doi.org/10.1038/nn1923>
- Devan, B. D., McDonald, R. J., & White, N. M. (1999). Effects of medial and lateral caudate-putamen lesions on place- and cue- guided behaviors in the water maze: Relation to thigmotaxis. *Behavioural Brain Research*. [https://doi.org/10.1016/S0166-4328\(98\)00107-7](https://doi.org/10.1016/S0166-4328(98)00107-7)
- Devan, B. D., & White, N. M. (1999). Parallel information processing in the dorsal striatum: Relation to hippocampal function. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.19-07-02789.1999>
- Diamantaki, M., Frey, M., Berens, P., Preston-Ferrer, P., & Burgalossi, A. (2016). Sparse activity of identified dentate granule cells during spatial exploration. *ELife*. <https://doi.org/10.7554/elife.20252>
- Diamantaki, M., Frey, M., Preston-Ferrer, P., & Burgalossi, A. (2016). Priming spatial activity by single-cell stimulation in the dentate gyrus of freely moving rats. *Current Biology*. <https://doi.org/10.1016/j.cub.2015.12.053>
- Diba, K., & Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, *10*(10), 1241–1242. <https://doi.org/10.1038/nn1961>

- Dragoi, G., & Buzsáki, G. (2006). Temporal Encoding of Place Sequences by Hippocampal Cell Assemblies. *Neuron*, *50*(1), 145–157.
<https://doi.org/10.1016/j.neuron.2006.02.023>
- Dragoi, G., & Tonegawa, S. (2011). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*. <https://doi.org/10.1038/nature09633>
- Dragoi, G., & Tonegawa, S. (2013). Distinct preplay of multiple novel spatial experiences in the rat. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1306031110>
- Dudman, J. T., & Gerfen, C. R. (2015). The Basal Ganglia. In *The Rat Nervous System: Fourth Edition*. <https://doi.org/10.1016/B978-0-12-374245-2.00017-6>
- Dupret, D., O'Neill, J., Pleydell-Bouverie, B., & Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*. <https://doi.org/10.1038/nn.2599>
- Edwards, S., Whisler, K. N., Fuller, D. C., Orsulak, P. J., & Self, D. W. (2007). Addiction-related alterations in D1 and D2 dopamine receptor behavioral responses following chronic cocaine self-administration. *Neuropsychopharmacology*.
<https://doi.org/10.1038/sj.npp.1301062>
- Ego-Stengel, V., & Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, *20*(1), 1–10.
<https://doi.org/10.1002/hipo.20707>
- Eichenbaum, H. (2000). A cortical–hippocampal system for declarative memory. *Nature Reviews Neuroscience*. <https://doi.org/10.1038/35036213>
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*.
<https://doi.org/10.1038/nature01964>
- Eliav, T., Geva-Sagiv, M., Yartsev, M. M., Finkelstein, A., Rubin, A., Las, L., & Ulanovsky, N. (2018). Nonoscillatory Phase Coding and Synchronization in the Bat Hippocampal Formation. *Cell*. <https://doi.org/10.1016/j.cell.2018.09.017>
- Endres, T., & Fendt, M. (2008). Inactivation of the lateral septum blocks fox odor-induced fear behavior. *NeuroReport*. <https://doi.org/10.1097/WNR.0b013e3282fb78d9>
- Evans, R. W. (1997). Design Guidelines for Shielding Effectiveness, Current Carrying Capability, and the Enhancement of Conductivity of Composite Materials. *Contractor*.

- Everitt, B. J., & Robbins, T. W. (2005). Neural systems of reinforcement for drug addiction: From actions to habits to compulsion. In *Nature Neuroscience*.
<https://doi.org/10.1038/nn1579>
- Feder, R., & Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. Part II. Hippocampal slow waves and theta cell firing during bar pressing and other behaviors. *Experimental Neurology*.
[https://doi.org/10.1016/0014-4886\(73\)90291-4](https://doi.org/10.1016/0014-4886(73)90291-4)
- Feldon, J., Rawlins, J. N. P., & Gray, J. A. (1982). Effects of lateral and medial septal lesions on response suppression maintained by response-contingent and response-independent shock. *Physiological Psychology*. <https://doi.org/10.3758/BF03327018>
- Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/1.1.1>
- Fernández-Ruiz, A., Oliva, A., de Oliveira, E. F., Rocha-Almeida, F., Tingley, D., & Buzsáki, G. (2019). Long-duration hippocampal sharp wave ripples improve memory. *Science*. <https://doi.org/10.1126/science.aax0758>
- Finnegan, R., & Becker, S. (2015). Neurogenesis paradoxically decreases both pattern separation and memory interference. *Frontiers in Systems Neuroscience*.
<https://doi.org/10.3389/fnsys.2015.00136>
- Floresco, S. B., Todd, C. L., & Grace, A. A. (2001). Glutamatergic afferents from the hippocampus to the nucleus accumbens regulate activity of ventral tegmental area dopamine neurons. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.21-13-04915.2001>
- Foster, D. J., & Wilson, M. A. (2006). Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084), 680–683.
<https://doi.org/10.1038/nature04587>
- Foster, D. J., & Wilson, M. A. (2007). Hippocampal theta sequences. *Hippocampus*.
<https://doi.org/10.1002/hipo.20345>
- Fox, S. E., Wolfson, S., & Ranck, J. B. (1986). Hippocampal theta rhythm and the firing of neurons in walking and urethane anesthetized rats. *Experimental Brain Research*.
<https://doi.org/10.1007/BF00236028>
- Frankland, P. W., & Bontempi, B. (2005). The organization of recent and remote memories. *Nature Reviews Neuroscience*. <https://doi.org/10.1038/nrn1607>
- Frankland, P. W., O'Brien, C., Ohno, M., Kirkwood, A., & Silva, A. J. (2001). α -CaMKII-dependent plasticity in the cortex is required for permanent memory. *Nature*.
<https://doi.org/10.1038/35077089>

- Frey, U., Matthies, H., Reymann, K. G., & Matthies, H. (1991). The effect of dopaminergic D1 receptor blockade during tetanization on the expression of long-term potentiation in the rat CA1 region in vitro. *Neuroscience Letters*. [https://doi.org/10.1016/0304-3940\(91\)90732-9](https://doi.org/10.1016/0304-3940(91)90732-9)
- Frey, U., Schroeder, H., & Matthies, H. (1990). Dopaminergic antagonists prevent long-term maintenance of posttetanic LTP in the CA1 region of rat hippocampal slices. *Brain Research*. [https://doi.org/10.1016/0006-8993\(90\)91578-5](https://doi.org/10.1016/0006-8993(90)91578-5)
- Freyja Ólafsdóttir, H., Barry, C., Saleem, A. B., Hassabis, D., & Spiers, H. J. (2015). Hippocampal place cells construct reward related sequences through unexplored space. *ELife*. <https://doi.org/10.7554/eLife.06063>
- Gangarossa, G., Longueville, S., De Bundel, D., Perroy, J., Hervé, D., Girault, J. A., & Valjent, E. (2012). Characterization of dopamine D1 and D2 receptor-expressing neurons in the mouse hippocampus. *Hippocampus*. <https://doi.org/10.1002/hipo.22044>
- Gasbarri, A., Sulli, A., Innocenzi, R., Pacitti, C., & Brioni, J. D. (1996). Spatial memory impairment induced by lesion of the mesohippocampal dopaminergic system in the rat. *Neuroscience*. [https://doi.org/10.1016/S0306-4522\(96\)00202-3](https://doi.org/10.1016/S0306-4522(96)00202-3)
- Gasbarri, A., Verney, C., Innocenzi, R., Campana, E., & Pacitti, C. (1994). Mesolimbic dopaminergic neurons innervating the hippocampal formation in the rat: a combined retrograde tracing and immunohistochemical study. *Brain Research*. [https://doi.org/10.1016/0006-8993\(94\)90512-6](https://doi.org/10.1016/0006-8993(94)90512-6)
- Gasbarri, A., Packard, M. G., Campana, E., & Pacitti, C. (1994). Anterograde and retrograde tracing of projections from the ventral tegmental area to the hippocampal formation in the rat. *Brain Research Bulletin*. [https://doi.org/10.1016/0361-9230\(94\)90288-7](https://doi.org/10.1016/0361-9230(94)90288-7)
- Gasbarri, A., Sulli, A., & Packard, M. G. (1997). The dopaminergic mesencephalic projections to the hippocampal formation in the rat. In *Progress in Neuro-Psychopharmacology and Biological Psychiatry*. [https://doi.org/10.1016/S0278-5846\(96\)00157-1](https://doi.org/10.1016/S0278-5846(96)00157-1)
- Gauthier, J. L., & Tank, D. W. (2018). A Dedicated Population for Reward Coding in the Hippocampus. *Neuron*. <https://doi.org/10.1016/j.neuron.2018.06.008>
- Gilbert, P. E., Kesner, R. P., & Lee, I. (2001). Dissociating hippocampal subregions: A double dissociation between dentate gyrus and CA1. *Hippocampus*. <https://doi.org/10.1002/hipo.1077>

- Gingrich, J. A., Andersen, P. H., Tiberi, M., Mestikawy, S. El, Jorgensen, P. N., Fremeau, R. T., & Caron, M. G. (1992). Identification, characterization, and molecular cloning of a novel transporter-like protein localized to the central nervous system. *FEBS Letters*. [https://doi.org/10.1016/0014-5793\(92\)80917-6](https://doi.org/10.1016/0014-5793(92)80917-6)
- Gingrich, J. A., Dearry, A., Falardeau, P., Bates, M. D., Fremeau, R. T., & Caron, M. G. (1992). Location and molecular cloning of D1 dopamine receptor. *Neurochemistry International*. [https://doi.org/10.1016/0197-0186\(92\)90204-5](https://doi.org/10.1016/0197-0186(92)90204-5)
- Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., & Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, *12*(10), 1222–1223. <https://doi.org/10.1038/nn.2384>
- Girardeau, G., & Zugaro, M. (2011). Hippocampal ripples and memory consolidation. In *Current Opinion in Neurobiology* (Vol. 21, Issue 3, pp. 452–459). <https://doi.org/10.1016/j.conb.2011.02.005>
- Glykos, V., & Fujisawa, S. (2016). A fine-timescale investigation of ventral 1:00 - 5:00 PM tegmental area neuronal signalling in working memory. *Society for Neuroscience*.
- Goekint, M., Bos, I., Heyman, E., Meeusen, R., Michotte, Y., & Sarre, S. (2012). Acute running stimulates hippocampal dopaminergic neurotransmission in rats, but has no influence on brain-derived neurotrophic factor. *Journal of Applied Physiology*. <https://doi.org/10.1152/jappphysiol.00306.2011>
- Gold, A. E., & Kesner, R. P. (2005). The role of the CA3 subregion of the dorsal hippocampus in spatial pattern completion in the rat. *Hippocampus*. <https://doi.org/10.1002/hipo.20103>
- Goldsmith, S. K., & Joyce, J. N. (1994). Dopamine D2 receptor expression in hippocampus and parahippocampal cortex of rat, cat, and human in relation to tyrosine hydroxylase-immunoreactive fibers. *Hippocampus*. <https://doi.org/10.1002/hipo.450040318>
- Gomperts, S. N., Kloosterman, F., & Wilson, M. A. (2015). VTA neurons coordinate with the hippocampal reactivation of spatial experience. *ELife*, *4*. <https://doi.org/10.7554/elife.05360>
- Gonzales, K. K., & Smith, Y. (2015). Cholinergic interneurons in the dorsal and ventral striatum: Anatomical and functional considerations in normal and diseased conditions. *Annals of the New York Academy of Sciences*. <https://doi.org/10.1111/nyas.12762>
- Goodroe, S. C., Starnes, J., & Brown, T. I. (2018). The Complex Nature of Hippocampal-Striatal Interactions in Spatial Navigation. *Frontiers in Human Neuroscience*. <https://doi.org/10.3389/fnhum.2018.00250>

- GoodSmith, D., Lee, H., Neunuebel, J. P., Song, H., & Knierim, J. J. (2019). Dentate Gyrus Mossy Cells Share a Role in Pattern Separation with Dentate Granule Cells and Proximal CA3 Pyramidal Cells. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0940-19.2019>
- Goutagny, R., Jackson, J., & Williams, S. (2009). Self-generated theta oscillations in the hippocampus. *Nature Neuroscience*. <https://doi.org/10.1038/nn.2440>
- Gribkoff, V. K., & Ashe, J. H. (1984). Modulation by dopamine of population spikes in area CA1 hippocampal neurons elicited by paired stimulus pulses. *Cellular and Molecular Neurobiology*. <https://doi.org/10.1007/BF00711003>
- Gribkoff, V. K., & Ashe, J. H. (1984). Modulation by dopamine of population responses and cell membrane properties of hippocampal CA1 neurons in vitro. *Brain Research*. [https://doi.org/10.1016/0006-8993\(84\)90768-6](https://doi.org/10.1016/0006-8993(84)90768-6)
- Gridchyn, I., Schoenenberger, P., O'Neill, J., & Csicsvari, J. (2020). Assembly-Specific Disruption of Hippocampal Replay Leads to Selective Memory Deficit. *Neuron*. <https://doi.org/10.1016/j.neuron.2020.01.021>
- Grosmark, A. D., & Buzsáki, G. (2016). Diversity in neural firing dynamics supports both rigid and learned hippocampal sequences. *Science*. <https://doi.org/10.1126/science.aad1935>
- Gupta, A. S., Van Der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2012). Segmentation of spatial experience by hippocampal theta sequences. *Nature Neuroscience*. <https://doi.org/10.1038/nn.3138>
- Gupta, A. S., van der Meer, M. A. A., Touretzky, D. S., & Redish, A. D. (2010). Hippocampal Replay Is Not a Simple Function of Experience. *Neuron*. <https://doi.org/10.1016/j.neuron.2010.01.034>
- Hamid, A. A., Pettibone, J. R., Mabrouk, O. S., Hetrick, V. L., Schmidt, R., Vander Weele, C. M., Kennedy, R. T., Aragona, B. J., & Berke, J. D. (2015). Mesolimbic dopamine signals the value of work. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4173>
- Hargreaves, E. L., Rao, G., Lee, I., & Knierim, J. J. (2005). Neuroscience: Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science*. <https://doi.org/10.1126/science.1110449>
- Hart, A. S., Rutledge, R. B., Glimcher, P. W., & Phillips, P. E. M. (2014). Phasic dopamine release in the rat nucleus accumbens symmetrically encodes a reward prediction error term. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2489-13.2014>

- Hart, E. E., Blair, G. J., O'Dell, T. J., Blair, H. T., & Izquierdo, A. (2019). Anterior cingulate cortex activity regulates effort-based decision making. *BioRxiv*. <https://doi.org/10.1101/792069>
- Hartley, T., Lever, C., Burgess, N., & O'Keefe, J. (2014). Space in the brain: How the hippocampal formation supports spatial cognition. In *Philosophical Transactions of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rstb.2012.0510>
- Healy, D. (1989). Neuroleptics and psychic indifference: a review. *Journal of the Royal Society of Medicine*, 82(10), 615–619.
- Hedrick, K. R., & Zhang, K. (2016). Megamap: Flexible representation of a large space embedded with nonspatial information by a hippocampal attractor network. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.00856.2015>
- Heien, M. L. A. V., Johnson, M. A., & Wightman, R. M. (2004). Resolving neurotransmitters detected by fast-scan cyclic voltammetry. *Analytical Chemistry*. <https://doi.org/10.1021/ac0491509>
- Heien, M. L. A. V., Khan, A. S., Ariansen, J. L., Cheer, J. F., Phillips, P. E. M., Wassum, K. M., & Wightman, R. M. (2005). Real-time measurement of dopamine fluctuations after cocaine in the brain of behaving rats. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.0504657102>
- Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D., & Buzsáki, G. (2000). Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.2000.84.1.390>
- Hicks, L. H. (1964). Effects of Overtraining on Acquisition and Reversal of Place and Response Learning. *Psychological Reports*. <https://doi.org/10.2466/pr0.1964.15.2.459>
- Hinman, J. R., Penley, S. C., Long, L. L., Escabí, M. A., & Chrobak, J. J. (2011). Septotemporal variation in dynamics of theta: Speed and habituation. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.00837.2010>
- Hipólito, L., Martí-Prats, L., Sánchez-Catalán, M. J., Polache, A., & Granero, L. (2011). Induction of conditioned place preference and dopamine release by salsolinol in posterior VTA of rats: Involvement of I-opioid receptors. *Neurochemistry International*, 59, 559–562. <https://doi.org/10.1016/j.neuint.2011.04.014>
- Hok, V., Save, E., Lenck-Santini, P. P., & Poucet, B. (2005). Coding for spatial goals in the prelimbic/infralimbic area of the rat frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.0407332102>

- Hok, V., Lenck-Santini, P. P., Roux, S., Save, E., Muller, R. U., & Poucet, B. (2007). Goal-related activity in hippocampal place cells. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2864-06.2007>
- HOLLISTER, L. E., EIKENBERRY, D. T., & RAFFEL, S. (1960). Chlorpromazine in nonpsychotic patients with pulmonary tuberculosis. *The American Review of Respiratory Disease*. <https://doi.org/10.1164/arrd.1960.81.4.562>
- Howe, A., De Guzman, R., Blair, G., & Blair, H. (2016). Place cells in the septohippocampal nucleus of freely behaving rats. *Society for Neuroscience*, Poster 356.26 / JJJ24. <https://www.abstractsonline.com/pp8/index.html#!/4071/presentation/16052>
- Howe, A. G., & Levy, W. B. (2007). A hippocampal model predicts a fluctuating phase transition when learning certain trace conditioning paradigms. *Cognitive Neurodynamics*, 1(2). <https://doi.org/10.1007/s11571-006-9012-7>
- Howe, M. W., Tierney, P. L., Sandberg, S. G., Phillips, P. E. M., & Graybiel, A. M. (2013). Prolonged dopamine signalling in striatum signals proximity and value of distant rewards. *Nature*. <https://doi.org/10.1038/nature12475>
- Huxter, J. R., Senior, T. J., Allen, K., & Csicsvari, J. (2008). Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. *Nature Neuroscience*. <https://doi.org/10.1038/nn.2106>
- Isomura, Y., Sirota, A., Özen, S., Montgomery, S., Mizuseki, K., Henze, D. A., & Buzsáki, G. (2006). Integration and Segregation of Activity in Entorhinal-Hippocampal Subregions by Neocortical Slow Oscillations. *Neuron*. <https://doi.org/10.1016/j.neuron.2006.10.023>
- Isomura, Y., Takekawa, T., Harukuni, R., Handa, T., Aizawa, H., Takada, M., & Fukai, T. (2013). Reward-modulated motor information in identified striatum neurons. *Journal of Neuroscience*, 33(25), 10209–10220. <https://doi.org/10.1523/JNEUROSCI.0381-13.2013>
- Jackson, J. C., Johnson, A., & Redish, A. D. (2006). Hippocampal sharp waves and reactivation during awake states depend on repeated sequential experience. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 26(48), 12415–12426. <https://doi.org/10.1523/JNEUROSCI.4118-06.2006>
- Jadhav, S. P. P., Rothschild, G., Roumis, D. K. K., & Frank, L. M. M. (2016). Coordinated Excitation and Inhibition of Prefrontal Ensembles during Awake Hippocampal Sharp-Wave Ripple Events. *Neuron*. <https://doi.org/10.1016/j.neuron.2016.02.010>
- Jadhav, S. P., Kemere, C., German, P. W., & Frank, L. M. (2012). Awake hippocampal sharp-wave ripples support spatial memory. *Science*, 336(6087), 1454–1458. <https://doi.org/10.1126/science.1217230>

- Jay, T. M., & Witter, M. P. (1991). Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *Journal of Comparative Neurology*. <https://doi.org/10.1002/cne.903130404>
- Jeewajee, A., Barry, C., Douchamps, V., Manson, D., Lever, C., & Burgess, N. (2014). Theta phase precession of grid and place cell firing in open environments. *Philosophical Transactions of the Royal Society B: Biological Sciences*. <https://doi.org/10.1098/rstb.2012.0532>
- Jensen, O., & Lisman, J. E. (2000). Position reconstruction from an ensemble of hippocampal place cells: Contribution of theta phase coding. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.2000.83.5.2602>
- Ji, D., & Wilson, M. A. (2007). Coordinated memory replay in the visual cortex and hippocampus during sleep. *Nature Neuroscience*. <https://doi.org/10.1038/nn1825>
- Johnson, A., & Redish, A. D. (2007). Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point. *Journal of Neuroscience*, 27(45), 12176–12189. <https://doi.org/10.1523/JNEUROSCI.3761-07.2007>
- Johnston, S. T., Shtrahman, M., Parylak, S., Gonçalves, J. T., & Gage, F. H. (2016). Paradox of pattern separation and adult neurogenesis: A dual role for new neurons balancing memory resolution and robustness. *Neurobiology of Learning and Memory*. <https://doi.org/10.1016/j.nlm.2015.10.013>
- Jong, W. L., Woon, R. K., Woong, S., & Min, W. J. (2009). Role of dentate gyrus in aligning internal spatial map to external landmark. *Learning and Memory*. <https://doi.org/10.1101/lm.1483709>
- Jung, M. W., & McNaughton, B. L. (1993). Spatial selectivity of unit activity in the hippocampal granular layer. *Hippocampus*. <https://doi.org/10.1002/hipo.450030209>
- Kamondi, A., Acsády, L., Wang, X. J., & Buzsáki, G. (1998). Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: Activity-dependent phase-precession of action potentials. *Hippocampus*. [https://doi.org/10.1002/\(SICI\)1098-1063\(1998\)8:3<244::AID-HIPO7>3.0.CO;2-J](https://doi.org/10.1002/(SICI)1098-1063(1998)8:3<244::AID-HIPO7>3.0.CO;2-J)
- Karlsson, M. P., & Frank, L. M. (2008). Network dynamics underlying the formation of sparse, informative representations in the hippocampus. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.4261-08.2008>
- Karlsson, M. P., & Frank, L. M. (2009). Awake replay of remote experiences in the hippocampus. *Nature Neuroscience*, 12(7), 913–918. <https://doi.org/10.1038/nn.2344>

- Kay, K., Chung, J. E., Sosa, M., Schor, J. S., Karlsson, M. P., Larkin, M. C., Liu, D. F., & Frank, L. M. (2020). Constant Sub-second Cycling between Representations of Possible Futures in the Hippocampus. *Cell*, *180*(3), 552-567.e25. <https://doi.org/10.1016/j.cell.2020.01.014>
- Keinath, A. T., Nieto-Posadas, A., Robinson, J. C., & Brandon, M. P. (2020). DG–CA3 circuitry mediates hippocampal representations of latent information. *Nature Communications*. <https://doi.org/10.1038/s41467-020-16825-1>
- Keithley, R. B., & Wightman, R. M. (2011). Assessing principal component regression prediction of neurochemicals detected with fast-scan cyclic voltammetry. *ACS Chemical Neuroscience*. <https://doi.org/10.1021/cn200035u>
- Kempadoo, K. A., Mosharov, E. V., Choi, S. J., Sulzer, D., & Kandel, E. R. (2016). Dopamine release from the locus coeruleus to the dorsal hippocampus promotes spatial learning and memory. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1616515114>
- Kempter, R., Leibold, C., Buzsáki, G., Diba, K., & Schmidt, R. (2012). Quantifying circular-linear associations: Hippocampal phase precession. *Journal of Neuroscience Methods*. <https://doi.org/10.1016/j.jneumeth.2012.03.007>
- Khan, Z. U., Gutiérrez, A., Martín, R., Peafiel, A., Rivera, A., & De La Calle, A. (2000). Dopamine D5 receptors of rat and human brain. *Neuroscience*. [https://doi.org/10.1016/S0306-4522\(00\)00274-8](https://doi.org/10.1016/S0306-4522(00)00274-8)
- Khodagholy, D., Gelineas, J. N., & Buzsáki, G. (2017). Learning-enhanced coupling between ripple oscillations in association cortices and hippocampus. *Science*. <https://doi.org/10.1126/science.aan6203>
- Kim, J. J., Rison, R. A., & Fanselow, M. S. (1993). Effects of amygdala, hippocampus, and periaqueductal gray lesions on short- and long-term contextual fear. *Behavioral Neuroscience*. <https://doi.org/10.1037/0735-7044.107.6.1093>
- Kitamura, T., Ogawa, S. K., Roy, D. S., Okuyama, T., Morrissey, M. D., Smith, L. M., Redondo, R. L., & Tonegawa, S. (2017). Engrams and circuits crucial for systems consolidation of a memory. *Science*, *356*(6333), 73–78. <https://doi.org/10.1126/science.aam6808>
- Kjelstrup, K. B., Solstad, T., Brun, V. H., Hafting, T., Leutgeb, S., Witter, M. P., Moser, E. I., & Moser, M. B. (2008). Finite scale of spatial representation in the hippocampus. *Science*. <https://doi.org/10.1126/science.1157086>
- Klausberger, T., Magill, P. J., Márton, L. F., Roberts, J. D. B., Cobden, P. M., Buzsáki, G., & Somogyi, P. (2003). Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature*. <https://doi.org/10.1038/nature01374>

- Knierim, J. J. (2015). The hippocampus. *Current Biology*.
<https://doi.org/10.1016/j.cub.2015.10.049>
- Knierim, J. J., & Neunuebel, J. P. (2016). Tracking the flow of hippocampal computation: Pattern separation, pattern completion, and attractor dynamics. *Neurobiology of Learning and Memory*. <https://doi.org/10.1016/j.nlm.2015.10.008>
- Knierim, J. J., & Zhang, K. (2012). Attractor Dynamics of Spatially Correlated Neural Activity in the Limbic System. *Annual Review of Neuroscience*.
<https://doi.org/10.1146/annurev-neuro-062111-150351>
- Kubota, Y., Liu, J., Hu, D., DeCoteau, W. E., Eden, U. T., Smith, A. C., & Graybiel, A. M. (2009). Stable encoding of task structure coexists with flexible coding of task events in sensorimotor striatum. *Journal of Neurophysiology*.
<https://doi.org/10.1152/jn.00522.2009>
- Kudrimoti, H. S., Barnes, C. A., & McNaughton, B. L. (1999). Reactivation of hippocampal cell assemblies: Effects of behavioral state, experience, and EEG dynamics. *Journal of Neuroscience*, 19(10), 4090–4101. <https://doi.org/10.1523/jneurosci.19-10-04090.1999>
- Laatikainen, L. M., Sharp, T., Bannerman, D. M., Harrison, P. J., & Tunbridge, E. M. (2012). Modulation of hippocampal dopamine metabolism and hippocampal-dependent cognitive function by catechol-O-methyltransferase inhibition. *Journal of Psychopharmacology*. <https://doi.org/10.1177/0269881112454228>
- Laurier, L. G., Corrigall, W. A., & George, S. R. (1994). Dopamine receptor density, sensitivity and mRNA levels are altered following self-administration of cocaine in the rat. *Brain Research*. [https://doi.org/10.1016/0006-8993\(94\)90255-0](https://doi.org/10.1016/0006-8993(94)90255-0)
- Lazarov, N. E., Schmidt, U., Wanner, I., & Pilgrim, C. (1998). Mapping of D1 dopamine receptor mRNA by non-radioactive in situ hybridization. *Histochemistry and Cell Biology*. <https://doi.org/10.1007/s004180050227>
- Lee, A. K., & Wilson, M. A. (2002). Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6), 1183–1194. [https://doi.org/10.1016/S0896-6273\(02\)01096-6](https://doi.org/10.1016/S0896-6273(02)01096-6)
- Lee, H., Kang, E., GoodSmith, D., Yoon, D. Y., Song, H., Knierim, J. J., Ming, G. L., & Christian, K. M. (2015). DISC1-mediated dysregulation of adult hippocampal neurogenesis in rats. *Frontiers in Systems Neuroscience*.
<https://doi.org/10.3389/fnsys.2015.00093>
- Leutgeb, J. K., Leutgeb, S., Moser, M. B., & Moser, E. I. (2007). Pattern separation in the dentate gyrus and CA3 of the hippocampus. *Science*.
<https://doi.org/10.1126/science.1135801>

- Leutgeb, S., & Mizumori, S. J. Y. (2002). Context-specific spatial representations by lateral septal cells. *Neuroscience*. [https://doi.org/10.1016/S0306-4522\(02\)00101-X](https://doi.org/10.1016/S0306-4522(02)00101-X)
- Leutgeb, S., Leutgeb, J. K., Treves, A., Moser, M. B., & Moser, E. I. (2004). Distinct ensemble codes in hippocampal areas CA3 and CA1. *Science*. <https://doi.org/10.1126/science.1100265>
- Leutgeb, S., & Mizumori, S. J. Y. (1999). Excitotoxic septal lesions result in spatial memory deficits and altered flexibility of hippocampal single-unit representations. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.19-15-06661.1999>
- Lever, C., Wills, T., Cacucci, F., Burgess, N., & O'Keefe, J. (2002). Long-term plasticity in hippocampal place-cell representation of environmental geometry. *Nature*. <https://doi.org/10.1038/416090a>
- Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*. [https://doi.org/10.1002/\(SICI\)1098-1063\(1996\)6:6<579::AID-HIPO3>3.0.CO;2-C](https://doi.org/10.1002/(SICI)1098-1063(1996)6:6<579::AID-HIPO3>3.0.CO;2-C)
- Lisman, J. E., & Grace, A. A. (2005). The hippocampal-VTA loop: Controlling the entry of information into long-term memory. In *Neuron*. <https://doi.org/10.1016/j.neuron.2005.05.002>
- Lisman, J., & Morris, R. G. M. (2001). Why is the cortex a slow learner? *Nature*. <https://doi.org/10.1038/35077185>
- Liu, X., Ramirez, S., Pang, P. T., Puryear, C. B., Govindarajan, A., Deisseroth, K., & Tonegawa, S. (2012). Optogenetic stimulation of a hippocampal engram activates fear memory recall. *Nature*. <https://doi.org/10.1038/nature11028>
- Logothetis, N. K., Eschenko, O., Murayama, Y., Augath, M., Steudel, T., Evrard, H. C., Besserve, M., & Oeltermann, A. (2012). Hippocampal-cortical interaction during periods of subcortical silence. *Nature*. <https://doi.org/10.1038/nature11618>
- Logothetis, N. K. (2015). Neural-Event-Triggered fMRI of large-scale neural networks. In *Current Opinion in Neurobiology*. <https://doi.org/10.1016/j.conb.2014.11.009>
- Lu Zhou, T., Tamura, R., Kuriwaki, J., & Ono, T. (1999). Comparison of medial and lateral septal neuron activity during performance of spatial tasks in rats. *Hippocampus*. [https://doi.org/10.1002/\(SICI\)1098-1063\(1999\)9:3<220::AID-HIPO3>3.0.CO;2-E](https://doi.org/10.1002/(SICI)1098-1063(1999)9:3<220::AID-HIPO3>3.0.CO;2-E)
- Lubenov, E. V., & Siapas, A. G. (2009). Hippocampal theta oscillations are travelling waves. *Nature*, 459(7246), 534–539. <https://doi.org/10.1038/nature08010>

- Luo, A. H., Tahsili-Fahadan, P., Wise, R. A., Lupica, C. R., & Aston-Jones, G. (2011). Linking context with reward: A functional circuit from hippocampal CA3 to ventral tegmental area. *Science*, 333(6040), 353–357. <https://doi.org/10.1126/science.1204622>
- M'Harzi, M., & Jarrard, L. E. (1992). Effects of medial and lateral septal lesions on acquisition of a place and cue radial maze task. *Behavioural Brain Research*. [https://doi.org/10.1016/S0166-4328\(05\)80160-3](https://doi.org/10.1016/S0166-4328(05)80160-3)
- Madar, A. D., Ewell, L. A., & Jones, M. V. (2019). Pattern separation of spiketrains in hippocampal neurons. *Scientific Reports*. <https://doi.org/10.1038/s41598-019-41503-8>
- Markus, E. J., Barnes, C. A., McNaughton, B. L., Gladden, V. L., & Skaggs, W. E. (1994). Spatial information content and reliability of hippocampal CA1 neurons: Effects of visual input. *Hippocampus*. <https://doi.org/10.1002/hipo.450040404>
- Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*. <https://doi.org/10.1098/rstb.1971.0078>
- Marrone, D. F., Satvat, E., Odintsova, I. V., & Gheidi, A. (2014). Dissociation of spatial representations within hippocampal region CA3. *Hippocampus*. <https://doi.org/10.1002/hipo.22367>
- Martig, A. K., & Mizumori, S. J. Y. (2011). Ventral tegmental area disruption selectively affects CA1/CA2 but not CA3 place fields during a differential reward working memory task. *Hippocampus*. <https://doi.org/10.1002/hipo.20734>
- McClelland, J. L., & Goddard, N. H. (1996). Considerations arising from a complementary learning systems perspective on hippocampus and neocortex. *Hippocampus*. [https://doi.org/10.1002/\(SICI\)1098-1063\(1996\)6:6<654::AID-HIPO8>3.0.CO;2-G](https://doi.org/10.1002/(SICI)1098-1063(1996)6:6<654::AID-HIPO8>3.0.CO;2-G)
- McHugh, T. J., Jones, M. W., Quinn, J. J., Balthasar, N., Coppari, R., Elmquist, J. K., Lowell, B. B., Fanselow, M. S., Wilson, M. A., & Tonegawa, S. (2007). Dentate gyrus NMDA receptors mediate rapid pattern separation in the hippocampal network. *Science*. <https://doi.org/10.1126/science.1140263>
- McNamara, C. G., Tejero-Cantero, Á., Trouche, S., Campo-Urriza, N., & Dupret, D. (2014). Dopaminergic neurons promote hippocampal reactivation and spatial memory persistence. *Nature Neuroscience*. <https://doi.org/10.1038/nn.3843>
- Menard, J., & Treit, D. (1996). Lateral and Medial Septal Lesions Reduce Anxiety in the Plus-Maze and Probe-Burying Tests. *Physiology and Behavior*, 60(3), 845–853.

- Miettinen, R., Hajszan, T., Riedel, A., Szigeti-Buck, K., & Leranth, C. (2012). Estimation of the total number of hippocampal CA1 pyramidal neurons: New methodology applied to helpless rats. *Journal of Neuroscience Methods*. <https://doi.org/10.1016/j.jneumeth.2011.12.017>
- Miki, T., Satriotomo, I., Li, H. P., Matsumoto, Y., Gu, H., Yokoyama, T., Lee, K. Y., Bedi, K. S., & Takeuchi, Y. (2005). Application of the physical disector to the central nervous system: Estimation of the total number of neurons in subdivisions of the rat hippocampus. *Anatomical Science International*. <https://doi.org/10.1111/j.1447-073x.2005.00121.x>
- Mishkin, M. (1978). Memory in monkeys severely impaired by combined but not by separate removal of amygdala and hippocampus. *Nature*. <https://doi.org/10.1038/273297a0>
- Mizumori, S. J. Y., McNaughton, B. L., Barnes, C. A., & Fox, K. B. (1989). Preserved spatial coding in hippocampal CA1 pyramidal cells during reversible suppression of CA3c output: Evidence for pattern completion in hippocampus. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.09-11-03915.1989>
- Mizuseki, K., Royer, S., Diba, K., & Buzsáki, G. (2012). Activity dynamics and behavioral correlates of CA3 and CA1 hippocampal pyramidal neurons. *Hippocampus*. <https://doi.org/10.1002/hipo.22002>
- Mohebi, A., & Berke, J. D. (2020). Dopamine release drives motivation, independently from dopamine cell firing. In *Neuropsychopharmacology*. <https://doi.org/10.1038/s41386-019-0492-7>
- Mohebi, A., Pettibone, J. R., Hamid, A. A., Wong, J. M. T., Vinson, L. T., Patriarchi, T., Tian, L., Kennedy, R. T., & Berke, J. D. (2019). Dissociable dopamine dynamics for learning and motivation. *Nature*. <https://doi.org/10.1038/s41586-019-1235-y>
- Monaco, J. D., De Guzman, R. M., Blair, H. T., & Zhang, K. (2019). Spatial synchronization codes from coupled rate-phase neurons. *PLoS Computational Biology*, *15*(1). <https://doi.org/10.1371/journal.pcbi.1006741>
- Montague, P. R., Dayan, P., & Sejnowski, T. J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.16-05-01936.1996>
- Morris, R. G. M., Garrud, P., Rawlins, J. N. P., & O'Keefe, J. (1982). Place navigation impaired in rats with hippocampal lesions. *Nature*. <https://doi.org/10.1038/297681a0>
- Mou, X., Cheng, J., Yu, Y. S. W., Kee, S. E., & Ji, D. (2018). Comparing mouse and rat hippocampal place cell activities and firing sequences in the same environments. *Frontiers in Cellular Neuroscience*. <https://doi.org/10.3389/fncel.2018.00332>

- Mowrer, O. H. (1947). On the dual nature of learning—a re-interpretation of “conditioning” and “problem-solving.” *Harvard Educational Review*, 17, 102–148.
- Muller, R. (1996). A quarter of a century of place cells. *Neuron*.
[https://doi.org/10.1016/S0896-6273\(00\)80214-7](https://doi.org/10.1016/S0896-6273(00)80214-7)
- Nádasy, Z., Hirase, H., Czurkó, A., Csicsvari, J., & Buzsáki, G. (1999). Replay and time compression of recurring spike sequences in the hippocampus. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.19-21-09497.1999>
- Nakashiba, T., Buhl, D. L., McHugh, T. J., & Tonegawa, S. (2009). Hippocampal CA3 Output Is Crucial for Ripple-Associated Reactivation and Consolidation of Memory. *Neuron*, 62(6), 781–787. <https://doi.org/10.1016/j.neuron.2009.05.013>
- Nitzan, N., McKenzie, S., Beed, P., English, D. F., Oldani, S., Tukker, J. J., Buzsáki, G., & Schmitz, D. (2020). Propagation of hippocampal ripples to the neocortex by way of a subiculum-retrosplenial pathway. *Nature Communications*.
<https://doi.org/10.1038/s41467-020-15787-8>
- Niv, Y. (2013). Neuroscience: Dopamine ramps up. In *Nature*.
<https://doi.org/10.1038/500533a>
- O’Keefe, J., & Dostrovsky, J. (1971). The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Research*.
[https://doi.org/10.1016/0006-8993\(71\)90358-1](https://doi.org/10.1016/0006-8993(71)90358-1)
- O’Keefe, J. (1976). Place units in the hippocampus of the freely moving rat. *Experimental Neurology*. [https://doi.org/10.1016/0014-4886\(76\)90055-8](https://doi.org/10.1016/0014-4886(76)90055-8)
- O’Keefe, J., & Nadel, L. (1978). Hippocampus Physiology -The hippocampus as a cognitive map. *The Hippocampus as a Cognitive Map*.
- O’Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, 3(3), 317–330.
<https://doi.org/10.1002/hipo.450030307>
- O’Mara, S. (2005). The subiculum: What it does, what it might do, and what neuroanatomy has yet to tell us. In *Journal of Anatomy*. <https://doi.org/10.1111/j.1469-7580.2005.00446.x>
- O’Neill, J., Senior, T. J., Allen, K., Huxter, J. R., & Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*. <https://doi.org/10.1038/nn2037>

- O'Neill, J., Senior, T., & Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron*. <https://doi.org/10.1016/j.neuron.2005.10.037>
- O'Reilly, R. C., & McClelland, J. L. (1994). Hippocampal conjunctive encoding, storage, and recall: Avoiding a trade-off. *Hippocampus*. <https://doi.org/10.1002/hipo.450040605>
- O'Keefe J, N. L. (1978). *The hippocampus as a cognitive map*. Oxford University Press.
- Ólafsdóttir, H. F., Carpenter, F., & Barry, C. (2016). Coordinated grid and place cell replay during rest. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4291>
- Packard, M. G. (1999). Glutamate infused posttraining into the hippocampus or caudate-putamen differentially strengthens place and response learning. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.96.22.12881>
- Packard, M. G., & Goodman, J. (2016). Neurobiology of procedural learning in animals. In *The Curated Reference Collection in Neuroscience and Biobehavioral Psychology*. <https://doi.org/10.1016/B978-0-12-809324-5.21086-9>
- Packard, M. G., & McGaugh, J. L. (1996). Inactivation of hippocampus or caudate nucleus with lidocaine differentially affects expression of place and response learning. *Neurobiology of Learning and Memory*. <https://doi.org/10.1006/nlme.1996.0007>
- Pan, W. X., Schmidt, R., Wickens, J. R., & Hyland, B. I. (2005). Dopamine cells respond to predicted events during classical conditioning: Evidence for eligibility traces in the reward-learning network. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.1478-05.2005>
- Papale, A. E., Zielinski, M. C., Frank, L. M., Jadhav, S. P., & Redish, A. D. (2016). Interplay between Hippocampal Sharp-Wave-Ripple Events and Vicarious Trial and Error Behaviors in Decision Making. *Neuron*. <https://doi.org/10.1016/j.neuron.2016.10.028>
- Parfitt, G. M., Nguyen, R., Bang, J. Y., Aqrabawi, A. J., Tran, M. M., Seo, D. K., Richards, B. A., & Kim, J. C. (2017). Bidirectional Control of Anxiety-Related Behaviors in Mice: Role of Inputs Arising from the Ventral Hippocampus to the Lateral Septum and Medial Prefrontal Cortex. *Neuropsychopharmacology*. <https://doi.org/10.1038/npp.2017.56>
- Park, E. H., Dvorak, D., & Fenton, A. A. (2011). Ensemble place codes in hippocampus: CA1, CA3, and dentate gyrus place cells have multiple place fields in large environments. *PLoS ONE*. <https://doi.org/10.1371/journal.pone.0022349>

- Patel, J., Fujisawa, S., Berényi, A., Royer, S., & Buzsáki, G. (2012). Traveling Theta Waves along the Entire Septotemporal Axis of the Hippocampus. *Neuron*, 75(3), 410–417. <https://doi.org/10.1016/j.neuron.2012.07.015>
- Patriarchi, T., Cho, J. R., Merten, K., Howe, M. W., Marley, A., Xiong, W. H., Folk, R. W., Broussard, G. J., Liang, R., Jang, M. J., Zhong, H., Dombeck, D., von Zastrow, M., Nimmerjahn, A., Gradinaru, V., Williams, J. T., & Tian, L. (2018). Ultrafast neuronal imaging of dopamine dynamics with designed genetically encoded sensors. *Science*. <https://doi.org/10.1126/science.aat4422>
- Paxinos, G., & Charles Watson. (2007). The Rat Brain in Stereotaxic Coordinates Sixth Edition. In *Elsevier Academic Press*.
- Pfeiffer, B. E., & Foster, D. J. (2013). Hippocampal place-cell sequences depict future paths to remembered goals. *Nature*, 497(7447), 74–79. <https://doi.org/10.1038/nature12112>
- Pinder, R. M. (2007). Pathological gambling and dopamine agonists: A phenotype? In *Neuropsychiatric Disease and Treatment*. <https://doi.org/10.2147/ndt.2007.3.1.1>
- Pirritano, D., Plastino, M., Bosco, D., Gallelli, L., Siniscalchi, A., & De Sarro, G. (2014). Gambling disorder during dopamine replacement treatment in Parkinson's disease: A comprehensive review. In *BioMed Research International*. <https://doi.org/10.1155/2014/728038>
- Poucet, B., & Hok, V. (2017). Remembering goal locations. In *Current Opinion in Behavioral Sciences*. <https://doi.org/10.1016/j.cobeha.2017.06.003>
- Purves, D., Augustine, G., Fitzpatrick, D., Katz, L., LaMantia, A.-S., McNamara, J., & Williams, M. (2001). *Neuroscience*. 2nd edition. Sunderland (MA): Sinauer Associates; 2001. <https://www.ncbi.nlm.nih.gov/books/NBK10940/>
- Ragozzino, M. E., Ragozzino, K. E., Mizumori, S. J. Y., & Kesner, R. P. (2002). Role of the dorsomedial striatum in behavioral flexibility for response and visual cue discrimination learning. *Behavioral Neuroscience*. <https://doi.org/10.1037/0735-7044.116.1.105>
- Raisman, G. (1966). The connexions of the septum. *Brain*, 89(2), 317–348. <https://doi.org/10.1093/brain/89.2.317>
- Ramirez, S., Liu, X., Lin, P. A., Suh, J., Pignatelli, M., Redondo, R. L., Ryan, T. J., & Tonegawa, S. (2013). Creating a false memory in the hippocampus. *Science*. <https://doi.org/10.1126/science.1239073>

- Ramirez, S., Liu, X., MacDonald, C. J., Moffa, A., Zhou, J., Redondo, R. L., & Tonegawa, S. (2015). Activating positive memory engrams suppresses depression-like behaviour. *Nature*. <https://doi.org/10.1038/nature14514>
- Ranck, J. B. (1973). Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. Part I. Behavioral correlates and firing repertoires. *Experimental Neurology*. [https://doi.org/10.1016/0014-4886\(73\)90290-2](https://doi.org/10.1016/0014-4886(73)90290-2)
- Rangel, A., Camerer, C., & Montague, P. R. (2008). A framework for studying the neurobiology of value-based decision making. *Nature Reviews Neuroscience*, 9, 545–556. <https://doi.org/10.1038/nrn2357>
- Rawlins, J. N. P., & Olton, D. S. (1982). The septo-hippocampal system and cognitive mapping. *Behavioural Brain Research*. [https://doi.org/10.1016/0166-4328\(82\)90039-0](https://doi.org/10.1016/0166-4328(82)90039-0)
- Redish, A. D. (2016). Vicarious trial and error. In *Nature Reviews Neuroscience*. <https://doi.org/10.1038/nrn.2015.30>
- Redish, A. D. D. (1999). Beyond the cognitive map: From place cells to episodic memory. In *Cambridge, MA: MIT Press*. The MIT Press.
- Redondo, R. L., Kim, J., Arons, A. L., Ramirez, S., Liu, X., & Tonegawa, S. (2014). Bidirectional switch of the valence associated with a hippocampal contextual memory engram. *Nature*. <https://doi.org/10.1038/nature13725>
- Regier, P. S., Amemiya, S., & Redish, A. D. (2015). Hippocampus and subregions of the dorsal striatum respond differently to a behavioral strategy change on a spatial navigation task. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.00189.2015>
- Reis, D. G., Scopinho, A. A., Guimarães, F. S., Corrêa, F. M. A., & Resstel, L. B. M. (2010). Involvement of the lateral septal area in the expression of fear conditioning to context. *Learning and Memory*. <https://doi.org/10.1101/lm.1534710>
- Restle, F. (1957). Discrimination of cues in mazes: A resolution of the “place-vs.-response” question. *Psychological Review*. <https://doi.org/10.1037/h0040678>
- Retailleau, A., & Morris, G. (2018). Spatial Rule Learning and Corresponding CA1 Place Cell Reorientation Depend on Local Dopamine Release. *Current Biology*. <https://doi.org/10.1016/j.cub.2018.01.081>
- Rich, P. D., Liaw, H. P., & Lee, A. K. (2014). Large environments reveal the statistical structure governing hippocampal representations. *Science*. <https://doi.org/10.1126/science.1255635>

- Richard, G. R., Titiz, A., Tyler, A., Holmes, G. L., Scott, R. C., & Lenck-Santini, P. P. (2013). Speed modulation of hippocampal theta frequency correlates with spatial memory performance. *Hippocampus*. <https://doi.org/10.1002/hipo.22164>
- Risold, P. Y., & Swanson, L. W. (1997). Connections of the rat lateral septal complex. In *Brain Research Reviews* (Vol. 24, Issues 2–3, pp. 115–195). [https://doi.org/10.1016/S0165-0173\(97\)00009-X](https://doi.org/10.1016/S0165-0173(97)00009-X)
- Ritchie, B. F., Aeschliman, B., & Pierce, P. (1950). Studies in spatial learning. VIII. Place performance and the acquisition of place dispositions. *Journal of Comparative and Physiological Psychology*. <https://doi.org/10.1037/h0055224>
- Robinson, J., Manseau, F., Ducharme, G., Amilhon, B., Vigneault, E., El Mestikawy, S., & Williams, S. (2016). Optogenetic activation of septal glutamatergic neurons drive hippocampal theta rhythms. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2141-15.2016>
- Roham, M., Daberkow, D. P., Ramsson, E. S., Covey, D. P., Pakdeeronachit, S., Garris, P. A., & Mohseni, P. (2008). A wireless IC for wide-range neurochemical monitoring using amperometry and fast-scan cyclic voltammetry. *IEEE Transactions on Biomedical Circuits and Systems*. <https://doi.org/10.1109/TBCAS.2008.918282>
- Rolls, E. T. (1996). A theory of hippocampal function in memory. *Hippocampus*. [https://doi.org/10.1002/\(SICI\)1098-1063\(1996\)6:6<601::AID-HIPO5>3.0.CO;2-J](https://doi.org/10.1002/(SICI)1098-1063(1996)6:6<601::AID-HIPO5>3.0.CO;2-J)
- Rolls, E. T. (2010). A computational theory of episodic memory formation in the hippocampus. In *Behavioural Brain Research*. <https://doi.org/10.1016/j.bbr.2010.03.027>
- Rolls, E. T. (2013). The mechanisms for pattern completion and pattern separation in the hippocampus. *Frontiers in Systems Neuroscience*. <https://doi.org/10.3389/fnsys.2013.00074>
- Rosen, Z. B., Cheung, S., & Siegelbaum, S. A. (2015). Midbrain dopamine neurons bidirectionally regulate CA3-CA1 synaptic drive. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4152>
- Rothschild, G., Eban, E., & Frank, L. M. (2017). A cortical-hippocampal-cortical loop of information processing during memory consolidation. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4457>
- Roux, L., Hu, B., Eichler, R., Stark, E., & Buzsáki, G. (2017). Sharp wave ripples during learning stabilize the hippocampal spatial map. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4543>

- Rueda-Orozco, P. E., & Robbe, D. (2015). The striatum multiplexes contextual and kinematic information to constrain motor habits execution. *Nature Neuroscience*, 18(3), 453–462. <https://doi.org/10.1038/nn.3924>
- Saal, D., Dong, Y., Bonci, A., & Malenka, R. C. (2003). Drugs of abuse and stress trigger a common synaptic adaptation in dopamine neurons. *Neuron*. [https://doi.org/10.1016/S0896-6273\(03\)00021-7](https://doi.org/10.1016/S0896-6273(03)00021-7)
- Sahay, A., Scobie, K. N., Hill, A. S., O'Carroll, C. M., Kheirbek, M. A., Burghardt, N. S., Fenton, A. A., Dranovsky, A., & Hen, R. (2011). Increasing adult hippocampal neurogenesis is sufficient to improve pattern separation. *Nature*. <https://doi.org/10.1038/nature09817>
- Salamone, J. D., & Correa, M. (2012). The Mysterious Motivational Functions of Mesolimbic Dopamine. In *Neuron*. <https://doi.org/10.1016/j.neuron.2012.10.021>
- Scatton, B., Simon, H., Le Moal, M., & Bischoff, S. (1980). Origin of dopaminergic innervation of the rat hippocampal formation. *Neuroscience Letters*. [https://doi.org/10.1016/0304-3940\(80\)90314-6](https://doi.org/10.1016/0304-3940(80)90314-6)
- Schmidt, B., Hinman, J. R., Jacobson, T. K., Szkudlarek, E., Argraves, M., Escabí, M. A., & Markus, E. J. (2013). Dissociation between dorsal and ventral hippocampal theta oscillations during decision-making. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2915-12.2013>
- Schmidt, B., Papale, A., Redish, A. D., & Markus, E. J. (2013). Conflict between place and response navigation strategies: Effects on vicarious trial and error (VTE) behaviors. *Learning and Memory*. <https://doi.org/10.1101/lm.028753.112>
- Schmitt, K. C., & Reith, M. E. A. (2010). Regulation of the dopamine transporter: Aspects relevant to psychostimulant drugs of abuse. In *Annals of the New York Academy of Sciences*. <https://doi.org/10.1111/j.1749-6632.2009.05148.x>
- Schultz, W. (1986). Responses of midbrain dopamine neurons to behavioral trigger stimuli in the monkey. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.1986.56.5.1439>
- Schultz, W., Apicella, P., & Ljungberg, T. (1993). Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.13-03-00900.1993>
- Schultz, W., Dayan, P., & Montague, P. R. (1997). A neural substrate of prediction and reward. *Science*. <https://doi.org/10.1126/science.275.5306.1593>
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. In *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.1998.80.1.1>

- Schwerdt, H. N., Kim, M. J., Amemori, S., Homma, D., Yoshida, T., Shimazu, H., Yerramreddy, H., Karasan, E., Langer, R., Graybiel, A. M., & Cima, M. J. (2017). Subcellular probes for neurochemical recording from multiple brain sites. *Lab on a Chip*. <https://doi.org/10.1039/c6lc01398h>
- Schwerdt, H. N., Kim, M., Karasan, E., Amemori, S., Homma, D., Shimazu, H., Yoshida, T., Langer, R., Graybiel, A. M., & Cima, M. J. (2017). Subcellular electrode arrays for multisite recording of dopamine in vivo. *Proceedings of the IEEE International Conference on Micro Electro Mechanical Systems (MEMS)*. <https://doi.org/10.1109/MEMSYS.2017.7863465>
- Schwerdt, H. N., Shimazu, H., Amemori, K. ichi, Amemori, S., Tierney, P. L., Gibson, D. J., Hong, S., Yoshida, T., Langer, R., Cima, M. J., & Graybiel, A. M. (2017). Long-term dopamine neurochemical monitoring in primates. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.1713756114>
- Schwerdt, H. N., Zhang, E., Kim, M. J., Yoshida, T., Stanwicks, L., Amemori, S., Dagdeviren, H. E., Langer, R., Cima, M. J., & Graybiel, A. M. (2018). Cellular-scale probes enable stable chronic subsecond monitoring of dopamine neurochemicals in a rodent model. *Communications Biology*. <https://doi.org/10.1038/s42003-018-0147-y>
- Scoville, W. B., & Milner, B. (2000). Loss of recent memory after bilateral hippocampal lesions. 1957. *The Journal of Neuropsychiatry and Clinical Neurosciences*. <https://doi.org/10.1176/jnp.12.1.103>
- Senzai, Y., & Buzsáki, G. (2017). Physiological Properties and Behavioral Correlates of Hippocampal Granule Cells and Mossy Cells. *Neuron*. <https://doi.org/10.1016/j.neuron.2016.12.011>
- Shapiro, M. A., Chang, Y. L., Munson, S. K., Jacobson IV, C. E., Rodriguez, R. L., Skidmore, F. M., Okun, M. S., & Fernandez, H. H. (2007). The four as associated with pathological Parkinson disease gamblers: Anxiety, anger, age, and agonists. *Neuropsychiatric Disease and Treatment*. <https://doi.org/10.2147/ndt.2007.3.1.161>
- Sheremet, A., Burke, S. N., & Maurer, A. P. (2016). Movement enhances the nonlinearity of hippocampal theta. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.3564-15.2016>
- Shoham, S., O'Connor, D. H., & Segev, R. (2006). How silent is the brain: Is there a “dark matter” problem in neuroscience? In *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*. <https://doi.org/10.1007/s00359-006-0117-6>

- Siegel, A., Ohgami, S., & Edinger, H. (1975). Projections of the hippocampus to the septal area in the squirrel monkey. *Brain Research*. [https://doi.org/10.1016/0006-8993\(75\)90027-X](https://doi.org/10.1016/0006-8993(75)90027-X)
- Singer, A. C., Carr, M. F., Karlsson, M. P., & Frank, L. M. (2013). Hippocampal SWR Activity Predicts Correct Decisions during the Initial Learning of an Alternation Task. *Neuron*. <https://doi.org/10.1016/j.neuron.2013.01.027>
- Singer, A. C., & Frank, L. M. (2009). Rewarded Outcomes Enhance Reactivation of Experience in the Hippocampus. *Neuron*. <https://doi.org/10.1016/j.neuron.2009.11.016>
- Singewald, G. M., Rjabokon, A., Singewald, N., & Ebner, K. (2011). The modulatory role of the lateral septum on neuroendocrine and behavioral stress responses. *Neuropsychopharmacology*. <https://doi.org/10.1038/npp.2010.213>
- Sirota, A., Csicsvari, J., Buhl, D., & Buzsáki, G. (2003). Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.0437938100>
- Skaggs, W. E., & McNaughton, B. L. (1996). Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271(5257), 1870–1873. <https://doi.org/10.1126/science.271.5257.1870>
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., Barnes, C. A., WE, S., BL, M., MA, W., CA, B., Skaggs, W. E., McNaughton, B. L., Wilson, M. A., Barnes, C. A., WE, S., BL, M., MA, W., & CA, B. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6(2). [https://doi.org/10.1002/\(SICI\)1098-1063\(1996\)6:2<149::AID-HIPO6>3.0.CO;2-K](https://doi.org/10.1002/(SICI)1098-1063(1996)6:2<149::AID-HIPO6>3.0.CO;2-K)
- Skinner, D. M., Martin, G. M., Wright, S. L., Tomlin, J., Odintsova, I. V., Thorpe, C. M., Harley, C. W., & Marrone, D. F. (2014). Hippocampal spatial mapping and the acquisition of competing responses. *Hippocampus*. <https://doi.org/10.1002/hipo.22233>
- Sławińska, U., & Kasicki, S. (1998). The frequency of rat's hippocampal theta rhythm is related to the speed of locomotion. *Brain Research*. [https://doi.org/10.1016/S0006-8993\(98\)00390-4](https://doi.org/10.1016/S0006-8993(98)00390-4)
- Smith, K. S., & Graybiel, A. M. (2013). A dual operator view of habitual behavior reflecting cortical and striatal dynamics. *Neuron*. <https://doi.org/10.1016/j.neuron.2013.05.038>
- Smith, K. S., & Graybiel, A. M. (2016). Habit formation coincides with shifts in reinforcement representations in the sensorimotor striatum. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.00925.2015>

- Sosa, M., Joo, H. R., & Frank, L. M. (2020). Dorsal and Ventral Hippocampal Sharp-Wave Ripples Activate Distinct Nucleus Accumbens Networks. *Neuron*, *105*(4), 725-741.e8. <https://doi.org/10.1016/j.neuron.2019.11.022>
- Squire, L. R., & Alvarez, P. (1995). Retrograde amnesia and memory consolidation: a neurobiological perspective. *Current Opinion in Neurobiology*. [https://doi.org/10.1016/0959-4388\(95\)80023-9](https://doi.org/10.1016/0959-4388(95)80023-9)
- Stanzione, P., Calabresi, P., Mercuri, N., & Bernardi, G. (1984). Dopamine modulates CA1 hippocampal neurons by elevating the threshold for spike generation: An in vitro study. *Neuroscience*. [https://doi.org/10.1016/0306-4522\(84\)90291-4](https://doi.org/10.1016/0306-4522(84)90291-4)
- Stellar, J. R., & Stellar, E. (1985). The Neurobiology of Motivation and Reward. In *The Neurobiology of Motivation and Reward*. <https://doi.org/10.1007/978-1-4615-8032-4>
- Stringer, K. G., Martin, G. M., & Skinner, D. M. (2005). The effects of hippocampal lesions on response, direction, and place learning in rats. *Behavioral Neuroscience*. <https://doi.org/10.1037/0735-7044.119.4.946>
- Stuber, G. D., Wise, R. A., & Author, N. N. (2016). Lateral Hypothalamic Circuits for Feeding and Reward HHS Public Access Author manuscript. *Nat Neurosci*. <https://doi.org/10.1038/nn.4220>
- Sun, F., Zeng, J., Jing, M., Zhou, J., Feng, J., Owen, S., Luo, Y., Li, F., Yamaguchi, T., Yong, Z., Gao, Y., Peng, W., Wang, L., Zhang, S., Du, J., Lin, D., Xu, M., Kreitzer, A., Cui, G., & Li, Y. (2018). A genetically-encoded fluorescent sensor enables rapid and specific detection of dopamine in flies, fish, and mice. *BioRxiv*. <https://doi.org/10.1101/332528>
- Sutton, R. S., Barto, A.G., Barto, A. G., & Barto, A.G. (1998). Reinforcement Learning: An Introduction. In *IEEE Transactions on Neural Networks*. MIT Press. <https://doi.org/10.1109/tnn.1998.712192>
- Suzuki, S. S., & Smith, G. K. (1988). Spontaneous EEG spikes in the normal hippocampus. III. Relations to evoked potentials. *Electroencephalography and Clinical Neurophysiology*. [https://doi.org/10.1016/0013-4694\(88\)90166-6](https://doi.org/10.1016/0013-4694(88)90166-6)
- Suzuki, S. S., & Smith, G. K. (1987). Spontaneous EEG spikes in the normal hippocampus. I. Behavioral correlates, laminar profiles and bilateral synchrony. *Electroencephalography and Clinical Neurophysiology*. [https://doi.org/10.1016/0013-4694\(87\)90123-4](https://doi.org/10.1016/0013-4694(87)90123-4)
- Suzuki, S. S., & Smith, G. K. (1988). Spontaneous EEG spikes in the normal hippocampus. IV. Effects of medial septum and entorhinal cortex lesions. *Electroencephalography and Clinical Neurophysiology*. [https://doi.org/10.1016/0013-4694\(88\)90197-6](https://doi.org/10.1016/0013-4694(88)90197-6)

- Swanson, L. W. (1982). The projections of the ventral tegmental area and adjacent regions: A combined fluorescent retrograde tracer and immunofluorescence study in the rat. *Brain Research Bulletin*. [https://doi.org/10.1016/0361-9230\(82\)90145-9](https://doi.org/10.1016/0361-9230(82)90145-9)
- Swanson, L. W., & Cowan, W. M. (1979). The connections of the septal region in the rat. *Journal of Comparative Neurology*. <https://doi.org/10.1002/cne.901860408>
- Swanson, L. W., & Cowan, W. M. (1975). Hippocampo-hypothalamic connections: Origin in subicular cortex, not Ammon's horn. *Science*. <https://doi.org/10.1126/science.49928>
- Swanson, L. W., Sawchenko, P. E., & Cowan, W. M. (1981). Evidence for collateral projections by neurons in Ammon's horn, the dentate gyrus, and the subiculum: A multiple retrograde labeling study in the rat. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.01-05-00548.1981>
- Sweeney, P., & Yang, Y. (2015). An excitatory ventral hippocampus to lateral septum circuit that suppresses feeding. *Nature Communications*. <https://doi.org/10.1038/ncomms10188>
- Syed, E. C. J., Grima, L. L., Magill, P. J., Bogacz, R., Brown, P., & Walton, M. E. (2015). Action initiation shapes mesolimbic dopamine encoding of future rewards. *Nature Neuroscience*. <https://doi.org/10.1038/nn.4187>
- Takamura, Y., Tamura, R., Zhou, T. L., Kobayashi, T., Tran, A. H., Eifuku, S., & Ono, T. (2006). Spatial firing properties of lateral septal neurons. *Hippocampus*, 16(8), 635–644. <https://doi.org/10.1002/hipo.20196>
- Takeuchi, T., Duzkiewicz, A. J., Sonneborn, A., Spooner, P. A., Yamasaki, M., Watanabe, M., Smith, C. C., Fernández, G., Deisseroth, K., Greene, R. W., & Morris, R. G. M. (2016). Locus coeruleus and dopaminergic consolidation of everyday memory. *Nature*. <https://doi.org/10.1038/nature19325>
- Takmakov, P., McKinney, C. J., Carelli, R. M., & Wightman, R. M. (2011). Instrumentation for fast-scan cyclic voltammetry combined with electrophysiology for behavioral experiments in freely moving animals. *Review of Scientific Instruments*. <https://doi.org/10.1063/1.3610651>
- Tang, W., Shin, J. D., Frank, L. M., & Jadhav, S. P. (2017). Hippocampal-prefrontal reactivation during learning is stronger in awake compared with sleep states. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.2291-17.2017>
- Tang, Y., Benusiglio, D., Grinevich, V., & Lin, L. (2016). Distinct types of feeding related neurons in mouse hypothalamus. *Frontiers in Behavioral Neuroscience*. <https://doi.org/10.3389/fnbeh.2016.00091>

- Tashiro, A., Makino, H., & Gage, F. H. (2007). Experience-specific functional modification of the dentate gyrus through adult neurogenesis: A critical period during an immature stage. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.4941-06.2007>
- Taylor, K. K., Tanaka, K. Z., Reijmers, L. G., & Wiltgen, B. J. (2013). Reactivation of neural ensembles during the retrieval of recent and remote memory. *Current Biology*. <https://doi.org/10.1016/j.cub.2012.11.019>
- Tecuapetla, F., Jin, X., Lima, S. Q., & Costa, R. M. (2016). Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution. *Cell*, 166(3), 703–715. <https://doi.org/10.1016/j.cell.2016.06.032>
- Temprana, S. G., Mongiat, L. A., Yang, S. M., Trincherro, M. F., Alvarez, D. D., Kropff, E., Giacomini, D., Beltramone, N., Lanuza, G. M., & Schinder, A. F. (2015). Delayed Coupling to Feedback Inhibition during a Critical Period for the Integration of Adult-Born Granule Cells. *Neuron*. <https://doi.org/10.1016/j.neuron.2014.11.023>
- Teuber, H. L., Milner, B., & Vaughan, H. G. (1968). Persistent anterograde amnesia after stab wound of the basal brain. *Neuropsychologia*. [https://doi.org/10.1016/0028-3932\(68\)90025-0](https://doi.org/10.1016/0028-3932(68)90025-0)
- Thierry, A. M., Gioanni, Y., Dégénétais, E., & Glowinski, J. (2000). Hippocampo-prefrontal cortex pathway: Anatomical and electrophysiological characteristics. *Hippocampus*. [https://doi.org/10.1002/1098-1063\(2000\)10:4<411::AID-HIPO7>3.0.CO;2-A](https://doi.org/10.1002/1098-1063(2000)10:4<411::AID-HIPO7>3.0.CO;2-A)
- Thompson, L. T., & Best, P. J. (1989). Place cells and silent cells in the hippocampus of freely-behaving rats. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.09-07-02382.1989>
- Thorn, C. A., Atallah, H., Howe, M., & Graybiel, A. M. (2010). Differential Dynamics of Activity Changes in Dorsolateral and Dorsomedial Striatal Loops during Learning. *Neuron*. <https://doi.org/10.1016/j.neuron.2010.04.036>
- Threlfell, S., Lalic, T., Platt, N. J., Jennings, K. A., Deisseroth, K., & Cragg, S. J. (2012). Striatal dopamine release is triggered by synchronized activity in cholinergic interneurons. *Neuron*. <https://doi.org/10.1016/j.neuron.2012.04.038>
- Tiberi, M., Jarvie, K. R., Silvia, C., Falardeau, P., Gingrich, J. A., Godinot, N., Bertrand, L., Yang-Feng, T. L., Fremeau, R. T., & Caron, M. G. (1991). Cloning, molecular characterization, and chromosomal assignment of a gene encoding a second D1 dopamine receptor subtype: Differential expression pattern in rat brain compared with the D1A receptor. *Proceedings of the National Academy of Sciences of the United States of America*. <https://doi.org/10.1073/pnas.88.17.7491>

- Tingley, D., & Buzsáki, G. (2020). Routing of Hippocampal Ripples to Subcortical Structures via the Lateral Septum. *Neuron*, *105*(1), 138-149.e5. <https://doi.org/10.1016/j.neuron.2019.10.012>
- Tingley, D., & Buzsáki, G. (2018). Transformation of a Spatial Map across the Hippocampal-Lateral Septal Circuit. *Neuron*, *98*(6), 1229-1242.e5. <https://doi.org/10.1016/j.neuron.2018.04.028>
- Tippmann-Peikert, M., Park, J. G., Boeve, B. F., Shepard, J. W., & Silber, M. H. (2007). Pathologic gambling in patients with restless legs syndrome treated with dopaminergic agonists. *Neurology*. <https://doi.org/10.1212/01.wnl.0000252368.25106.b6>
- Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*. <https://doi.org/10.1126/science.1105370>
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological Review*. <https://doi.org/10.1037/h0061626>
- Tonegawa, S., Liu, X., Ramirez, S., & Redondo, R. (2015). Memory Engram Cells Have Come of Age. *Neuron*. <https://doi.org/10.1016/j.neuron.2015.08.002>
- Traub, R. D., Bibbig, A., LeBeau, F. E. N., Buhl, E. H., & Whittington, M. A. (2004). Cellular mechanisms of neuronal population oscillations in the hippocampus in vitro. *Annual Review of Neuroscience*. <https://doi.org/10.1146/annurev.neuro.27.070203.144303>
- Treves, A., & Rolls, E. T. (1994). Computational analysis of the role of the hippocampus in memory. *Hippocampus*. <https://doi.org/10.1002/hipo.450040319>
- Trouche, S., Koren, V., Doig, N. M., Ellender, T. J., El-Gaby, M., Lopes-dos-Santos, V., Reeve, H. M., Perestenko, P. V., Garas, F. N., Magill, P. J., Sharott, A., & Dupret, D. (2019). A Hippocampus-Accumbens Tripartite Neuronal Motif Guides Appetitive Memory in Space. *Cell*. <https://doi.org/10.1016/j.cell.2018.12.037>
- Tsien, J. Z., Huerta, P. T., & Tonegawa, S. (1996). The essential role of hippocampal CA1 NMDA receptor-dependent synaptic plasticity in spatial memory. *Cell*. [https://doi.org/10.1016/S0092-8674\(00\)81827-9](https://doi.org/10.1016/S0092-8674(00)81827-9)
- Van Der Meer, M. A. A., & Redish, A. D. (2011). Theta phase precession in rat ventral striatum links place and reward information. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.4869-10.2011>

- Van Der Meer, M., Kurth-Nelson, Z., & Redish, A. D. (2012). Information processing in decision-making systems. In *Neuroscientist*.
<https://doi.org/10.1177/1073858411435128>
- van Dijk, M. T., & Fenton, A. A. (2018). On How the Dentate Gyrus Contributes to Memory Discrimination. *Neuron*. <https://doi.org/10.1016/j.neuron.2018.04.018>
- Van Strien, N. M., Cappaert, N. L. M., & Witter, M. P. (2009). The anatomy of memory: An interactive overview of the parahippocampal- hippocampal network. *Nature Reviews Neuroscience*. <https://doi.org/10.1038/nrn2614>
- Vandecasteele, M., Varga, V., Berényi, A., Papp, E., Barthó, P., Venance, L., Freund, T. F., & Buzsáki, G. (2014). Optogenetic activation of septal cholinergic neurons suppresses sharp wave ripples and enhances theta oscillations in the hippocampus. *Proceedings of the National Academy of Sciences of the United States of America*.
<https://doi.org/10.1073/pnas.1411233111>
- Vanderwolf, C. H. (1969). Hippocampal electrical activity and voluntary movement in the rat. *Electroencephalography and Clinical Neurophysiology*, 26(4), 407–418.
[https://doi.org/10.1016/0013-4694\(69\)90092-3](https://doi.org/10.1016/0013-4694(69)90092-3)
- Vazdarjanova, A., & Guzowski, J. F. (2004). Differences in hippocampal neuronal population responses to modifications of an environmental context: Evidence for distinct, yet complementary, functions of CA3 and CA1 ensembles. *Journal of Neuroscience*. <https://doi.org/10.1523/JNEUROSCI.0350-04.2004>
- Verney, C., Baulac, M., Berger, B., Alvarez, C., Vigny, A., & Helle, K. B. (1985). Morphological evidence for a dopaminergic terminal field in the hippocampal formation of young and adult rat. *Neuroscience*. [https://doi.org/10.1016/0306-4522\(85\)90275-1](https://doi.org/10.1016/0306-4522(85)90275-1)
- Voon, V., Thomsen, T., Miyasaki, J. M., De Souza, M., Shafro, A., Fox, S. H., Duff-Canning, S., Lang, A. E., & Zurowski, M. (2007). Factors associated with dopaminergic drug-related pathological gambling in Parkinson disease. *Archives of Neurology*. <https://doi.org/10.1001/archneur.64.2.212>
- W. E. Skaggs, B. L. McNaughton, K. M. Gothard, E. J. M. (1993). An Information-Theoretic Approach to Deciphering the Hippocampal Code. *Proceedings of the IEEE*.
<https://doi.org/10.1109/PROC.1977.10559>
- Wang, D. V., & Ikemoto, S. (2016). Coordinated interaction between hippocampal sharp-wave ripples and anterior cingulate unit activity. *Journal of Neuroscience*.
<https://doi.org/10.1523/JNEUROSCI.1042-16.2016>
- Wang, S.-H., & Morris, R. G. M. (2010). Hippocampal-Neocortical Interactions in Memory Formation, Consolidation, and Reconsolidation. *Annual Review of Psychology*.
<https://doi.org/10.1146/annurev.psych.093008.100523>

- West, M. J., Slomianka, L., & Gundersen, H. J. G. (1991). Unbiased stereological estimation of the total number of neurons in the subdivisions of the rat hippocampus using the optical fractionator. *The Anatomical Record*.
<https://doi.org/10.1002/ar.1092310411>
- Whishaw, I. Q., & Vanderwolf, C. H. (1973). Hippocampal EEG and behavior: Change in amplitude and frequency of RSA (Theta rhythm) associated with spontaneous and learned movement patterns in rats and cats. *Behavioral Biology*.
[https://doi.org/10.1016/S0091-6773\(73\)80041-0](https://doi.org/10.1016/S0091-6773(73)80041-0)
- Wightman, R. M., Heien, M. L. A. V., Wassum, K. M., Sombers, L. A., Aragona, B. J., Khan, A. S., Ariansen, J. L., Cheer, J. F., Phillips, P. E. M., & Carelli, R. M. (2007). Dopamine release is heterogeneous within microenvironments of the rat nucleus accumbens. *European Journal of Neuroscience*. <https://doi.org/10.1111/j.1460-9568.2007.05772.x>
- Wikenheiser, A. M., & Redish, A. D. (2015). *Hippocampal Sequences and the Cognitive Map*. https://doi.org/10.1007/978-1-4939-1969-7_5
- Wikenheiser, A. M., & Redish, A. D. (2015). Hippocampal theta sequences reflect current goals. *Nature Neuroscience*, *18*(2), 289–294. <https://doi.org/10.1038/nn.3909>
- Wilber, A. A., Skelin, I., Wu, W., & McNaughton, B. L. (2017). Laminar Organization of Encoding and Memory Reactivation in the Parietal Cortex. *Neuron*.
<https://doi.org/10.1016/j.neuron.2017.08.033>
- Willshaw, D. J., Dayan, P., & Morris, R. G. M. (2015). Memory, modelling and Marr: A commentary on Marr (1971) 'Simple memory: A theory of archicortex.' In *Philosophical Transactions of the Royal Society B: Biological Sciences*.
<https://doi.org/10.1098/rstb.2014.0383>
- Wilson, M. A., & McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, *265*(5172), 676–679.
<https://doi.org/10.1126/science.8036517>
- Wilson, M. A., & McNaughton, B. L. (1993). Dynamics of the hippocampal ensemble code for space. *Science*, *261*(5124), 1055–1058. <https://doi.org/10.1126/science.8351520>
- Wilson, M. A., & Tonegawa, S. (1997). Synaptic plasticity, place cells and spatial memory: Study with second generation knockouts. *Trends in Neurosciences*.
[https://doi.org/10.1016/S0166-2236\(96\)01023-5](https://doi.org/10.1016/S0166-2236(96)01023-5)
- Wirtshafter, H. S., & Wilson, M. A. (2019). Locomotor and Hippocampal Processing Converge in the Lateral Septum. *Current Biology*, *29*(19), 3177-3192.e3.
<https://doi.org/10.1016/j.cub.2019.07.089>

- Wirtshafter, H. S., & Wilson, M. A. (2020). Title: Differences in reward biased spatial representations in the lateral septum and hippocampus. *ELife*.
<https://doi.org/10.7554/eLife.55252>
- Witter, M. P., & Groenewegen, H. J. (1990). The subiculum: Cytoarchitectonically a simple structure, but hodologically complex. *Progress in Brain Research*.
[https://doi.org/10.1016/S0079-6123\(08\)61240-6](https://doi.org/10.1016/S0079-6123(08)61240-6)
- Wixted, J. T., & Squire, L. R. (2010). The role of the human hippocampus in familiarity-based and recollection-based recognition memory. In *Behavioural Brain Research*.
<https://doi.org/10.1016/j.bbr.2010.04.020>
- Wong, L. C., Wang, L., D'Amour, J. A., Yumita, T., Chen, G., Yamaguchi, T., Chang, B. C., Bernstein, H., You, X., Feng, J. E., Froemke, R. C., & Lin, D. (2016). Effective Modulation of Male Aggression through Lateral Septum to Medial Hypothalamus Projection. *Current Biology*. <https://doi.org/10.1016/j.cub.2015.12.065>
- Wu, C. T., Haggerty, D., Kemere, C., & Ji, D. (2017). Hippocampal awake replay in fear memory retrieval. *Nature Neuroscience*, 20(4), 571–580.
<https://doi.org/10.1038/nn.4507>
- Wyss, J. M., Swanson, L. W., & Cowan, W. M. (1979). A study of subcortical afferents to the hippocampal formation in the rat. *Neuroscience*. [https://doi.org/10.1016/0306-4522\(79\)90124-6](https://doi.org/10.1016/0306-4522(79)90124-6)
- Yang, K., Broussard, J. I., Levine, A. T., Jenson, D., Arenkiel, B. R., & Dani, J. A. (2017). Dopamine receptor activity participates in hippocampal synaptic plasticity associated with novel object recognition. *European Journal of Neuroscience*.
<https://doi.org/10.1111/ejn.13406>
- Yartsev, M. M., & Ulanovsky, N. (2013). Representation of three-dimensional space in the hippocampus of flying bats. *Science*. <https://doi.org/10.1126/science.1235338>
- Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*. <https://doi.org/10.1016/j.tins.2011.06.006>
- Yin, H. H., & Knowlton, B. J. (2004). Contributions of striatal subregions to place and response learning. *Learning and Memory*. <https://doi.org/10.1101/lm.81004>
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2006). Inactivation of dorsolateral striatum enhances sensitivity to changes in the action-outcome contingency in instrumental conditioning. *Behavioural Brain Research*. <https://doi.org/10.1016/j.bbr.2005.07.012>
- Yin, H. H., Knowlton, B. J., & Balleine, B. W. (2004). Lesions of dorsolateral striatum preserve outcome expectancy but disrupt habit formation in instrumental learning. *European Journal of Neuroscience*. <https://doi.org/10.1111/j.1460-9568.2004.03095.x>

- Yin, H. H., Ostlund, S. B., Knowlton, B. J., & Balleine, B. W. (2005). The role of the dorsomedial striatum in instrumental conditioning. *European Journal of Neuroscience*. <https://doi.org/10.1111/j.1460-9568.2005.04218.x>
- Yu, J. Y., & Frank, L. M. (2015). Hippocampal-cortical interaction in decision making. In *Neurobiology of Learning and Memory* (Vol. 117, pp. 34–41). Academic Press Inc. <https://doi.org/10.1016/j.nlm.2014.02.002>
- Zeitler, M., Fries, P., & Gielen, S. (2006). Assessing neuronal coherence with single-unit, multi-unit, and local field potentials. *Neural Computation*, 18(9), 2256–2281. <https://doi.org/10.1162/neco.2006.18.9.2256>
- Zhang, K., Ginzburg, I., McNaughton, B. L., & Sejnowski, T. J. (1998). Interpreting neuronal population activity by reconstruction: Unified framework with application to hippocampal place cells. *Journal of Neurophysiology*. <https://doi.org/10.1152/jn.1998.79.2.1017>
- Ziv, Y., Burns, L. D., Cocker, E. D., Hamel, E. O., Ghosh, K. K., Kitch, L. J., Gamal, A. El, & Schnitzer, M. J. (2013). Long-term dynamics of CA1 hippocampal place codes. *Nature Neuroscience*. <https://doi.org/10.1038/nn.3329>
- Zola-Morgan, S., Squire, L. R., & Amaral, D. G. (1986). Human amnesia and the medial temporal region: Enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus. *Journal of Neuroscience*. <https://doi.org/10.1523/jneurosci.06-10-02950.1986>
- Zutshi, I., Brandon, M. P., Fu, M. L., Donegan, M. L., Leutgeb, J. K., & Leutgeb, S. (2018). Hippocampal Neural Circuits Respond to Optogenetic Pacing of Theta Frequencies by Generating Accelerated Oscillation Frequencies. *Current Biology*. <https://doi.org/10.1016/j.cub.2018.02.061>