

Of Memories and Ripples

Functional and Mechanistic Aspects of Memory Sequences During Hippocampal Ripples

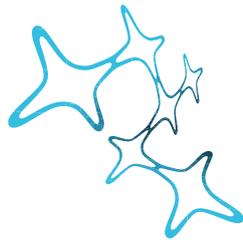
ÁLVARO TEJERO-CANTERO



Dissertation at the
Graduate School of Systemic Neurosciences
at the Ludwig-Maximilians-Universität München

Of Memories and Ripples

*Functional and Mechanistic Aspects
of Memory Sequences During Hippocampal Ripples*



Graduate School of
Systemic Neurosciences

LMU Munich

Dissertation at the
Graduate School of Systemic Neurosciences
at the Ludwig-Maximilians-Universität München

submitted by
ÁLVARO TEJERO-CANTERO

from Burgos, Spain

Oxford,
May the 9th, 2012

1st Supervisor/Reviewer Prof. Dr. Christian Leibold

2nd Supervisor/Reviewer PD Dr. Stefan Glasauer

Date of PhD Defense July the 25th, 2012

Of Memories and Ripples:

Functional and Mechanistic Aspects of Memory Sequences During Hippocampal Ripples

Copyright of the Text

- © Álvaro Tejero-Cantero, licensed under CC-BY-NC-SA (see below).
excerpts in Maier et al., 2011: © Elsevier Inc.
excerpts in Kammerer et al., 2012: © Springer Science + Business Media LLC.

Copyright of the Figures

<i>Figure numbers</i>	<i>Copyright holder (special rights)</i>
1.1, 1.2, 1.4, 1.5, 1.6, 2.7, 3.5, 3.7, 4.8.	indicated in respective Figure captions, (check each copyright holder for rights)
1.7, 2.1-2.4, 2.6, 2.9-2.16, 2.19, 2.22- 2.27, 3.6, 3.8-3.10, 3.18, 3.19-3.22, 3.25, 3.30, 3.31.	Elsevier Inc. (Maier et al., 2011), standard restrictive copyright
4.16-4.20, 4.22, 4.23.	Springer Science + Business Media LLC (Kammerer et al., 2012), standard restrictive copyright
All others.	Álvaro Tejero-Cantero, copyleft (CC-BY-NC-SA)

The CC-BY-NC-SA License

Full name. Creative Commons Attribution—NonCommercial—Share Alike 3.0

Summary and legal details. <http://creativecommons.org/licenses/by-nc-sa/3.0/>

a Ricardo,

hermano
compañero de juegos
amigo

Abstract

The hippocampus is one of the regions in the mammalian brain that is associated with memory of events in their spatiotemporal context. Sequences of neuronal activity in the hippocampus are the chief candidate for a neurophysiological correlate of such contextual, or episodic memory. Simultaneously to replaying these behaviorally-related activity sequences, the hippocampus engages in a powerful and fast oscillation known as sharp-wave ripples (SWR). Ripples in turn participate in a brain-wide pattern of activity and may orchestrate the local strengthening of memories and their broadcasting to the cortex. In this Thesis, both memory sequences and ripple oscillations are studied in the light of the unifying hypothesis that the coordinated activation of a neuronal assembly represents an individual memory item in the sequences, and is at the same time responsible for the individual cycles in the oscillations. To test the hypothesis, we investigated SWR in vitro and in vivo in the mouse, using intracellular recordings of currents in CA1 pyramidal cells referenced to the local field potential. Expanding current hypotheses on SWR generation, we found powerful, well ripple-locked and spatially pervasive but CA1-local excitatory inputs, indicative of presynaptic assemblies of CA1 principal neurons. Combining a novel peeling reconstruction algorithm for synaptic currents with recordings at different holding potentials, we could for the first time unravel individual synaptic contributions during ripples. Analysis of the strikingly precise timing of currents demonstrated that inhibition aligns its phase to excitation over the course of a ripple. We carried on the dissection of ripples to the theoretical domain by incorporating the effect of inhibition into a mean field model of sequence replay. Using this model, we inquired what are the neuronal assembly size and inhibitory feedback strength that maximize the capacity of a hippocampal network to store memories, so that those memories can be successfully retrieved during ripple episodes. We found that a linearly coupled inhibitory population indeed helps increase storage capacity by dynamically stabilizing replay in an oscillatory manner for lower assembly sizes than in absence of inhibition. The findings about the temporal structure of neuronal activation during ripples complement our experimental observations. Collectively, they offer new insights on the physiology and function of sharp-wave ripples, paving the way for an integrated, continuous-time model of large networks of sparsely connected neurons that replay activity sequences concomitant to transient ensemble oscillations.

Table of contents

Abstract	iii
List of Tables	ix
List of Figures	xi
Contributions	xiii

I Introduction

Prologue	2
1 Hippocampal Memory of Sequences	4
1.1 A Grammar for Memory in the Brain	4
1.2 Traces of Spatial Experience in the Hippocampus	9
1.2.1 Rat in a Linear Track: Place Cells, Reactivation and Ripples	10
1.2.2 Rat in a Linear Track: Assemblies, Oscillations And Sequences	12
1.3 Anatomical Context	14
1.3.1 The Hippocampus in the Brain	14
1.3.2 The Hippocampal Transverse Slice: Cells and Pathways	17
1.3.3 The CA1 Area of Hippocampus	19
1.3.4 The CA1 Pyramidal Cell	22
1.4 Physiology of Sharp-Wave Ripples	23
1.4.1 The Interpretation and Content of Hippocampal LFP	23
1.4.2 Sharp-Wave Ripples in the Local Field Potential	26
1.5 Research Goals Revisited	31

II Materials

2 Experimental Approach to Study SWR	34
2.1 In Vivo, Head-Fixed Parallel Recordings	36
2.2 An In Vitro Model for Investigating SWR	39
2.3 Electrotonic Isolation of Currents	41
2.3.1 Reversal Potentials	42
2.3.2 AMPAR-Mediated Currents	44
2.3.3 Timing of Spikes	48
2.3.4 Higher Holding Potentials	49
2.4 Pharmacological Isolation of Excitatory Currents	51
2.4.1 DIDS Block of Inhibition for Voltage Clamp Experiments	51
2.4.2 DNDS Block of Inhibition for Current Clamp Experiments	54
2.5 SWR in a CA1 Minislice	56

III Results

3 On Sharp-Wave Ripples	60
3.1 Amplitude-Based Large-Event Detection	61
3.2 Average Shapes of CPSCs Triggered to SWR Peaks	63
3.2.1 Average Amplitudes	63
3.2.2 Average Deconvolutions	65
3.3 Atomic Intracellular Events: Spontaneous PSCs	67
3.4 Steep Slopes	69
3.4.1 Comparison of Upward and Downward Steepest Slopes	69
3.4.2 Rhythmicity of Input: Inter-Slope Intervals	71
3.5 Locking of CPSCs to the Extracellular Ripple	73
3.5.1 Coherence	74
3.5.2 Slope-Triggered Averages of SWRs	75
3.5.3 Vector Strength of Steep Slopes	78
3.5.4 Concluding Remarks	80
3.6 Peeling Reconstruction of Synaptic Inputs	82
3.6.1 Position of the Problem	82
3.6.2 Anatomy of a Peeling Reconstruction Step	83
3.6.3 Peeling Reconstruction in Action	85
3.6.4 Comparison With Spontaneous PSCs and Validation	88
3.6.5 Detection by Deconvolution	90
3.7 Unweaving Single Post-Synaptic Currents from CPSCs	93
3.7.1 Onset Timing in Comparison to Timing of Steep Slopes	93
3.7.2 In-Ripple PSC Rates, Amplitudes and Kinetics	96
3.7.3 Cycle-Wise PSC Features	98
3.7.4 PSC Kinetics Are Stable Over the Course of the Ripple	101
3.7.5 Phase Drift Aligns Excitation and Inhibition During the Ripple	101
4 On Sequence Memory	104
4.1 Sequence Replay With Binary Neurons	105
4.1.1 Assumptions of the LK2006 Model	105
4.1.2 Imprinting	110
4.1.3 Capacity	112
4.1.4 Successful Replay as a Quality Criterion	113
4.1.5 Inhibitory Population	114
4.1.6 Cellular Simulations	115
4.1.7 Dynamical Regimes	118
4.2 Construction of a Mean-Field Model	120
4.2.1 Markovian Dynamics and Binomial Statistics	120
4.2.2 Mean-Field Model for Activation Averages in the Case $Q \rightarrow \infty$	122
4.2.3 Synaptic Input as a Random Process and the Optimal Detector	122
4.2.4 Synaptic Input Distributions From the Afferent Populations	123
4.3 Phase-Space Tools to Analyze Replay	132
4.3.1 The Mean-Field Rule Fully Specified	132
4.3.2 Receiver-Operating Characteristics	133
4.3.3 Mean Field Dynamics in the Phase Plane	135
4.3.4 Bayes-Optimal Threshold	139
4.4 Inhibition Increases Capacity for Sequence Replay	142
4.4.1 Instantaneous Inhibition	142
4.4.2 Dynamic Inhibition	147

4.4.3 Recurrent Inhibition	150
IV Discussion	
5 Ripples, Assemblies, and Sequences	154
5.1 Biophysical Mechanisms of Sharp-Wave Ripples	155
5.1.1 Executive Summary	155
5.1.2 Excitation and Inhibition	157
5.1.3 Origin of Ripple-Coherent Spikes	159
5.1.4 Emergence of the SWR Phenomenon	161
5.1.5 Discussion of Methods	163
5.1.6 Assemblies	166
5.1.7 Summary	167
5.2 Theoretical Model of Sequence Replay	168
5.2.1 Executive Summary	168
5.2.2 Replay Benefits From an Adaptive Threshold	169
5.2.3 Static Inhibition and the Optimal Threshold	169
5.2.4 Dynamic Inhibition and Oscillations in the Gamma Range	169
5.2.5 Transient Retrieval	170
5.2.6 Capacity Bounds Depend on Model Assumptions	171
5.2.7 More Sophisticated Models Needed to Formulate Predictions	171
5.3 General Conclusion and Outlook	172
Appendix A Electrophysiology Methods	174
A.1 Animal Surgery and Electrophysiology In Vivo	174
A.2 In Vitro Experimental Procedures	175
Appendix B Participation in Associations	180
B.1 Average and Variance of the Potentiation State	180
B.2 Correlations of Potentiation are Linear in Coding Ratio	181
Bibliography	182
Tools	190
Acknowledgements	192
Eidesstattliche Erklärung	194

List of Tables

Hippocampal Rhythmical and Nonrhythmical LFP patterns in CA	25
Summary of Main Electrophysiological Recording Conditions	35
Summary of Additional Electrophysiological Recording Conditions	36
Comparison of SWR Incidence and Ripple Frequency	57
Numbers of SWR Events Detected per Cell and Condition	62
Resources Required by Cellular Simulations	118
Firing Probabilities of Presynaptic Neurons According to Group	124
Morphological Connection Probabilities	124
Potential Probabilities	125
Probabilities of Spike Transmission	126
Accuracy of Gaussian Approximations to the Binomials	127
Means and Variances of Synaptic Input to On, Off and Inh Populations	131
Parameters Used for the Exploration of Recurrent Inhibition	150
Composition of Intracellular Solutions	176

List of Figures

Place Cells, Replay and Sharp-Wave Ripple Complexes	11
From Neocortex to Allocortex in the Macaque Brain	15
The Dissected Mouse Hippocampus	16
Orthogonal Sections of Mouse Brain: Views of the Hippocampus	17
Hippocampal Formation and Parahippocampal Cortices of the Rat	20
Passive Integration of Compartment-Specific Inputs to CA1 Pyramidal Cell	22
LFP Reflects Intracellular Currents	25
Hippocampal Recordings During LFP Ripples in Head-Fixed, Awake Mice	36
Statistics of LFP Ripples	37
Intracellular Current-Clamp Recordings in Head-Fixed, Awake Mice	37
Current-Clamp Recordings at Different Potentials	38
Time Course of Membrane Potential During Ripples	38
Intracellular Voltage-Clamp Recordings in Head-Fixed, Awake Mice	40
Layer Profile of Sharp Wave-Ripples in a Mouse Slice	41
Experimental Confirmation of Inhibitory and Excitatory Reversal Potential	43
Phasic Currents in CA1 Pyramidal Neurons During SWR In Vitro	44
Membrane Potential is Rippled During Field SWR	45
Coherence of Extracellular Ripple Activity in CA1	46
Parallel Recordings in Voltage Clamp	47
Correlation of cPSCs Across Simultaneously Recorded Principal Cells	47
Coherence of cPSCs Across Simultaneously Recorded Principal Cells	48
Spikes Are Entrained by Rippled Current Injection	49
SWR-Concurrent Spikes of Slightly Depolarized CA1 Cells Lock to Ripple	50
Example K _{Glu} 1 –45 mV	50
Example Cs _{Glu} –6 mV	50
Intracellular Blockade of Inhibition	52
Example DIDS –45mV	53
Example DIDS –60mV	53
Intracellular Blockade of Inhibition Reveals Strong Excitatory Contribution to cPSCs	54
DNDS Suppresses GABA _A Currents but Does Not Affect Spiking	55
DNDS Effectively Blocks Stimulus-Induced Fast IPSCs	55
Current Clamp Recordings of Two CA1 Pyramidal Neurons With DNDS	56
Spiking is Locked to Ripples Under Intracellular Block of Inhibition	56
Ripple-Coherent Excitatory Currents Persist in CA1 Minislices	57
SWR Peak-Triggered CPSC Averages	64
SWR Peak-Triggered CPSC Averages (DIDS Condition and its Control)	65
SWR Peak-Triggered Deconvolution Averages	65
SWR Peak-Triggered Deconvolution Averages (DIDS Condition and its Control)	66
Simulation of Synaptic Kernel for AMPA Currents	67
Kinetics of Spontaneous PSCs	69
Shape of CPSCs in Two Extreme Scenarios	70
Steepest Slopes of Excitatory Currents	71
Inter-Slope-Intervals in the Excitatory Conditions: Pooled Statistics	72

Inter-Slope-Intervals in the Excitatory Conditions: Single-Cell Statistics	72
Coherence Between SWRs and Their Concurrent CPSCs	75
Coherence Between SWR and Their Preceding CPSC	76
Averages of SWRs Triggered by Steep Slopes in Concurrent CPSCs	76
SWR Peak-Triggered Averages of SWR: Minislices Have Longer SWR	77
Averages of SWR Triggered by Steep Slopes in Preceding CPSC	77
Phase Locking of Steepest cPSC Slopes to Concurrent SWR	79
Phase Locking of Steepest Onsets to Subsequent SWR	81
Principle of Operation of Peeling Reconstruction	84
Example: Peeling Reconstruction of a CEPSC	86
Example: Peeling Reconstruction for an Entire Recording	88
Validation: Fit Parameters Compared to Spontaneous PSCs	89
Summary of the Peeling Reconstruction Analysis (KGlu –66)	94
Summary of the Peeling Reconstruction Analysis (CsGlu –6)	95
Steep Slopes vs. Peeling Reconstruction for Detection of PSCs	96
EPSC Event Rates, Amplitudes and Kinetics Across Cells	97
IPSC Event Rates, Amplitudes and Kinetics Across Cells	98
Amplitude Distribution of EPSCs Depends on Ripple Cycle (KGlu – 66)	99
Amplitude Distribution of IPSCs Depends on Ripple Cycle (CsGlu – 6)	100
In-Ripple PSC Time Constants	101
Relative Shifting of Excitatory vs. Inhibitory Phases during a Ripple	102
Phase Drift at Population Level	102
Combinatorial Neural Code for Sensory Percepts	106
Association of Memory Traces	107
Activation Variables	110
Populations of On and Off Neurons	110
Storage of an Association in the Synaptic Matrix	111
Effective Synaptic Network Induced by an Association	112
Dynamical Regimes in Phase Space	118
Three Asymptotic Dynamical Regimes in Cellular Simulations	119
Three Dynamical Regimes in Cellular Simulations With Inhibition	119
Reduced Connectivity of the Excitatory Network	121
Synaptic Input and Involvement in Associations	129
To Fire or Not to Fire?	132
Receiver Operating Characteristic (ROC Curves) for the Neuron Detector	134
Overview of Nullclines and Fixed Points of the Mean-Field Model	136
The Sigmoid Function Has Three Roughly Linear Regimes	137
Effect of the Firing Threshold on the Nullclines and Fixed Points	138
Phase Diagram	139
The Optimal Detector Threshold is Linear in the Network Activity	141
Threshold Coupling to Activity Along the Escape Trajectories	142
Instantaneous Inhibition	143
Replay-Level Threshold Normalization	144
Dynamical Feedback Inhibition	148
Inhibition is Tuned Up at the Edge of Stability	149
Phase Diagram Examples in Presence of Recurrent Inhibition	151
Summary of Research Presented in this Thesis	154
The Squared Coefficient of Variation is Linear in f Under Sparse Coding	181

Contributions

This Thesis presents work carried out in collaboration and published peer-reviewed journals (see List of Publications below). Here we list the contributions of third parties whose inclusion was essential to the integrity of the work and identify those that have appeared in the aforementioned publications.

Synaptic Basis of Hippocampal Ripples

All materials reported in Chapter 2 and the results in Chapter 3 with exception of cycle-dependent statistics of ripple-associated post-synaptic currents have been published in Maier et al., 2011 (reference **1** below). Excerpts therein have been reused in Chapters 2, 3 and 5 (discussion) in accordance with the copyright policy of Elsevier.

Experiments were carried out at the laboratories of Dietmar Schmitz (in vitro) and James F.A. Poulet (in vivo). Data analyses presented in Chapter 2 were carried out by Nikolaus Maier. Data analyses in Chapter 3 were carried out by Álvaro Tejero-Cantero, with exception of the manual postselection of EPSCs for EPSC fit statistics, which was done by Christian Leibold.

Mean-field Model of Sequential Memory

Most of the results reported in Chapter 4 have been published in abridged form in Kammerer et al., 2012 (reference **2** below). Excerpts therein have been reused in Chapters 4 and 5 (discussion) by permission from Springer.

Results in that Chapter were obtained jointly by Álvaro Tejero-Cantero and Axel Kammerer, building on a model by Christian Leibold and Richard Kempter (2006) and supervised by Christian Leibold. Exceptions are the numerical code for the verification of mean-field results with cellular simulations, which was written by Axel Kammerer and the Bayes-optimal threshold derivation, established by Álvaro Tejero-Cantero.

List of Own Publications

1. Nikolaus Maier*, Álvaro Tejero-Cantero*, Anja L. Dorn, Jochen Winterer, Prateep S. Beed, Genela Morris, Richard Kempter, James F.A. Poulet*, Christian Leibold*, Dietmar Schmitz*: *Coherent Phasic Excitation during Hippocampal Ripples*. *Neuron*, **2011**, 72, 137-152.
2. Axel Kammerer*, Álvaro Tejero-Cantero*, Christian Leibold: *Inhibition Enhances Memory Capacity: Optimal Feedback, Transient Replay and Oscillations*. *Journal of Computational Neuroscience*, **2012**, 1-12.

★ equal contribution.

Part I

Introduction

Prologue

How does our brain store memories? Does it ever exhaust its capacity for the storage of impressions? Does it risk becoming *full*? How does it retrieve what it saves? And how can it so vividly render the recollections of past events, made as they are of sounds, odours, images, sensations in general that are so rich and deeply intertwined? These questions have fascinated philosophers and scientists for centuries. In view of the breathtaking complexity of the biological machinery that are brains, it may never be possible to offer an explanation in simple terms—where all the ingredients are spelled out at the same deep level of detail and univocal causal paths are delineated. Notwithstanding the complexity of the global picture, the scientific community has managed to elucidate anatomical and physiological aspects of the brains *structure*, as well as facets of its *functional* output during learning, motor and sensory tasks.

Seldom are function and structure as intricately interlocked in a system as they are in the brain: cell types, architecture of interconnections, chemical diversity, nouns and verbs encoded by exquisite electrical patterns; everything in the brain seems delicately crafted for a specific function, region by region, system by system. For how could such a rich palette of structural expression be accounted for otherwise; one that displays continuity across species in the course of evolution yet allows for variation to meet needs imposed by the changing physical and social environment of individuals? We observe everywhere in the brain how local architecture enables specific functions, as rapid reaction, recollection of the past, integration of impressions from the senses, comparison and assessment. At the same time, we continuously confront evidence that the details of these tasks, even in minute timescales, are able to exquisitely sculpt the hardware that subserves their realization. It is not only the force of evolution that makes brains adapted, it is also the use of brains over a lifetime that hones them to machines capable of everything that we respect and admire, everything that we fear and avoid.

This Thesis swings between form and function as it tries to offer new knowledge about the workings of event memory. We ask here two related questions about memories and tackle them with rather different methods, as form and function relay each other at the center of our interest. First, we want to know how already stored memories are reenacted in the brain, be it for conscious recollection or as an unconscious maintenance process. Some of the most striking reenactments of past experience are best visible in the hippocampus—a brain region, as it starts to pulsate electrically in a very fast and well synchronized oscillation. How are neurons coordinated so that they participate with the required precise timing, and what do individual cell classes—instruments—play in the symphony at each stage during the reenactment of memories? That is the setting for our first question, which we tackle with a descriptive spirit—how *are* things in Nature?

Memories unfold in time, one impression succeeding another so as to build stories, sequences of discernable events. Arguably memories would lose much of their interest if they could not be woven into the fabric of our existence—by recalling past sequences we predict what to expect and decide if it is worth our pursuing one or other course of action. How does one impression give way to the next during the reenactment of a story? Following ideas developed from the 1970s about how memories are associated in pairs, and pairs among them to create storylines, we enquire how many of those storylines can fit into a neural network, and if the diversity of neural substrates—form—could not help increase that capacity—function. That is our second question, that we confront from a more *normative* viewpoint—how *should* simple memory circuits be designed to fulfill their putative function?

To restate both questions with the necessary precision so that they can guide our investigation, we shall introduce in **Chapter 1** the notions of neuronal assemblies, plasticity and storylines as sequences—function. Coming down from these abstract concepts to a concrete realization in the brain, we shall present the anatomical and physiological traits of the hippocampus that enable them—form.

We then shall head on to our first question on the timing of neuronal activity during fast hippocampal oscillations. This question we address by extensive analysis of electrophysiological signals recorded from both the extracellular space and principal cells in slices of mouse hippocampus. The presentation is divided in two Chapters. **Chapter 2** is concerned with the recordings by our collaborators in Berlin and how they relate to previous characterizations of fast hippocampal oscillations, and is mainly about methods. The next two Chapters present the main results of this Thesis. **Chapter 3** discusses our data analysis work based on state-of-the-art techniques in the time domain. The results on the relative timing of excitatory and inhibitory inputs to pyramidal cells in vitro will contribute towards the understanding of the inner workings of the fast hippocampal oscillation.

In **Chapter 4** we turn to the capacity of the hippocampus to hold sequences of memories and how to improve upon it. There, we extend the mathematical model of sequence memory by Leibold and Kempter (2006) in order to gain insight about the role of inhibition in the replay of memories. With this biologically inspired model, we find under what conditions an inhibitory population may stabilize the replay of sequences and thus enhance the function of the hippocampus as a memory network. Ultimately, these lines shall converge, pushed by the collective effort, towards answering the overarching question of how biological networks operate in time to store and recover as many sequences of memories as possible in as little brain as feasible. In the last **Chapter 5** we critically appraise our contribution in light of the published literature and discuss the limitations of our experimental and analytical methods of interrogation. It is our hope that those limitations will be as eloquently portrayed as their modest achievements.

Chapter 1

Hippocampal Memory of Sequences

In the first two Sections in this Chapter we intend to present our research questions with as little technical language as possible: we only assume knowledge of what neurons are and how they transmit electrical signals across synapses. The reader familiar with basic neuroscience and the Hebbian notions of assemblies and plasticity can skim through the first Section; if additionally acquainted with the classical place cell reactivation experiments she may skip the second altogether and head straight to Section 1.3. There we present the anatomy of the hippocampus in its brain context, zoom down to the region of our interest (CA1) and there concentrate on the pyramidal cell, which will be our telescope of sorts to investigate the network activity. We complement the introduction of the physical substrate with a description of one of its physiological modes of expression: fast oscillations. These fast oscillations could subserve the replay of memories of sequences of events.

We are at that point ready to reformulate our research question in precise terms, which we do in Section 1.5.

1.1 A Grammar for Memory in the Brain

This Section discusses a few foundational concepts about neural network realizations of sequence memory, without particularizing on *what* is the content of those sequences or *where* in the brain they occur. We aim to establish a common ground with the reader around the abstract concepts of memory, neuronal assemblies, synchrony, network rhythms and plasticity to set the context for further discussion particularizing on *spatial memory* as realized in the *hippocampus*. We shall only need to know about the hippocampus that it is a brain structure found in vertebrates that has been widely recognized in modern times to subserve important memory functions and to play a central role in a number of human brain disorders, most notably epilepsy.

Memories This work seeks to provide insight on a particular aspect of the general theme of memory. We are interested in understanding the unique features of evolved, versus engineered, memory architectures. Memory is expressed in psychophysical tests as a correlation between past experience and present action, often measured as performance in a well-defined task. Memory is diverse in several ways, and in humans numerous distinctions are made according to e.g. how long it lasts (long vs. short-term), whether the subject is conscious of demanding its recall (explicit vs.

implicit), or what kind of tasks it helps perform (motor programmes, reflex reactions, forecasting and decision-making or the solution of abstract problems demanding a relational understanding of the world). Most tasks require several of these “clear-cut” memory types. Each variant of memory attracts considerable attention from brain scientists often with the ultimate goal of helping overcome human memory deficits and cognitive disorders.

Memory systems An enormous body of research has been directed at understanding each of the types of memory enumerated above. One wants to know where, when, and how is the imprint of the experience established and how is it read out when action is due or recollection is commanded. All along the way the reductionistic approach demands dissection of these broad answers down to the biophysical and biochemical mechanisms involved: *what* changes in the electrical activity in memory-related areas during storage, maintenance and recollection or *how* is the dialectic relationship between the ephemeral electrical expression of memory and its more stable neuro-physical substrate: the number and types of cells, their structural connectivity, the chemical state inside them and on their membranes, and particularly where they come in contact forming synapses; all this and more is in the limelight in memory research. The answers feed back into the functional classification of memory and inform it structurally, in a dialogue that is aimed—perhaps too ambitiously in an evolved setting—at delineating precise boundaries for *memory systems* responsible of *memory functions*.

Memory traces in neuronal activity Any contribution to the daunting task of understanding memory is faced with its enormous diversity: diversity of behavioral expressions and performances, diversity across species, across brain subsystems in a single model organism, and also necessary diversity of the levels of description involved (Craver, 2007). Clearly, focus is needed to characterize any one facet, at the likely expense of perspective. Here is the line of reasoning leading to our research question.

The most elementary changes that can be related to memory are at the level of synapses and even smaller structures. It is very difficult, however, to produce a picture of those changes leading to network-level understanding, for several reasons:

- the number of synapses is exceedingly large, scaling as the square of the number of neurons,
- their size is minute,
- their detailed state is stipulated by concentrations of chemicals and spatial distributions of membrane features changing across several time scales.

All this makes it very hard to acquire data in order to characterize memory in populations of synapses, but also to use it for modelling in search of generalization beyond of what is measurable. Alternatively, one can address memory indirectly through the firing activity of the cells hosting the synapses, or even their membrane potential. Whilst still a very ambitious program, this line of attack (electrophysiology of populations of neurons) has already produced valuable insights — of which a few will be recounted below. If synapses change to reflect experience, they are probably informed to do so by electrical activity of neurons as transduced from sensory organs; and their modification is reflected back on the electrical activity transmitted across neurons.

The electrophysiological dynamics of a network must thereby indirectly bear the trace of the memory that was encoded as physicochemical changes in the synapses. We thus frame this work under the wider research question: can we discover traces of memories in the activity patterns of groups of neurons?

Neuronal assemblies Several or even many neurons must be involved in the expression of each of those *memory traces*, for otherwise, with unreliable, biological units, the failure to fire of a single or a handful of neurons would threaten survival. Their activity must be concerted in time, because downstream of them, other neurons assume the task of processing the memory trace in a cognitive setting and eventually cause motor action. These other neurons are known to lose track of what was communicated to their synapses as fast as current leaks out of their membranes i.e. in only a few dozen milliseconds, and so they demand a certain degree of simultaneity in their inputs to do their part in the computation. In addition to this biophysical constraint, the expression of memory traces is subject to a behavioral constraint—where memories are instrumental for survival, sequences of complex scenes must be recalled promptly, which requires well-timed activity. We come thus to the concept of a *neuronal assembly*: a set of neurons that by sheer number, or good synchronization, or adequate doses of both, manage to make themselves heard in the immensity of the brain network. Following the path delineated by Hebb (1949), we hypothesize that memory traces must be visible as synchronous activation of neurons, or neuronal assemblies (see Harris, 2005 for review).

Synchronized discharge of assemblies and brain rhythms There is in the previous sentence an implicit assumption: that *activation* (firing) is *the* phenomenon because it can be relayed to another cell via synaptic communication, and that the subthreshold dynamics is only the detail of the mechanism leading to it. This may be an oversimplification promoted by our measurement instruments: intracellular recordings are limited to a handful of neurons at a time and are thus insufficient to address the network-level significance of subthreshold communication. Two main such channels of subthreshold communication have been suggested: direct electrical contact between cell membranes via *gap junctions* or coupling through volume conduction in the extracellular milieu (*ectopic*), which has the trait of being dependent on proximity (metric, rather than topological neighbourhood).

What is meant by *synchronous*? In the brain, the nearest closest analog of a clock is provided by activity waves at the scale of brain regions. Our blunt instruments, in this case extracellular electrodes, reveal only low frequency (<500 Hz) local oscillations of the potential of the electric field in the extracellular milieu (*local field potentials*; LFPs), but their mere existence and prominence against the noise strongly suggests already two preconditions. First, the arrangement of the currents generating the potential is geometrically orderly. Otherwise, e.g. with cells placed totally at random, sources would cancel and the LFP would be flat in average. Second, that those cells act in timely consonance (the capacity of the extracellular milieu being negligible to produce time integration by itself). That consonance must be at the scale of the variation of the potentials, which is the oscillation cycle. Our previous demand of synchrony of neuronal assemblies in the timescale of a neuron membrane time constant is thereby complemented with the knowledge that large groups of neurons must be firing at oscil-

lation frequency for the oscillation to occur at all, and that it is likely that postsynaptic targets sense that oscillation as increased or decreased input at certain times in the oscillation. To what extent the assemblies that express memory directly generate the oscillation, or whether they simply participate in it by tuning their firing to a certain stage of a cycle to match an increased or decreased sensitivity of their postsynaptic targets remains open.

Despite the well-founded expectations raised by the observations above, based on the local field potential, the evidence for the existence of assemblies is still scarce. Since firing of putatively assembly-participating cells is generally very sparse, it is difficult to see them simultaneously active when only few cells (under 200) in a brain area can be recorded at the same time.

Fire together, wire together We have noted above that from the point of view of an observer, whether internal (a neuron) or external (a multielectrode array monitoring brain activity), memories should be visible as tightly firing assemblies of neurons. There is still one more argument to the importance of timing and cooperativity: the establishment of assemblies (as opposed to their readout, discussed above) is dependent on synaptic changes that require quasi-simultaneous but ordered firing. Two questions have to be solved: *a*) how do cells representative of the same memory item bind together into an assembly (where order does not matter) and *b*) how do they relay their activity to the next assembly. It turns out that the second, seemingly more complex question may help solve the first. Consider an assembly formed in response to a sensory stimulus (i.e., an assembly of cells tuned to a certain stimulus feature). Other cells down the line will be driven by the extensive connectivity of this original, ad-hoc assembly. These cells, which fire on short order, may or may not be structurally connected. If they are close enough to be connected, they will fire in rapid succession (allowing for the randomness of axonal delays). One might produce synaptic release, the other (postsynaptic) may experience backpropagating action potentials, i.e. an elevation of membrane potential at the dendrites. These two requirements, when manifested in a window of about 5 to 40 ms, can strengthen the connections between neurons, a phenomenon that has been called synaptic time-dependent plasticity (STDP). By this expedient, cells in the target protoassembly are programmed to fire jointly above and beyond the stimulus-driven synchronization induced by their afferents: the repetition of synchronous stimulation strengthens the connections in the direction of information flow. This form of directed plasticity usually goes by the name of Hebbian plasticity (Hebb, 1949). A succeeding sensory input may be encoded into the same assembly, if it is quick enough, because convergence of sensory input to the memory network will make the cells there establish an assembly, regardless of the different sensory modalities of the inputs. Or it could be assigned a different assembly of the memory network. The directionality of synaptic connections will promote the firing of the second assembly by the first. This is the hypothesized mechanism for the formation of an association of two items. Now presentation of the first sensory stimulus will trigger activation of the second, at least to a degree that makes it stand out for a readout system. Note that it is the simultaneity of the inputs that decides whether the sensory experience is bound into a single assembly with predominantly symmetric connections within, or is rather split in several such assemblies, with

asymmetric connectivity among them in order of activation. The asymmetric structure of chemical synaptic communication, and the conjunctive character of synaptic plasticity (requiring pre and postsynaptic activation) are here the requisites for this mechanism to operate the formation of associations in a memory network.

Sequences of assemblies The associations formed by succeeding sensory impressions may be woven together into storylines represented by neuronal assemblies activated in sequence. A central concern when contemplating such a cue-target/cue-target/cue-... scenario is robustness in the face of noise. Memory architectures that can complete a defective pattern, i.e. recover a target from a damaged or incomplete cue (Heerden, 1963; Willshaw et al., 1969; Marr, 1971), are instrumental in this setting, even if they have to trade off capacity for robustness.

Inhibitory cooling The mechanism just described depends only on excitatory connections between neurons, i.e. there is no need for cells that upon firing discourage or block the firing of their postsynaptic partners. However, the brain contains such neurons; they are a numerous minority (about one tenth to one fifth, depending on the region) and they seem to have intricate designs resulting in complex brain state-dependent firing patterns. How to account for their presence? The natural tendency of theorists to ignore them has been solidly backed by the relative absence of data on these cells as compared to their excitatory counterparts. Indeed, inhibitory neurons are not only scarcer, but are also difficult to characterize extracellularly from spike shapes, and to address intracellularly due to their generally smaller somatic size. Often, they have been assigned a dynamical homeostasis role. Two general characteristics contributed to this view: the short range of most of their projections (“local interneurons”) and the seemingly lesser complexity of their synapses, which are directly on the dendritic shafts of the target cell, in contrast to the membrane specializations that form excitatory inputs (“spines”, Yuste, 2011). Further, supporting the need for a homeostatic brake, excitatory-only recurrent networks have the tendency to explode into epileptic-like states of whole-network activation. Inhibitory interneurons with no particular selectivity in their projections would sense the overall activity of the network and scale it down, introducing a stabilizing feedback loop. This regulatory role we describe and quantify in our model for sequence retrieval from a memory network, which is presented in Chapter 4.

... or fading from the future? There have been theoretical proposals as well that embed interneurons at the heart of memory sequences (Rabinovich et al., 2001). Establishing mixed excitatory-inhibitory assemblies and asymmetric inhibitory connections (stronger from the target assembly to the cue assembly than in the natural order of activation cue-assembly) provides another mechanism for relaying activity, one that has been suggested to underlie how locusts reliably represent complex odours by sequences of neuronal activity unfolding with time (Wehr and Laurent, 1996). Activity grows in the initial assembly and is progressively relayed to another ensemble, whose interneurons shut down the incoming activity before the second assembly itself is silenced by the interneurons in the third assembly, etc. In this setting, oscillations increase the combinatorial space for expression of a *single* (but complex) memory item by adding to identity the use of time, thus enabling a so-called spatiotemporal

code. While such an architecture based on asymmetry of interneuronal connectivity is also liable to enable transient assemblies that represent distinct memory items, no evidence has been found of its existence in the hippocampus.

Consolidation of assemblies It is natural to ask if there may be a process to help assemblies related to behavioral events stand out from opportunistic, coincidental, random assemblies, or noise. Such a process would offer the advantage of demanding less repetitions of the behavioral cues to establish the memory. This could indirectly enable the memorization of longer sequences, which those are less frequently presented. Following the mechanistic picture we have suggested above, assemblies have to fire in order to weld into a sequence, however now, for consolidation purposes the sensory input is absent by definition. Therefore, for a brain network to be able to engage into consolidation, it must be possible for it to operate in a hyperexcitable state when sensory input is absent. In contrast, when sensory input is active, network-level hyperexcitability is undesirable to the extent that activation of anything but the strongest-driven assemblies would result into confusion.

Broadly, a systems-level state switch could be implemented either by bathing a network in a neuromodulatory agent, or by a network-wide excitatory signal, or by the release of an inhibitory block. Which mechanism is used defines the timescale of the network switch. While sensory deprivation during sleep seems ideally suited for consolidation, the pattern of network activity believed to be linked with consolidation has also been observed in awake immobility, including very brief pauses during movement. This observation would speak for consolidation not to be dependent on neuromodulation, which is generally slow-acting.

In the excitable state, random (uncued) activation of ensembles of neurons would occasionally have enough overlap with a recently formed assembly so as to kickstart the sequence reproduction process. Naturally, premature termination, or a branching-off to an unrelated sequence are possible if the connections were weakly established. This “lifting up from the noise” that is suggested to underlie local consolidation in the brain area where sensory inputs first converge would work just as well for systems-level consolidation, i.e. for sending information about the sequenced firing to other brain areas (Battaglia et al., 2011). Particularly well-established cue assemblies would start their sequences from the random excitable state correspondingly often, and their targets outside of this memory network would mirror the sequences as described above, eventually enabling the transfer or copy of the memory to the remote location.

1.2 Traces of Spatial Experience in the Hippocampus

In the previous Section we have introduced the notions of memory, neuronal assemblies, time defined by brain rhythms, sequences of activity and Hebbian plasticity. Here we present these concepts in practice, taking as a motivational guide the most thoroughly researched memory system (the hippocampus), the current paradigmatic experimental realization of episodic memory (reactivation of cells that code for space), and the model systems of choice for its investigation (rodents).

How these elements are woven together is at best illustrated by recounting one of the now classical experiments on sequence memory.

1.2.1 Rat in a Linear Track: Place Cells, Reactivation and Ripples

We describe next the reactivation experiment by Diba and Buzsáki (2007). The behavior of the animal, and the brain activity in form of oscillations and neuronal action potentials are illustrated in Figure 1.1.

Behaviour and LFP rhythms In the study by Diba and Buzsáki, a rat ran along a linear track for water reward at either end. The track used for this kind of experiments is itself rather featureless, but the rat can see the laboratory environment above and around, which provides visual landmarks in addition to other sensory cues that human observers do not easily perceive. A set of electrodes had been implanted in the brain of the rat in order to monitor the electrical activity in the hippocampus. When the rat was running an oscillatory pattern became apparent. It was the so-called theta rhythm, a slow (5-10Hz) oscillation of the local field with a somewhat triangular wave shape. The theta rhythm (θ activity) is associated with states of movement, where it shows oscillation frequency increasing with locomotion speed, as well as with shallow (REM) sleep phases. While not running (i.e. when grooming, resting or when in deep, non-REM sleep), a different brain pattern sets in: most of it seems unstructured but it is irregularly punctuated up to once per second by sudden, acute deflections called *sharp waves*. A feature of these sharp waves is that they display a fast superimposed oscillation, called a *ripple*. The joint phenomenon is denominated *sharp-wave ripple* (SWR) and often referred indistinctly by either of its two components as a sharp wave or a ripple.

Neuronal activity The oscillatory rhythms just introduced, theta and sharp waves, are expressed in the electrode recordings depending on the behavioral state. What do individual cells do during them? It is in the high-frequency components of the signal (above 500 Hz) where the neuronal firing is visible as blips or spikes. The availability of several close-by electrodes in the probe permits one to compare and triangulate the high-frequency signals as seen from slightly different vantage points, and assign to each of them a putative “source identity”. The inference of which kind of neuron is behind one of these clusters, or units, relies in often heuristic assessments of its firing phenotype—both the thin, but not structureless, shape of the spike, and the rate and pattern of firing, usually as seen in the autocorrelation of spike times.

Place cells and the theta rhythm Some of the identified units fire preferentially at particular spots of the linear track, called place fields by analogy with other sensory receptive fields. These fields are established after only a few repetitions, are stable over the course of the experience, and fade out a few days later. They represent a neural correlate of a higher-order sensory construct: “location” (involving vision, olfaction, proprioception...). A spatial trajectory of an animal is thus represented by several of these *place cells* firing in sequence. This has been confirmed in several studies where the position has been reconstructed with probabilistic methods from the neural activity alone, once the meaning of the spikes was known (by observing both positions and spikes for some time).

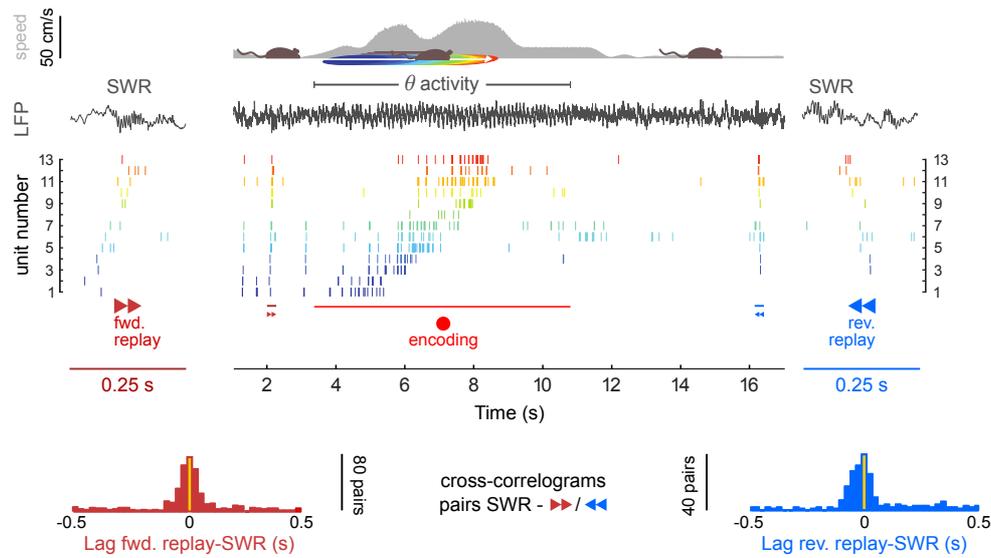


Figure 1.1. Place Cells, Replay and Sharp-Wave Ripple Complexes. A rat runs along a linear track while the electrical activity in region CA1 of its hippocampus is recorded by an array of electrodes. **Behaviour.** Speed is shown by a filled grey curve at the top panel. The rat is practically immobile before and after running. **LFP.** An example trace from one electrode is shown in dark grey. A rhythmic oscillation at 4-8 Hz (θ) is visible during movement; elsewhere the LFP is irregularly punctuated by Sharp-Wave Ripples (SWR; two examples shown magnified at the flanks). **Units.** Units have been ordered according to their firing during movement and their spikes are color-coded according to the region of the environment (*place field*) in which they are active (see schematic of rat movement at top). **Replay.** Two example instants at which the majority of cells activate quasi-simultaneously have been marked \blacktriangleright and \blacktriangleleft ; they are shown expanded at the left and right side, respectively. The magnification shows a firing order that is correlated resp. anticorrelated with that during subsequent resp. precedent movement, representative of many runs. Replay occurs coincident with fast (ripple) oscillations in the LFP. **SWRs.** The synchronization of both forward and reverse replay events extends to the whole dataset, as established by the crosscorrelograms at bottom. Adapted with permission from Diba and Buzsáki (2007) and Carr et al. (2011).

Reactivation and ripples These “conjunctive sensory units” (units that activate upon reception of the appropriate multimodal sensory stimuli) also discharge occasionally in absence of sensory input. Usually, they fire very sparsely and with no apparent pattern. But during hippocampal ripples that appear irregularly during awake immobility or slow wave sleep (SWS), they activate spontaneously forming short bursts (1/10th of a second) in an order that is highly correlated to the order in which they were activated during the previous experience. This *reverberation* of the run along the track is repeated time and again both in forward and reverse order. Forward reverberations, or reactivations, tend to happen before the animal starts running on the given direction, and are speculated to play a role in planning and anticipation of events, whereas reverse reactivations overwhelmingly correlate with rest at the target location and offer a potential solution to the *credit assignment problem*, i.e. finding out which of the past actions were essential for bringing about reward. Forward reactivation is presumed to rely on synaptic mechanisms. Reverse replay, following experience rather immediately, has been suggested to rely on an activity-based mechanism: proximate-field place cells remain more excitable than distant-field cells upon arrival of a post-movement rising excitation wave (SWR), thus advantaging replay in the reverse order (Buzsáki, 1989).

The reactivation of spatial experiences in experiments with awake, behaving animals motivates our interest in the more mechanistic aspects of the architecture of memory networks for retrieval (Chapter 4) and the synaptic underpinnings of the retrieval-related rhythm of SWR (Chapters 2 and 3). We direct the reader further interested in spatial reactivation to the recent reviews by Buhry et al. (2011) and Carr et al. (2011) for discussion of aspects such as two-dimensional environments, reactivation during other behavioral states, functional value, concurrent reactivation of assemblies and correlated LFP activity in other brain areas, reward and motivation, etc. In the next Section we concentrate on the concepts of neuronal assemblies and oscillatory rhythms (defined in Section 1.1), important for the memory of sequences in the light of the example experiment just described.

1.2.2 Rat in a Linear Track: Assemblies, Oscillations And Sequences

In order to frame our research question, we shall now recapitulate the conceptual elements that converge in this story. First, there are two dominant electrical patterns: theta and irregular sharp-wave activity hosting fast ripples. Second, the firing of excitatory neurons during theta is tuned to a particular position—a neural correlate of a high-level physical variable. Third, a sufficient number of cells manifest this spatial tuning so that even capturing only a few permitted Diba and Buszáki to show a reasonable coverage of the linear track. Fourth, the same cells that fire in short succession during the behavior fire at other times, either before running or after. Fifth, when they fire, they follow the sequence expressed during running, or its reverse, with high significance. And sixth: they precisely choose the irregularly occurring sharp-waves, outside of theta, for their coordinated reactivation.

The abstract concepts outlined in Section 1.1 can be matched to some extent with a concrete realization in this experiment

1. **memory** — place cells represent a neural correlate of a physical variable; the experiment does not establish, however, that this correlate has a functional value. Although this important issue has been tackled in later experiments, e.g. Dupret et al. (2010), the leap from reactivation to memory should be taken in this work as a working hypothesis, not an established fact.
2. **memory systems** — place cells are found in a particular region of the cortex, the hippocampus; a region that is anatomically and cytoarchitecturally extremely well defined and that is known from lesion studies from rodents to humans to be critically concerned with the establishment of episodic memories.
3. **memory traces** — the sequential activation of neurons as the rat progresses on the track is reinstated later during e.g. immobility. There is presumably a biophysical structural substrate, likely synaptic, that supports this well-timed reactivation, but we cannot investigate it in large scale; we content ourselves with its imprint or *trace* in the spiking activity of the small fraction of the reactivating cells that we can measure in an electrophysiological experiment (note that the term *trace* is often used for the durable expression, perhaps synaptic, of the memory, for which we will use rather simply the term *memory*).

4. ***assemblies*** — one may suggest that cells that fire together in the short temporal framework of one sharp wave (about 0.1 s) form an assembly. Our working definition of assembly relied, however, on the representation bearing an atomic unit of content, a single memory item. The successive positions, at the resolution of one place field, do map better to our concept of unitary, structureless representation than the whole track as represented by the sequence. It would be uneconomic to encode each possible trajectory in one assembly. Additionally, studies in two dimensional environments have shown that trajectories that pass through a certain place field do show firing of the corresponding place cell. We take this not as evidence against assemblies, but as indication that they exist at a finer temporal level of expression.
5. ***rhythm-synchronized assembly expression*** — recordings with current multi-electrode arrays allow to sample at most up to two hundred neurons, of which forty percent will be place cells, only a fraction of which, in turn, will be active in a given environment. If assemblies exist over very short time scales, the current numbers of simultaneously recorded neurons are insufficient to prove it. In spite of that, here is our hypothesis in this Thesis, in mnemonic form:

one spatial location (memory item) \sim one assembly \sim one ripple cycle

6. ***consolidation*** — the shortened time-frame for reactivation of behavioral sequences has been proposed to be well adapted to currently known mechanisms of plasticity (long-term plasticity, Buzsáki et al., 1992). The ordered reactivation matches the needs of STDP for establishing synaptic flows forward in time (for planning) and backward (for valuation, or credit assignment).
7. ***broadcasting*** — beyond the immediate consolidation of the path as a reinforced sequence of firing cells in the hippocampus, there may be the need of sending away the information. Two functional needs could be served by broadcasting this information to other areas in the cortex. First, retrieval in the hippocampus could cue sensory areas to reinstate the full flavor of the experience (“mental time travel”, Carr et al., 2011). Second, it could subserve consolidation of the memories in longer-term storage areas of the cortex (McClelland et al., 1995). That memories become with time independent of the hippocampus has been observed in lesion studies, most notably in humans. It is not excluded that the replay of memories to reinstate the experience could have as a side effect their becoming independent of the hippocampus. What is interesting in this context is the coincidence of hippocampal reactivation with sharp-wave ripples: ripples constitute a very strong collective activation of the entire hippocampal network and thus are ideally suited to function as carrier waves to hippocampal target areas. Sleep is ideally suited for separating learning and consolidation and so prevent reinstatement to result in hallucinations (Diekelmann and Born, 2010).

1.3 Anatomical Context

In this Section we present the main anatomical and physiological characteristics of the hippocampus relevant to our work. We place the hippocampus in its wider anatomical context in the brain in Section 1.3.1. We discuss its internal structure of layers and pathways in Section 1.3.2, and home in subsequently on its CA1 subfield (Sec. 1.3.3) and, further into it, on the CA1 pyramidal cell (Sec. 1.3.4).

When referring to abbreviations in any of the figures, we will use special type, e.g. **CAM**, which stands for *Cornu Ammonis*.

1.3.1 The Hippocampus in the Brain

The hippocampus is found in all mammals with a distinctive architecture that changes little with phylogenetic development. There are however morphological variations of the basic elongated and curved theme from species to species (Insausti, 1993, Fig. 2; compare macaque staining in Figure 1.2 with mouse in 1.4 and rat in 1.5).

We zoom out from the area of the rat brain where the recordings shown in Figure 1.1 of the previous Section were obtained. Cornu Ammonis 1, or CA1, is a field of the hippocampus proper, which is formed by CA1, CA2, and CA3. Together with the U- or V-shaped dentate gyrus (**DG**), these three fields conform the hippocampus^{1.1}. The hippocampus and the nearby structures subiculum (**SB**), presubiculum (**PSB**), parasubiculum (**PASB**), and entorhinal cortex (**ENT**) constitute the hippocampal formation (**HF**). Histologically, the hippocampal formation transitions from folded archicortex (ancient cortex) to the modern neocortex in the medial temporal lobe (e.g. region 36 of the perirhinal cortex; Figure 1.2), and is conventionally ascribed to the limbic system. The limbic system forms a medial ring around the corpus callosum and the hilus of the hemisphere. It is on the *limbus*, or border, of the medial surface of the lobe and comprises several cortices. The inner ones (olfactory, amygdala, septum and hippocampus) are simpler in construction, whereas the outer ones (subgenual, cingulate, retrosplenial and parahippocampal) look more modern in architecture. While the cortices in the parahippocampal gyrus (entorhinal and presubicular regions) are structured in six layers much like neocortex, some of these layers have few cells, which led to the denomination of schizocortex or “split cortex”. The term limbic system is an important keyword for communicating about subcortical brain areas in the cerebrum, but lately anatomists tend to eschew its use and refer to the hippocampal formation as an independent subcortical structure (Amaral and Lavenex, 2007).

Figure 1.2 illustrates the aforementioned spatial transition from the phylogenetically more modern neocortex, convolved at the right, to the older hippocampus, housed in a fold at the left (labeled **CAM** and **DG**). The basic layout of cells and fiber pathways of the hippocampal formation is the same across mammals (see Section 1.3.2). In rodents, however, the shape of the hippocampus is somewhat simpler than in primates and of more homogeneous thickness, and it takes up a much larger fraction of the brain. Each of the hippocampi resembles a curved tube. Both are joined

1.1. Some authors limit *hippocampus* to the cytoarchitecturally similar Cornu Ammonis fields and consider the DG as part of the larger *hippocampal formation*, which is not cytoarchitecturally homogeneous but defined by predominantly unidirectional pathways threading through the areas involved.

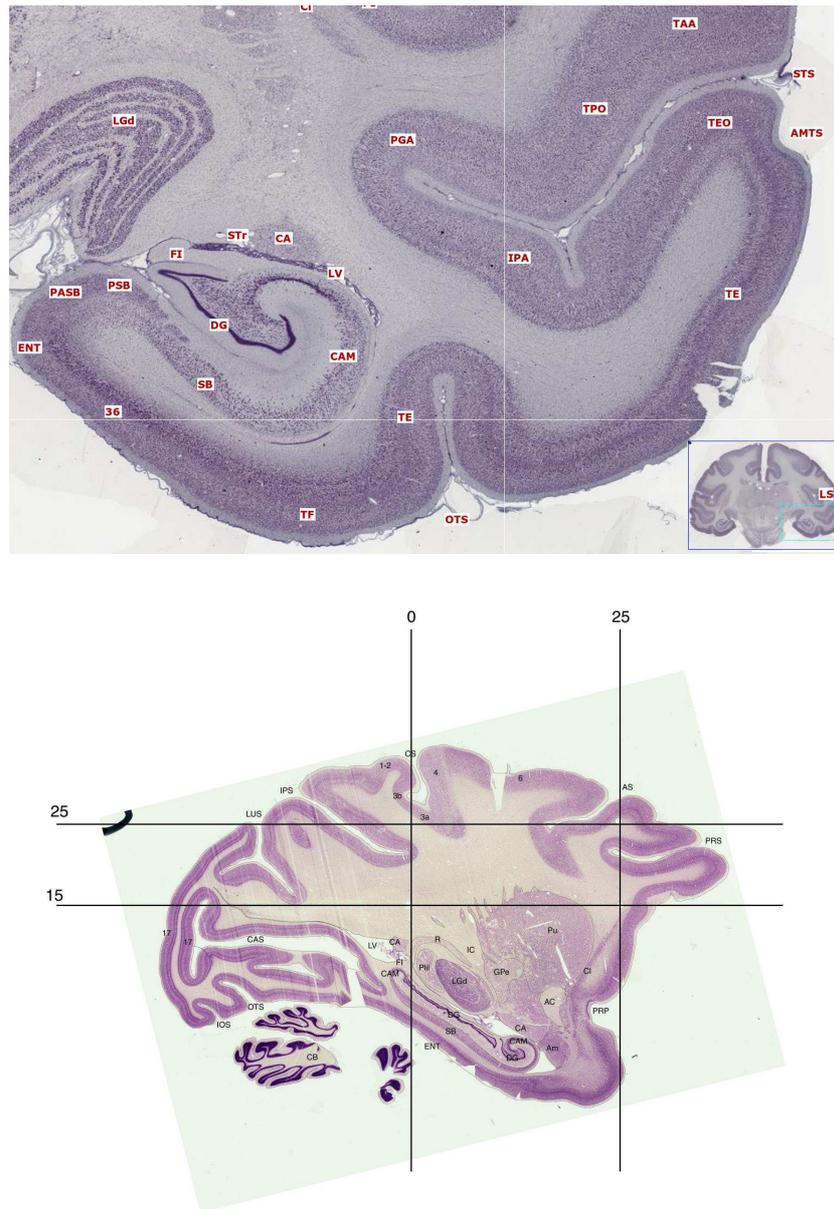


Figure 1.2. From Neocortex to Allocortex in the Macaque Brain. Nissl staining of coronal (left) and sagittal sections (right) of the macaque brain. The six-layer cortex goes over to a three layered architecture in the hippocampus through a transition zone. **Top.** Coronal section shows right temporal (inset shows location), the area beneath the lateral sulcus (**LS**; inset). From right to left, starting at the occito-temporal sulcus in the middle of the image (**OTS**): temporal neocortex (**TE**), area TF of parahippocampal cortex (**TF**), and cortical area 36 (**36**). Entering the hippocampal formation: entorhinal cortex (**ENT**), parasubiculum (**PASB**), presubiculum (**PSB**), subiculum (**SB**), Cornu Ammonis (**CAM**) and dentate gyrus (**DG**); the latter two conforming the hippocampus proper). Adjacent to the hippocampus: lateral ventricle (**LV**), caudate nucleus (**CA**), stria terminalis (**STr**) and fimbria (**FI**). Above: dorsal lateral geniculate nucleus (**LGd**). **Bottom.** Sagittal section showing the elongated shape of the hippocampal formation and its posterior location. The cerebellum (**CB**) is the salient caudal structure. Notice again the difference in neuronal layering between neocortex and allocortical or archicortical regions. **Technique and source.** Nissl stainings here and in Figs. 1.3, 1.4 label the whole neuronal population (and glia) but is less apt for morphology than Golgi silver, which penetrates only a few neurons whose morphology can be well studied. Images courtesy of brainmaps.org (Mikula et al., 2007).

at their upper ends by the hippocampal commissure that crosses the midline under the anterior corpus callosum (Fig. 1.3). The axis of the tube is denominated septotemporal because the hippocampus connects with the medial septum at its inner, dorsal end (right panel, Figure 1.3). The radial direction is quite structured and explained in terms of a transverse axis with positions ranging from DG-proximal to DG-distal, and a depth axis orthogonal to it. The structure of the slices is quite similar across the septotemporal axis, which, paired with the fact that the cylinder is quite twisted in three dimensions, leads to coronal, sagittal and horizontal sections resembling warped versions of each other (Figure 1.4).

The degree of self-similarity of parallel slices, especially when taken transverse to the septotemporal axis together with the strong, cohesive longitudinal connectivity has led to hypotheses that, under some regimes, it may support parallel processing. Despite the existence of extensive connectivity along the septotemporal axis with defined spatial patterns of efference and afference, a connectivity which is forcibly lost upon sectioning, acute slices still display several network phenomena seen *in vivo* with remarkable fidelity; a prime example to be exploited in this work is oscillations. Refer to Wu et al. (2005b) for a brief round-up of *in vivo* and *in vitro* literature supporting this so-called lamellar hypothesis.

Slices for experimental work are generally below 0.5 mm in thickness (but see Wu et al., 2005b) and maintained in a heavily gassed solution (usually 95% O₂, 5% CO₂) akin to cerebrospinal fluid (ACSF; artificial cerebrospinal fluid), because profuse oxygenation is critical to tissue survival. The richness of electrophysiological phenomena supported by the hippocampal slice of rodent tissue has spawned a whole field of electrophysiology in the slice, allowing not only studies at the cellular/sub-cellular level (e.g. synapses and plasticity, firing patterns) but also of microcircuits (local connectivity) and mesocircuits (e.g. population activity involving several hippocampal fields). The degree of detail with which cellular and local circuit physiology is accessible in the slice, especially via long intracellular recordings, has facilitated these extraordinary advances.

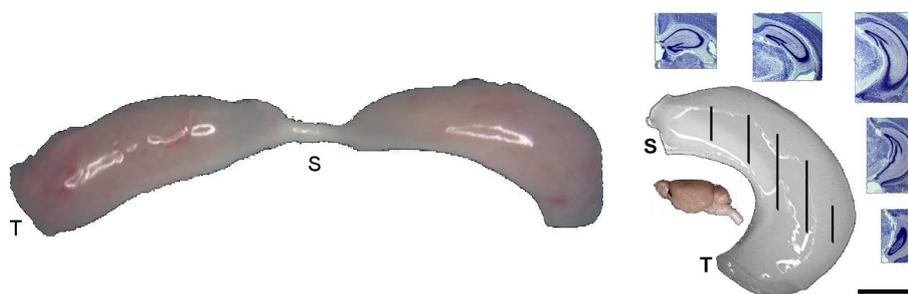


Figure 1.3. The Dissected Mouse Hippocampus. **Left.** Top view of both fresh hippocampi including subiculum (strain DBA 2/J, weight of mouse 21g; weight both hippocampi ~25 mg, whole brain ~400 mg). Hippocampi communicate by an extension of the fornix called the fimbria; the pathway is called the hippocampal commissure. Adapted from the Mouse Brain Gross Anatomy Atlas (Mouse Brain Library project). **Right.** Left hippocampus from the side (miniature of mouse brain inset at left for reference). Nissl stainings in insets at right reflect successive coronal sections along the septotemporal axis **S-T** at locations marked by vertical lines. Calibration is 1 mm (for main image). Adapted from Lu et al. (2001) by permission.

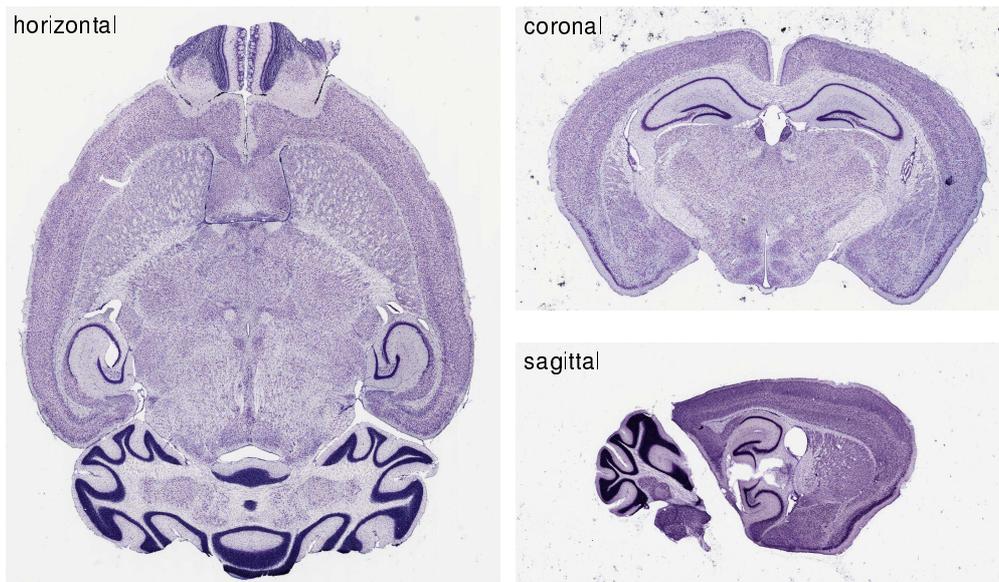


Figure 1.4. Orthogonal Sections of Mouse Brain: Views of the Hippocampus. Horizontal (from above), coronal (from front) and sagittal (from right side) whole-brain Nissl-stained sections. The hippocampus is easily identified as the darkest stained structure after the cerebellum, which is at the back of the brain (not visible in coronal). The dentate gyrus is visible in all three as the smaller, U or V-shaped, densely stained cell strip, interlocking with the Cornu Ammonis, which is larger and of gentler curvature. Lamination is much less marked than in nonhuman primate brains (Fig. 1.2). Images courtesy of brainmaps.org (Mikula et al., 2007).

Septo-temporal patterns of connectivity The wealth of pathways and subtle gradients of connectivity between transverse slices does not lend itself very well to verbal description. Instead, a database where queries can be filtered by position of source and target in 3D (septo–temporal position, proximal–distal location along the transverse axis and superficial–deep across strata, see below) is more appropriate. The effort of compiling that database for the rat hippocampus has recently been carried out and the results presented in tabular form (Strien et al., 2009).^{1,2} In this Thesis we study phenomena that are readily measurable on the slice and thus we focus on longitudinally local circuits. For these two reasons, we direct the interested reader to the publication by Strien et al. for connectivity across the slice planes. We shall now move on to the intraslice connectivity, cell types and catalog of oscillations.

1.3.2 The Hippocampal Transverse Slice: Cells and Pathways

In this Section we describe how both the cells in the slice and their connections are organized according to layers. We mainly followed Amaral and Lavenex (2007), Strien et al. (2009), and Jones and McHugh (2011), together with some of the references therein to prepare a summary that focuses in the main features that are important to understand the work we present in later Chapters.

^{1,2} This database of connectivity is regularly updated; notices are posted at <http://temporal-lobe.com>.

Hippocampal layering and pathways in comparison to cortex The hippocampus is organized differently from cortex, which in most mammals is arranged in modular columns with six layers. Instead, the old cortex in the hippocampus has three visible layers and intrahippocampal connectivity is not reciprocal but rather progressive, mainly embedded in a closed loop from layer II to layers V/VI of entorhinal cortex through the hippocampus and subiculum.

Neurons More than 80% of cells in cortex and hippocampus are classified as excitatory, with a roughly similar number of excitatory synapses. Principal cells show some differences between hippocampal subfields and are rather homogeneous inside each of them. Inhibitory cells, however, show a staggering variety in their location, their afference and efference patterns, and their physiology.

Radial organization in layers Starting in the radial frame from the deep layers we find first what is called in the DG the *hilus* and in Cornu Ammonis the *alveus* and *stratum oriens*. It is populated mainly by incoming and outgoing axons and interneurons, and in CA by the basal dendrites of pyramidal cells. Next towards the surface is the principal cell layer, which also contains some interneurons. In DG this layer is denominated, after its characteristic principal cells, the *granule cell layer* (*stratum granulare*) and correspondingly in Cornu Ammonis and subiculum, *pyramidal layer* (*stratum pyramidale*). The most superficial layer is called *molecular layer* in the DG and in the subiculum. The molecular layer displays some substructure in Cornu Ammonis, depending on the subfield. In CA3 this *stratum moleculare* subdivides into three strata: *lucidum* (layer of afference of the DG), *radiatum* (hosting the apical dendrites of the pyramidal cells just below) and *lacunosum-moleculare*. In CA2 and CA1 stratum lucidum is missing. Region CA2 is small and has been little studied in the rodent literature from a functional point of view; it has neurons similar to CA3 but unlike CA3 it does not receive projections from the DG, it is more resistant to epilepsy-induced cell loss and less suitable for LTP (see Jones and McHugh, 2011 for review).

Longitudinal topography The entorhinal cortex connects its superficial to its deep layers through the hippocampus. The main pathway is comprised of four glutamatergic synaptic relays at principal cells. The three first hops have been dubbed collectively the *trisynaptic loop*, although as such, it is not a closed loop. First, granule cells of the dentate gyrus receive excitation from the medial EC, layer II (1st synapse) through the *perforant path*. They in turn project to CA3 pyramidal cells via a synaptically strong pathway, made of so-called *mossy fibers* (2nd synapse). Some more lateral layer II EC fibers bypass DG and directly impinge on CA3 principal cells. Finally, CA3 pyramidal cells send axon collaterals (*Schaffer collaterals*) to CA1 cells (3rd synapse). These connect back to the deep layer V of entorhinal cortex, directly or through the subiculum. Some connectivity from EC layer III pyramidal cells to CA1 neurons (*temporoammonic pathway*) proceeds in exactly the opposite direction (*backprojections*) to distal superficial dendrites of CA1 pyramidal cells, but does not progress further than CA1. CA3 cells project back to the DG hilus and molecular layer; this pathway is stronger at temporal levels, but in general sparse. There are also reported backprojections from CA1 to CA3 and subiculum to CA1, but they are not as well known as those enumerated above (Strien et al., 2009).

The closing of the entorhinal-entorhinal loop is provided spot-on the EC layer of origin (II) by the projections from CA1-innervated pre and parasubiculum to layers II and III respectively.

Subfield connectivity These cross-field pathways are complemented by in-field connectivity, such as the granule cell excitatory and inhibitory feedback loop with mossy cells and hilar interneurons respectively in the hilus, or the important recurrent, associative connectivity of CA3 pyramidal cells. CA1 field connectivity will be discussed further on.

Extrahippocampal connections: cortex The superficial layers of the EC receive sensory information from the neocortex via the parahippocampal gyrus and or the perirhinal cortex and routes it to the hippocampus. Conversely, the deep layers of EC route information from CA1 back to the neocortex through the same areas. Dorsal hippocampal CA1 cells project directly to retrosplenial and perirhinal cortices (A35 and A36 in the Figure), and receive inputs from the perirhinal and postrhinal cortices in stratum lacunosum moleculare; at the mid septo-temporal level they rather project to the medial prefrontal cortex. The dentate gyrus receives direct input from pre and parasubiculum as well.

Other extrahippocampal targets Besides this tight embedding in a loop with the neocortex via the entorhinal cortex, there are a few other extrahippocampal pathways of note. CA3 projects through the commissure to the contralateral hippocampus (CA1 does as well, but much more sparsely). These commissural fibers are far more relevant in rodents than in primates. CA1, CA3 and subiculum project to the lateral septum; the CA1 and CA3 collaterals are routed through the fornix. The ventral hippocampus projects to the olfactory bulb, the anterior olfactory nucleus, and to nuclei in the amygdala and striatum.

Other extrahippocampal inputs Amygdala, hypothalamus, thalamus, brainstem and basal forebrain all project to the hippocampus. Brainstem inputs of special importance are the dopaminergic fibers from the ventral tegmental area to CA1, because they can modulate synaptic transmission. The basal forebrain and brainstem inputs may modulate the main intrahippocampal excitatory pathway according to sleep/wake cycles or novelty. Acetylcholine (ACh) from septal nuclei, in particular, has been associated with modulation of the CA3→CA1→EC→neocortex pathway (including CA3 intrinsic connections). Reduced ACh levels during quiet states, including sleep, would reinforce auto-associative activity and output to the cortex, thus providing a neuromodulatory substrate to the consolidation phase of the two-stage model of memory consolidation (Hasselmo, 1999).

1.3.3 The CA1 Area of Hippocampus

We have seen above that CA1 is at the epicenter of the hippocampus as a connectivity hub, receiving input directly and through DG and CA3 from the entorhinal layers that gather information from the cortex (III, II), but also from those that project it back (V, VI) and in direct dialogue also with other cortical and subcortical areas.

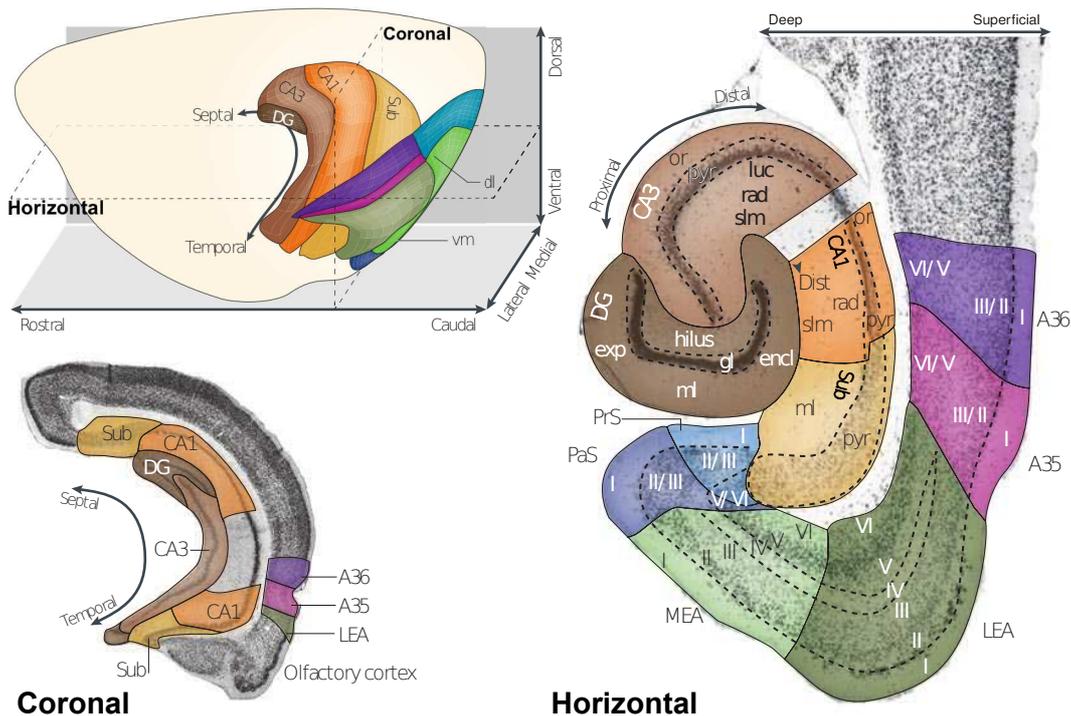


Figure 1.5. Hippocampal Formation and Parahippocampal Cortices of the Rat. Coronal and horizontal caudal cross-sections of the rat's cerebrum are presented with reference to the 3D schematic at top left. **Reference frames.** Schematic shows brain-wide frame (straight arrows). The septotemporal axis is called dorsoventral by some authors. The horizontal section is almost transverse to this axis. The transverse frame is spanned by the proximodistal axis which is parallel to the hippocampal cell layer and starts from the dentate gyrus and, orthogonally to it, the radial or superficial-deep axis. **Fields.** Dentate gyrus (DG; dark brown), CA3 (medium brown), CA2 (not indicated), CA1 (orange), subiculum (Sub; yellow), pre-subiculum (PrS; medium blue), parasubiculum (PaS; dark blue) and entorhinal cortex (medial, MEA in light green and lateral, LEA in dark green) constitute the hippocampal formation. The entorhinal cortex is described by dorsolateral to ventromedial (dl, vm in schematic respectively) and rostrocaudal axes. The perirhinal cortex consists of areas A35 (pink), A36 (purple) and the postrhinal cortex (blue-green; only visible in schematic). **Strata.** In the horizontal section cortical layers are numbered I-VI from superficial to deep. In hippocampus, gl stands for granule cell layer; luc for stratum lucidum; ml for molecular layer; or, stratum oriens; pyr, pyramidal cell layer; rad, stratum radiatum; slm, stratum lacunosum-moleculare. Adapted from Strien et al. (2009) with permission.

The measurements from CA1 in linear track replay experiments *in vivo* (see brief account in Section 1.2.1) also contributed to our fascination for the area: pyramidal cells there seem to code for space, and the electrical activity measured in CA1 shows a variety of oscillations in different time scales which are clearly correlated with behavioral state and undoubtedly play a role in setting the stage for the coding of episodic memory. We find reasons thus to examine in more detail the inhabitants of CA1: pyramidal cells and a rich zoo of interneurons, highlighting those aspects of their morphology and connectivity that will be most relevant later.

The pyramidal cell layer is packed in the rat with 4-5 cell bodies and includes some basket cells as well. Cell body counts estimate a population of a quarter to half a million cells in this layer of the rat hippocampus (West et al., 1991), slightly above

the corresponding numbers for CA3. In contrast to CA3, CA1 has recurrent loops of restricted septotemporal scope. This is a precious anatomical observation for any slice model in which observations could be explained by recurrence, as the case is ours. These recurrent connections are made in stratum oriens, sometimes *en passant* towards the subiculum (Deuchars and Thomson, 1996).

Schaffer collaterals The CA3 intrinsic output pathway synapses strongly within CA3 itself before targeting CA1 both at the level of the stratum radiatum and oriens. Schaffer collaterals are the strongest input to CA1 and have been extensively investigated, not least because the right electrical stimulation of the fiber bundle can produce long-term potentiation (LTP). Through this pathway a single CA1 neuron may receive thousands of contacts from CA3 cells. The topography of the connection follows the principle of reverse distance: distal CA3 projects to proximal CA1 and proximal CA3 to distal CA1 (to further complicate things, there are differences as well in the synaptic target layers according to proximodistal location in CA1).

Temporoammonic pathway Axons from the entorhinal cortex impinge on the distal apical dendrites of pyramidal cells.

Subiculum and EC The projections from CA1 to the subiculum follow again the same reverse topographic pattern as Schaffer collaterals from CA3 to CA1, and they respect the seemingly columnar organization of the target region.

Interneurons Interneurons have generally a confined axonal scope and use γ -aminobutyric acid (GABA), hence their longer denomination: local inhibitory interneurons. The majority of them lack dendritic spines. A testimony to the likely functional importance of the otherwise relatively few interneurons is how their dendritic arborizations and axonal plexa are clearly confined to certain strata, to the point of being a definitory geometric feature as important as their genotype. One speaks for example of bistratified cells because their axons arborize in the two strata above and below stratum pyramidale, or of oriens-lacunosum moleculare (O-LM) cells, because their dendrites are in oriens along with the cellular body, while their axon makes contacts only in lacunosum-moleculare. The study of interneuronal firing patterns in vivo is relatively recent (Klausberger et al., 2003). Since interneuronal firing patterns are often as specific to behavioral states as their projections are to layers, and since specific interneuron classes have privileged control over the spiking output of pyramidal cells because of their afference sites at the soma or directly on the axon, the extension of the present knowledge from just a few units in the anesthetized condition to populations in the behaving animal will surely shed abundant light on hippocampal function.

A final important feature of the interneuronal population at large is that there is abundant recurrent connectivity, not only by chemical synapses but also through gap junctions; this has led even to rescue Golgi's *syncytium* as a characterization of the resulting networks (Amitai et al., 2002). This connectivity is highly interneuron-class specific.

Note that CA1 was initially named *regio superior* in the terminology of Ramón y Cajal and later Blackstad (CA3, correspondingly, was *regio inferior*); we have seen this terminology employed in the literature produced well into the 1970s.

1.3.4 The CA1 Pyramidal Cell

In this Thesis we use pyramidal cells as a telescope to peek into network activity. We do so by listening to the synaptic inputs they receive, both from interneurons and from other pyramidal neurons. Here we enumerate some properties of CA1 pyramidal cells that are relevant for this observer role.

Pyramidal cells are the dominant excitatory cell type in the mammalian cortex. Across structures in the forebrain such as hippocampus, amygdala or cerebral cortex they share an asymmetric, pyramidal-shaped soma and two distinct dendritic arborizations (see Fig. 1 in Spruston, 2008 for a morphological comparison of neocortical vs. hippocampal CA1 and CA3 pyramidal cells and Fig. 1.6 below for a CA1 example). Towards the base of the pyramid, the basal dendritic tuft branches off the several dendrites that emanate from the soma. The apical dendritic tuft is built as ramifications of one single apical dendrite connected to the soma, with sometimes a second principal branch emanating from this principal dendrite close to the soma. The basal tuft is more compact than the apical one, covering perhaps 1/4 to 1/3 of the extension of the cell away from the soma.

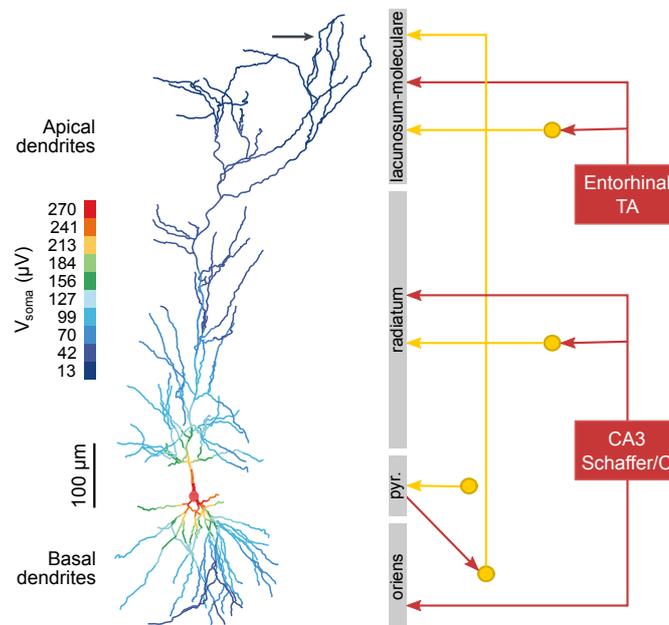


Figure 1.6. Passive Integration of Compartment-Specific Inputs to CA1 Pyramidal Cell.

Cell reconstruction and schematic of afference patterns. The axon (not shown) emerges from the basal end of the soma and projects through the alveus chiefly to subiculum and secondarily to parasubiculum and entorhinal layer V/VI and through the fimbria to a variety of extrahippocampal targets depending on the septotemporal location of the cell. **Left.** The voltage seen at the soma is predicted by cable theory (Rall, 1967) to strongly attenuate according to where the conductance is injected into the cell (here a factor 20 between distal inputs and somatic inputs). Color overlay on the morphology of a CA1 cell represents the somatic voltage change that would be induced by an example conductance of 0.3 nS at each potential synaptic site of injection. **Right.** Schematic of main excitatory input from nearby CA3 and entorhinal cortex showing some of the local interneuronal circuits. Interneurons are represented by a yellow disc positioned at the stratum where their soma is located. TA: temporoammonic pathway from layer III of the entorhinal cortex; Schaffer/C: inputs from ipsilateral CA3 (Schaffer collaterals) and contralateral CA3 (commissural afferents). Attenuation heat map adapted by permission from Spruston, 2008.

Cell bodies of CA1 pyramidal cells are smaller than those of CA3 cells. CA1 pyramids have cortical-like spines on which asymmetric synapses of excitatory character are formed, but they do not show the massive thorny excrescences characteristic of the DG afference onto CA3 pyramidal cells. Inhibitory input forms symmetric synapses on the shafts. This anatomical difference helps to quantify the compartmental preference of different types of inputs. The stratification of inputs is as follows: excitation is predominantly distributed on the dendritic tree while inhibition via GABA targets also the soma and axon; CA3 input is concentrated on the near apical and basal dendrites whereas the distal apical dendrites receive input from the entorhinal cortex and the thalamus in anatomically almost equal proportion. Thus near-somatic input is more local (including interneurons and recurrence) and distal input is more extrinsic. Due to the considerable electrotonic length of the cell, a strong coalition of distal inputs would be needed to match more proximal synaptic afference of equivalent intensity at the source. The question of how diverse sources of input contribute towards the output is still open, not least because a non-homogeneous distribution of conductances (active or passive) could compensate for distance (“synaptic conductance scaling”, Magee and Cook, 2000), and because integration could be behavioral-state dependent. However, any mechanism of compensation based on active conductances will be largely absent in the voltage clamp recording configuration, where the cell voltage is held constant.

The principal excitatory neurotransmitter is glutamate, and synapses contain variable amounts of NMDA, AMPA and metabotropic receptors. Of consequence for our theoretical model below, synapses have been found that express NMDA receptors but no functional AMPA receptors. These have been dubbed *silent* synapses. Their repeated activation may lead to the installation of AMPARs, one of the proposed cellular mechanisms for long term potentiation (LTP). Both AMPA and GABA-A synapses have fast kinetics, with current tail-offs estimated at around 2 and 4 ms respectively.

An extensive review of pyramidal cells concentrated on CA1 and cortical neurons in rat was published recently by Spruston (2008). We refer the reader to that publication for a discussion of plasticity.

1.4 Physiology of Sharp-Wave Ripples

Having sketched the most salient anatomical features of the hippocampus, we change our bearings and look at the physiological pattern at the center of our interest: sharp-wave ripples. Before diving into their electrographic signature and the associated unit activity, we devote a few paragraphs to the singular characteristics of hippocampal LFPs, and to other hippocampal patterns of activity that are apparent in them beside ripples.

1.4.1 The Interpretation and Content of Hippocampal LFP

We have seen in the previous Section that the local hippocampal tissue is a) laminated in the radial direction, both for the location of cell bodies and their connections, b) self-similar in the transverse or septotemporal direction and c) self-similar along the

longitudinal or proximodistal direction. This frame is locally orthogonal and defines two spatial axes along which the layout of neural hardware varies little (transverse and longitudinal) and one in which it varies in a structured way (radial). The main consequence is that, in the hippocampus, one can obtain information about the activities of neural populations from local field potentials (LFPs).

LFPs reflect synaptic currents Extracellular electrodes will record similar signals when displaced in the two homogeneous axes and changing signals when recording in the radial direction, the change in current polarity reflecting the new relative positions of current sources and sinks as the electrode goes deeper in the radial direction. The fast components (>500 Hz) will betray the charge sinks and sources due to spiking activity, whereas the slow components reflect well the underlying synaptic currents (Johnston and Wu, 1994). Most synapses being associated with pyramidal cells, and these having to their advantage when “voting” for the LFP that they are distributed neatly in a compact layer, it is rather likely that the LFP is heavily influenced by synaptic currents at the pyramidal dendrites (be they of inhibitory or excitatory nature, see Ellender and Paulsen, 2010). However, the attenuation scheme applicable for cable conduction (Fig. 1.6) is not the same as for the volumetric current displacements intervening in the formation of LFP signals, whereby the relation of LFPs with e.g. currents observed at the soma is ill-specified as both weight differently distal contributions (see Fig. 1.7).

The laminated structure of the hippocampus along its radial axis permits investigation of the radial structure of currents by a technique known as current-source density analysis (CSD). CSD profiles are produced by obtaining derivatives of the measured potential along the radial axis with the additional assumptions that the surrounding medium is resistive, uniform and isotropic. CSDs require multiple measurement points along the superficial-deep axis and thus is most frequently performed with shank electrodes. Another, closely related application of radial voltage profiles is to find LFP “generators” by applying independent component analysis (Makarova, 2011).

Oscillations are visible in the hippocampus The open geometry of the pyramidal cell layer and the stratification of inputs around it prevent cancellations that would result from random or spherically symmetric configurations; if neurons engage in coordinated activity, this becomes readily visible extracellularly. Not only that: the LFP is well-suited for observing fast-paced oscillations, since it roughly reflects synaptic currents and not membrane potentials. Membrane potentials are slowed down with respect to the currents that generate them by the membrane leak current. With a characteristic time constant of e.g. 15 ms, fast synaptic phenomena on the scale of e.g. 5 ms will be seen smeared when measuring transmembrane potentials. The LFP will reveal them, amidst other signals. The importance of LFPs in the assessment of neural oscillations is thus undergirded not only by historical but also by structural reasons.

Oscillations in several frequency bands The hippocampus supports behaviorally-dependent oscillations in several frequency ranges (Buzsáki and Draguhn, 2004). In fact, with the exception of the intervals between ripple oscillations during large irregular activity (LIA) epochs, and the transitional and uncommon slow irregular activity (SIA), the hippocampus is always oscillating, and often displays nested oscillations. Table 1.1 summarizes the most important LFP patterns and how they combine.

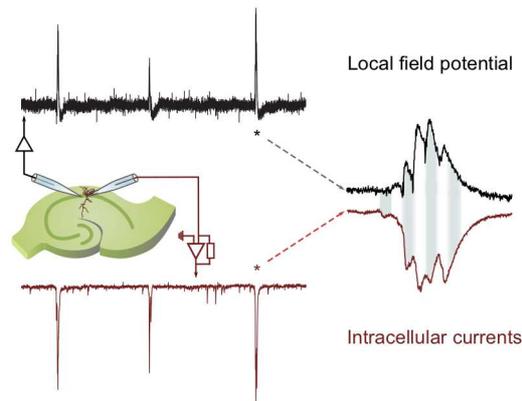


Figure 1.7. LFP Reflects Intracellular Currents. A simultaneous measurement from the dataset used in this Thesis evinces the parallelism between the local field potential recorded by an extracellular electrode (black) and the intracellular currents (dark red), in this case acquired at the reversal potential of inhibition in the voltage clamp configuration. The coordination is particularly striking in this example of a strong oscillatory event (sharp-wave ripple; see below); not every ripple from every recorded pyramidal cell will sample the overall synaptic activity detected by the LFP electrode with such fidelity.

f. band (Hz)	name (abbr.)		LIA	δ	θ	γ	RPL
—	large irregular activity (LIA)		LIA				
1.5–4	Delta (δ)		δ				
4–10	Theta (θ)		θ	×			
30–80	Gamma (γ)		γ	×	×		
120–200	Ripples / fast oscillation (RPL)		RPL	×	×		

Table 1.1. Hippocampal Rhythmical and Nonrhythmical LFP patterns in CA. LFP patterns have characteristic frequency bands and may appear embedded one in another. **Left.** Denominations and typical frequency bands of hippocampal CA LFP patterns. **Right.** Compatibility matrix of LFP patterns. A marked entry (×) represents the possibility of finding instances of the row pattern embedded in the column pattern. The table is composed for mouse and rat in vivo after O’Keefe (2007), complemented by observations of ripples during exploration (O’Neill et al., 2006). We have omitted the slow irregular activity pattern, that is associated to behavioral transitions and less well investigated, as well as the beta oscillation (10-20 Hz) that appears in EC or DG but not CA.

Plasticity linked to theta and gamma A hint to the functional importance of hippocampal oscillations is provided by the in vitro investigation of plasticity with stimulation protocols: the modulation of theta activity, with its phase dynamics (“theta precession”) has been suggested to be the in vivo nearest equivalent of the tetanic stimulation that generates LTP. For its part, gamma oscillations, which are believed to be orchestrated by gabaergic inhibitory “hub” neurons, render pyramidal firing most effective in the ~ 10 ms windows where inhibition is at its lowest, thereby setting the stage for another form of plasticity —spike time-dependent plasticity (STDP)—, which requires presynaptic to be followed by postsynaptic activation in 10-20 ms. By this, gamma would bind populations into assemblies.

Both gamma and theta have been dissected into at least two types — gamma, according to frequency into fast and slow gamma, and theta according to its pharmacological resistance to cholinergic antagonists, into atropine-sensitive and insensitive forms. The investigation of the functional correlates of these subtypes is still ongoing. For our part, we shall not dwell on the behavioral-state dependence of the diverse rhythms, their function, or their mechanisms. Suffice it to add, in preparation for the next Section, that theta is the rhythm present during exploration while spatial memories are being built. Indeed, the firing patterns of cells during theta can effectively be used to train a Bayesian probabilistic model of the animal's position that can subsequently be used to recover position it from neural activity alone (Davidson et al., 2009). Such “theta firing patterns” re-emerge later during the faster ripple oscillation, to which we now turn.

1.4.2 Sharp-Wave Ripples in the Local Field Potential

During about one tenth of a second, between 50,000 and 100,000 hippocampal neurons (a 10-20% of the total in the rat hippocampus) organize an intricate neuronal dance propagating from CA3 through CA1 to the deep entorhinal layers and out of the hippocampus; a burst that qualifies as the most synchronous pattern of activity in the brain (Chrobak and Buzsáki, 1994) and is well visible in the LFP signal. In area CA1 this pattern is expressed as a violent high-frequency oscillation (~ 200 Hz) in the pyramidal layer, dubbed ripple, that coexists with large-amplitude LFP deflections, conforming “sharp-wave ripple complexes” (SWR, Buzsáki, 1986; Buzsáki et al., 1992, Ylinen et al., 1995). In this Section we present some of the accumulated knowledge about SWR: where, when, for how long and how many times they happen, how they are affected by common drugs, how single-neuron activity correlates to the ripple oscillations and, finally, how SWR may ultimately relate to performance in memory tasks.

1.4.2.1 Behavioral Correlates of SWR

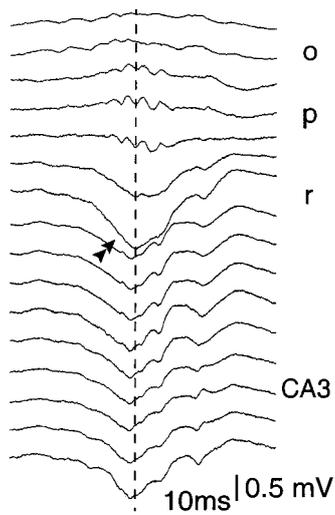
SWR have been studied predominantly during the deep phase of sleep (slow-wave sleep, SWS) and awake immobility (O'Keefe and Nadel, 1978; O'Keefe, 2007). In recent years, a phenomenon with much the same LFP signature and pattern of unit activation has been reported to occur embedded in theta oscillations during movement or brief pauses (< 2.4 s, O'Neill et al., 2006). These movement-related SWR have been named “exploratory SWR” (eSWR) to distinguish them from the original immobility SWR (iSWR) and sleep SWR (sSWR).

1.4.2.2 Spatiotemporal Expression

The following notes refer to the *in vivo* situation in the rat, except where specified.

Duration and recurrence SWR occur irregularly at 0.01 to 3 Hz and last between 30 and 120 ms (Buzsáki, 1986). There is a tendency for SWR events to bunch together in 100-ms separated doublets (Buzsáki et al., 2003; mouse) or even longer chains reaching up to 700 ms (Davidson et al., 2009).

The electrographic signature of SWR is dependent on the radial coordinate The denomination “sharp-wave ripple” reflects the concurrent expression of two patterns visible in the local field potential (LFP) at different radial levels (layers):



SWR in CA1 of the rat during SWS
(Csicsvari et al., 1999a, with permission).

sharp wave. a biphasic deflection with a first, larger, negative phase followed by a positive upstroke. The name is due to the sharp peak that stands out of the noisy baseline when observed at low time resolution; in fact the hippocampal pattern described as large irregular activity (LIA) is characterized by irregular activity of low amplitude punctuated by large sharp waves. Sharp waves are most prominent in stratum radiatum, the apical dendritic layer (r in the Figure), and reverse polarity slightly above pyramidale (p), becoming positive in oriens (o).

ripple. a fast oscillation of 120 to 200 Hz frequency, whose amplitude is modulated by the sharp wave. and is most prominent in stratum pyramidale.

The SWR depth profile is so characteristic, that it was used as one of the landmarks to ascertain that eSWRs had indeed the physiological features of sleep/immobility SWRs and were not a different kind of network phenomenon (O'Neill et al., 2006, Fig. 2).

Proximo-distal axis: SWR involve the whole hippocampo-subicular region SWRs are measured from CA3 through CA1 up to the subiculum and the entorhinal cortex, where they show a lower ripple frequency and appear with a time lag. (Chrobak and Buzsáki, 1994; Chrobak and Buzsáki, 1996). It is interesting and important for systems consolidation theories that SWRs propagate to the output structures of the hippocampus but never invade the input, superficial layers of EC (II/III).

Inside the hippocampus, while a sizable CA3 burst predates the CA1 phenomenon, the fine structure of both oscillations is not phase-coherent. Any site in CA1 can be an emergence site; propagation is towards the output structures (Csicsvari et al., 2000) and preserves coherence along the longitudinal axis of CA1, at least for the most powerful ripples.

Bilateral dimension of SWR SWRs are present simultaneously in both hippocampi, but they are not phase-coherent between e.g. the ipsi- and contralateral CA1 (Buzsáki, 1996; rat; Buzsáki et al., 2003; mouse).

1.4.2.3 Unit Activity

The subthreshold potential of pyramidal cells is affected in an oscillatory manner by the ripple. On this topic we shall have much to say in further Chapters; let us therefore direct our attention now to the spiking output of pyramidal cells and interneurons.

Firing rates Both pyramidal cells and interneurons elevate their firing rates during SWR. The individual unit firing rates vary considerably (see scatterplot in Csicsvari et al., 1999b, Fig. 4) but, in average, pyramidal cell rates are boosted by a factor six or seven to about 7 Hz and inhibitory interneuron rates ramp up by a

factor three or four to about 40 Hz. Probabilities of unit participation in any one SWR are at 5% for pyramids and about 20% for interneurons. More powerful ripples recruit additional CA1 pyramidal cells more effectively than interneurons (Csicsvari et al., 1999a).

Multi-unit synchrony and field coherence Another observation linking unit activity and ripple is that stronger synchrony of multi-unit activity associates with more widespread spatial coherence (Csicsvari et al., 2000). The same authors report that unit activity correlates better to the local field at $\leq 300 \mu\text{m}$. In a previous paper Csicsvari et al. (1999b) showed that summed interneuronal activity correlated similarly well to the wave shapes of theta and ripples. Pyramidal cell activity correlated similarly well to ripple wave shapes, but considerably worse to theta.

Locking of unit activity to ripples The high individual rates of interneurons, that sometimes reach up to 200Hz during ripples (Ylinen et al., 1995) allow to obtain statistics of locking to LFP ripple at the single-unit level. Single interneuronal units fire phase-locked to ripple undulations. By comparison, pyramidal spikes are scarce (often just one per ripple). But when pooling them across units, one observes strong phase modulation as well. Concretely, the 'contrast' of the phase modulation reaches a value of 3 both for principal cells and interneurons in the pyramidal layer. Interneurons in alveus/oriens are even better phase locked, with modulation contrast values close to 5. These levels roughly double those obtained from analyzing spike-to-theta-rhythm coupling (Csicsvari et al., 1999b).

In-ripple phases The earlier impression from extracellular recordings that subsets of interneurons (singled out by soma location) show different ripple-associated firing patterns was confirmed in studies combining juxtacellular recordings with intracellular labeling (Klausberger et al., 2003). Ripples were accompanied by increased spiking of some basket cells as well as bistratified and trilaminar interneurons. Firing probability was reduced in oriens-lacunosum moleculare and axo-axonic cells. These results are based on low cell counts and anesthetized animals; ongoing experiments in the awake behaving rat will likely sharpen the picture and highlight some differences.

1.4.2.4 In Vivo vs. In Vitro Ripples

SWRs in the slice are more regular and uniform in size than SWRs in vivo. Additionally, the activity in CA3 concurrent to the CA1 ripple is different in both settings. While in vivo CA3 seems to send a multitude of small amplitude excitatory inputs from distributed sources, in vitro the activity in CA3 is larger and ripple-patterned. This is perhaps a consequence of the fact that slicing severs the transversally extended recurrent axonal plexa of CA3 cells.

1.4.2.5 Initiation and the Role of CA3

CA3 oscillations concurrent with CA1 ripples are generally of a lower frequency (80-140Hz); they are often classified as "fast gamma" and the reason of these two different oscillating modes is attributed to the different resonance properties of both areas, when considered as voltage-controlled oscillators (Sullivan et al., 2011). Poor locking of units to CA3 oscillations in vivo (Ylinen et al., 1995) seems to indicate that the CA1 rhythm emerges in response to CA3 activity but does not carry over the details of activity there (but see Both et al., 2008 for contradicting evidence from an in vitro

stimulation paradigm). The specific currents generating the CA1 oscillation are still a subject of debate. Currents generating SWR in CA1 are at the center of interest of this work. At present it is widely believed that a volley of unspecific CA3 Schaffer collateral input generates the CA1 sharp wave which is shortly accompanied by ripples mainly engendered by the local inhibitory populations. This view is supported by the existence of interneurons with tight inhibitory control over the pyramidal cell soma, such as parvalbumin- and cholecystokinin-positive interneurons, and their high firing rates with strong ripple phase modulation.

An *in vitro* stimulation paradigm (Behrens et al., 2005) has shown that LTP-inducing protocols can enhance SWR expression *in vitro*. The induction of SWR turned out to be NMDA-receptor dependent. However SWRs that were induced in presence of NMDA transmission were expressed in its absence just fine. Low-frequency stimulation, a protocol that is known to induce long-term depression also reduced SWR expression. Importantly, in this experiment the sequence of excitation and inhibition was maintained during the development of SWR in individual neurons; this is in support that behind the generation of SWR are stable distributions of synaptic weights.

1.4.2.6 Pharmacological Sensitivity of Ripples

Ripples are suppressed by several gap-junction blockers (Ylinen et al. 1995; Maier et al., 2003; Behrens et al., 2005). These so-called uncoupling agents are rather unspecific drugs. Some of them facilitate the disruption of the ripple without affecting the sharp wave^{1,3} (it is concretely the case of halothane *in vivo*). These observations and modeling studies, together with anatomical work have spawned the idea that the requisite synchronization for sharp waves could be established by the networks defined by the axons of pyramidal neurons, connected by gap junctions. The relative role and sometimes the very existence of this form of synchronization, relative to synaptic coordination or even ectopic influences remains under discussion in the community

Pharmacological blockade of slow receptors NMDA and GABA_B failed to affect SWRs (Maier et al., 2003); but impairment of fast excitation with the AMPA-receptor antagonist CNQX extinguished them completely (Maier et al., 2003; Behrens et al., 2005; Wu et al. 2005a). In rats anesthetized with ketamine, Ylinen et al. (1995) observed a reduction in ripple frequency. The principal effect of ketamine (besides affecting NMDA receptors) is to increase GABA_A currents thus tilting the balance in favor of inhibition.

Maier et al. (2003) applied gabazine and bicuculline to remove GABA_AR-mediated inhibition in the slice; this caused the emergence of epileptiform bursts that still contained a 200 Hz component. To avoid those bursts, Nimmrich et al. (2005) tested *in vitro* the effect of gabazine on CA1 tissue alone (a “minislice”; the bursts stay with the CA3 remnants). They observed that in this disinhibited preparations SWR oscillations persisted if the dendrites were excited by puffs of KCl. The authors interpret these results as establishing the need for the tonic components of excitation and inhibition to initiate the ripples, but by elimination they attribute the fast-frequency component to gap-junctional communication.

1.3. Differential disruption of sharp wave and ripple is an effect that has been only recently reported for synaptic transmission by the use of the GABA_AR agonists zolpidem and diazepam (see Koniaris et al., 2011 for details on the dose-dependent effects).

What do we make out of this confusing, sometimes contradicting evidence? Both application of pharmacological agents and the use of anaesthetics *in vivo* have shown that slow synaptic communication is unessential for ripple expression. The extent of the involvement of fast excitation, fast inhibition and gap junctions is much less clear, and probably all three have a role in at least one of the different stages of triggering the sharp wave, eliciting the ripple embedded in it and terminating the phenomenon without it becoming an epileptic burst.

1.4.2.7 Functional Correlates of SWR

The widely accepted *consolidation* theory about the functional significance of SWR is based upon

- i. the correlations of unit activity observed during ripple oscillations with fore-going or forthcoming activity patterns of spatial nature,
- ii. the emerging link between SWR and performance, and
- iii. the powerful expression of SWR along the hippocampal exit pathway, together with the correlation of ripple oscillations with well-defined brain states in other parts of the brain.

These pieces of evidence combined undergird the idea that SWR contain episodic information that, by repetition, is solidified or consolidated in hippocampal networks and eventually learned by target cortical networks as well. The temporal details of this systemic consolidation are unclear, although it seems that the hippocampal “version” of the memory is less resistant over long time spans^{1,4}. Let us briefly revisit these three mainstays of the consolidation theory.

The sequence replay literature, of which the experiment by portrayed in Figure 1.1 is an example, has demonstrated a link between place cell activity during exploration and sharp-wave content. The observation of episodes of remote replay, whereby the animal reenacts trajectories in other sections of the environment, or other environments altogether, can have mixed interpretations. On the one hand, remote replay of local trajectories may be important for decision-making and underlines the active memory role of SWR. On the other hand, replay of remote trajectories in other environments points at SWR subserving a memory maintenance (memory-passive) role, i.e. as a framework for memories to be strengthened during sleep and waking. To dissect which role is predominant, perhaps in a behaviorally dependent fashion, more experiments will be necessary.

These roles of SWR have begun to be investigated by assessing performance in behavioral tasks. Rats learning three hidden positions in search for food found those goals more easily after sleep in proportion to how much the patterns of activity coding for the goals had been reactivated during sleep ripples (Dupret et al., 2010, Fig. 5c). Another line of investigation is direct disruption (using electrical stimulation) of SWRs between learning and test in a memory task (Ego-Stengel and Wilson, 2010; Girardeau et al., 2009). Probing the significance of eSWR *in vivo* will be the next step

1.4. Recently, though, one group observed that recall of remote contextual fear memories can be dependent on hippocampal CA1 indefinitely. According to Goshen et al. (2011), previously reported detachment of recall capability from hippocampus (via lesions or pharmacological inhibition) may be actually a compensation phenomenon that requires some time to develop.

and likely involve the use of fast stimulation such as offered by optogenetic addressing of specific neuronal populations.

Finally, neural activity in other parts of the brain such as the prefrontal cortex, V1 in the visual cortex, or ventral striatum has been shown to correlate to the expression of hippocampal SWR, in particular slow oscillations in the cortex (see Buhry et al., 2011 and Battaglia et al., 2011 for references and review). This evidence builds upon the generally acknowledged role of the hippocampus in the establishment of contextual memories and as a non permanent store to support the general idea that the electrographic correlate of hippocampal-driven systems consolidation is the SWR.

1.5 Research Goals Revisited

At the beginning we presented our interest in neuronal assemblies of coactive neurons as the physical implementation of the memory trace. We would like to pursue that topic at two levels of investigation, where we have identified respective knowledge gaps. First, we wish to know if these assemblies exist. For that we need to study neural activity potentially representative of the expression of memory traces. We need *data* to inform our mental models. In view of the now established sequential reactivation of place-coding cells during SWR, these oscillations provide an excellent framework to address the more general question of assembly expression. SWR are a very fast phenomenon, and thus we require exquisite time resolution for their investigation: we need electrophysiological recordings. Multiunit extracellular recordings offer one avenue to investigate the problem, and indeed they have provided strong hints already as to the existence of assemblies and how they might operate in time. Another route, intracellular recordings of currents, has the advantage of fixing the readout: we want to see assemblies in action from the point of view of the unit of coding, the pyramidal cell. Along the way we will look very closely at the network states that facilitate the timely expression of coordinated activity: we wish to know how inhibition and excitation play the tune of memory at the millisecond scale — gleaning this information from intracellular recordings, we aim at improving the state of knowledge about the synaptic basis of sharp-wave ripples.

Second, the underpinnings of sequence memory must be better understood from a functional perspective. We need a *model* to inform further acquisition of data. How many memories can be resident in a neural network so that the storylines in which they are embedded replay successfully end-to-end? What is the optimal size for coactive assemblies from this theoretical, functional perspective? Is there a role for inhibitory activity? how would it develop in time? What factors influence the overall memory capacity for sequences of a neuronal network of excitatory and inhibitory neurons? We mean to answer these questions by incorporating as much as we learn from the data collected for our first endeavor: the model must be biologically constrained, yet tractable so that it can formulate predictions that enrich our hypothesis-making for further data collection.

In summary, the objective of this Thesis is to better understand how the brain stores and retrieves memories, and that at two levels — first, how it should do, provided with a minimum repertoire of expression capabilities, in order to maximize its capacity and, second, how does it organize the physiology of its cells as building blocks to effectively make it possible.

Part II

Materials

Chapter 2

Experimental Approach to Study SWR

In this Chapter we describe the electrophysiological recordings acquired to investigate the synaptic basis of sharp wave-ripples (SWR). All experiments reported here and most of the analyses have been done by collaborators in our publication *Coherent Phasic Excitation during Hippocampal Ripple* (Maier, Tejero-Cantero, Dornn, Winterer, Beed, Morris, Kempter, Poulet, Leibold, and Schmitz, 2011). The figures here have been edited (sometimes substantially) with respect to the paper for clarity of exposition in a longer format.

Our general approach to observe the activity of the network during hippocampal Sharp Wave ripples was to record simultaneously from two electrodes. One electrode captured the variations of the local field potential (LFP), which are believed to mainly reflect local synaptic activity. A second electrode sampled in whole-cell patch either the currents incoming into a nearby cell (voltage clamp configuration) or the variations in its membrane potential (current clamp).

The field activity measured with the extracellular electrode can neither be traced back to the individual synapses responsible for it, nor to the neurons which they connect. By comparison, the target cell of the synaptic activity revealed by the intracellular electrode is known, although we still ignore the identity of the afferents that induce it. This intracellular electrode can operate in two modes. In voltage clamp we do not know if the measured current is enough to fire the cell, whereas in current clamp the spiking is apparent but little can be said about the excitatory or inhibitory composition of the input, and thus of the types of cells that are active. In addition, in the more physiological current-clamp configuration synaptic input is convolved with the leak conductance kernel, which has a time constant of dozens of milliseconds and thus smears faster details (compare Fig. 2.9 to Fig. 2.10 below). In contrast, the voltage clamp configuration allows to follow the very fast currents present during the extracellular SWR, which is why we use it. On occasion, we added a second intracellular or extracellular electrode to determine the spatial extent of the oscillatory synchronization.

We start the Chapter by providing qualitative evidence of the existence of fast excitatory currents incoming on to CA1 pyramidal cells in vivo in the awake mouse. The first **Section (2.1)** shows the features of the local field potential and whole-cell recordings from a head-fixed mouse compared against observations from the literature.

In order to better quantify the phenomena observed in vivo, we move in the next Section to an in vitro slice model of SWR (Maier et al., 2009) that allows optically tar-

geted cell patches, fast pharmacology wash-in and low-noise recordings with a stable baseline. We present the chief features of the experimental model in **Section 2.2**. To disentangle excitatory from inhibitory currents we employed electrotonic isolation. The procedure and a few recordings of currents during SWR are discussed in **Section 2.3**, where we also address the question of their spatial extents in the slice .

As advanced in the Introduction, the present state of knowledge about the synaptic basis of ripples demanded that we focus our attention on the *excitatory* synaptic input. To verify that the currents observed under electrotonic isolation of excitation are indeed exclusively excitatory, we additionally blocked chloride channels from inside the cell. We describe the corresponding approaches for voltage and current clamp in **Section 2.4**. The locking of CA1 pyramidal cells discharges to the extracellular firing is substantiated there.

description (all figures in mV)	Solution (Table A.1)	No. cells	Notes	Figures in Chapter 2
				2.1, 2.2 (lfp)
in vivo cc		12	head-fixed	2.3, 2.4, 2.5
in vivo vc ~ -70	KGlu 0	8	$E_{Cl^-} \simeq -70$	2.6
vc -66 KGlu	KGlu 1	8	$E_{Cl^-} \simeq -67.3$	2.9
vc -45 KGlu	KGlu 1	8	recorded 2nd	2.17
vc -6 CsGlu	CsGlu 2	6	$E_{AMPA} \simeq -6.5$ mV	2.18
vc -66 KGlu_{CA1}	KGlu 1	5	CA1 minislice	2.27
vc -60 KGlu	KGlu 1	7	DIDS control	2.19
vc -45 KGlu	KGlu 1	5		2.22
vc -60 DIDS	DIDS 3	7	Pharma. block of inhibition	2.21, 2.22
vc -45 DIDS	DIDS 3	5		2.20
cc ctrl. DNDS	KGlu 4 ₋	5/7	DNDS control	2.23, 2.24
cc DNDS	KGlu 4 ₊	5/7	Pharma. block. inh. for cc	+2.25, 2.26

Table 2.1. Summary of Main Electrophysiological Recording Conditions. Experimental conditions employed in this Thesis for the assessment of SWR. In all conditions an extracellular electrode was present; we report here in the first column only additional extra or intracellular electrodes, together with their recording mode, voltage clamp —vc— or current clamp —cc— and holding potential, if applicable. The name of the condition is highlighted in bold. The reversal potential of chloride for cells recorded with the solution KGlu 0 is about -70 mV; that for those recorded with solution KGlu 1 is estimated to be -67.3 . The reversal for AMPAR-mediated excitation with the solution CsGlu 2 is about -6.5 mV (Fig. 2.8). Consecutive colored rows indicate that the same cells were recorded under successively different experimental conditions (holding potential and/or pipette solution), in the order given in the table.

Finally, we present in **Section 2.5** a deafferented preparation of CA1 (*minislice condition*), which allowed us to investigate what is strictly local to CA1 in the generation of ripples versus what is dependent on afference from CA3 or subiculum.

We sum up the experimental conditions introduced in this Chapter in Table 2.1. Further analyses of rhythmicity of currents, locking to the extracellular ripple and inhibitory-excitatory timing relations are deferred to the next **Chapter 3**.

description of condition	No. cells	Notes	Figures (this Chapter)
cc synth. ripple	6	current injection	2.15
cc spont. ripple	7	depolarized	2.16
cc example		~ -80 mV	2.10
2 nd LFP pipette	—	~ 50 - 500 μ m	2.11
$2 \times$ vc -66 mV	2×20	~ 40 - 580 μ m	2.12, 2.13, 2.14
cc example	1	-80 mV baseline	2.10.

Table 2.2. Summary of Additional Electrophysiological Recording Conditions. Continuation of Table 2.1. Columns, units and conventions as described in caption to that table.

2.1 In Vivo, Head-Fixed Parallel Recordings

To directly study synaptic inputs onto CA1 pyramidal neurons during SWRs in vivo, without confounding effects of anesthesia, we used a recently established approach where mice were habituated to head restraint in the recording setup over several days (Crochet and Petersen, 2006; Harvey et al., 2009; Margrie et al., 2002). During quiet wakefulness, LFP recordings in the area CA1 revealed SWRs comparable to those observed in chronically implanted animals (e.g., Buzsáki et al., 2003; Fig. 2.1).

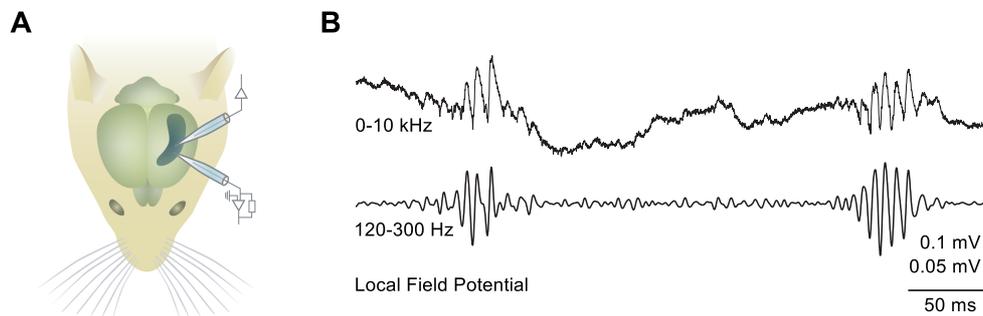


Figure 2.1. Hippocampal Recordings During LFP Ripples in Head-Fixed, Awake Mice.

A. Schematic of the recording configuration with an intracellular and extracellular electrode.

B. Example recording from the extracellular electrode in CA1. Two SWR events are visible in the unfiltered (top) and 120-300 Hz bandpass-filtered local field potential (bottom).

The average incidence of SWR events was 0.14 ± 0.02 Hz. The frequency at which their power peaked was in the range from 127 to 147 Hz and averaged 136.3 ± 0.2 Hz for the 1,288 events from 16 mice that were analyzed. Their duration was between 12.1 ms and 156.6 ms with a mean of 65.0 ± 0.7 ms (Fig. 2.2 C). These values are to some extent dependent on the details of the detection procedure, which is described in Appendix A.1 (*SWR detection*).

Next, we combined LFP recordings with simultaneous whole-cell recordings and stainings from nearby CA1 pyramidal neurons (Fig. 2.3). Six neurons out of 19 attempted were successfully stained. They all revealed the typical morphology of CA1 principal cells (Figs. 2.3 A and 2.6 A).

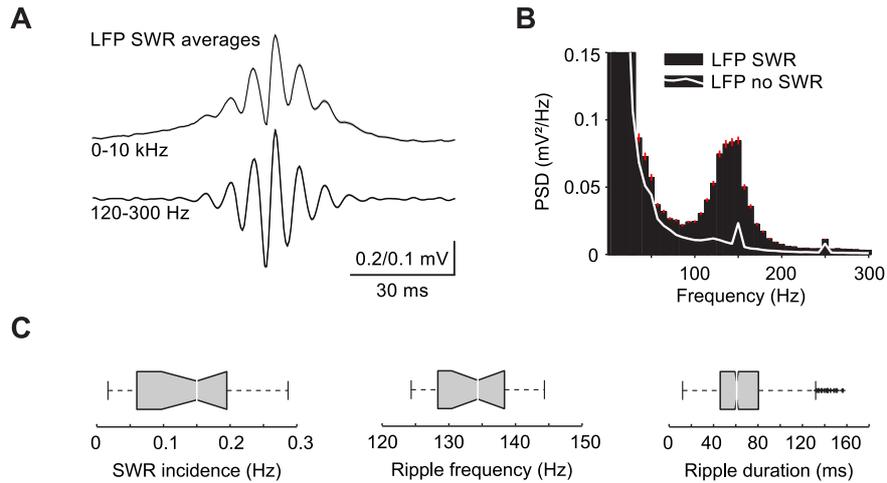


Figure 2.2. Statistics of LFP Ripples. **A.** Grand averages (\pm SEM, grey lines) of 1,288 SWR events from 16 mice. **B.** Averaged power spectral density (PSD) of all identified ripples and of eventless epochs preceding ripples (white line). **C.** SWR incidence, ripple peak frequency and duration from all recorded animals summarized with notched boxplots. The central line represents the median of the respective distribution; the notched area the 95% confidence interval about it. The box extends from the 25% quartile to the 75% quartile. The whisker ends signal the maximum and minimum value that are not farther than 1.5 times the box width from the respective box edge. Little crosses mark values that do not fit within that range.

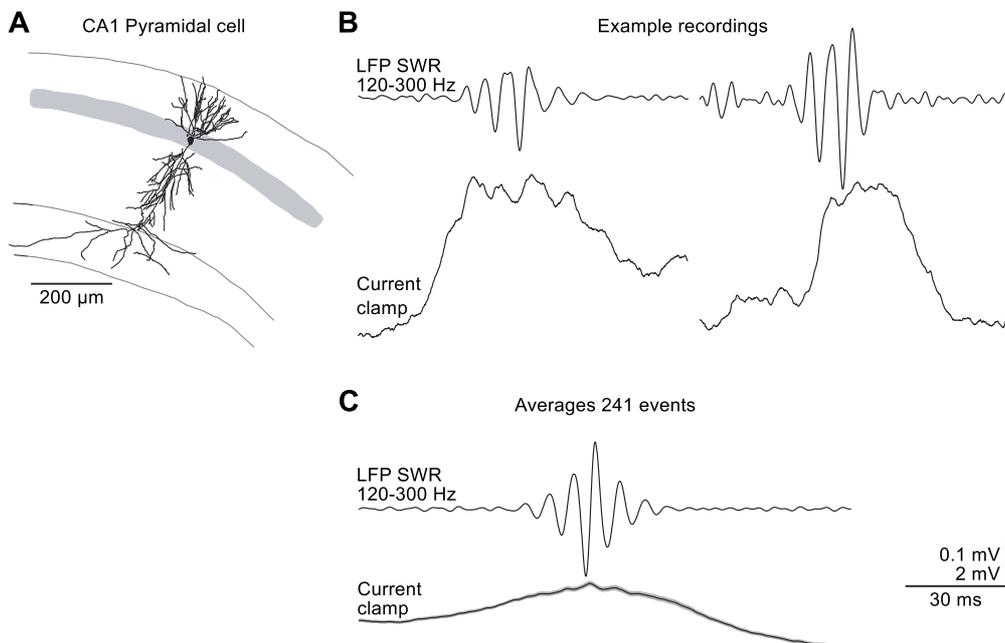


Figure 2.3. Intracellular Current-Clamp Recordings in Head-Fixed, Awake Mice. **A.** Anatomical reconstruction of a CA1 pyramidal cell. **B.** Examples of rippled membrane potentials from this cell recorded in current clamp at the time of the extracellular LFP ripple, shown filtered above as a reference. **C.** Averages (\pm SEM in grey) of 241 membrane potential traces from 12 cells in 10 mice at the bottom during LFP ripples, shown filtered above. Note the preserved ripple signature in the average of the intracellular signal.

Whole-cell recordings during ripples in current clamp showed that synaptic input during SWRs was frequently superimposed with fast ripple-associated voltage fluctuations (Figs. 2.3 B and C; 241 events from 12 cells with average resting potential of -61.8 ± 1.4 mV and action potential amplitude of 47.1 ± 3.5 mV). The course of the membrane potentials at the time of the SWR was biphasic: a jagged depolarization was observed first, in synchrony with the extracellular ripple, followed by a rather tonic hyperpolarization (single examples in Fig. 2.4; summary of polarity in Fig. 2.5).

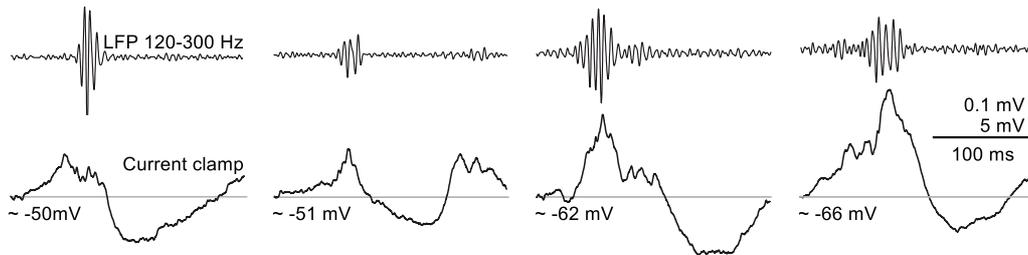


Figure 2.4. Current-Clamp Recordings at Different Potentials. Example events recorded in current clamp at different initial membrane potentials (bottom row) reveal ripple-associated de- and hyperpolarizing synaptic inputs in coincidence with the LFP ripple, which is shown above for reference filtered between 120 and 300 Hz.

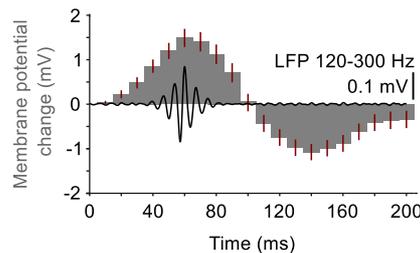


Figure 2.5. Time Course of Membrane Potential During Ripples. Histogram of changes of membrane potentials binned to a 10 ms time windows (means \pm SEM from 241 events in 12 cells). The black trace averages the concurrent extracellular signal in the ripple band (120-300 Hz). Each intracellular event was referred to the average voltage during the 60 ms preceding the LFP ripple peak (“0” change). On average, ripple-associated depolarization is followed by pronounced hyperpolarization.

We also observed the rippling of the synaptic input in our recordings in voltage clamp mode. At a voltage close to the reversal potential of Cl^- (-70 mV) ripple-associated inward currents were jagged. Figure 2.6 B shows representative examples; The undulation is washed out to a great extent upon averaging all 421 events from 8 cells (Fig. 2.6 C).

The biophysical interpretation and quantitative analysis of in vivo whole-cell voltage-clamp recordings is complicated:

- i. The extracellular ion concentrations are not known with sufficient precision to determine the reversal potential for Cl^- . Hence the cell cannot be clamped at a potential that isolates excitation at the single-cell level for its analysis during SWR (see Section 2.3 and Equation 2.1 below for an outline of the intended technique, called electrotonic isolation).

- ii. Even with a hypothetically good estimation of the reversal potential of Cl^- , the high series resistance of in vivo recordings (Lee et al., 2006; Margrie et al., 2002) and voltage-clamp errors (Williams and Mitchell, 2008) make difficult to determine both the polarity and the timing of fast synaptic currents, in particular if they arise from distal synapses.

We therefore turned to a previously established in vitro model of hippocampal SWRs where these problems are alleviated and the recordings have enough stability of baseline to envisage the application of advanced analysis techniques in the time domain.

2.2 An In Vitro Model for Investigating SWR

The difficulties to hold a CA1 cell clamped at a potential such that the driving force for either the inhibitory, or the excitatory synaptic contributions is greatly diminished (in order to isolate the other) led us naturally to consider an in-vitro approach on a murine slice preparation. The experimental model is described at length in Maier et al. (2009). There, the previous approach reported in Maier et al. (2003) is improved upon by establishing reliable SWR recordings in a submerged chamber, thus enabling optically guided whole-cell recordings or even Ca^{2+} imaging. Recordings on such slices have established that the layered profile, propagation pattern, frequency of occurrence, ripple frequency and sensitivity to gap-junction blockers of sharp waves and ripples are in qualitative agreement with observations in vivo (Buzsáki, 1986; Ylinen et al., 1995; Csicsvari et al., 2000; see also Section 1.4 in previous Chapter).

Sharp waves vs. ripples vs. sharp-wave ripples Throughout this work we refer to sharp-wave ripples as a unitary, not composite phenomenon. This is just a convenience of language applicable to the recordings presented here, all of which were made in the pyramidal layer of CA1 and rather consistently showed fast fluctuations (“ripples”) patterning a slow LFP deflection (“sharp wave”). The terms sharp wave and ripple should however not be conflated in general. Concomitant to pyramidal layer rippled sharp waves (SWR), large sharp waves with less or no ripple, are visible at other depths (see Fig. 2.7 and discussion below). Conversely, ripple-like oscillations have been observed in absence of an underlying sharp wave. Finally, bath application of gap-junction blockers such as halothane (Ylinen et al., 1995), carbenoxolone and octanol (Maier et al., 2003) suppress efficiently the ripple without eliminating the sharp wave. For all these reasons, the term *SWR* in this work is meant only as a phenomenological description of a jagged slow deflection in the field, and will be used interchangeably with both the designations *sharp wave* and *ripple*, without implying any statement about a common mechanism of both phenomena. We note, though, that the detection algorithm employed here for SWR events in the main results chapter (Chapter 3) is amplitude-based and therefore emphasizes the sharp waves. This is in contrast to the in vivo literature, where the customary approach is rather to detect-by-ripple because the baseline is not stable enough (compare Figs. 2.1 B above and 2.9 A below), and because the same electrodes that detect SWR are often used to detect unit activity, which is more abundant in the cell body layer.

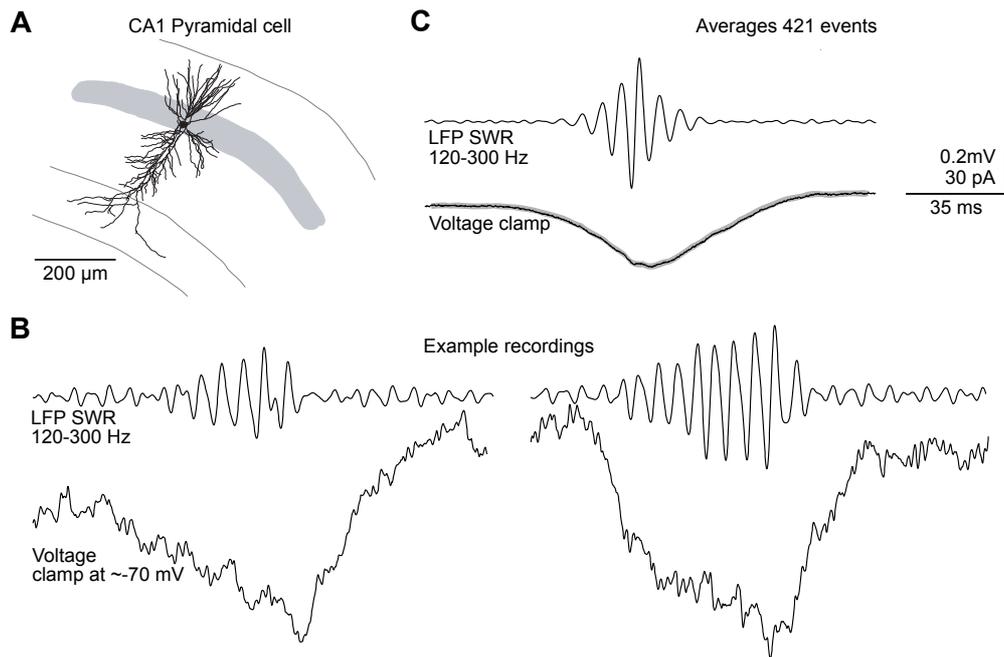


Figure 2.6. Intracellular Voltage-Clamp Recordings in Head-Fixed, Awake Mice. Recordings obtained from CA1 pyramidal neurons during LFP ripples. **A.** Anatomical reconstruction of a CA1 pyramidal neuron in a different animal to that in Fig. 2.3. **B.** Ripple-associated voltage-clamp data and the simultaneously recorded LFP. Holding potential was close to the resting membrane potential (-70 mV) of the cell. **C.** Averages (\pm SEM, grey) of synaptic currents concurrent to LFP ripples (top; filtered). A total of 421 events from 8 cells in 5 mice were included.

Layer profile The expression of sharp waves and ripples along the depth axis of the hippocampus (orthogonal to the pyramidal layer) provides evidence about the synaptic origin of the phenomenon insofar the layer preference of the afferents is known.

Figure 2.7 displays the situation in our in vitro model (Maier et al., 2009). Ripple amplitude is maximum in the pyramidal layer. It decreases sharply away from it towards stratum radiatum (rad.) and oriens. In comparison, in the rat in vivo (Ylinen et al., 1995) ripples are also concentrated in the pyramidal layer but much less visible in the dendritic layers. In vivo, underlying sharp waves are also less spread across layers than in the mouse in vitro, showing maximum amplitude in the stratum radiatum and in stratum pyramidale. Phase reversal of both sharp waves and ripples occurs in our preparation below the pyramidal cell layer, much as in the in vivo situation (Buzsáki et al., 1983, Fig. 14, Ylinen et al., 1995, Fig. 2). The standard interpretation of the differential expression of ripple oscillations and sharp wave deflections is that the latter reflect compound EPSPs in the stratum radiatum of CA1 induced by the Schaffer collateral/commisural afference whereas ripples are generated in the pyramidal cell layer (Buzsáki et al., 1992). Their patterning would be due to perisomatic-targeting interneurons (Ylinen et al., 1995) that have been found to fire locked to the extracellular ripple (Klausberger et al., 2003). Elucidation of the cellular mechanisms behind both phenomena is much aided by the capability to resolve fast contributions of specified kind, excitatory or inhibitory; indeed, more direct evidence than afforded by spiking patterns is provided by examining the currents that cells *see* during the SWR. To this task we now turn.

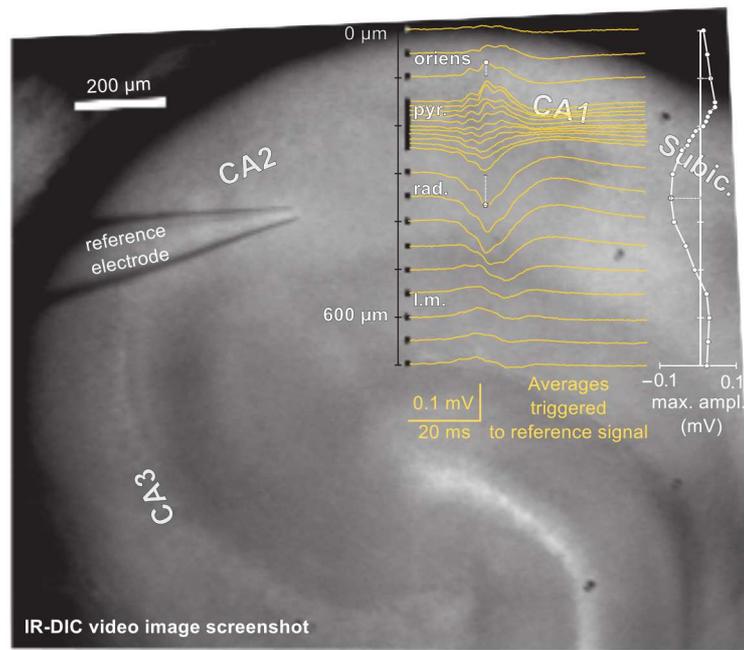


Figure 2.7. Layer Profile of Sharp Wave-Ripples in a Mouse Slice. Example laminar profile recorded on a slice of mouse hippocampus (background image) by Nikolaus Maier. Averages of 20 LFP recordings during sharp waves are shown in yellow, next to the corresponding electrode positions (black dots). These averages were triggered to the signal from the reference electrode, positioned in CA1 stratum radiatum close to CA2. The maximum amplitude deflection (positive or negative, depending on the trace) is summarized by the white voltage profile to the right. Further experiments in $n = 8$ slices confirm the amplitude reversal at roughly $200 \mu\text{m}$ down from st. alveus (not shown). Rad.: stratum radiatum; pyr.: stratum pyramidale; l.m., stratum lacunosum moleculare; Subic.: Subiculum. Modified from Fig. 5A in Maier et al. (2009), with permission per CC-BY license.

2.3 Electrotonic Isolation of Currents

This is Ohm's law for the postsynaptic current due to the flow of one or more species of ions crossing a population of channels with conductance g and ensemble reversal potential E when the potential of the membrane patch is V_m :

$$I_{\text{synaptic}} = g(V_m - E).$$

The linearity of synaptic current with driving force $V_m - E$ holds with good approximation if the responsible conductances are not strongly voltage dependent and if the ratio of ion concentrations at both sides of the membrane is small (Hille, 2001, Chp. 13). Currents flowing through several channel types combine additively. For example if excitatory resp. inhibitory currents are mostly due each to one conductance only:

$$I_{\text{synaptic}} = g_{\text{exc}}(V_m - E_{\text{exc}}) + g_{\text{inh}}(V_m - E_{\text{inh}}). \quad (2.1)$$

This is the case in our preparation, where excitation is due mostly to AMPA receptor (AMPA)-mediated currents, activated by the neurotransmitter glutamate, and inhibition is mostly due to GABA receptor (GABA)-mediated currents activated by the

neurotransmitter GABA (γ -aminobutyric acid). We did not see evidence of significant slow currents at the holding potentials of -66 , -45 , and -6 mV, such as those mediated by NMDA receptors or GABA_B receptors, which have time constants on the order of the hundreds of milliseconds.

From Eq. 2.1, two strategies become apparent to isolate either of the two contributing currents — make either of the factors g or $V_m - E$ zero. The first is termed *pharmacological isolation* (more on it below); the second necessitates control of V_m to nullify the electromotive force induced by the gradient of potential from the inside to the outside of the cell and is denominated *electrotonic isolation*. Electrotonic isolation is thus an electrophysiological technique consisting in holding a cell clamped close to the the reversal potential corresponding to the ion or ions flowing through a given population of channels. The electrical driving force due to the gradient between the interior of the cell and the extracellular milieu is thus so diminished that currents due to those ions become, at least in theory, unobservable compared to other current flows.

Glutamate receptors and others such as nicotinic acetylcholine or serotonin are channels that let through both Na⁺ (with reversal around $+50$ mV) and K⁺ (-70 mV) in similar proportions, such that the combined reversal potential lies somewhere between -10 and 0 mV. At that potential the efflux of K⁺ is exactly compensated by the influx of Na⁺. Below that the Na⁺ inflow dominates, and above it is the outflow of potassium, thus creating a hyperpolarizing current. With the cell in voltage clamp held at V_m close to $E_{exc} \simeq -5$ mV, even if presynaptic action increases the conductance g_{exc} at the postsynaptic membrane, it will be opening doors nobody wants to go through. For its part, fast inhibitory synaptic transmission in CA1 takes place through channels specific to Cl⁻ that open when the neurotransmitter GABA docks at GABA_A receptors, with reversal $E_{inh} \simeq -65$ mV. Knowledge that the reversal potentials of AMPAR- and GABA_AR-mediated transmission are respectively well above and well below the firing threshold of about -40 mV justifies retrospectively the use of the labels *exc* and *inh* in Eq. 2.1: activation of the first will increase likelihood of postsynaptic discharge by raising membrane potential closer to firing threshold and, conversely, that of the latter will lower it away from threshold, thus effectively inhibiting spiking.

In the following, we first estimate the reversal potentials applicable in our preparation (Section 2.3.1). With knowledge of the reversal potential for inhibition we then present in the next Section 2.3.2 some features of the excitatory currents concurrent to the LFP ripple as observed in our in vitro model — their frequency in the ripple range, their coherence to the extracellular ripple and across most of CA1. Next (Section 2.3.3) we examine current clamp data to check if spikes are more likely to happen at a certain phase of the ripple. Finally, we show in Section 2.3.4 mixed excitatory-inhibitory currents and solely inhibitory currents made visible by raising the holding potential.

2.3.1 Reversal Potentials

To determine reversal potentials (Fig. 2.8) we quantified the current flow into the cell by the amplitude of PSCs triggered by electrical stimulation. Bath application of synaptic transmission blockers for AMPA, NMDA and GABA_B ensured pharmacological isolation of GABA_AR-mediated inhibition, whereas blockers of NMDA, GABA_A

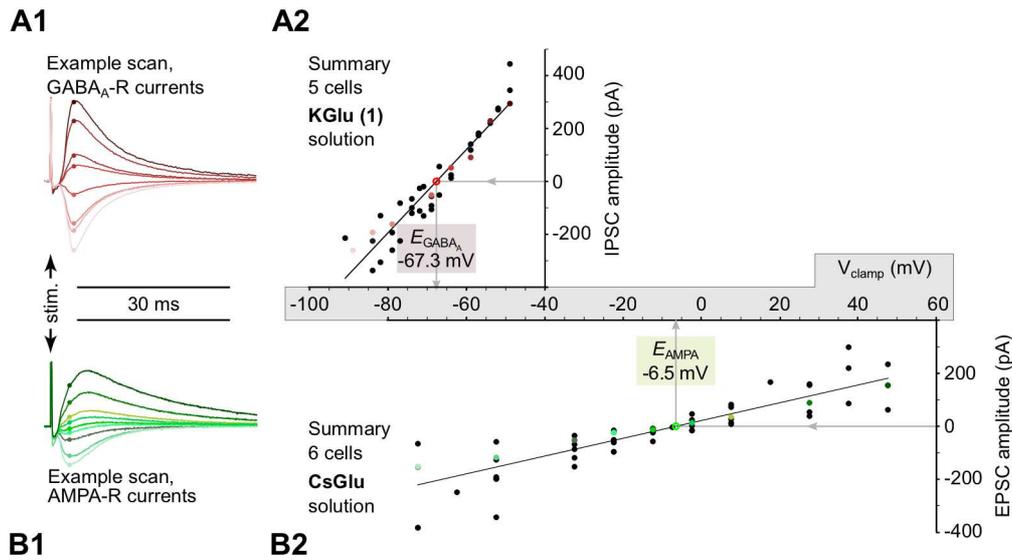


Figure 2.8. Experimental Confirmation of Inhibitory and Excitatory Reversal Potential.

A1. Example currents due to electrically-evoked synaptic transmission via GABA_A receptors. Bath application of 10-20 μM NBQX, 50 μM D-APV and, for some cells, 1 μM CGP blocked AMPA, NMDA and GABA_B currents respectively. The patch electrode solution was **KGlu 1** (see Table A.1 in the Appendix). Each trace averages ten events corresponding to the liquid junction-corrected holding potential indicated to the right at A2 by the abscissa of the corresponding colored disc. **A2.** Linear regression of PSC amplitudes for a total of five cells against holding potential. PSC amplitudes were determined at the time of the minimum of current at the lowest holding potential (indicated by dots to the left at A1) and baseline-corrected for the offset due to holding potential changes. Holding potential corrected by 14 mV for liquid junction offset. **B1.** Example currents due to electrically-evoked synaptic transmission via AMPA receptors. Bath application of 2 μM gabazine, 50 μM D-APV and 1 μM CGP blocked GABA_A, NMDA, and GABA_B currents respectively. The patch electrode solution was **CsGlu** (see Table A.1). All else as in A1. **B2.** Linear regression of PSC amplitudes for a total of six cells analogous to analysis in A2.

and GABA_B isolated excitatory AMPAR-mediated excitation. The protocol consisted, then, in triggering PSCs at a range of holding potentials for the cell (panels A1 and B1) and measuring their amplitude (A2, B2).

The amplitude of the evoked PSCs depends linearly on the driving force according to Ohm's law (Eq. 2.1); linearity was observed both for GABA-mediated inhibition (Fig. 2.8 A2) and AMPA-mediated excitation (B2). The reversal potential was obtained by applying linear regression to stimulation data pooled from several cells. The holding potential at which the amplitude of the current vanished was -67.3 mV for inhibitory currents and -6.5 mV for excitatory currents. At -66 mV, all excitatory currents are inward by definition, while inhibitory inputs, if present, should display very small amplitudes due to the low Cl^- driving force, and be outward directed (influx of negative charges is equivalent to outflow of positive ions).

It turns out that neither glutamatergic NMDA receptor-mediated currents, nor gabaergic GABA_B receptor-mediated ones play a significant role in our experimental scenario (see Appendix A).

2.3.2 AMPAR-Mediated Currents

Using an established in vitro model of hippocampal SWRs (Maier et al., 2009; schematic in Fig. 2.9 A), we recorded first the extracellular LFP and currents from cells held at their reversal potential of inhibition (-67.3 mV, Fig 2.8). We analyzed the polarity, rhythmicity, coherence with the concurrent LFP signal, and spatial coherence across CA1 of these putatively excitatory currents in a series of experiments that we introduce in this Section. Additional analyses of rhythmicity, coherence and relation to the timing of inhibitory currents are presented in the next Chapter 3.

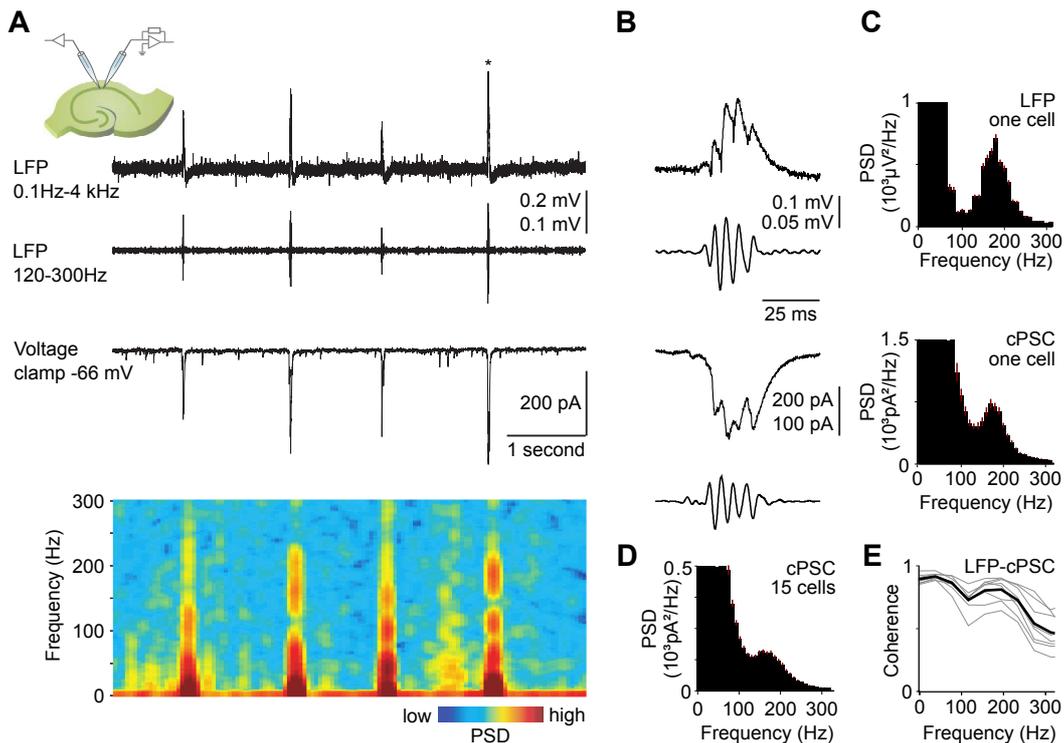


Figure 2.9. Phasic Currents in CA1 Pyramidal Neurons During SWR In Vitro. **A.** In hippocampal slices, extracellular recordings were combined with close-by patch-clamp recordings from CA1 pyramidal cells (diagram at top left). Extracellular SWRs in CA1 (first trace), filtered version (second), and simultaneously recorded PSCs (third; cell clamped at -66 mV) show the systematic co-occurrence of LFP ripples and compound PSCs (cPSCs). At bottom the spectrogram of the current in A shows periods of enhanced power at ~ 120 - 250 Hz coincident with cPSCs. **B.** Magnification of the SWR marked with * in A and its simultaneous cPSC. Filtered versions (120-300 Hz) are shown below each. **C.** Average of power spectral densities of 133 LFP ripples (top) and respective cPSCs (bottom) from the recording in A; both spectra show a distinct component at ~ 200 Hz. **D.** Average PSD of 1,137 PSC segments during ripples (15 cells) peaks at ~ 195 Hz. **E.** Coherence between LFP and cPSC signals peaks at ripple frequency (~ 200 Hz; 8 cells in 6 slices).

Sharp waves were found to spontaneously occur in the extracellular signal at a rate of 0.77 ± 0.05 Hz (average from recordings in 28 slices). They bore ~ 200 Hz ripples, similar to ripples observed in vivo (Fig. 2.9 B, top). SWRs were found to originate in CA3, propagate through CA1 to the subiculum and exhibit a polarity inversion below *stratum radiatum* just like their in vivo homologues (Buzsáki, 1986). The signal from

the patched cells displayed large-amplitude PSCs simultaneous to extracellular SWRs. Closer inspection revealed compound bursts of postsynaptic currents (cPSCs) with a distinct frequency matching the dominating frequency of LFP ripples (Fig. 2.9 B, bottom). In average, the peak cPSC frequency lay close to 200 Hz (Figs. 2.9 D 1, 137 events from 15 cells: average 194 ± 6 Hz and range 160-240 Hz). It was, however, more variable even within the same cell than the peak LFP frequency during ripples (compare histograms of 133 events from one cell in Fig. 2.9 C). This is evidenced by the spectrogram at the bottom of Figure 2.9 A as well, where even consecutive cPSCs display noticeable differences of spectral composition. Spectral content of SWRs was analyzed in 100 ms stretches of raw data centered on the SWR peak using the Fast Fourier Transform (FFT). Frequency resolution of the resulting power spectral density (PSD) plots was 9.98 Hz.

A similar fast frequency component was observed for postsynaptic potentials in the current-clamp configuration (Fig. 2.10; using the same KGlu-based intracellular solution, KGlu 1). The rippled character of the postsynaptic membrane potential was thus confirmed in the more physiological condition of current clamp. Figure 2.10 eloquently displays however the disadvantage of current clamp for the analysis of fast synaptic activity: the membrane time constant on the order of tens of ms smears the individual synaptic events, whose rhythmic arrival is only highlighted by high-pass filtering. High-pass filtering enhances the fluctuations of a time series and is similar to performing a derivative, which is a substantial part of inverting the convolution with the synaptic kernel of the leak conductance (see Chapter 3 for a precise formulation).

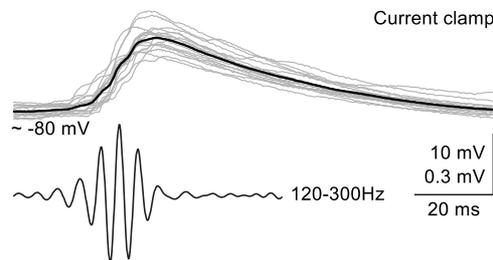


Figure 2.10. Membrane Potential is Rippled During Field SWR. At top, single (grey) and averaged (black) ripple-associated voltage traces from a CA1 pyramidal cell recorded in the current-clamp configuration; the recording was performed at resting membrane potential. Below, band pass-filtered version of the membrane voltage average. Note rhythmic fluctuations at ~ 5 ms in the signals.

Finally, we determined the coherence of cPSC and field ripple oscillations. In eight simultaneous whole-cell/LFP recordings, we observed a peak of coherence at ~ 200 Hz (Fig. 2.9 E), indicating preferential coordination of both oscillations at the ripple frequency.

2.3.2.1 Spatial Coherence of Ripples Across CA1

Ripples in the LFP in vivo are known to be synchronized in the average over most of CA1 (Chrobak and Buzsáki, 1996, Fig. 1E, rat), with delays due to the propagation

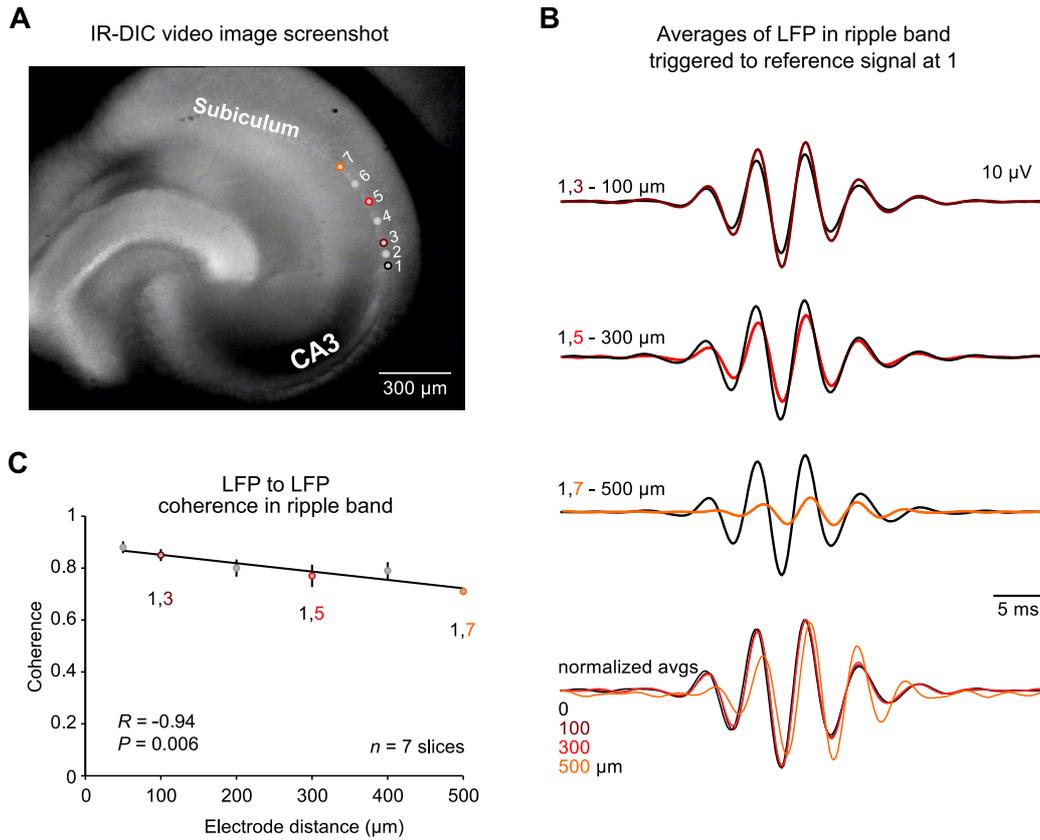


Figure 2.11. Coherence of Extracellular Ripple Activity in CA1. **A.** Screenshot of an example experiment showing electrode positions. **B.** Examples of ripple-band averages (120–300 Hz) with the electrode at increasingly distant positions (colors as in A). At bottom all four traces are aligned to facilitate comparison. **C.** LFP-to-LFP peak coherence in the ripple band as a function of distance, summarizing dual recordings in seven slices.

of the field oscillation through the tissue towards the subiculum. We verified that this was the case also in our preparation, using triggered averages and coherence. The coherence $C_{xy}(f)$ of two signals x and y is defined as their normalized cross-spectral density (Shiavi, 2007, Chp. 9). For a given frequency f , the coherence is higher the more stable the phase difference of the signals x and y ,

$$C_{xy}(f) \equiv \frac{|P_{xy}(f)|^2}{P_x(f)P_y(f)}. \quad (2.2)$$

Coherence can be used to quantify synchrony established in specific frequency bands. See Section 3.5 in the next Chapter for details.

Figure 2.11 shows the coherence of the extracellular signal during ripples as a function of distance. One LFP electrode was left at position 1 and provided the reference signal. A second electrode was moved in steps and measured the distant signal (positions 2...7, panel A). The extracellular SWR coherence was compared between positions. A statistically significant decline of coherence was observed, probably due to the phase shift (panel B), but overall coherence remained still high for distant recordings (panel C).

To demonstrate the synchrony of inputs to cells constituting the local network, we examined how the observed single-cell-to-ripple coherence (Fig. 2.9 E) extends to the network level.

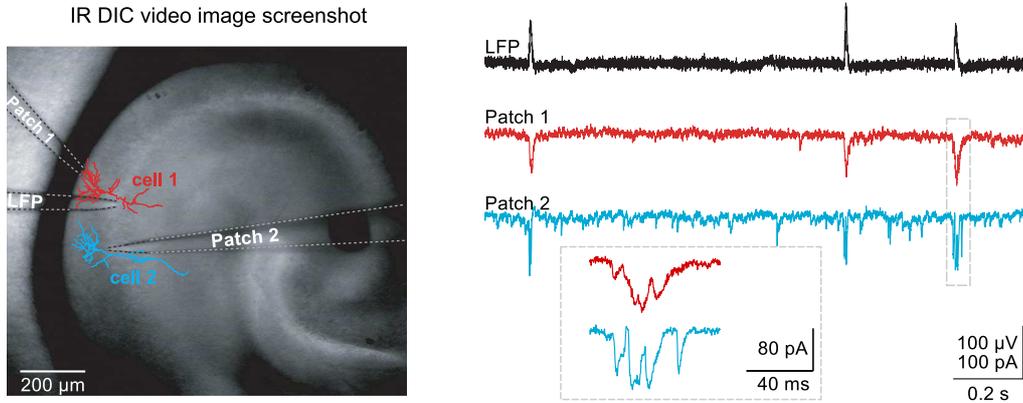


Figure 2.12. Parallel Recordings in Voltage Clamp. **Left.** Reconstructions of pyramidal cells (blue and red) overlaid on an infrared differential interference-contrast (IR DIC) video image of a hippocampal slice. *LFP*, *Patch 1* and *Patch 2* indicate the electrodes' arrangement. **Right.** Simultaneous recording from both cells and LFP. The inset magnifies the boxed events.

If ripple-locked cPSCs indeed represent signatures of neuronal population oscillations, we would expect synchrony of cPSCs across multiple cells in the local network, and cell-to-cell input coherence should extend over a considerable distance. We tested this hypothesis in dual pyramidal cell recordings (Fig. 2.12). Consistent with inputs from a synchronized network during SWRs, cPSCs were correlated (Fig. 2.13 A).

The amplitudes of cross-correlation peaks decreased with distance and their time

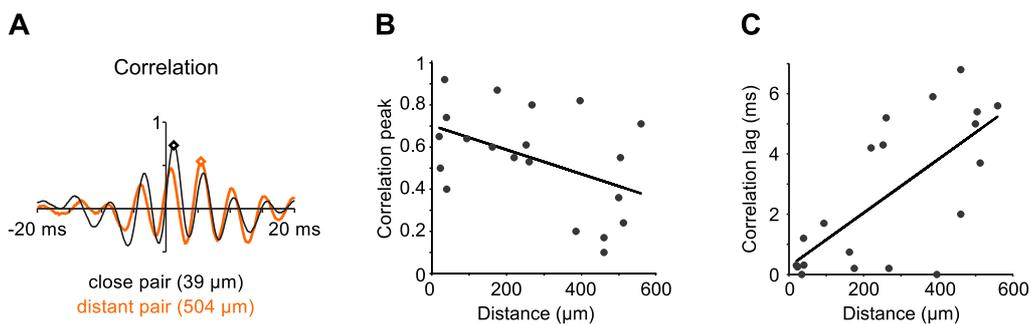


Figure 2.13. Correlation of cPSCs Across Simultaneously Recorded Principal Cells. **A.** Cross-correlations during ripples of two example dual recordings at different cell distances. Peak correlation value of close (black diamond) vs. distant cell pair (orange diamond) was 0.74 vs. 0.55, and respective lags 1.2 ms vs. 5.4 ms. **B.** Peak correlation value decreases as a function of distance between recorded neurons ($n = 20$ dual recordings; $R = -0.46$; $P = 0.04$). Cells located in proximal CA1 were taken as reference. **C.** Lag of correlation peak increases with distance between recorded neurons ($n = 20$ dual recordings; $R = 0.7$; $P = 0.0006$).

lags increased (Fig. 2.13 B, C with 2,132 SWR-associated cPSCs from a total of 20 dual pyramidal cell recordings). Both observations are in agreement with a spread of SWR activity from proximal to distal sites in CA1 with respect to CA3.

In addition, we found that ripple-associated cPSCs in pairs of pyramidal neurons were phase-coherent, as demonstrated by coherence maxima in the ripple frequency range (Fig. 2.14 A). Cell-to-cell coherence of cPSCs insignificantly decreased with increased spatial separation (Fig. 2.14 B; $R = -0.26$, $P = 0.26$). In line, comparison of cPSC coherence in close by ($<100 \mu\text{m}$) vs. distant ($450\text{--}580 \mu\text{m}$) neuron pairs revealed no significant difference (Fig. 2.14 C; $P = 0.39$; rank-sum test). Collectively, these results on dual principal cell recordings confirm that ripple-locked cPSCs are indeed signatures of population oscillations.

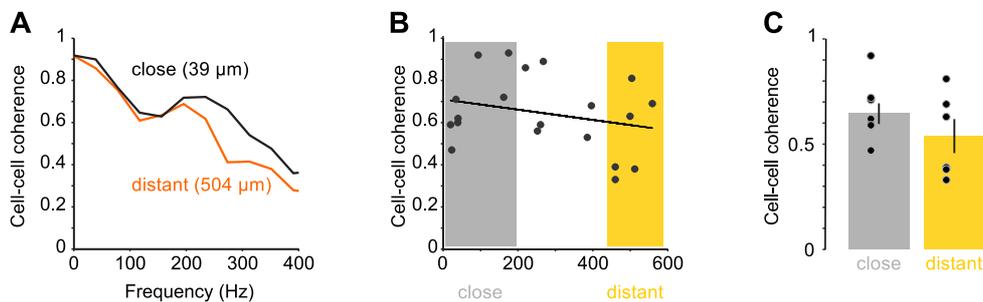


Figure 2.14. Coherence of cPSCs Across Simultaneously Recorded Principal Cells. **A.** CPSC coherence from close vs. distant dual recordings (same pairs as in Fig. 2.13 A). **B.** Summary of peak coherence values in the 120–300 Hz range as a function of distance within CA1. Linear regression shows a modest negative correlation ($n = 20$ simultaneous recordings; $R = -0.26$; $P = 0.26$). **C.** Mean peak coherence values in close ($<200 \mu\text{m}$) and distant ($>450 \mu\text{m}$) recordings are similar (colors as in B; mean peak coherences: close: 0.65 ± 0.05 , $n = 8$; distant: 0.54 ± 0.08 , $n = 6$; $P = 0.41$, Mann-Whitney test). See also Fig. 2.11.

2.3.3 Timing of Spikes

The precedent observation that excitatory PSCs are phasic and ripple-locked raised the question whether they could account for the timing of action potentials in target CA1 principal neurons. We conducted two analyses in current clamp whose results were confirmed later under pharmacological block of inhibition (see Section 2.4.2).

First, we injected a synthetic excitatory current (green line in Fig. 2.15) with rhythmicity and individual PSC kinetics similar to those observed during spontaneous SWRs. The single beats of the ripple there are of only 50 pA amplitude compared to up to 400 pA observed in measurements, but they add to a baseline current that brings the otherwise extremely sparse-firing pyramidal cells closer to their spiking threshold. This phasic input considerably enhanced firing probability (Fig. 2.15 A) and the resulting spikes were significantly locked to it with a high vector strength value of 0.96. The set of experiments demonstrated that a current waveform resembling the input observed during ripples can indeed account for the occurrence and the timing of action potentials.

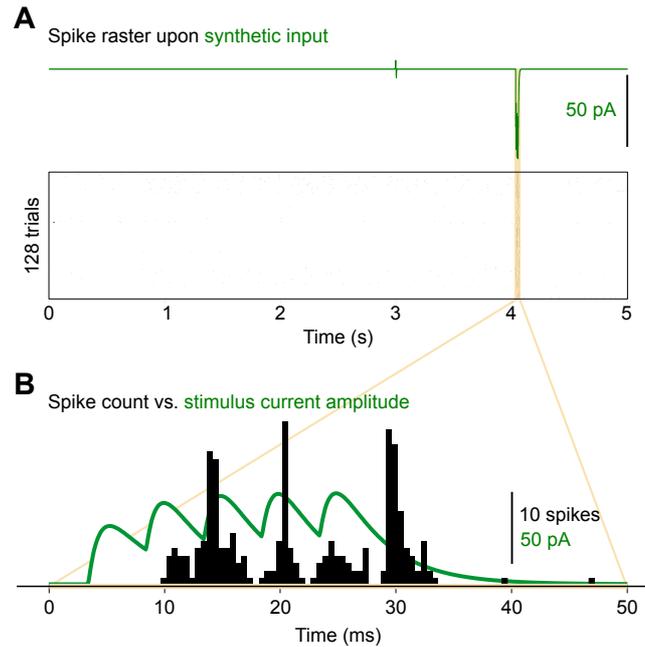


Figure 2.15. Spikes Are Entrained by Rippled Current Injection. **A.** Spike raster from six cells held close to their firing threshold in 128 trials (bottom; 10, 20, 20, 22, 26, and 30 trials, respectively). A brief current injection (highlighted at top) enhanced spiking. **B.** Histogram of the spike count during the stimulation window highlighted in A overlaid on top of the amplitude of the injected current (green).

In a second set of experiments, we analyzed action potential timing during spontaneous SWRs, measuring the membrane potential variations in current clamp at a slightly depolarized value (Fig. 2.16). Since CA1 pyramidal cells have a resting potential close to -80 mV (Fig. 2.10), depolarization brings them closer to the reversal of Cl^- (-67.3 mV) and helps minimize confounding inhibitory input. It also enhances the extremely low firing probability of these cells, which, in our preparation fire only sporadically during ripples (as compared to e.g. Böhner et al., 2011), presumably a sign of the strength with which inhibition holds them back or maybe an offshoot of the many excitatory afferents cut when preparing slices.

The scarcity in the firing was indeed a problem for the collection of data: out of 27 cells investigated, only seven concentrated 72% of the spikes and showed significant locking of action potentials with respect to the ripple in an interval of 30 ms centered to the peak of the ripple-band-filtered LFP ($P < 0.05$, Rayleigh test; mean AP phase -56°). The cell-wise variability in locking strength of action potentials to field ripples hints at a high variability of the strength of recruitment into the synchronous neuronal assembly, in line with recent research (Böhner et al., 2011).

2.3.4 Higher Holding Potentials

The intracellular pipette solution K_{Glu} 1 employed so far has a reversal potential of -67 mV, whence no inhibitory currents (outward, positive-signed traces) should be observed. Holding the cell membrane at a more depolarized potential allows to verify that there exist simultaneous inhibitory currents and that they are rhythmic as well,

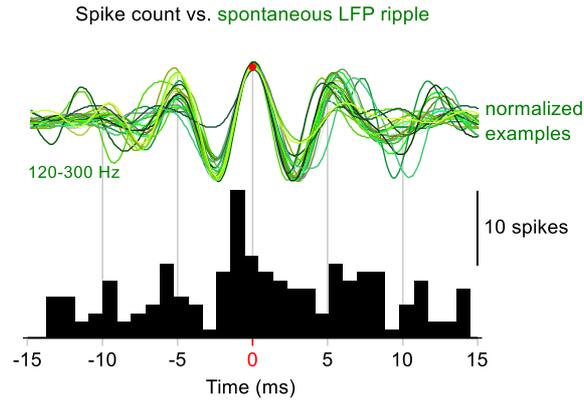


Figure 2.16. SWR-Concurrent Spikes of Slightly Depolarized CA1 Cells Lock to Ripple.

as expected from the literature. For eight cells held first at -66 mV we continued the recordings a few minutes in each case at -45 mV. As can be seen in the examples of Figure 2.17, cPSCs tend to be of mixed polarity, with an overall inward current at the beginning followed by an outward current.

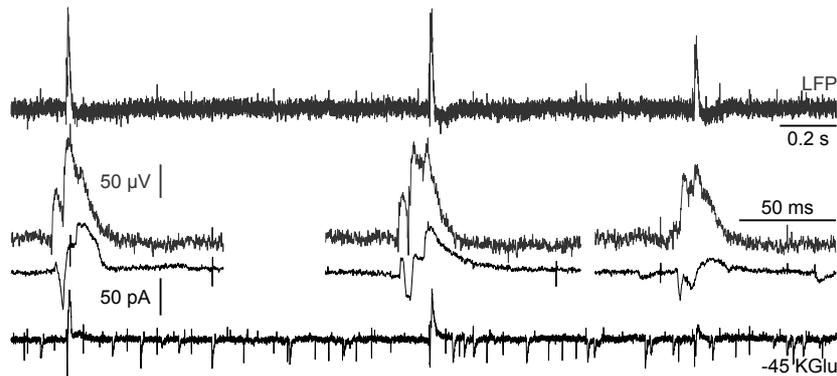


Figure 2.17. Example KGlu 1 -45 mV. Three-second sections of simultaneous LFP (top) and current recordings (bottom) with intracellular solution KGlu 1 at a holding potential of -45 mV. Three individual SWR/cPSC events zoomed-in as indicated by scale bars.

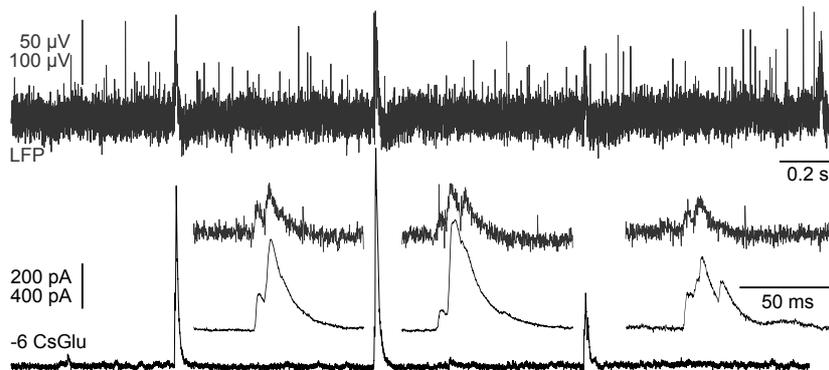


Figure 2.18. Example CsGlu -6 mV. Three-second sections of simultaneous LFP (top) and current recordings at -6 mV holding potential (bottom) with intracellular solution CsGlu (reversal potential of excitation -6.5 mV; see Fig. 2.8 B2). SWRs are regularly associated with compound IPSCs (cIPSCs).

It is not possible to hold the cell at the much higher potentials needed to eliminate the driving force for excitation with our potassium-based solution. We therefore turned to a solution based on cesium gluconate (CsGlu; Appendix Table A.1). On all cells recorded, currents were consistently outward, rippled and of large magnitude (Fig. 2.18).

2.4 Pharmacological Isolation of Excitatory Currents

The results presented so far support the statement that excitatory currents, and hence presynaptic spiking activity of pyramidal cells, is rhythmic with ripple frequency and well coordinated with the extracellular ripple. But we may enquire if electrotonic isolation of e.g. excitatory currents by voltage-clamping the cell at its soma is enough to warrant such conclusions, or if we are rather seeing currents that are inward, yes, but correspond to inhibitory inputs in disguise. Two reasonable objections can be raised to the power of electronic isolation alone:

- voltage-clamp of the soma does not ensure control of the membrane potential over the whole membrane. In an spatially extended hippocampal CA1 pyramidal cell, control over the farthest branches of the dendritic tree is likely to be deficient. If the dendrites of a cell with reversal potential of inhibition around -67 mV are not held at that potential but left somewhat free e.g. at a lower intermediate potential closer to the resting potential, a driving force will arise that in the presence of high enough inhibitory conductances will result in inward currents activated by inhibitory presynaptic partners.
- individual cells might deviate from the statistically determined reversal potential (Fig. 2.8). The holding potential will be off from the real reversal and thus a driving force may arise for inhibitory currents.

To further experimentally corroborate our hypothesis of the existence of ripple-coherent excitatory PSCs, we sought to directly investigate excitation during ripples by blocking inhibition pharmacologically. First, we addressed the question whether the SWR-associated currents remained inward and ripple-rhythmic when inhibition was blocked pharmacologically. We then went to the current clamp configuration to confirm that pyramidal cells could be fired by excitatory currents with ripple-coordinated timing.

2.4.1 DIDS Block of Inhibition for Voltage Clamp Experiments

Bath-application of GABA_A receptor antagonists is experimentally inappropriate because they not only block inhibitory PSCs, but also disrupt SWRs as a collective network phenomenon (Ellender et al., 2010; Maier et al., 2003; Nimrich et al., 2005). We therefore blocked GABAergic synaptic inputs at the single-cell level by perfusing cells intracellularly with a pipette solution containing 4,4'-diisothiocyanostilbene-2,2'-disulfonic acid (DIDS) and cesium fluoride (CsF; the detailed composition is available in the Section A.2 of the Appendix). This approach was first described in Nelson et al. (1994) where it was applied both in vivo and in vitro to rat visual cortical neurons. Inhibitory currents were largely blocked in the whole range of tested voltages between -90 and 20 mV. Henceforth the condition using CsF-DIDS will be referred to in short as *DIDS condition*.

To demonstrate the reliability of DIDS recordings, we first tested the mechanics of the re-patch procedure. Employing the same pipette solution as before in Section 2.3.2 (“control”) in two successive whole-cell patches of the same cell, we verified that synaptic stimulation produced the same inhibitory postsynaptic current (Fig. 2.19 A). We then set up a new re-patch experiment, this time using CsF-DIDS for the second patch. Inhibitory PSCs evoked by fiber stimulation were successfully blocked when cells were perfused with the DIDS-based pipette solution (Fig. 2.19 B; antagonists of AMPAR- and NMDAR-mediated synaptic transmission applied in the bath).

We also recorded currents evoked by UV-flash-triggered photolysis of ‘caged’ GABA with control solution. Following re-patching of the same cells with CsF-DIDS and repeated ‘un-caging’ of GABA, we indeed observed blockade of postsynaptic GABA currents (Fig. 2.19 C).

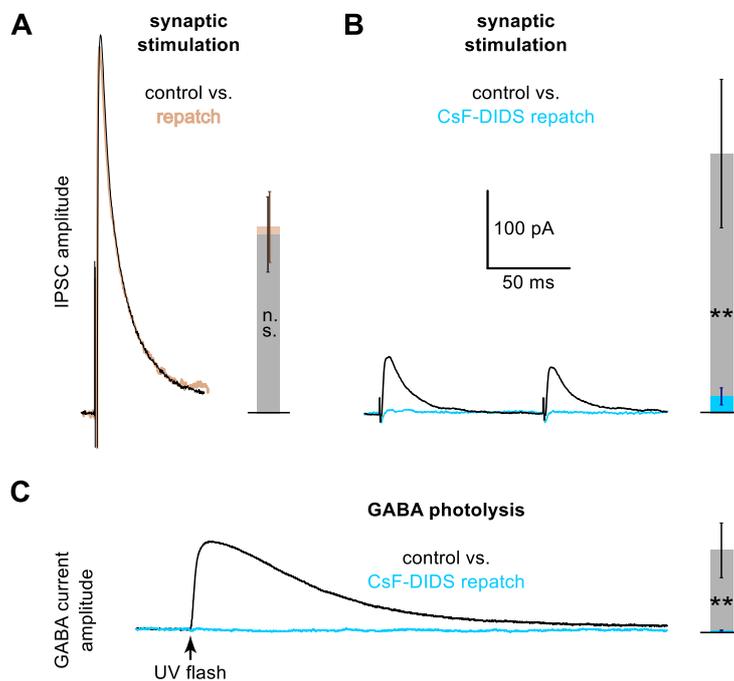


Figure 2.19. Intracellular Blockade of Inhibition. **A.** Control experiment showing the reliability of re-patch recordings. Short electrical pulses were extracellularly delivered to elicit IPSCs; their amplitudes were compared before (black) and after (brown) re-patching the cells with identical intracellular solution (20 μM NBQX and 50 μM D-APV in the bath solution). Bars show insignificant difference of IPSC average amplitudes of first patch and re-patch. **B.** Stimulus-evoked IPSCs are blocked after intracellular perfusion with CsF-DIDS. Control (black): 329 ± 95 pA; re-patch (blue): 21 ± 11 pA ($n = 6$; $P = 0.002$, rank-sum test). Cells were clamped at -59 mV. **C.** Postsynaptic current after GABA uncaging by a brief UV-flash without (black; control) and with intracellular blockade of inhibition (blue; 1 mM CsF-DIDS perfusion; 10 μM NBQX and 50 μM D-APV in extracellular medium; clamping potential, -59 mV). Control: 104 ± 34 pA; re-patch: 3.2 ± 1.3 pA ($n = 5$, $P = 0.008$, rank-sum test).

Subsequently, we used this approach to test our hypothesis of oscillation-locked excitatory inputs during ripples. We applied CsF-DIDS in re-patches of seven cells after having collected a sufficient number of ripple-associated cPSCs under control conditions close to the potential of Cl^- -reversal. In line with our hypothesis, ripple-

associated fast synaptic inputs indeed persisted in the re-patch recording with disrupted GABA_A -R mediated synaptic transmission (Fig. 2.20 at -45 mV, Fig. 2.21 at -60 mV, Fig. 2.22 comparing an example event at -60 mV with one recorded in re-patch using control solution K Glu 1).

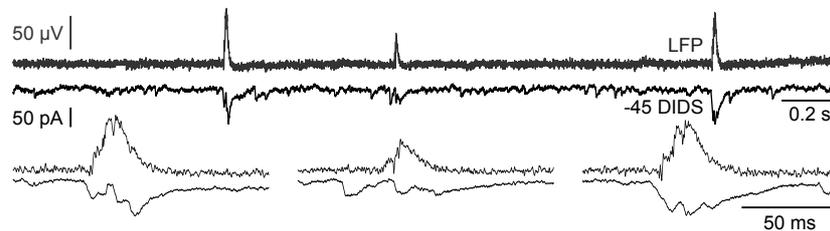


Figure 2.20. Example DIDS -45 mV. Three-second sections of simultaneous LFP (top) and current recordings at -45 mV holding potential (below) with intracellular solution CsF-DIDS. SWR-associated currents are systematically inward and rippled as in the three zoomed examples at bottom.

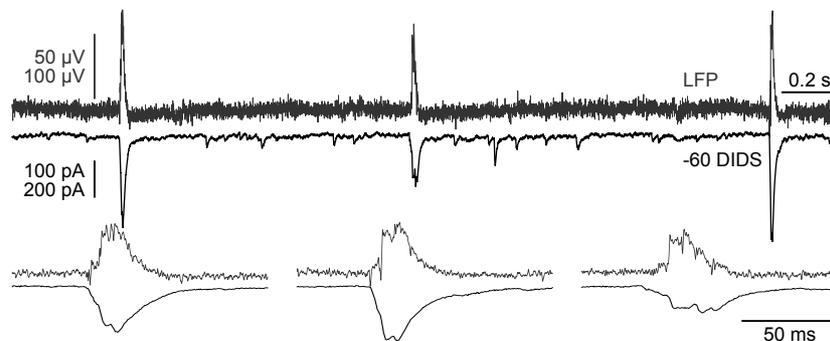


Figure 2.21. Example DIDS -60 mV. Three-second sections of simultaneous LFP (top) and current recordings at -60 mV holding potential (below) with intracellular solution CsF-DIDS. SWR-associated currents are systematically inward and rippled as in the three zoomed examples at bottom.

It is likely that voltage clamp at the reversal potential of inhibition can alone reduce or eliminate most inhibitory synaptic currents in our model. Furthermore, the prevailing line of argument in the literature for inhibition as the main organizer of the ripple always pointed at perisomatic targetting cells, and especially basket cells, as the main candidate for the orchestration of the fast oscillations (Ylinen et al., 1995; Klausberger et al., 2003; Taxidis et al., 2011). Indeed, ultrastructural anatomical studies in rat CA1 (Megas et al., 2001) have established that symmetric synapses characteristic of inhibitory transmission dominate at and proximal to the soma both in the axon initial segment and thick proximal apical dendrites (40% of inhibition is perisomatic; but inhibitory synapses are, in total, only 3% by number; afference to soma is 98% inhibitory, and to proximal basal dendrites, 48%), whereas the afference to more distal radiatum and to oriens (main targets of Schaffer collaterals) is predominantly onto dendritic spines and excitatory. Although the fraction of inhibitory synapses over the total in lacunosum-moleculare (sometimes to spines) increases again to 14-17% (Megas et al., 2001, Fig. 2), the layer receives mostly entorhinal or sub-cortical inputs which are absent in our hippocampal slice.

In view of the above, the majority of inhibitory inputs, at least as judged from number of synaptic contacts, would be well controlled by voltage-clamping the cell's soma at the reversal potential of GABA_AR-mediated transmission.

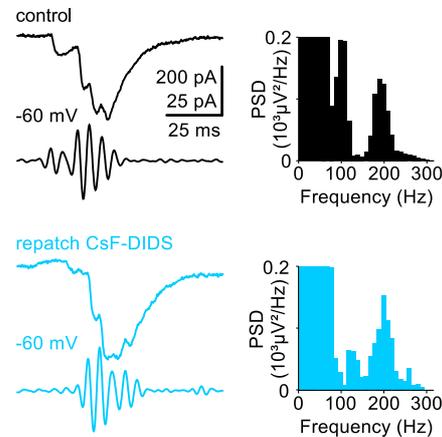


Figure 2.22. Intracellular Blockade of Inhibition Reveals Strong Excitatory Contribution to cPSCs. **Left.** Ripple-associated PSCs in CA1 principal cells remain stable (>15 min) following intracellular blockade of inhibition. Top: cPSCs in control (black) and re-patched with CsF-DIDS (blue). Bottom: band pass-filtered (120–300 Hz) cPSCs. **Right.** PSDs of the individual events at left.

Further, if far-out dendrites would be making relevant inhibitory contributions, one would expect them more attenuated than input close to the soma (unless larger distal conductances compensate for that, see Magee and Cook, 2000) and significantly slower in their kinetics (bar compensation mechanisms e.g. via an inhomogeneous distribution of dendritic ion channels Magee, 1999; in our voltage clamp data, active conductances do not play a role, however). If the decay time constant of IPSCs extends much beyond the typical cycle length of the ripple, at ~ 5 ms, their ability to pattern the ripple-associated currents and thus precisely time spikes will be severely challenged. Those hypothetical inhibitory dendritic inputs should also stand out as slow spontaneous PSCs in the electrotonically controlled condition and be missing in recordings with additional intracellular pharmacological block, a situation that we never observed. We will show that IPSC kinetics during ripples can be as fast as those of EPSCs and that only a few events have decays beyond 10 ms (Fig. 3.26).

In the next Chapter we will considerably extend our analysis of currents in the DIDS condition and show that their main characteristics are retained when the pharmacological block is absent if the driving force for inhibition is weak.

2.4.2 DNDS Block of Inhibition for Current Clamp Experiments

To further substantiate our finding that spikes can be timed by rippled excitatory currents, we also analyzed action potential timing during blockade of inhibition at the single-cell level. For that, we applied DNDS, which blocks GABA_AR-mediated inhibition from the intracellular side without changing action potential firing, fol-

lowing Dudek and Friedlander, 1996 (Figs. 2.23 and 2.24). Using this approach we found that action potentials were locked to ripples (Figs. 2.25 and 2.26; $\log_{10} P$ values between -54.3 and -1.6 ; $n = 1,119$ spikes associated with 1,564 SWRs; 7 cells). IPSCs appeared as inward in this experiment, because the reversal potential was set to about -35 mV in order to have a small, but inward, driving force when recording APs (DNDS seemed to work best when the Cl current was inward). To allow for the cell to express spikes, V_{clamp} was set in the range from -65 to -45 mV.

Together, these experiments demonstrate that the ripple-locked excitatory inputs remaining after block of inhibition can effectively regulate the spike timing of target principal neurons.

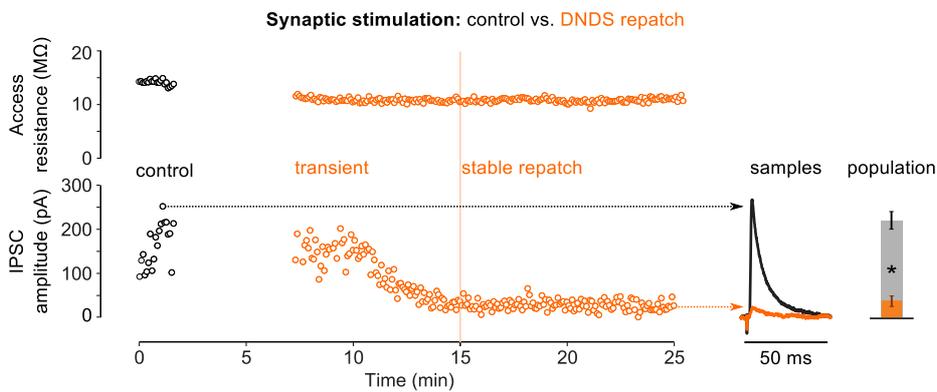


Figure 2.24. DNDS Effectively Blocks Stimulus-Induced Fast IPSCs. **Left.** Representative experiment to assess the potency of DNDS to block IPSCs. Upon a brief control period (black) in which IPSCs were elicited with synaptic stimulation the electrode was withdrawn (AMPA, NMDA, and GABA_B receptors blocked). The same cell was re-patched (orange) with pipette solution supplemented with 500 μM DNDS. After $\sim 5 - 10$ min IPSCs gradually decreased in amplitude (bottom panel); the cell's access resistance was monitored during the experiment (top). **Right.** Sample traces and population statistics comparing control IPSC amplitudes with those under DNDS after 15 minutes or longer (stable repatch): 234.9 ± 21.4 pA vs. 37.9 ± 12.9 pA respectively ($P = 0.0006$, two-tailed rank-sum test; $n = 7$ repatches).

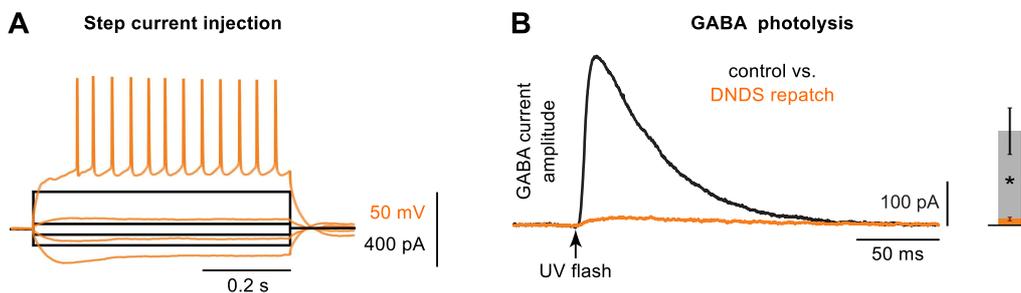


Figure 2.23. DNDS Suppresses GABA_A Currents but Does Not Affect Spiking. **A.** Input/response example in a CA1 pyramidal cell perfused with DNDS-based K-Gluconate solution. De- and hyperpolarizing current steps were applied at $+280$ pA, ± 40 pA and -120 pA. **B.** GABA uncaging experiments. UV flash-evoked GABA currents were compared in the absence (black) and presence (orange) of 500 μM DNDS. On average, DNDS reduced GABA currents by $93.8 \pm 3.5\%$ (control: 210.4 ± 52.4 pA vs. DNDS: 12.4 ± 4.3 pA; $P = 0.008$, two-tailed ranksum test; $n = 5$ repatch experiments).

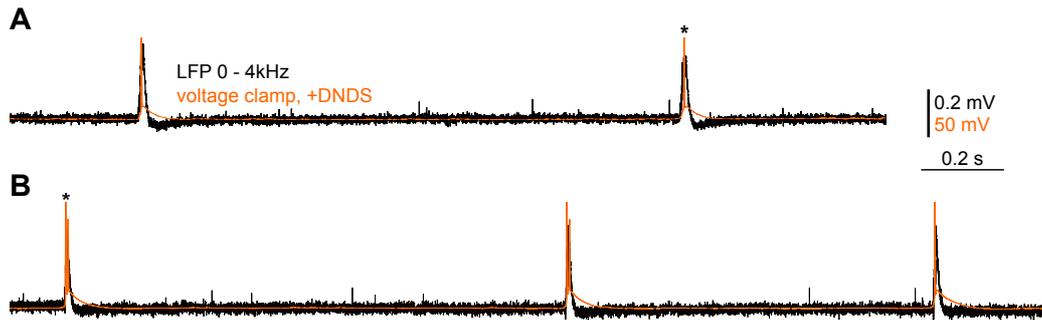


Figure 2.25. Current Clamp Recordings of Two CA1 Pyramidal Neurons With DNDS. Cells were held at slight depolarization to enhance firing probability. Representative examples from two cells (A and B) display single or double spikes during SWRs. Marked SWRs (asterisk) are shown larger in Fig. 2.26.

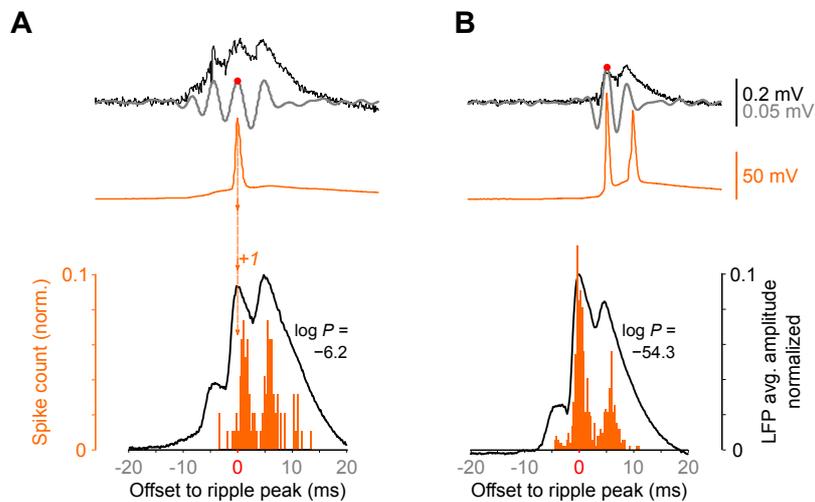


Figure 2.26. Spiking is Locked to Ripples Under Intracellular Block of Inhibition. Example SWRs from Fig. 2.25 and spike-to-LFP locking statistics from the two cells therein. **Top.** Wide-band LFP (black) superimposed with ripple band-filtered LFP (grey) and LFP ripple maximum (red disc); below: intracellular voltage trace. **Bottom.** Spike-time histograms (orange) of the two example cells demonstrate phase-locking of spikes and LFP SWRs (black; averaged and normalized). The $\log_{10}P$ values indicate the significance levels determined by the Rayleigh test for the uniformity of a circular distribution. Logarithms of significances in five other recordings were -1.6 , -1.7 , -2.6 , -11.4 and -17.1 . The vector strength when pooling all spikes across cells was 0.38 with phase 15° and $\log_{10}P = -76$. See also Figs. 2.15 and 2.16.

2.5 SWR in a CA1 Minislice

Several studies (Maier et al., 2003; Nimrich et al., 2005) have shown ripples in the extracellular potential measured in a deafferented slice of CA1 mouse hippocampus. To complement these studies we asked whether ripple-coherent cPSCs represent the spiking output of CA3 pyramidal neurons (Both et al., 2008) or whether they are generated locally within the CA1 network. We used *minislices* where area CA1 was isolated from the adjacent CA3 and subiculum (Fig. 2.27 A, C). In eight CA1 minislices, SWRs occurred at a rate of 0.46 ± 0.09 Hz (median: 0.46 Hz; range: 0.13

Hz to 0.93 Hz; Fig. 2.27 B). The frequency of the peak power in the ripple frequency bands across these events was 213.1 ± 6.6 Hz on average (median: 215 Hz; range: 175 Hz to 235 Hz; Fig. 2.27 B, PSD at right).

	Figure	Sample size	SWR incidence (Hz)	Ripple frequency (Hz)
in vivo	2.2	16 mice	0.14 ± 0.02	136.3 ± 0.2 Hz
full slice	2.9	28 slices	0.77 ± 0.05	194 ± 6 Hz
minislice	2.27	8 slices	0.46 ± 0.09	213.1 ± 6.6 Hz.

Table 2.3. Comparison of SWR Incidence and Ripple Frequency. Statistics of SWR incidence and extracellular ripple frequency in our three experimental models for hippocampal ripples in CA1.

To test whether ripple-coherent cPSCs survived in the isolated area CA1, we again recorded from principal neurons voltage-clamped close to the reversal potential of Cl^- (-66 mV). SWRs in CA1 minislices were indeed accompanied by phasic inward currents at ripple frequency coherent with the SWR (Fig. 2.27 D, 725 cPSCs from 5 cells).

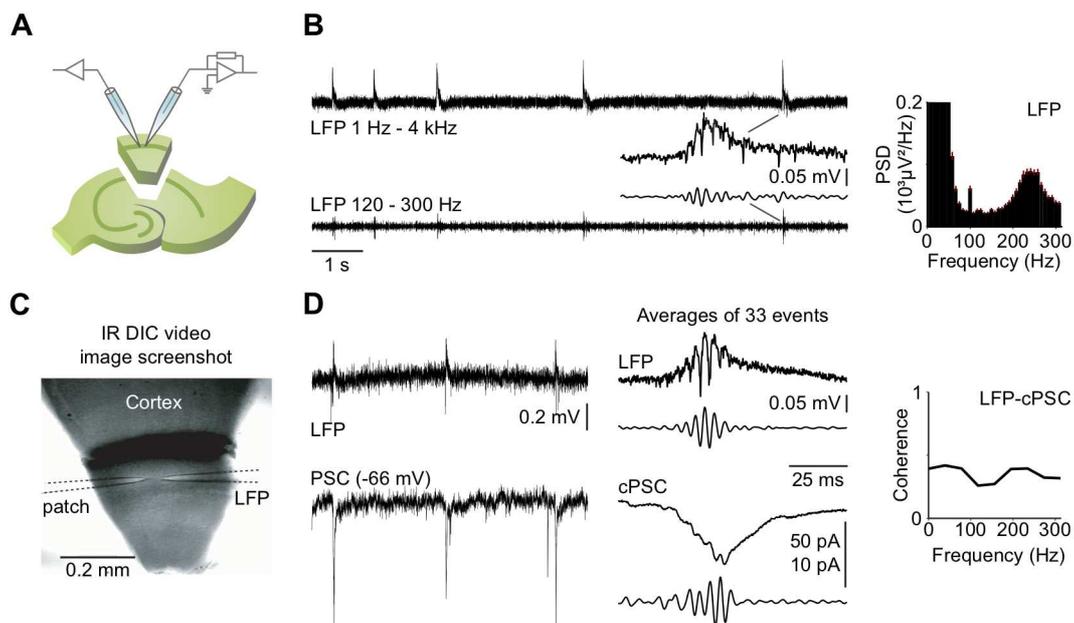


Figure 2.27. Ripple-Coherent Excitatory Currents Persist in CA1 Minislices. **A.** Schematic of the CA1 minislice approach. **B.** (Left) CA1 minislices do generate SWRs at slow pace. Top: SWRs recorded in *stratum pyramidale*. Middle: magnified event. Bottom: Filtered version of the above, as indicated. (Right) Averaged PSD plot from this experiment (153 individual spectra). Note the peak at ripple frequency. **C.** (Left) In a different CA1 minislice, PSCs were recorded from a principal neuron in parallel with the LFP. The video image screenshot displays the CA1 minislice and the electrode arrangement. **D.** (Left) Example LFP-SWRs and associated PSCs. Middle: 33 LFP- and PSC events (1st and 3rd from top) were averaged triggered on the LFP ripple maximum. Respective filtered versions below. (Right) Coherence of LFP and cPSCs from this experiment reveals a peak at ripple frequency (~ 200 Hz).

Part III

Results

Chapter 3

On Sharp-Wave Ripples

In this Chapter we present results on ripple-associated excitatory and inhibitory currents in vitro. We refer the reader to the previous Chapter 2 for a description of the electrophysiological techniques used to obtain the data as well as for a first acquaintance with the measurements. The analytical methods used are introduced right next to the results, because often both are delicately interdependent.

The reader in search of a quick, graphical summary is advised to consult directly Figures 3.16 (PSCs lock to LFP ripple), 3.21 (in-ripple PSCs are like spontaneous), 3.22/3.23 (peeling reconstruction in voltage clamp works like a spike detector), 3.27/3.28 (PSC amplitude depends on ripple cycle) and 3.31 (inhibition aligns with excitation over the course of ripples).

We start out the Chapter in **Section 3.1** by presenting our amplitude-based SWR detection algorithm. Such a detection strategy is only possible because our in vitro recordings exhibit high signal-to-noise ratio, and a rather stable baseline.

The extraordinary regularity of ripple-associated currents makes it possible to assess how much a cell participates of the ripple phenomenon at the network level (or how good is our experimental access to it), simply by averaging the compound post-synaptic currents (cPSCs) triggered to the SWR peak. Another relatively simple analysis that does not require delving into the fine structure of the oscillation is an average of deconvolutions (a hybrid between averaging amplitudes and averaging derivatives that will be introduced below). With these two analyses we obtain the overall polarity of amplitudes and slopes (**Section 3.2**), thus helping quantify the impressions from the data snapshots that introduced the raw datasets to the reader in Figures 2.9, 2.17, 2.18, 2.20, 2.21 and 2.27 of the previous Chapter. The triggered averages supply also a first rough indication of the timing relationship between excitation and inhibition.

Key to our understanding of in-ripple post-synaptic currents (PSCs) is an evaluation of the kinetics of the out-of-ripple events. We often will refer to those, with slight abuse of language, as *spontaneous* PSCs or sPSCs. **Section 3.3** discusses their properties, which are used to inform the reconstruction approach later in Section 3.6.

The realization that spontaneous PSCs rise much faster than they decay naturally suggests to use the steepest slopes of the ripple-associated currents as indicators of the onset of PSCs. While only approximate, and biased in ways that will be made apparent below, this technique makes few assumptions and is adequate for the assessment of the rhythmicity in the intracellular current, providing a complementary view to that of spectral analyses. In contrast to spectral techniques, which weigh ripple events in averages according to their amplitude, local slope statistics treat large and small cPSCs on an equal footing. Also in **Section 3.4** we obtain cell-wise evidence that the E/IPSC isolation approaches translate into steeper downward resp. upward slopes, thus supporting our confidence on the successful electrotonic and pharmacological isolation of excitatory and inhibitory conductances.

So far we have dwelt on the intracellularly measured signal, the compound post-synaptic current. How does it coordinate with that other thermometer of network activity that is the LFP? In **Section 3.5** we start with traditional spectral methods by examining the **coherence** of both signals. Next we adopt the steepest slopes in a cPSC as proxies for PSC onsets. This allows to produce two locking measures: “onset” (steep slope)-triggered averages of the extracellular SWR, and vector strength of such onsets in the ripple.

Having a well-established prior about the time course of single, spontaneous PSCs we ask next (**Section 3.6**) if it can be used to reveal the individual synaptic contributions that make up a cPSC. A number of assumptions are needed for the resulting iterative procedure, which *peels* off earlier events in the cPSC to uncover later arrivals. This approach allows one to obtain onsets, amplitudes and the rise and decay kinetics of individual PSCs in the ripple. In the final **Section (3.7)** we exploit the peeling reconstruction method just presented to obtain raster plots of presynaptic spikes and histograms of their rhythmic arrival, enhancing the results based on steep slopes. We confirm that their kinetics match those observed in off-ripple events. An unsolved question about the generation of SWRs concerns the relative timing of excitatory and inhibitory spikes during the ripple. We show evidence, based on our *in vitro* data that, during the ripple, there is a progressive synchronization of the inhibitory and excitatory phases. We further ask how the PSC features change over the course of the ripple. The resulting data on ripple-cycle dependence of amplitudes will inform an estimation of the excitatory and inhibitory conductances, which is a key ingredient of a continuous-time model for sequence replay. Such an extension of the discrete-time model that we present in **Chapter 4** would allow to link the observed currents to the expected spiking dynamics.

Recording conditions and cell coloring Recording conditions are listed in Table 2.1 (the composition of the solutions used is available in Appendix A). They are grouped into datasets that we will informally designate excitatory, DIDS, minislice, and inhibitory (Table 3.1). Cells are numbered from 0 to 33; missing numbers correspond to failed recordings, e.g. the LFP SWR signal was absent or other experimental or storage errors. In order to provide visual help for the reader comparing results from different analyses, we have assigned a unique color to each cell within a recording condition, both in text and graphics.

To complement the plots presented in this Chapter, Table 3.1 summarizes the numbers of events detected by cell and condition and subsequently used for analysis.

3.1 Amplitude-Based Large-Event Detection

We describe here the method used to detect SWRs. Because it is generally applicable to the detection of large signals, and especially cPSCs, we describe it as a large-event detection algorithm despite it being applied here only on the extracellular trace.

Detection of events was on the extracellular trace Across this work the SWR was taken as the temporal framework against which the currents were examined. Although the extracellular SWR went almost always hand in hand with prominent complex currents, so that we could have used either for detection of the combined event, we relied

excitatory		Solution	Holding V	0	1	2	3	4	5	6	7
	KGlu 1	-66 mV		143	102	201	69	208	116	136	110
		-45 mV		142	211	193	113	157	120	154	129

DIDS		Solution	Holding V	8	9	10	11	13	14	15
	Control	-45 mV		27	19	42	×	12	×	28
		-60 mV		51	34	36	34	19	11	39
	DIDS	-45 mV		96	32	53	×	31	×	119
		-60 mV		94	90	53	101	51	37	36

minislice		Solution	Holding V	16	19	20	21	22
	KGlu 1	-66 mV		54	80	232	124	71

inhibitory		Solution	Holding V	27	28	29	30	31	33
	Cs Glu	-6 mV		42	118	138	198	230	135

Table 3.1. Numbers of SWR Events Detected per Cell and Condition. Nicknames for the overall datasets are given in sans serif for each table. Cell pairs that were recorded from the same slice (consecutively, the LFP electrode staying in place) share background color in the table header, e.g. 10 and 11. A cross indicated that no recording was made for that condition.

exclusively on detection of the SWR to locate the cPSC, and we referred all in-ripple times in both electrodes to the SWR peak.

SWRs were selected using an amplitude-based criterion SWR detection was performed on 4-100 Hz band pass-filtered extracellular traces (2nd order zero-phase, acausal Butterworth filter). Their amplitudes were tallied, and the histogram was fitted with a Gaussian that was dominated by the eventless epochs of small amplitude. The tails provided us with an expected frequency of rare amplitudes. We found the threshold as the lowest amplitude which appeared 100 times more often than expected from the Gaussian fit of amplitudes. Two thresholds were at hand — one for negative, another for positive rare amplitudes. Exploiting the often biphasic character of SWRs, both were used to detect events. This feature is less interesting in the version finally selected, that incorporates a somewhat aggressive high-pass of 4Hz because compensation (to achieve zero mean) of any peak will create artificial surrounding negative peaks. The maximum in a (-75, 75) ms “search window” around each over-threshold point was selected as representative for it, and the total time over any of the thresholds in that range was stored. Only maxima with at least 7 ms of over-threshold amplitudes around them were retained.

SWRs were characterized by the time of the peak As a last step, the characteristic “SWR peak” reference time was obtained from the unfiltered trace. This made comparable the results of different algorithms e.g. using different filter bands, and

offered a very precise reference point. However, as we discovered on later data once the algorithm was tested on an initial dataset, it had the drawback of not being robust against extremely high-frequency amplitude blips (3-5kHz, 1 or 2 samples generally). This is cured by instead selecting maxima after low-passing at 1 kHz. Such an edge frequency still allows for a precise enough identification of the peak for ~ 200 Hz phenomena and is more robust. To confirm we reran our analyses with this setting and obtained a slight quantitative improvement, as a few spurious blips were now correctly ignored. However, the figures presented here correspond to the unfiltered maximum finding.

Events were checked manually We examined the raw traces with the marked events as well as scatterplots of cPSC amplitudes concurrent to SWR amplitudes and detected two cells in the minislice dataset where the low signal to noise ratio would have demanded specialized detection settings. For them the histogram of SWR amplitudes is clearly bimodal. Upon analyzing the low-amplitude SWRs we concluded that they were fake events i.e. the LFP was not rippled, and removed those by hand.^{3.1}

Results are lower bounds The algorithm outlined above strives to avoid a bias towards high-amplitude SWR. Since a roughly linear correlation between SWR amplitude and concurrent cPSC amplitude is observed (not shown) our results would certainly be embellished by such a selection bias. We preferred, however, to run the risk of washing out some features rather than biasing our sample of SWRs to stress the point. Most of the metrics of rhythmicity and locking in this work are expected to improve with a more lax approach that includes only the high-amplitude SWR and thus the predominantly high-amplitude associated cPSCs.

3.2 Average Shapes of CPSCs Triggered to SWR Peaks

This Section provides an overview of the polarity, size and ripple modulation of the currents through their cell-wise amplitude and deconvolution averages triggered to the SWR peak in the LFP. Both facilitate interpretation of later, more derivative analyses, by enabling to discern what features are particular to one cell or to one condition, and what is common across them.

3.2.1 Average Amplitudes

Figures 3.1 and 3.2 (for DIDS-related conditions) display the averages of traces, cell by cell for each condition separately.

There is considerable cell-to-cell variability Condition -66 KGlu (full slice) at top left in Figure 3.1 shows already to what extent the current amplitudes are variable across cells; note that it is genuinely a cell effect and not a slice effect: while both cell 1 and cell 2 share slice and show similar (low) amplitudes, another pair of cells that share slice (5 and 6) show starkly differing amplitudes. Also the polarity of the

3.1. Cells affected by low-amplitude false alarm SWR clusters are number 19 and 22. We used manual cutoffs of $200 \mu\text{V}$, which filters out 122 of the 222 candidates from cell 19, and $35 \mu\text{V}$ which rejects 45 of 118 candidates initially detected for cell 22).

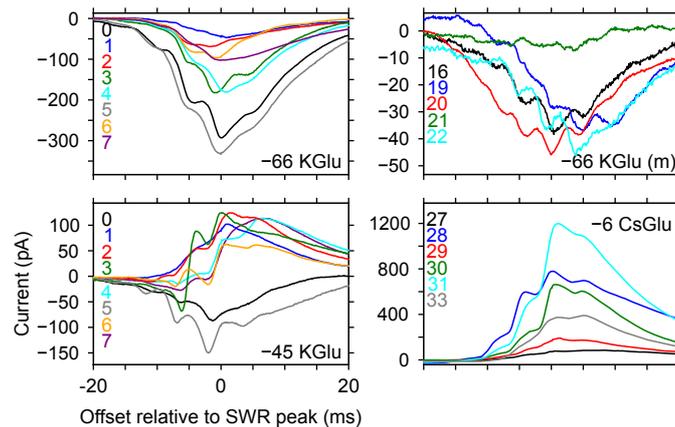


Figure 3.1. SWR Peak-Triggered CPSC Averages. Conditions are indicated in each panel.

currents reacts differently to a change in the holding potential: whereas cell 5 maintains an average inward current profile at -45 mV, the average cPSC of 6 betrays predominantly outward currents. Since cells were phenotypically checked for firing characteristics of pyramidal cells and sometimes reconstructed via biocytin filling, and since the pyramidal layer in the mouse contains about 3 or 4 densely packed cell bodies, it is likely that this variability is intrinsic to pyramidal cells and not due to e.g. a misplacement of the electrode or patching an interneuron.

Excitation hits first during the ripple Both -45 KGLu conditions (cells 0 to 7 and DIDS control recordings of cells 8 to 15 in Figure 3.2) illustrate that responses to an elevation of the holding potential are cell dependent: some maintain an inward polarity, some go over to strictly outward currents and finally a subset of *mixed* cells show time-dependent polarity (3, 4, 7 and 15). The details of this behavior are informative of the evolution of the balance of excitation to inhibition during the ripple. Mixed cells exhibit always first inward currents, followed by outward currents, both in the average and in all of the many individual cPSCs that we examined by eye. A quantitative evaluation of the ongoing balance of excitation and inhibition from this data is not possible without further, cell-wise measurements of reversal potentials (including those of AMPA-R currents) and conductances.

Inhibitory currents are large and tail off more slowly In the -6 CsGlu condition we observe particularly large currents, peaking at around $1 \mu\text{A}$. Their average decays, although no more modulated than those recorded at -66 KGLu are slower, pointing at a perhaps longer decay time constant of gabaergic conductances (Section 3.3). Note that recordings at -45 mV for cells that reverse polarity totally or partially also show this slow tail-off, which suggests that the decay of those cPSCs is dominated by inhibitory PSCs with their perhaps slower kinetics.

Currents in minislices are appreciably smaller SWRs observed in the minislice condition -66 KGLu (*m*) have small amplitudes, which complicates extraction from the background signal (see Fig. 2.27 in Chapter 2 for examples and Section 3.1 above for discussion). The average currents in the minislice cells 16 to 22 are shallower as well, and their signal to noise ratio lower as evinced by the jittery averages at similar number of events (compare cell 20 with 232 vs. cell 0 at -66 mV with 143 events, see Table 3.1).

DIDS succeeds at blocking outward currents Compare the panels -45 DIDS and -60 DIDS with their respective controls (Fig. 3.2). Upon raising the holding potential the driving force for excitation is reduced. This results in the lower current amplitudes visible at -45 DIDS ($\sim \div 2$). But currents stay inward. In the control recordings only cells 10 and 13 stay inward, with reduced amplitudes. All the others present outward currents.

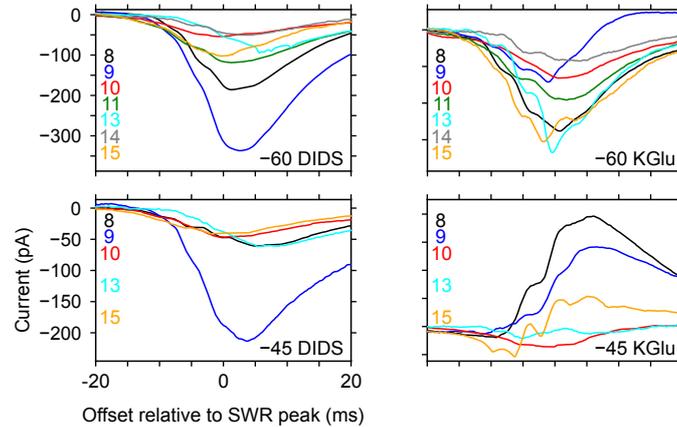


Figure 3.2. SWR Peak-Triggered CPSC Averages (DIDS Condition and its Control). Conditions are indicated in each panel. Cells 11 and 14 were not recorded at -45 mV.

3.2.2 Average Deconvolutions

Here we observe the ripple modulation of the currents adding the refinement of a sharpened view of the oscillation as offered by deconvolution with an exponential kernel (see Section 3.6 for an in-depth discussion and application of the idea).

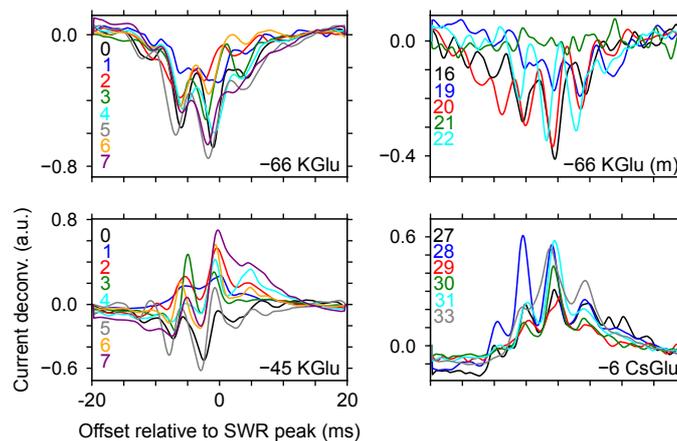


Figure 3.3. SWR Peak-Triggered Deconvolution Averages.

Why deconvolution? From a practical point of view deconvolution with an exponentially decaying causal kernel permits to sharpen the time expression of phenomena that have a fast rise and to whose decay an exponential function can be fitted with some success. This class of phenomena includes single synaptic events measured postsynaptically, which, we believe, make up the cPSCs that we observe.

Such an exponential deconvolution (short for deconvolution with an exponential causal kernel) is nothing else than a weighted sum of the amplitude of the signal and its derivative: $\text{deconv}(s) = s + \tau ds/dt$. The weight of the derivative τ is the assumed typical decay constant of the signal components (PSCs). Since we could expect PSCs to have different characteristic decay constants depending on whether they are excitatory or inhibitory, what cell they are recorded in or even the phase of the ripple oscillation at which they arrive, any chosen τ for such a large set is a compromise. Here (Figs. 3.3 and 3.4 below) we have used $\tau = 5$ ms.

If τ is large, the deconvolution is dominated by the derivative, which is noisy because it emphasizes variations in the low and in the high-amplitudes equally. If τ is very small, the deconvolution is just the original signal. In between, using the an adequate value of τ has the effect of *sharpening* the signal without amplifying too much the noise. If the spikes that generate PSCs would be instantaneous, and PSCs would be decaying exponentials with time constant τ , deconvolution would recover the presynaptic spike train from the postsynaptic superposition of PSCs. To the extent that all three assumptions are only partially accurate (a spike lasts for an appreciable 1-2 ms, PSCs are not simple exponentials and we do not know their exact kinetics, which furthermore, vary from one to the next), the result is a collection of peaks indicative of the times of the presynaptic spikes. Those peaks are more spikey than in the unprocessed cPSC, facilitating the analysis of timing that we envisage here.

The averages in the Figures are preceded by filtering and normalization. Smoothing the high-frequency oscillations (>400 Hz) eliminates jitter in the derivative due to amplitude noise, and does not compromise our analysis which is centered on the sub-250 Hz part of the spectrum. Normalization places putative synaptic events (here revealed by peaks) from large and small events on the same footing before the averaging. Overall, the deconvolutions succeed generally at reducing the support of the cPSC and turn the putative synaptic events into rather symmetric peaks, thus confirming the order of magnitude of $\tau \simeq 5$ ms as representative of decay times.

DIDS ripple modulation The Figures emphasize the 5 ms rhythmicity of the cPSCs. Despite the fact that cPSCs recorded under DIDS are generally less well modulated (compare amplitude average in Fig. 3.2; but see also non-DIDS cells 1, 7 and 33) deconvolution reveals that they are ripple-rhythmic as well.

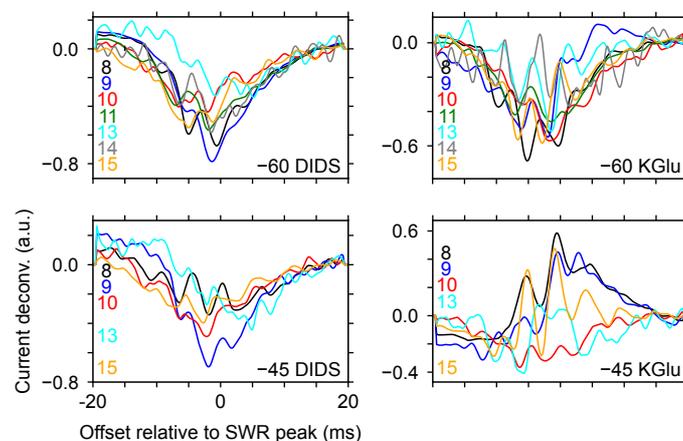


Figure 3.4. SWR Peak-Triggered Deconvolution Averages (DIDS Condition and its Control). See Fig. 3.3 for the processing steps. Cells 11 and 14 were not recorded at -45 mV.

Timing across cells Sharper peaks means easier visual inspection of oscillatory coordination. Jitter is specially low for outward currents in the mixed (-45 mV) and inhibitory conditions (-6 mV); Figure 3.3.

Time of cPSC peak Inward currents reach their maxima briefly in anticipation of the SWR peak used for reference when averaging; outward current peaks seem slightly delayed, almost overlapping with the SWR peak. See Figure 3.28 (bottom row) for a confirmation of this observation at the level of individual PSCs.

3.3 Atomic Intracellular Events: Spontaneous PSCs

Before diving into the details of the complex postsynaptic currents that appear during SWR it is convenient to equip ourselves with some knowledge of the kinetics of the much simpler currents that dominate the inter-SWR time.

Kinetics of currents: from empirical fits to theoretical modelling The time development of postsynaptic currents and potentials is the result of a sequence of biological processes involving rather complex chemical pathways in both the presynaptic and postsynaptic sites, diffusion in the cleft and other associated processes. Figure 3.5 sketches the transduction process at a synaptic site and highlights the postsynaptic current (orange), which is the center of our interest when measuring in voltage clamp.

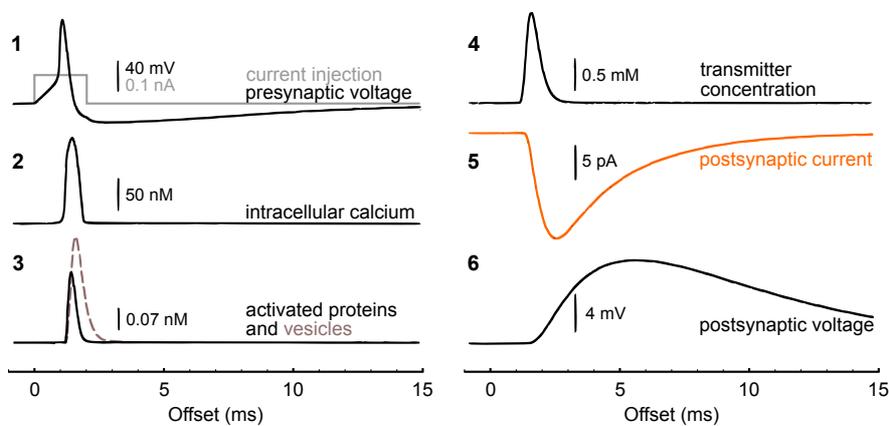


Figure 3.5. Simulation of Synaptic Kernel for AMPA Currents. Traces represent different stages (1-6) of the transduction process from a presynaptic spike to a postsynaptic depolarization via intermediate chemical signals for AMPAR-gated ion channels. **1.** Presynaptic current injection and action potential elicited by voltage-gated ion channels. **2.** Intracellular Ca^{2+} concentration rises in the presynaptic terminal with a high-threshold calcium current providing a transient calcium influx and an active calcium pump in charge of removal. **3.** Ca^{2+} ions bind to a protein and produce an active form (relative concentration shown by solid line) which promotes the exocytosis of synaptic vesicles (dashed line). **4.** Concentration of transmitter in the synaptic cleft. The transmitter binds to AMPA/kainate receptor-ionophor complexes and opens the associated ion channels. **5.** Postsynaptic current produced by the gating of those receptors according to a six-state kinetic scheme characterized from single channel recordings in the locust muscle ($E_{\text{rev}} = 0$ mV, $\bar{g}_{\text{AMPA}} = 0.1$ nS). **6.** Excitatory postsynaptic potential in response to transmitter release ($E_{\text{leak}} = -70$ mV). Adapted by permission from Destexhe et al. (1994). Please refer to their Fig. 4 and Eq. 25 for details of the model.

A function of the form $((t - t_0)/\tau) e^{1-(t-t_0)/\tau}$, $t > 0$ was originally proposed by Rall (1967) as an empirical fit to measured synaptic signals and has been widely used since for modeling currents in computational models of neuronal networks (Gerstner and Kistler, 2002). This function belongs in the general class of weighted sums of decaying exponentials. More recently those have been shown to constitute the general solution of Markov kinetic schemes for the activation of ion channels (when these are analytically solvable, i.e the transmitter release is pulsed in the case of synaptic conductances). These schemes posit time-independent probabilities of state transition for the gating variables that represent channels, yet in their full generality are rich enough that their descriptive power spans the range of voltage-gated, synaptic and neuromodulator-induced currents (Destexhe et al., 1994). The following slightly more general form of the original alpha function is a good empirical compromise between matching the waveform of spontaneous PSCs in our recordings and introducing too many parameters^{3,2}. It results from a second-order kinetic scheme with two closed forms of the receptor, far from saturation, and the announced simplification of instantaneous transmitter release (see Destexhe et al., 1994, Appendix B):

$$\alpha(t; A, t_0, \tau_r, \tau_d) = AN(\tau_r, \tau_d) \Theta(t - t_0) (e^{-(t-t_0)/\tau_d} - e^{-(t-t_0)/\tau_r}). \quad (3.1)$$

The factor $\Theta(t - t_0)$ is the Heaviside step function (0 for $t < t_0$, 1 elsewhere). The normalization factor is found from the constraint $\alpha(t_{\text{peak}}) = A$ to be:

$$N(\tau_r, \tau_d) = (\tau_r/\tau_d)^{\tau_r/(\tau_d-\tau_r)} - (\tau_r/\tau_d)^{\tau_d/(\tau_d-\tau_r)}.$$

The peak amplitude is then A . The time offset t_0 , also called *onset* or arrival time, represents the delay between the presynaptic spike and the start of the postsynaptic current stroke. The time course of the current is divided into a *rise phase* from baseline amplitude at the onset t_0 , to A at peak and a *decay phase* from peak amplitude back to zero as $t \rightarrow \infty$. The decay phase is well characterized by τ_d with little influence from τ_r , while the time of the rise depends more substantially on both τ_r and τ_d . The classic, one- τ alpha function is recovered in the limit $\tau_d \rightarrow \tau_r$ but it enforces a relation between rise and decay time that is too constraining and fits our sPSCs noticeably less well.

Identification of non SWR-associated PSCs and their kinetics EPSCs outside SWRs (“spontaneous EPSCs”) were detected among the -66 KGLu 1 cells as steep downward slopes (top 5% of all maximal slopes on the 0.5–400Hz filtered intracellular trace of each cell). A total of 1,000 events from 5 cells were postselected by eye (by Christian Leibold) to exclude those that are too small to be distinguished from noise and also to avoid multiple events where a second PSC arrives during the tail of the first.

The fit was performed with the double exponential from above (Eq. 3.1). The average time constants obtained from the fits were $\tau_r = 1.70 \pm 0.04$ ms and $\tau_d = 4.04 \pm 0.08$ ms. Here the PSC detection algorithm is based on slopes, which are influenced by both time constants. To highlight the separation of timescales used by the algorithm, we plot in Figure 3.6 B histograms of durations of rises and decays (20-80% and 80-20% of maximum amplitude) rather than time constants.

3.2. Usually, more sophisticated (i.e. parameter-rich) functions are fitted when the details of the synaptic interactions are the focus of research *and* the availability of a stimulation paradigm allows to obtain a good characterization of the current by averaging over many repetitions. Here we aim at fitting currents without stimulation, outside the SWR. Below (Section 3.6) the scenario becomes more adverse as we attempt fitting the overlapping SWR-associated PSC bursts.

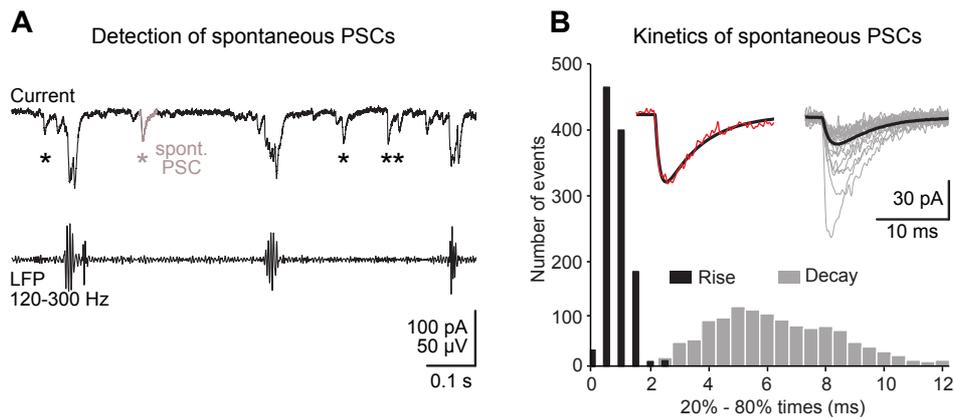


Figure 3.6. Kinetics of Spontaneous PSCs. **A.** Spontaneous, non ripple-associated PSCs (as those marked by asterisks in the upper right trace; highlighted example in brown) were extracted from currents measured in cells 0-7 at -66 mV. **B.** Histograms of 20%–80% rise- and 80%–20% decay times of these non-ripple PSCs (0.83 ± 0.01 ms and 6.66 ± 0.09 ms). *Left inset:* To determine rise and decay times of PSCs (red), a double exponential (black; Eq. 3.1) was fitted; *right inset:* example PSCs (grey) and their average (black).

These kinetics derived from spontaneous EPSCs (not associated with ripples) were fast enough to account for excitatory currents in ripple-associated cPSCs (Fig. 3.6 B). When comparing these values to the kinetics of AMPAR-mediated currents reported in the literature it is important to remember that our in vitro preparation was recorded at a relatively high 32°C , while the body temperature of mice oscillates between 37 and 39°C .

At an initial stage of this work we conducted a similar exercise on a preliminary dataset recorded from cells held at -40 mV whose spontaneous events were directed outward, resulting on comparable average rise kinetics and somewhat slower (1 to 2 ms) average decays.

3.4 Steep Slopes

In this Section we exploit in two ways the fact that PSCs are characterized by a steep onset phase followed by a gentler decay (see Fig. 3.6 for the separation of timescales). First, we reason that depending on the dominant type of input, upward or downward slopes should be the prevalent among the steepest; we verify that our putative excitatory currents indeed contain steeper downward than upward currents, as expected. Then we take advantage that the initial sharp deflection of a PSC can be used as a proxy for the onset itself to construct a rhythmicity measure that is not sensible to amplitude, by contrast to spectral techniques. We test that again on our K_{Glu} 1 and DIDS data.

3.4.1 Comparison of Upward and Downward Steepest Slopes

Figure 3.7 shows two extreme scenarios for incoming currents during ripples that could result in similarly looking but actually different shapes of the current observed at an intermediate reversal potential between that of excitation and that of inhibition. At potentials below the reversal potential of excitatory synaptic transmission, excitatory

(inward) currents within cPSCs should display downward slopes (onsets) steeper than their upward slopes (decays). At the same time, if the potential is above the reversal potential of inhibitory GABAergic transmission, putative *inhibitory* (outward) currents should display steeper upward slopes (onsets) followed by shallower downward slopes (decays).

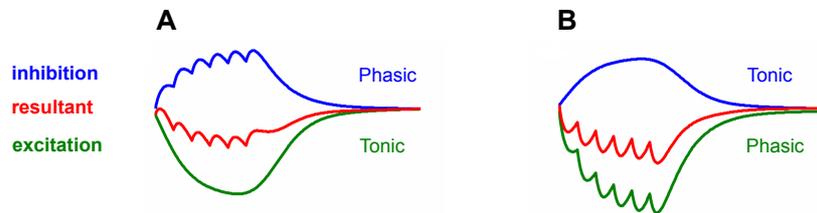


Figure 3.7. Shape of CPSCs in Two Extreme Scenarios. **A.** Phasic, high-frequency (~ 200 Hz) inhibitory current input during a slower, sharp wave-associated excitatory wave. **B.** Fast ~ 200 Hz trains of excitatory current inputs during a slow inhibitory signal. Red traces represent the resultant current that would be observed at an intermediate potential between the reversals of each of the two components. Curves courtesy of Richard Kempter.

We hypothesize that for our recordings at potentials close to the reversal of inhibition (with or without the supplementary help of the pharmacological isolation agent DIDS; see Section 2.4.1) the observed downward slopes will be steeper than the upward ones, in accordance with the green trace in scenario B of Figure 3.7. Indeed, at potentials below the reversal potential of excitatory synaptic transmission, *excitatory* currents within cPSCs are inward and should thus display downward slopes (rises) steeper than their upward slopes (decays). In addition, at the potential we have chosen, putative inhibitory outward currents should display only small amplitudes, due to the small driving force for Cl^- . From the kinetic information obtained from isolated PSCs, downward slopes are expected to be steeper even when several time-shifted PSCs overlap.

Detection of steep slopes Steep slopes were searched in windows of 80 ms around cPSC peaks using a low-pass version of the current signal (Butterworth zero-phase filter of passband 0.5–400Hz and second order). This is the black trace in Figure 3.8 A. Minima of the time derivative of this smoothed current (in grey in the Figure) indicate downward slopes, and maxima indicate upward slopes. Slopes are filtered according to whether they belong in the top decile or top quartile by value in the local cPSC. The distributions of values of the derivatives at such extremal points are shown as cumulative histograms for electrotonic isolation alone (Fig. 3.8 B) and with the supplementary aid of pharmacology (Fig. 3.8 C).

Downward slopes at the reversal of inhibition are steeper We analyzed first the slopes within cPSCs in the eight cells recorded with K₂Glu 1 at -66 mV (1,085 cPSCs in total). In line with the expectations from the EPSC kinetics, we found that downward slopes were steeper than upward slopes (Fig. 3.8 B): the analysis furnished slope averages (\pm SEM) of 21.1 ± 0.3 pA/ms vs. 13.1 ± 0.1 pA/ms (top 25%) and 35.7 ± 0.5 pA/ms vs. 18.9 ± 0.2 pA/ms for 10% steepest downward vs. upward slopes in individual cPSCs, with a P value of 1.6×10^{-178} according to the Kolmogorov-Smirnov (KS) test.

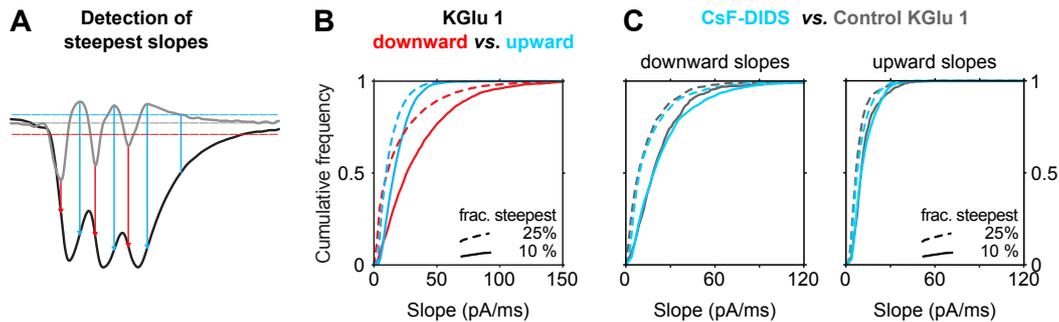


Figure 3.8. Steepest Slopes of Excitatory Currents. Cumulative distributions of downward and upward 25% (dashed line) and 10% (continuous) steepest slopes during cPSCs. **A.** Detection procedure of steepest slopes in an exemplary cPSC. The derivative (grey) of the 400 Hz-low pass-filtered current trace (black) is compared to the 10% threshold levels of steepest slopes (horizontal blue line: upward; red: downward). Vertical lines mark the times of identified slopes. **B.** Distributions of upward and downward slopes from all cPSCs in cells under electrotonic isolation of excitation for two steepness thresholds: top 10% (solid) and top 25%, (dashed). **C.** Same as B but comparing slopes in control condition (black) and after CsF-DIDS perfusion (blue). Two separate panels have been used for downward (left) and upward (right) slopes.

... also with DIDS solution We again analyzed downward and upward slopes of putative EPSCs and compared their values before (“control”) and following perfusion of the cells with CsF-DIDS. Downward cPSC slopes were still steeper than upward slopes following intracellular block of inhibition (control: 24.3 ± 0.8 pA/ms; $n = 224$ cPSCs; CsF-DIDS: 26.6 ± 0.7 pA/ms; $n = 462$ cPSCs; 7 re-patched cells; $P = 0.1$; KS test), while upward slopes were slightly enhanced (control: 12.9 ± 0.3 pA/ms; CsF-DIDS: 13.9 ± 0.2 pA/ms; Fig. 3.8 C; $P < 0.0001$; Kolmogorov-Smirnov test).

Note that in our set of cells with pharmacological isolation via CsF-DIDS, slopes were consistently gentler (about 25 pA/ms in average) than in the electrotonic isolation dataset KGLu 1 (about 35 pA/ms; top steepness decile). Since the steepness was conserved upon repatching control solution with CsF-DIDS, the difference is attributable not to any effect of DIDS but to different cPSC characteristics in both set of cells. SWR peak-triggered averages of cPSCs (Figs. 3.1 top left and 3.2 top row) show similar temporal extension of currents in both cases, whereupon it must be concluded that the different mean amplitude is the cause of the observed difference.

3.4.2 Rhythmicity of Input: Inter-Slope Intervals

We showed above (Fig. 3.6 B) that the kinetics derived from spontaneous EPSCs (not associated with ripples) were fast enough to account for excitatory currents in ripple-associated cPSCs. Working under the hypothesis that currents during SWR reflect compound bursts of PSCs, of which we choose to observe the excitatory components by extinguishing inhibitory conductances, we realized that rhythmicity could be gleaned from the intervals between successive downward slopes. Indeed, Figure 3.8 A above shows by example that steep slopes appear only slightly, and, due to the quasi-linear character of the PSC rise, predictably, shifted with respect to the real onset of the PSC. As long as we are concerned with time differences and the inhomogeneity of the time shift remains in the sub millisecond domain, this discrepancy is of no consequence.

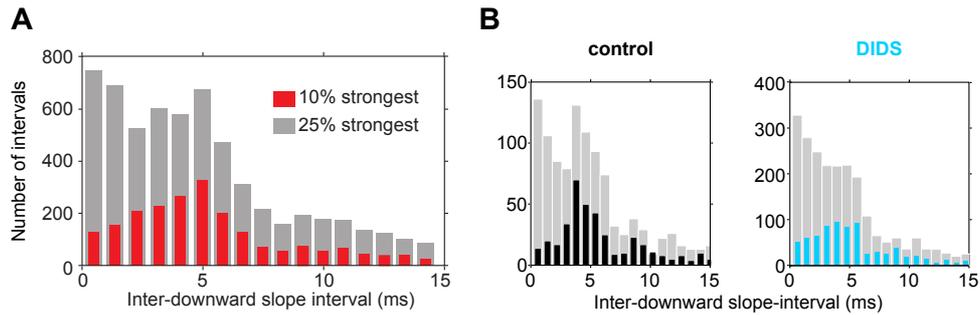


Figure 3.9. Inter-Slope-Intervals in the Excitatory Conditions: Pooled Statistics. **A.** Distributions of interdownward slope intervals for two steepness thresholds (25% in grey). Both peak at 5 ms, indicating ripple frequency modulation of intracellular currents. **B.** Inter-downward slope-intervals detected in cPSCs in control (left, 7.4 ± 0.2 ms, mean \pm SEM; for 10% steepest slopes in black) and CsF-DIDS (right, 6.3 ± 0.1 ms; blue). There are less intervals in the control recordings because those were kept short in order short to increase the success rate of the re-patch with Cs-DIDS.

To quantify the rhythmicity of fast network input during cPSCs, we calculated the distributions of inter downward steep slope intervals and observed that they peaked at 4–5 ms, consistent with ripple frequency. We found these numbers both under control conditions and after CsF-DIDS administration (Fig. 3.9; see Fig. 3.10 for single-cell data). They were robust for a range of low-pass frequencies (400–600 Hz) and fractions of selected steepest slopes (1–25%). Filtering out higher frequencies is required to avoid too many local extrema of the derivatives being detected for what is in practice the same onset.

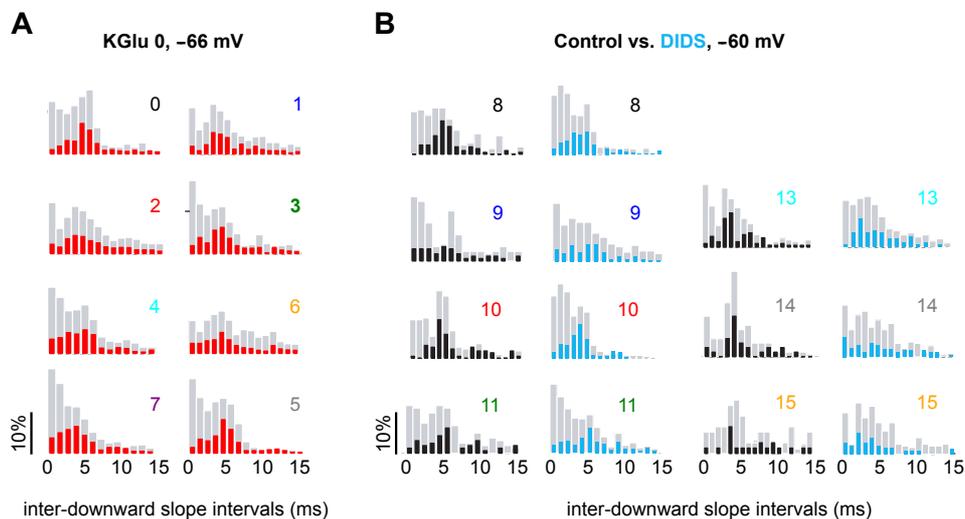


Figure 3.10. Inter-Slope-Intervals in the Excitatory Conditions: Single-Cell Statistics. **A.** Inter downward-slope interval distributions of ripple-associated cPSCs recorded at -66 mV (i.e., at Cl^- reversal potential, revealing excitatory synaptic input) from eight single-cell-and-LFP recordings; all histograms demonstrate consistent peaks at ~ 5 ms. Black and grey distributions represent 10% and 25% of steepest slopes respectively. See Fig. 3.9 for data pooled across cells. **B.** Single cell inter-downward slope-interval histograms for the seven cells with supplementary pharmacological isolation. Black bars correspond to control, and blue histograms to CsF-DIDS condition for the 10% steepest slopes in individual cPSCs; grey histograms indicate values for 25% steepest slopes in individual cPSCs.

Taken together, these results derived from experimentally blocking the somatic postsynaptic action of GABAergic inputs corroborate our hypothesis that ripples are accompanied by a strong oscillation-coherent phasic excitatory component.

Spectral methods We found that, for our data, Fourier spectral methods were less well adapted for the quantification of fast structure of cPSCs. This was in part due to the short duration of the events and the disproportionate weight assigned to events with higher amplitude oscillations. Additionally, the temporal fine structure differences between events are particularly problematic for analyses in the frequency domain, since for high frequencies f a temporal jitter Δt is amplified into frequency jitter $\Delta f = -f^2 \Delta t$ owing to the reciprocal relation $f = 1/t$.

Prior and further work After implementing PSC detection based on slopes we found a previous application of a similar idea. Compte et al. (2008) detected slopes on *voltage* traces by thresholding them after passing a differentiator filter (a combination of a derivative and a filter, in their case <200 Hz because they were interested in the gamma range). Importantly, they set thresholds based on the count of E/IPSCs that each hypothetical threshold would produce, selecting on a cell-by-cell basis one that leads to diminishing EPSC and increasing IPSC counts with increasing depolarization (their Fig. 5E). Methods based on thresholding a filtered derivative have several limitations. First, their time resolution is set by the filter that is so necessary to moderate the increased noise inherent to taking derivatives. If events succeed rapidly, they may compensate in amplitude and thus escape detection, or if they have the same sign, they may coalesce into a waveform that does not show the two original slope extrema after filtering. The method also requires similar kinetics of excitatory and inhibitory events, else an excitatory decay could pass for an inhibitory onset. A final caveat is that, when comparisons are involved, it should be borne in mind that the derivative will increase with increasing amplitude. Normalizing amplitudes should then prevent a bias from inhomogeneous detector sensitivity.

We avoided some of these pitfalls by not employing our method in mixed sign traces and referring values to the local slope statistics. Further work to characterize the regimes of applicability of these methods could start by providing detection rates (ROC curves) as a function of slope thresholds and filter passband from synthetic ground truth data obtained from neural simulations with artificial noise. A comparison with another, more involved, PSC detection method is the object of Section 3.7.1 and in particular Figure 3.24, where the counts detected events are histogrammed against time for several values of the steepness threshold.

3.5 Locking of CPSCs to the Extracellular Ripple

We provide three measures of temporal coordination between intra- and extracellular signals during SWRs:

- i. **coherence**, demonstrating synchronization of both signals at ripple frequency,

- ii. **slope-triggered SWR averages**, showing the alignment of the intracellular steep slopes with a particular phase of the extracellular ripple oscillation, and
- iii. **vector strength**, which quantifies the magnitude of the effect as a function of the threshold for slope selection, enables to evaluate significance of phase alignment and provides the relative phase PSC onset-SWR.

For each of them, a comparison is produced between the results for concurrent SWR and cPSC those for each cPSC and its *subsequent* SWR, intended as a control of how much in the coordination of oscillations is due to the stereotypical nature of the ripple in the slice. Such stereotypy is mainly limited to the extracellular signal; the intracellular one is much more variable. This can be readily appreciated by browsing them at a fast pace for which purpose we have created a video (www.minin.es/omr_video) of 1,000 SWR/cPSC events from the KGLu dataset at -66 mV .

3.5.1 Coherence

Coherence addresses two shortcomings of the quantification of synchronisation between time series by correlation of their amplitudes. The first is that one has to obtain a single correlation value, which implies computing correlations for an unspecified range of time lags and usually taking the highest value. The second is that correlation is not specific of a frequency band and offers no information about band-specific coordination. These deficiencies are tackled by another measure, called coherence.

Coherence is obtained from spectra. Given two time series x and y , one multiplies the Fourier transform of one, say x , with the complex conjugate of the spectrum of the other, y to get the cross-spectrum C_{xy} , and normalizes its square modulus by the product of the power of x and the power of y , $P_x P_y$:

$$\text{coh}_{x,y}(f) \equiv \frac{|C_{x,y}(f)|^2}{P_x(f)P_y(f)}.$$

The resulting measure^{3.3}, which depends on the frequency, only captures linear correlations between x and y and presupposes both time series to be stationary. It estimates the degree both of phase and amplitude dispersion at the frequency f , and ranges between 0 and 1. Perfect coherence at f implies stable phase difference (i.e. total phase lock) *and* amplitude ratio over the time considered. This mix-up of two features is undesirable, and that is a further reason to compute vector strength below (Section 3.5.3), which addresses phase locking alone. Larger amplitude epochs weigh proportionally more towards coherence. One would want this feature only whenever low amplitudes are close to the noise level (so as to deemphasize them), but not in general.

We calculated coherences of concurrent SWRs and cPSCs (Fig. 3.11) and of “shifted” SWRs (one event ahead) and cPSCs (Fig. 3.12) .

3.3. A more specific denomination for coherence is magnitude-squared coherence (Shiavi, 2007); the names coherency or squared coherency are also employed.

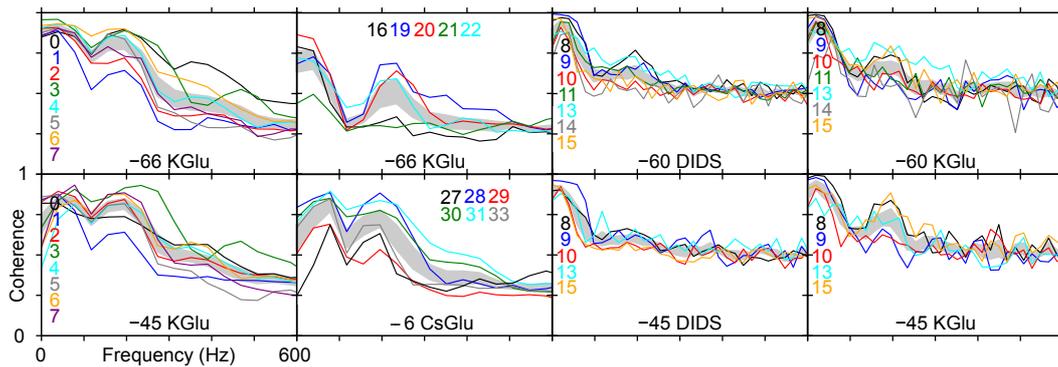


Figure 3.11. Coherence Between SWRs and Their Concurrent CPSCs. Top. Coherence of simultaneous cPSC and extracellular SWR. Symmetric windows of 128 ms around the SWR peak were used for the PSDs calculated with the Welch periodogram method with 256 FFT points. Averages across cells (\pm SEM) are shown in light grey under the colored curves.

Coherence achieves very high values with a clear peak at around 200 Hz and secondary peaks at higher frequencies. At -6 and -45 mV an additional 100 Hz peak stands out from the DC coherence, presumably due to a diminished correlation of amplitudes.

The cells in the DIDS dataset (four rightmost panels) show a poor 200 Hz peak, with the situation only slightly better in the control condition before the repatch. The additional wiggleness with respect to all the other datasets is due to the data acquisition at a reduced rate of 5 kHz instead of 10 kHz.

Some particularities of individual cells can be appreciated. For example cell 3 displays a very broad main peak and a well-defined secondary one close to 500 Hz. Two of the minislice cells (16 and 21) are nearly in the ripple band, whereas the others show a very clear-cut peak of slightly higher frequency matching well the higher frequency of SWRs in the minislice (see Table 2.3 in the previous Chapter). The minislice cells total 725 cPSCs and confirm that SWRs in CA1 minislices are indeed accompanied by phasic, coherent inward currents at ripple frequency.

Finally, the coherence profiles between out-of-sync SWRs and cPSCs (Fig. 3.12) are almost indistinguishable from the synchronous-event coherences, pointing towards a strong stereotypy of ripples in the slice. Indeed successive SWRs looked much more similar in the *in vitro* data than in the *in vivo* recordings.

3.5.2 Slope-Triggered Averages of SWRs

If cPSCs are phase-locked to the SWR, the oscillatory structure in the LFP during ripples should be revealed by the steep slopes in cPSCs. Slope-triggered averages (Fig. 3.13) were computed by averaging windows of 40 ms LFP data centered to the times of the steepest 10% slopes of each cPSC. CPSCs with only one steep slope were not considered. A total of 1,085 cPSCs were represented with 5,161 events, each cell having an average of 4 to 5 onsets per cPSC event. Ripple-modulation (at ~ 5 ms) of the resulting averages signals a consistent phase of the putative PSC onsets in the oscillation defined by the extracellular ripple. The presence of the bulk of the signal after the onset (0 ms offset) indicates a tendency for steeper slopes to concentrate at the beginning of the cPSC.

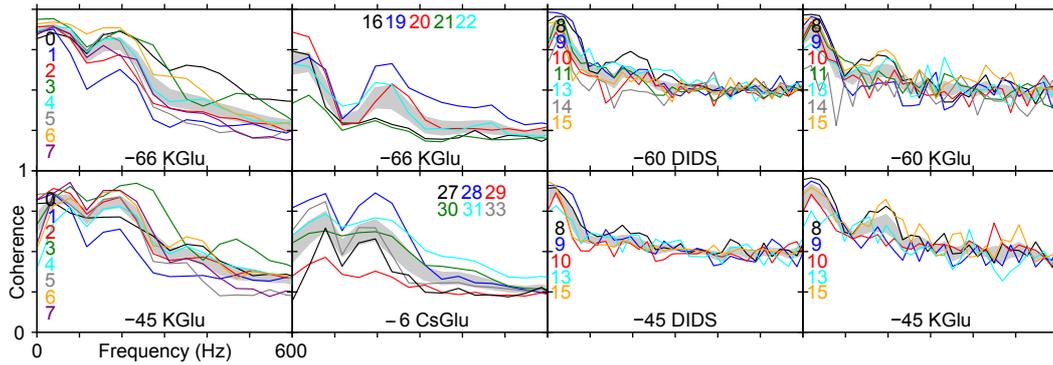


Figure 3.12. Coherence Between SWR and Their Preceding CPSC. Coherence between each cPSC and the *next* recorded SWR. See caption to Fig. 3.11 for the calculation details.

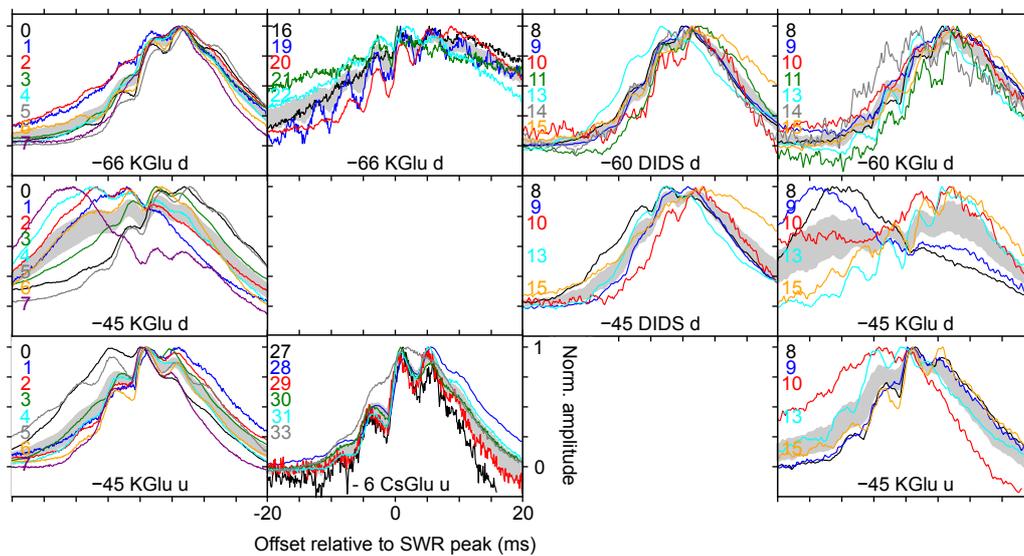


Figure 3.13. Averages of SWRs Triggered by Steep Slopes in Concurrent CPSCs. LFP segments of 40 ms around downward slopes were used. For holding potentials between the reversals of excitation and inhibition, averages are shown triggered both by downward and upward slopes (d, u). Jagged averages are due to low event count (e.g. DIDS control; rightmost column) or normalization of SWR amplitudes close to the noise (e.g. cells 27 and 29).

Polarity changes upon raising the holding potential to -45 mV (for the main dataset KGlU 1) are readily visible in that the slopes that trigger more rippled averages are no longer downward but upward; see for example cells 2, 4 and 7. Ripples are also visible in the DIDS averages, that do not shift upon raising the holding potential as some of their control counterparts do. This is just again DIDS blocking inhibitory currents.

Whereas most cell averages display about 5 cycles of SWR oscillation, two minislice cells (19 and 20) show extremely well defined averages over 7-8 cycles. To elucidate the cause, we looked in Figure 3.14 at the extracellular average waveform, triggering it to the SWR peak much as we did for cPSCs in Figures 3.1 and 3.2. The disparities of shape observed between 1 and 2 as well as 5 and 6 in the Figure give an idea of the variability of SWR events, since these cell pairs come each from a single slice and should therefore show roughly similar SWR wave shapes.

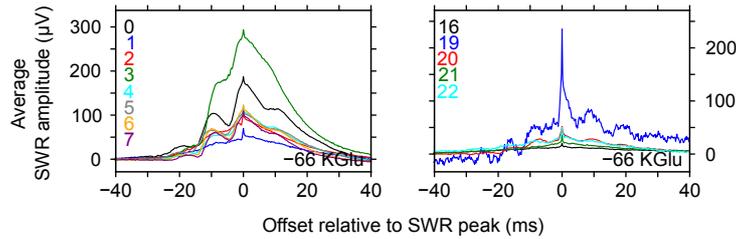


Figure 3.14. SWR Peak-Triggered Averages of SWR: Minislices Have Longer SWR.

SWR averages by cell in full slices vs. minislices. Whereas full-slice SWRs have at most five well-defined cycles, cells 19 and 20 in the minislice have more. Note that the high blip at trigger time (zero) is due to selecting the representative SWR zero as the peak of an unfiltered trace; when recalculated for SWR peaks selected from 1,000 Hz low-passed traces to eliminate this contamination the curve shapes were almost identical except for the central spike.

The analysis showed that not only minislice cPSCs (see Fig. 3.1), but also minislice SWRs in those cells were longer, in particular featuring two visible cycles after the peak instead of one as in the full slice. This observation highlights that both the LFP and the intracellular signal capture the network dynamics. It also begs for a more extensive study of SWR in minislices, one that can properly establish this observation statistically and look for the underlying mechanism. A preparation where the full slice is recorded first, then deafferented and recorded again would permit a clustering approach such as pioneered by Reichinnek et al. (2010). One caveat learned from this case, is that studies of coordination may only offer comparability where the yardstick itself (here the extracellular ripple) is similar enough across the comparanda.

Finally, note that here our control of averages of SWR triggered to the previous cPSC (Fig. 3.15) shows more sensitivity than in the case of coherence: some cells lose completely their patterning (e.g. 1, 2, 16, 29; seems to preferentially affect jagged averages, i.e. low-amplitude SWRs), and the others lose it partially.

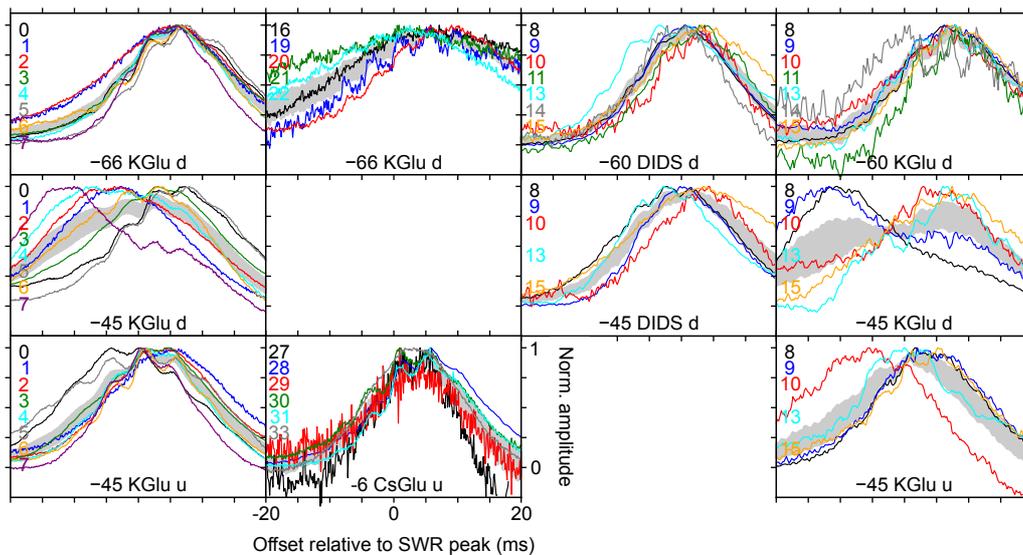


Figure 3.15. Averages of SWR Triggered by Steep Slopes in Preceding CPSC. Patterning of averages is deteriorated in comparison to the concurrent case (Fig. 3.13; see its caption for labels and methods.)

3.5.3 Vector Strength of Steep Slopes

The detailed time structure of the slope-triggered averages calculated in the previous Section hints at PSCs having a definite phase in the extracellular ripple during SWR. In this Section we obtain PSC phases without the confounding influence of amplitude by means of the vector strength, which has the additional advantage over coherence of being applicable to non stationary data.

3.5.3.1 Local Phase Estimation

What phase can be attached to each of the steep slopes in the extracellular ripple? Since the ripple is not a harmonic, sine-like wave, a procedure is needed first to assign phase angles to times in the ripple. Proposed algorithms for the task of local phase estimation in brain signals include wavelet-based ridge extraction (Roux et al., 2007), which also estimates local frequency, and the Hilbert transform, which computes only the phase. While in theory wavelet-based techniques have a promising potential for short-length signals such as ours, it is not clear how much difference there is in practice (Le Van Quyen et al., 2001).

A harmonic real signal such as $s(t) = A \cos \omega t$ is cast in complex form with an imaginary oscillating part delayed by $\pi/2$:

$$\bar{s}(t) = A e^{i\omega t} = \underbrace{A \cos \omega t}_{s(t)} + i \underbrace{A \sin(\omega t)}_{H[s](t)}.$$

The Hilbert transform H obtains the imaginary part $H[s(t)]$, representative of the phase, for every signal $s(t)$ (not only monochromatic) by applying a $-\pi/2$ phase shift to each of its frequency components:

$$\bar{s}(t) = s(t) + iH[s](t).$$

This is implemented by an integral transform (P.V. stands for principal value):

$$H[s](t) = \text{P.V.} \int_{-\infty}^{\infty} \frac{1}{\pi \tau} s(t - \tau) d\tau = h \star s = A(t) e^{i\varphi(t)}$$

The transform is equivalent to a convolution of the amplitude signal $s(t)$ with the Hilbert kernel $h(t) = (\pi t)^{-1}$. This kernel is nonlocal, i.e. it extends to $\pm\infty$, but its flanks decay rapidly, so that local information from the amplitude $s(t)$ dominates the calculated phase. The convolution theorem states that the Fourier transform of a convolution is the product of the transforms of the factors. Since the transform of the Hilbert kernel for positive frequencies is $\hat{h}(f) = -i = e^{-i\pi/2}$, the HT applies a phase shift uniformly to the signal at all frequencies. That the HT exists does not guarantee that it has a meaning; only when the signal frequencies are confined to a narrow band does it make sense to interpret $A(t)$ as its envelope amplitude and $\varphi(t)$ as its phase (with $d\varphi/dt$ matching the frequency of maximum power; see Boashash (1992).

3.5.3.2 Slope Phases in the Ripple

To assign phases with respect to the ripple component of the LFP, we applied a Hilbert transform on the extracellular potential filtered in the ripple band (120-300 Hz) with a Butterworth acausal filter of order two. Each putative PSC event detected by its extremal slope was assigned a Hilbert phase, φ_j . As an indicator for locking quality of these collected phases we use the vector strength (VS)

$$\text{VS}(\varphi) = \frac{1}{N} \left| \sum_{j=1}^N e^{i\varphi_j} \right|.$$

This quantity, also known as *mean resultant length* or *phase locking value* is an average of phasors and is normalized between 0 and 1.

Each of the panels in Figures 3.16 and 3.17 shows the vector strength as a function of the fraction of steepest slopes whose phases are examined, together with the average of those phases in a polar plot inset. The uppermost left panel of Figure 3.16, for instance, displays eight cells for which phase locking becomes better as steeper slopes are selected, which are likelier to be actual proxies of real PSCs. This monotonic behavior, which can be appreciated for all conditions and cells, constitutes an indication that steep slopes are indeed a viable feature to assess PSC timing. The values reached for e.g. 5% of steepest slopes (corresponding to the top three onsets per ripple) are considerably high, and the mean phase remains rather stable across cells (in the example at about 270°), suggesting that PSC times a dynamical feature of the SWR/cPSC phenomenon and not dependent on the details of the afference to each of the cells.

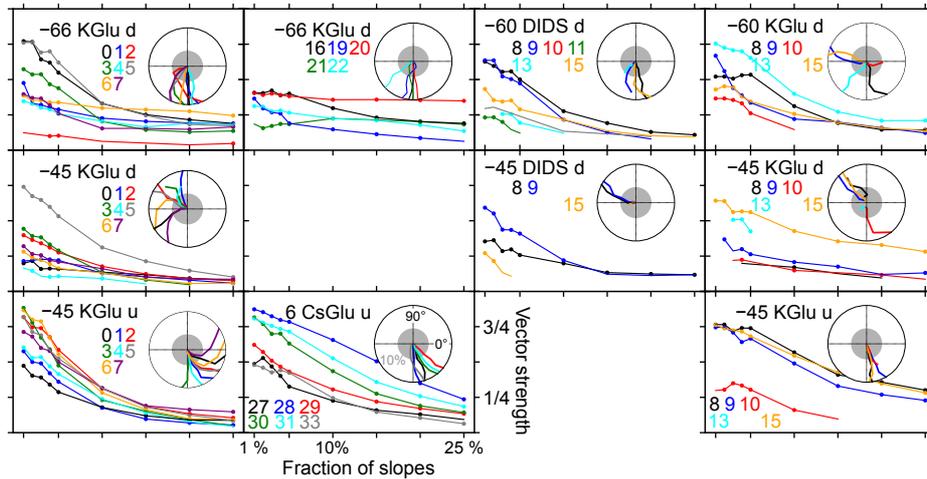


Figure 3.16. Phase Locking of Steepest cPSC Slopes to Concurrent SWR. Each panel represents one condition and one direction of steep slopes (d for downward and u for upward slopes). The vector strengths are shown as a function of the fraction of steepest slopes chosen as proxies for PSC onsets, in the range from 25% (right) to 1% (left). Polar plots insets show the corresponding phase averages, from 25% slopes at the periphery to 1% at the center through the mark of 10% (grey disc). All vector strengths and phases plotted are significant according to a Rayleigh test for the uniformity of a phase distribution with $P < 0.05$; Dots have been added on the line plot when the significance level increased so that $P < 0.01$.

Excitatory onsets The mean phase of putative EPSCs at 10% threshold is $-114 \pm 10^\circ$, which means that the steepest downward slopes lead the ripple cycle peak (at 0°) by about 1.5 ms. No statement can be made about the pyramidal presynaptic spikes that presumably cause them; see the Discussion in Chapter 5.

Minislices Cells in minislices show very significant locking despite generally lower vector strengths. Similarly, for cells recorded in minislices, the mean phase of putative EPSCs with respect to LFP ripples at 10% threshold is $-99 \pm 7^\circ$, comparable with those derived from intact slices ($-114 \pm 10^\circ$). This compatibility is additional evidence of a local origin of ripple-coherent excitatory PSCs within area CA1.

Inhibitory onsets Mean phase over cells at 10% threshold of steepest slopes is $-60 \pm 13^\circ$, corresponding to cIPSCs leading the LFP ripple cycle peak by ~ 0.8 ms.

Intermediate holding potentials Cells 0 to 7 held at -45 mV show mostly outward currents. Correspondingly, the phase of their upward slopes advances towards that of inhibitory currents in the CsGlu -6 dataset. The downward slopes are for most cells phase-opposed to the upper. This makes sense if one considers a rise time of about 1-2 ms and a decay of 2-3 ms: the detected downward slopes are, in reality, decays of the inhibitory PSCs.

DIDS dataset The pharmacologically isolated excitatory currents show significant locking with phases similar to those only electrotonically isolated, further increasing our confidence that the latter are mainly excitatory in nature. Control shows a very dispersed phase that is not easy to understand, because with the very little data available in hand, cells 15 and 9 display PSC onsets well ahead of those measured for other cells. At -45 mV the same qualitative effects are observable as commented above.

Finally, our control consisting of examining phases of PSCs not in the concurrent SWR but in the next one (Fig. 3.17) shows that the ability of vector strength to capture nonlinear correlations provides an additional finesse in the appreciation of locking, as compared to a technique like coherence that only appraises linear correlations. This enhanced sensitivity leads to most values of vector strength being strongly depressed and entire cells losing locking significance, even at the laxer $P < 0.05$ level.

3.5.4 Concluding Remarks

We have performed three analyses of synchronization between SWR and cPSCs: coherence, onset-triggered averages of SWRs and vector strengths of onsets in the oscillatory framework of the SWR. What did we learn?

- the striking stereotypy of the whole ripple phenomenon in the slice, revealed by comparing synchronization of simultaneous SWR-cPSC waveforms with a shifted control cannot be overlooked, but is less prominent when assessed with a technique that includes nonlinear correlations, such as vector strength. The question left in the air is, naturally, whether the assemblies underlying the ripple are changing with each event or the stereotypy means that a single set of cells is staging the whole show.

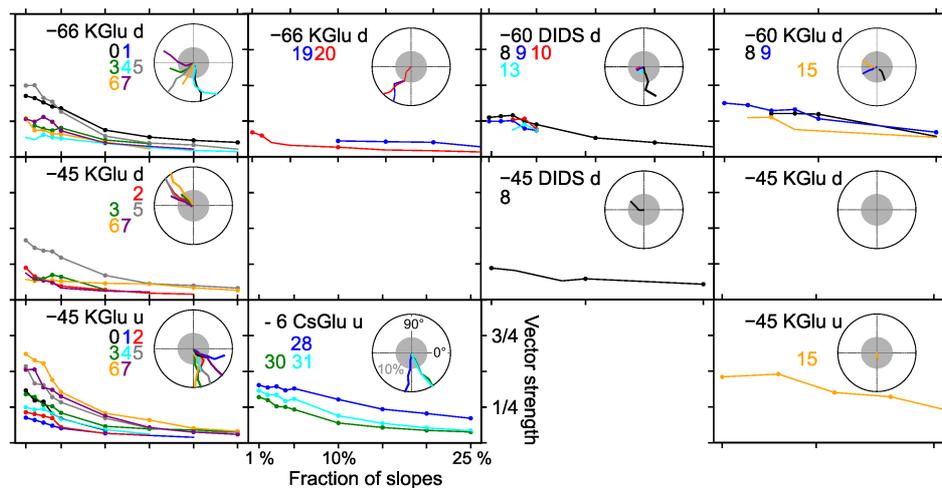


Figure 3.17. Phase Locking of Steepest Onsets to Subsequent SWR. Same as 3.16 but calculating the phase of the observed steep slopes not on the concurrent SWR but on the subsequent. Cell numbers for which no single value was significant have been deleted.

- minislice ripples have more cycles, both in their extra and intracellular expression. This fact necessitates a more solid statistical backing (here observed only in two cells, since the others had too low-amplitude SWRs).
- DIDS observations ranked lower in all synchronization measures, but all the parameters that could be extracted were in quantitative agreement with the more telling electrotonic-only results.
- the intermediate holding potentials are interesting for the assessment of the interplay of inhibition and excitation, but necessitate of a delicate and sophisticated data analysis. The most promising approach may be in comparing the composition of currents incoming to the *same cell* as the potential is raised to the most depolarized value that the cell will sustain (see Compte et al., 2008). Along these lines, a second patch could offer an additional reference measurement while the main cell is progressively depolarized.
- phases show a certain dispersion across cells but cluster in a rather stable fashion across conditions.
- inhibition slightly lags excitation, by about 1.5 ms. Both types of synaptic input seem to be exquisitely well timed to a similar extent. These findings are consistent with phase relationships derived from extracellular recordings in vivo by Ylinen et al. (1995), where unfortunately only one interneuron could be isolated (their Fig. 7).

It is time to go one step further and try to delve into the details of the excitatory-inhibitory interplay. The next Section presents the technique that will clear our way into the dissection of the timing and shape of individual PSCs, which is then presented in Section 3.7.

3.6 Peeling Reconstruction of Synaptic Inputs

In this Section we target the individual contributions that make up compound post-synaptic currents during ripples. We describe a novel algorithm for waveform reconstruction that is based on peeling fitted components off the remaining, unexplained traces. As these components are first detected by deconvolution, we refer to this peeling reconstruction algorithm also by the name of *peeling deconvolution*. The reader who prefers to find out first what results were attained with this approach first and then maybe come back here to learn about the details of the algorithm is encouraged to skip to the next Section.

3.6.1 Position of the Problem

Compound post-synaptic currents (cPSC) are presumably the result of the superposition in time of individual PSCs. The interval between PSCs during ripples is roughly similar to their decay kinetics (4-5 ms). This poses a primary detection problem and a secondary problem for characterization, since no PSC is visible in isolation. Detection is complicated because, as the ripple advances, the sum of the tails of all earlier PSCs becomes an ever increasing fraction of the signal to be analyzed. This sum of tails is at the simplest a sum of exponential decays with somewhat different characteristic times and offsets. It is the base upon which an incoming PSC is compounded—let us assume to our benefit, for the moment, that it is simply linearly accumulated. Onset slopes are systematically reduced by this mechanism more than decays are and the ability to detect incoming PSCs from their most defining kinetic characteristic—their fast onset—is gradually lost. Even if detection was not jeopardised and the resulting estimation of onset times was not biased, there remains the problem of how to adequately characterize the amplitude and kinetics of a PSC that is overlaid on an unknown mass of current from previous events. This is the characterization problem.

To address these challenges, we decided to take advantage of the causal structure of the problem: earlier events are *not* affected by later events, and there is an ever so short span of time between the first PSC and the second where it can be observed in pristine isolation. Do we know enough about the shapes of PSCs to be able to extrapolate what we do not see, the tail of the PSC that hides, as if it were the submerged mass of an iceberg below the waterline of succeeding synaptic activity? If so, we may subtract the entire first PSC (visible + “underwater”) from the rest of the trace and look at the remainder with confidence that it is *as if* the first event had never been there in the first place, an instance of the original problem upon which to re-iterate our peeling step.

We banked on a positive answer for the particular combination of signal to noise ratio, structure of noise, directionality of currents, available sampling ratio, kinetics of events, and rate of new events. It is important to underline that there is no guarantee that the relative success experienced with this combination can be reproduced for less benevolent settings. That said, the challenge before us was more complex than, for example, is the case for studies of synaptic facilitation. There, the experimenter is in control of the arrival times (she *defines* them through the stimulation protocol) and can obtain a very accurate characterization of kinetics (through repetition and averaging), which may even be constant throughout the synaptic sequence (Richardson and Silberberg, 2008). In those cases, only amplitudes remain to be estimated. Another recent peeling study (Grewe et al., 2010), aimed at uncovering the

spikes that underlie the very slow dye fluorescence activity associated with Ca^{2+} imaging, could rely on homogeneous kinetics *and* predictable amplitudes, plus known stimulation times (the concrete onset time still had to be fitted). However, the challenging demands of imaging severely limited sampling rate and signal to noise ratio in that study. Stimulation-related assumptions and homogeneity of kinetics are absent from our approach, which frequently succeeds at unmixing events that overlap even over their rise phase.

What alternative approaches did we consider? Perhaps the naivest consists in fitting the whole trace at once, using local optimization (to be quickly lost in local minima in a space of 25+ dimensions), or global techniques such as simulated annealing or genetic algorithms (which share the problem of estimating how many events are present, see the “detection problem” above). Since our base functions (exponentials) form an overcomplete basis (frame), matching pursuit (Mallat and Zhang, 1993) could be envisaged. As global fits above, estimation of the number of events to fit poses a problem, and the technique does not exploit the time structure of synaptic traces. We had disappointing results in a superficial exploration. Finally, factorial hidden Markov models (Ghahramani and Jordan, 1997) seem to offer a sufficiently general framework but are more complex to implement and potentially slow. At the time of this writing we became aware of a mathematically sophisticated approach from the imaging community (Vogelstein et al., 2010). Although we cannot offer insight on the details of their algorithm, the authors claim faster-than real-time performance, offer validation with synthetic data and do a round-up of alternative algorithms (under the homonymous headline, p. 3702), all of which commends its consideration by the interested reader.

To briefly recapitulate, the prescription of our finally chosen method is: detect the first PSC, reconstruct its full wave shape from the unadulterated visible fragment well into the sunken future, and peel that best guess away from the trace. Repeat until no more candidates left. In the next Section we first outline and then detail these steps.

3.6.2 Anatomy of a Peeling Reconstruction Step

Figure 3.18 summarizes the structure of a single iteration of the algorithm. Note that the references in this Section to maxima and minima should be understood in the context of a negative-going trace, i.e. a cEPSC.

Preprocessing We remove any local baseline shifts that survive a global highpass filter at 0.5 Hz (top panel). The next steps of the algorithm assume a zero baseline throughout. The detection step will be conducted on the smoothed trace that results from applying a low-pass Butterworth filter of order 2 at 400 Hz. The trace extends 30 ms before and after the SWR peak.

Detection As a first step in every iteration, we detect candidate events as peaks in the deconvolution with a single exponential kernel of decay constant $\tau_d = 4$ ms, equal to the average decay time constant of spontaneous events (Fig. 3.6; note that the histogram there reports 80%-20% amplitude decay times, not time constants). Such single-exponential deconvolution of a signal y is calculated as a linear combination of its amplitude and its derivative:

$$\text{deconv}(y) = y + \tau_d \frac{dy}{dt}. \quad (3.2)$$

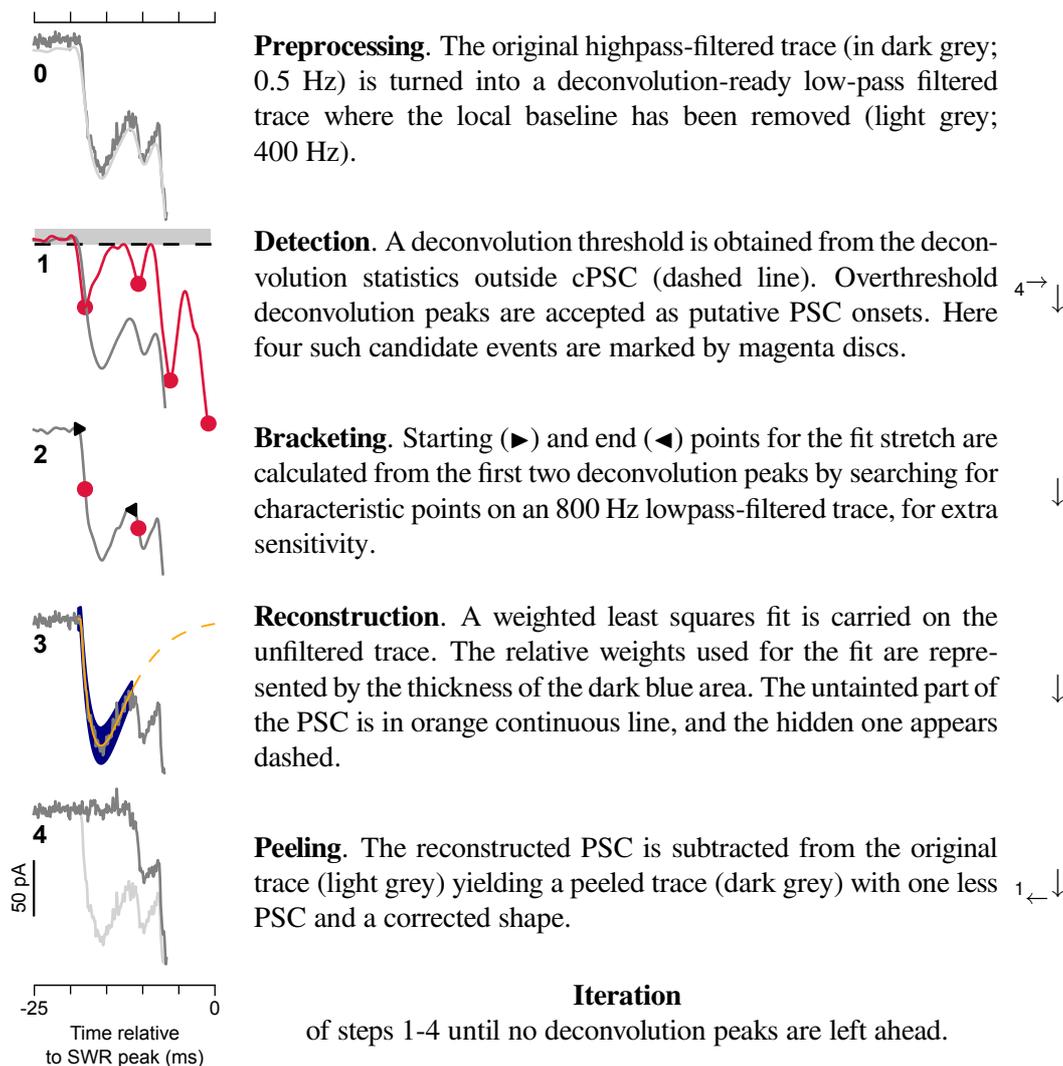


Figure 3.18. Principle of Operation of Peeling Reconstruction.

We discuss in Section 3.6.5 below the choice of this particular deconvolution operator. Deconvolution peaks are accepted as putative PSC onsets if they are larger than the mean plus four standard deviations of the deconvolution peaks of noise surrounding the event and if the 400 Hz lowpass-filtered trace has both a negative amplitude and a negative derivative there.

Bracketing Second, we identify the starting and end point for the one-PSC fit. As reference we use the first two accepted deconvolution peaks. The starting point is estimated as the zero crossing before the first deconvolution peak (preceding the peak by at most 3 ms) roughly indicating a maximum of the cPSC trace. The endpoint is estimated either as the first zero crossing of the trace after the first peak of the deconvolution or the last maximum before the second peak, whichever occurs first. Maxima and zeros are searched on the 800 Hz lowpass-filtered trace for increased resolution. These choices are adequate in most cases to provide the largest possible fitting window

that does not include contributions from upcoming events. Incorporating as much of the decay is critical to obtaining a good estimate in the reconstruction step.

onset doublets. If no local maximum can be identified between the first two deconvolution peaks, we label those intervals as *doublets* and handle them specifically. Indeed, an extremely rapid succession of synaptic inputs can lead to composite currents where the decay phase of the initial event is masked in its entirety by the rapid rise of a second, superposed event. Generally this results in traces that do not show any local extremum between the first two onsets but at most a concavity change.

Reconstruction Third, we run a locally-weighted constrained least-squares fit using the double exponential of Equation 3.1 that has proven to offer an accurate representation of spontaneous PSCs.

weights. A Gaussian weight function emphasizes the kink of the PSC around its minimum with an adaptive length constant equal to the distance of the minimum to the end of the fit or to its start, whichever is shorter but in any case between 1 and 10 ms. This procedure of locally weighting a regression is often used in machine learning, where the point of maximum interest, with the larger weight, receives the name query point and the decay constant of the weight function that of bandwidth (Ng, 2011). Here we extend upon that approach by making bandwidth opportunistically shrink or expand in dependence of the available information forward and backward of the turning point.

constraints. In order to avoid artificially long decays, we add a quadratically growing penalty to the fit error whenever a decay constant $\tau_d = \tau_*$ is exceeded. The penalty threshold τ_* was either 8 or 10 ms. Artificially long decays jeopardise the method, because they subtract too much amplitude from subsequent deconvolution peaks thus making them invisible. About only 5% of spontaneous decays are longer than 10 ms (compare Fig. 3.21 below). An additional sentinel system allows to plug functions that decide on whether a fitted PSC is suitable; for our dataset we applied only the requirement that PSCs must have at least 5 pA amplitude.

Peeling Finally, with the fit parameters obtained in the previous step we rebuild the PSC, including the previously missing tail, and subtract it from the unfiltered trace.

Onset doublets (see above) are handled separately by fitting them simultaneously with a sum of two component alpha functions. Only the first PSC fit of the sum is subtracted from the cPSC.

After subtraction of each PSC, we restart the iterative procedure by detecting deconvolution peaks after the endpoint of our previous fit and proceed as described above until no more events are left in the cPSC event window of $(-50, 100)$ ms around the SWR maximum.

3.6.3 Peeling Reconstruction in Action

The procedure explained in the previous Section can be applied to traces of any length. The robustness of the procedure is increased by re-running the detection stage after each peeling step, by the penalty imposed upon long decays together with the requirement of zero baseline, and by repeating the second fit of double fits.

3.6.3.1 Peeling a Single CPSC Event

Figure 3.19 shows a reconstruction example comprising seven PSCs, which provides an excuse to discuss some of the challenges of the problem and limitations of the solution.

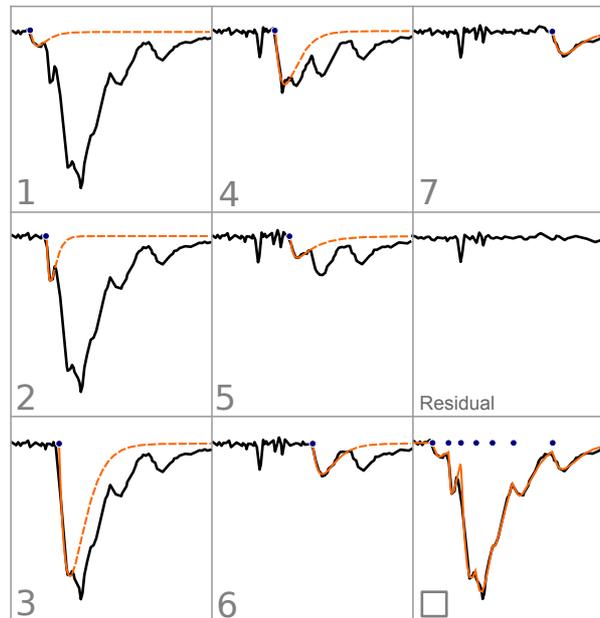


Figure 3.19. Example: Peeling Reconstruction of a CPSC. In each panel the leftmost PSC is fitted by a model function (dashed orange) and subtracted from the trace. Here 7 PSCs were fitted to the cPSC. The residual displays the fitting error consisting mainly of high-frequency components indicative of the limits of the model function. The putative fitted onsets of the PSCs are shown in blue as the fit progressed and displayed together in the final panel.

The initial two events are reasonably well fit; the extent of their decays enables a good estimation of the tail. The third event is challenging for a variety of reasons.

Smoothness of the onset Firstly, we see in the third event that the double exponential function does not capture well smooth biological onsets. The Heaviside factor in Equation 3.1 generates a strong kink of discontinuous derivative that will fail to represent the smoother increase of the biological signal. We considered, and discarded, two alternatives. The first was to employ a simple alpha function with only one time constant, te^{-t} . This has a smooth onset but its kinetics are not rich enough to represent our data, having only one parameter. The second was to use a more complex function, perhaps a sum of two error functions. This added parameters to a fit that already is rather unconstrained in view of the paucity of post-peak available data, and created problems of its own such as to how best glue the component functions. Therefore we went ahead with the double exponential. High-frequency leftovers such as the small peak left behind by the fit of the third event are due to this shortcoming of the model function, but are expected to influence only mildly the extrapolated time course of the event, by virtue of the emphasis on the area around the PSC peak provided by the weight function.

Super-events or presynaptic assemblies in action? Secondly, the third event is of extraordinary magnitude and is likely to be the result of two or more almost coincident PSCs. The slight change in slope that is visible in panel 2 is an indication, but it is not enough to force a change of concavity and thus not be picked up by the doublet-fitting branch of the algorithm, which would have anyways a hard time with as little decay information as is available in this case.

Veiled decays Finally, the amount of decay information available is very scarce. It points at a faster decay than predicted by the fitting algorithm. Here we have experienced that the model function is again the limitation: there is no way to produce very sharp turns at the maximum with a double exponential function.

Despite the above three challenges illustrated in the third event, the algorithm suggests a fit that does not compromise the remainder of the reconstruction process, so that all of the originally clearly visible events will be fitted. Indeed, a tail whose amplitude is overestimated can detract amplitude from succeeding small deconvolution maxima in the rest of the trace. A delicate balance is needed, because it is precisely tail subtraction that enables deconvolution to ignore noise-related peaks in later iterations that would have been above threshold at the start, especially at the end of the trace where several tails artificially enhance the amplitude of the noise.

The fourth event has either a noise blip or a small PSC riding at its peak, which the algorithm ignores. The fifth event is barely visible as tail bump in panel 1, yet it is fit with a rather long tail that might hide an undetected, rather small additional PSC (falling below the 5 pA amplitude threshold). For the sixth event, the bracketing routine suggests a slightly too late bracket close. The last event was almost unaffected by the bulk of the preceding compound current, and is fit unproblematically. The situation there is as for spontaneous PSCs, which are fit in Section 3.6.4 below.

3.6.3.2 Peeling a Population

We have seen that peeling reconstruction is not immune against a number of pitfalls. How does it fare at the statistical level? The answer to this question is very much complicated by the related question of what constitutes a good reconstruction. As can be deduced from the discussion above, each may have his own opinions, or trained eye, as to what is a missed PSC or an overestimated decay time. There is no ground truth to speak of, short of applying the method on stimulation data or on synthetic traces, which we did not have the time to do systematically. However, there are a few heuristic criteria that we may apply. Our code provides a pluggable architecture both for single PSC fits and whole cPSC reconstructions. Logical predicates (oracle functions) emit an accept/reject verdict that can be used in logical expressions. After extensive experimentation including constraints e.g. on the area above zero in the residual trace (its increment is a sign of a badly overshooting fit) we settled for just rejecting those events where the standard deviation was larger in the residual than in the original trace.

Using these admittedly relaxed criteria, which were complemented by extensive random inspection of single reconstructions, we found a reconstruction acceptance rate above 95% for all 14 cells analyzed, but number 4 (92.8%), totalling 1,053 excitatory and 848 inhibitory events. Figure 3.20 shows timing and amplitude information from all the reconstructions from cell 0 recorded at -66 mV.

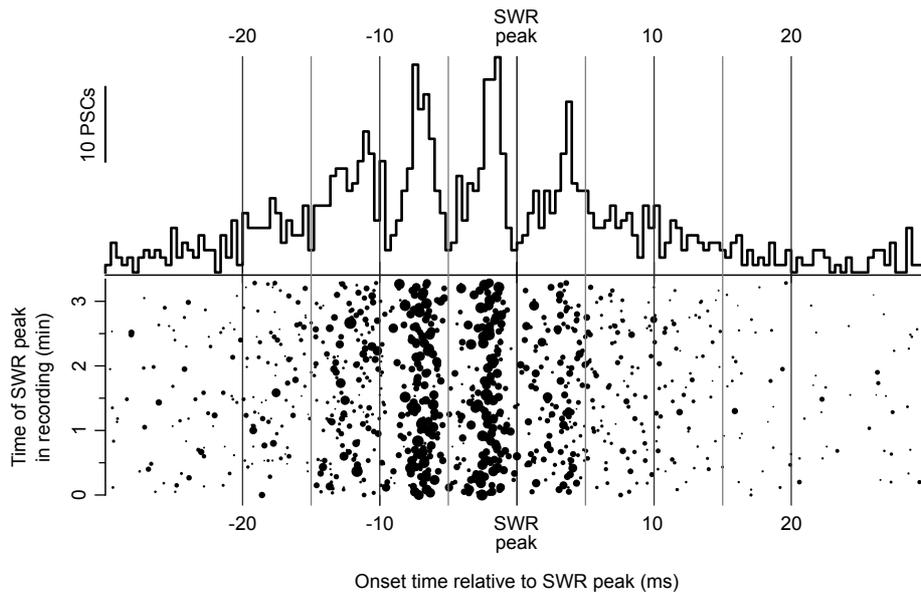


Figure 3.20. Example: Peeling Reconstruction for an Entire Recording. Reconstruction of all cPSCs detected for cell 0. The raster at bottom stacks all cPSCs measured in about 200 seconds of measurement, aligned by the SWR peak. PSC onsets are represented by discs of area proportional to the PSC amplitude. At top, histogram of arrival times demonstrating exquisite in-ripple synchronization in the millisecond range.

With this technology we did not attempt reconstructions of noisier data (such as from the minislice experiments). Neither did we try to fit the DIDS traces, whose lower sampling rate of just 5 points per millisecond complicates the task, adding to their less-well defined cPSC shapes. We also eschewed the harder problem of mixed currents where a downstroke could just as well be the onset of an EPSC or the decay of an IPSC, and currents can cancel, making it fundamentally ill posed. Notwithstanding, the strong separation of timescales between the rise and the decay phases of PSCs of any sign may still enable a further version of the present algorithm to fit such traces, especially if inward and outward PSC amplitudes are similar. Another future challenge lies in the deconvolution of voltage recordings obtained during ripples in the more natural configuration of voltage clamp.

3.6.4 Comparison With Spontaneous PSCs and Validation

The peeling reconstruction algorithm can be applied to sparse trains of synaptic events as well, such as those between cPSCs in our recordings. Whereas simpler algorithms also do the job, using peeling reconstruction avoids the need of manually purging events that ride on top of each other (as we did for Fig. 3.6). Figure 3.21 compares statistics of PSCs detected during ripples and outside of ripples from recordings of cell 0 at -66 mV. Data about other cells is presented below in the same format, in Figures 3.25 and 3.26.

Inter onset intervals during SWR peak sharply at ripple frequency The Figure shows first in the top left panel that the distribution of inter onset intervals of the noise is rather long-tailed and does not present any peaks. A more detailed study

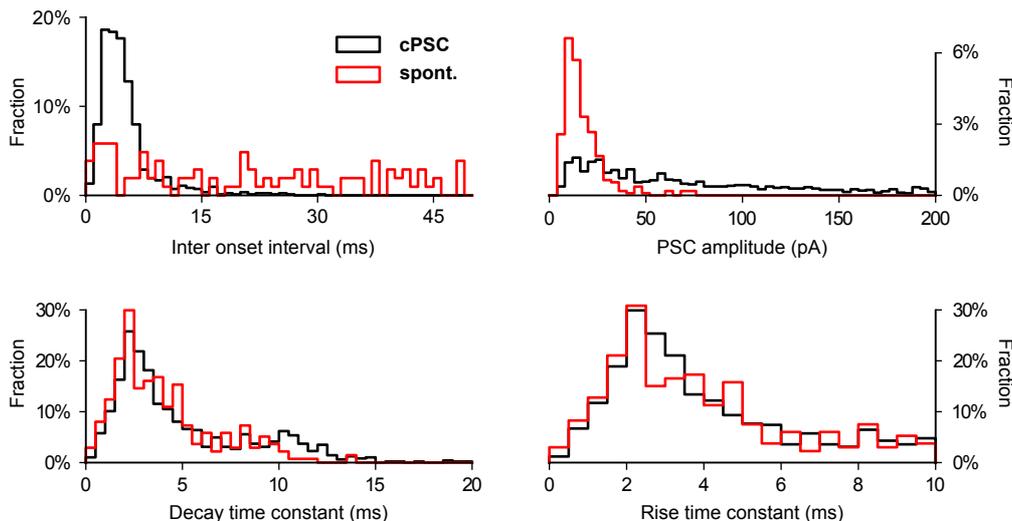


Figure 3.21. Validation: Fit Parameters Compared to Spontaneous PSCs. The epochs for the analysis of spontaneous events were the intervals $(-180, -120)$ ms with respect to each SWR peak, while in-ripple PSCs were fitted in the interval $(-30, 30)$ ms, as in all other reconstruction analyses. An independent assessment of the spontaneous PSC kinetics (Fig. 3.6) yielded $\tau_r = 1.70 \pm 0.04$ ms and $\tau_d = 4.04 \pm 0.08$ ms.

including longer time spans would be needed to assess its Poissonian character. By comparison, the in-ripple cPSCs present a clear peak at around 4 ms. This peak is even better defined than those in the corresponding inter-slope histograms (Figs. 3.9 and 3.10), and, importantly, both are clearer than those obtained with spectral methods (Fig. 2.9 C in the previous Chapter).

Ripple-embedded PSCs many times larger than spontaneous In the top right panel of Figure 3.21 we can appreciate that while amplitudes of the spontaneous PSCs are in the 20 pA range that we encountered in our previous analysis of Section 3.3, PSCs during ripples exhibit a long tail reaching well up to 10 times and more the size of spontaneous PSCs. This may mean that several PSCs from distinct presynaptic partners belonging in the same assembly are collated into one, so well-timed that they are inaccessible for the time resolution of peeling reconstruction on somatic voltage clamp data. Alternatively, excitatory conductances may just increase during ripples.

Exquisite accord of PSC kinetics in and out of ripples The bottom panels of Figure 3.21 show seemingly identical distributions of both kinetic constants for in-ripple vs. out-of-ripple PSCs. A caveat is that the peeling reconstruction algorithm *uses* as a guess both for detection (deconvolution kernel constant) and fitting (least squares guess) the known kinetics of EPSCs, so it is not entirely unbiased in its choice of time constants. However, the success of many reconstructions (low squared area of residual, and observed match to the original) and the ultimate freedom to fix time constants are strong hints that what we see is the underlying statistical distribution, and that it *is* the same for all EPSCs observed in the recording, whether ripple associated or not. The penalty function (here discouraging $\tau_d > 10$ ms) makes itself visible in the little bump of decay times. The corresponding curve for spontaneous PSCs shows a small mass of events beyond the 10 ms mark. This inexactitude is a price to pay to ensure stability in the progression of the reconstruction.

Robustness of peeling reconstruction The above results are, thus, consistent in every aspect with our previous investigation of in-ripple currents and spontaneous PSCs by independent methods. They thus add to our confidence in the performance of peeling reconstruction. We additionally stress-tested the robustness of the reconstruction algorithm on individual events and at the population level by changing structural parameters (bracketing strategy, use of a second pass to catch leftover, unfitted events), signal resolution (filter bands), detection constants (deconvolution time constant and threshold), fit constraints (penalty threshold and penalty function), weighting schemes (constant bandwidth vs. adaptive; flat weighting) and by jittering parameter guesses for each PSC. No relevant changes of the fit statistics were observed. Finally, at the population level, averages of measured currents and averages of their reconstructions agree (see Figs. 3.22 and 3.23 in the next Section), which does not guarantee correctness at the single-cPSC level but hints at absence of bias in the errors of the reconstructions.

Do synaptic inputs compound linearly during ripples? The main underlying assumption is that individual currents sum linearly, in fact that synapses work as linear time-invariant systems. While in current clamp the arrival of a synaptic event can change the voltage of the membrane and thus affect succeeding synaptic events (through active conductances), in voltage clamp this source of nonlinearity in the temporal summation is absent. Another source of nonlinearity would be the arrival of PSCs in short order at the same synaptic site. The currents could engage nonlinear synaptic mechanisms, such as vesicle depletion. Though in principle possible, the low firing rates of pyramidal cells during ripples (our own observations in vitro; generally below 10 Hz in vivo, Csicsvari et al., 1999b) coupled with the extensive arborization, make it unlikely in the case of excitatory afference. From our own data (amplitude histogram in Fig. 3.21), it can be argued (admittedly, making use of the linearity assumption), that the total number of events arriving within a cPSC in our preparation is at most 50, whereas intact (not in-vitro) pyramidal cells have dozens of thousands of synaptic terminals. In support, Cash and Yuste (1999) failed to see nonlinear summation due to reduction of driving force (does not apply here, under voltage clamp) or conductance shunt in hippocampal CA1 pyramidal cells, attributing the disagreement with the predictions of cable theory mostly to active, compensating conductances.

The arguments above do not establish, by far, that the currents *do* compound linearly (much less in the high-conductance state in vivo), after all, for example, why should they distribute across synaptic terminals in a non-clustered fashion?. On the other hand, if currents compounded nonlinearly, the process would be time- and membrane-site dependent, and it could show hysteretic effects due to e.g. chemicals being not cleared out fast enough. Analyses would be nearly impossible. We are thus bound to see how far we can get within the limits of the necessary assumption of linearity.

3.6.5 Detection by Deconvolution

Deconvolution is just the inverse of convolution. The convolution m of two real-valued functions of real argument f and g , denoted $m = f * g$, results from flipping one of the functions (i.e. $f(t) \rightarrow f(-t)$), shifting it by s ($f(-t) \rightarrow f(s-t)$) and evaluating the integral of the product with the other:

$$c(s) = \int_{-\infty}^{\infty} g(t)f(s-t) dt.$$

The convolution function $c(s)$ gives the area of overlap of g and flipped f for all values of the relative shift s .

The reason why the convolution operation is so important is that it models the process of acquisition of measurements through an instrument. The measurement v is generally *not* the underlying reality but the convolution of the instrument function f with that underlying reality, g . This is denoted as $c = f * g$. Deconvolution is the process of obtaining g from the measurements and some knowledge of the instrument, or transfer, function f . As an example, astronomical observations of stars usually show so-called diffraction spikes that make them look like crosses. Deconvolution can be used to remove those spikes by mathematically reversing the optical transformation applied by the telescope.

What instrument f are we deconvolving from the measured signal c , when we observe currents in voltage clamp? The requirements of helping detection of a postsynaptic current and observing the presynaptic spikes match in this case. For detection we desire to have as compact a deconvolution as possible — to avoid precisely the kind of mixing in which fast-paced currents incur. But we also know that, as long as the current is of synaptic origin, it is caused by an action potential elsewhere, which is already a very compact function, extending for at most 1-2 ms (vs. perhaps 10 for the current it ultimately generates on the postsynaptic end).

Thus it becomes of help to consider what is the synaptic transfer function, even though we have detection and reconstruction phases that are uncoupled, i.e. we can detect with a biologically unrealistic kernel and still fit with a plausible one. This is a complex subject because if we really want to go back all the way to the presynaptic spike, many processes are involved, from the presynaptic vesicle dynamics to the electronics of the amplifier. A good start was offered by the function in panel 5 of Figure 3.5. We adopted instead the somewhat simpler double exponential of Equation 3.1, which we denote here, for short, α . The question is, given that the current is the convolution of the synaptic kernel α and the spikes g , $c = \alpha * g$, how to obtain the spikes back from the current. The theory of Green functions, central to solving differential equations, offers the way to associate to a convolution kernel α its corresponding operator (synaptic operator) such that $\hat{\alpha}[c] = g$. For linear combinations of exponentials, only derivatives are involved. In particular, for our double exponential the first and second derivative of the current c are needed. Derivatives are a double and sharp-edged sword in signal analysis because they amplify the noise. They are often used together with filtering to compensate for that. Here we determined (by looking at the data) that the second derivative brought about by the double exponential kernel does not benefit the detection process and we thus selected for practical reasons the simpler synaptic operator presented above in Equation 3.2, which corresponds to the simple mono-exponential synaptic kernel $\alpha(t) = \Theta(t - t_0) e^{-(t-t_0)/\tau}$ and has only one derivative.

Analytic deconvolution operators are nice, but application of the convolution theorem allows to solve for g even if no analytic expression for $\hat{\alpha}$ can be found. The convolution theorem states

$$c = \alpha * g \quad \Leftrightarrow \quad F[c] = F[\alpha] \cdot F[g],$$

where \cdot is the usual product of functions and F is one of a number of integral transforms. These statements for functions carry over naturally into the discrete domain for time series. The Fourier transform is the most commonly used of the conformant transforms and one of special interest because it is fast to numerically compute, using the fast Fourier transform algorithm (FFT; $O(N \log N)$ instead of $O(N^2)$). The deconvolved signal g before application of the synaptic kernel α can thus always be found as

$$g = F^{-1} \left[\frac{F[c]}{F[\alpha]} \right].$$

Selection of the synaptic operator parameter The choice of τ can be, in fact, guided by the data. Richardson and Silberberg (2008) use a variational approach, whereby the deconvolution is calculated for several trial constants τ_t and tabulated. A good single-time-constant deconvolution of a multiexponential event does not retrieve a perfect impulse, but should tail-off to zero. Richardson and Silberberg exploit this requirement and select the τ_t minimizing the long-time tail of the deconvolved trace. This is not an option for us, as there is often very little event tail available, and we resort here to using the τ_d from prior knowledge of the kinetics of spontaneous events and leaving to the reconstruction step its fine-tuning, together with the other PSC parameters. Incidentally, it is also no option for them, but they can assume the kinetics to remain the same during the experiment and thus can always use the last event to find τ_d . Applying the variational procedure for doublets is conceivable, but the approach reaches its limit well below the number of tightly packed events observed in our traces.

Standardized evaluation of reconstruction algorithms We propose that future implementations are benchmarked with regard to the three following, non-dimensional characteristics. These suggestions can be tweaked to suit the characteristics of the particular family of synaptic kernels under use, but the important point is to allow systematic comparisons by creating classes of equivalence of the input with respect to the reconstruction process.

asymmetry. the ratio between the time constant dominating the rise and that dominating the decay, i.e. roughly, here, τ_r/τ_d . Asymmetry is critical to any detection procedure leaning on derivatives (such as deconvolution). Note that a double exponential (Eq. 3.1) with asymmetry 1, i.e. $\tau_d = \tau_r$ is still rather asymmetric, which begs for a better name or a better characterization (e.g. 20-80% amplitude variations can be calculated analytically in terms of the time constants).

burst rate. the ratio between the duration of the event and the inter-event time interval Δ , i.e. roughly, here, $(\tau_r + \tau_d)/\Delta$. The burst rate is an expression of the amount of shadowing of one event by the next and is key to the fitting of time constants, be it on the deconvolved trace or in the original one. In some cases it may be of interest to look rather at τ_r/Δ , since overlapping in the rise phase poses qualitatively harder challenges.

signal to noise ratio. the quotient between event amplitude and a measure of the noise level (such as its standard deviation σ , if it is a Gaussian process), i.e. roughly A/σ .

The dataset reported here presented often asymmetries 1, burst rates of 1 to 5 and variable signal-to-noise ratios reaching down to 1 (standard deviation of ~ 5 pA with the lowest admitted PSC size decreed to be 5 pA).

3.7 Unweaving Single Post-Synaptic Currents from CPSCs

In the previous Section we presented a reconstruction algorithm designed to find out the component PSCs that make up a cEPSC or a cIPSC. Here we apply it to data from all cells in the excitatory dataset KGlu -66 (0..7) and to those in the inhibitory dataset CsGlu -6 (27..33) in order to characterize the onset times, amplitude statistics and kinetics of in-ripple PSCs.

Recapitulation of the peeling method For the reader who may have skipped that Section, we summarize here how peeling reconstruction works. It starts by detecting the first PSC with a sophisticated version of steep slope search of Section 3.4 (that uses deconvolution). It then finds out where the event onset is, and where the current starts to be shadowed by a succeeding PSC. Next, it fits its shape. Finally, it subtracts the whole reconstructed shape (including the previously unseen tail) from the remainder of the trace. The second event in the cPSC, which originally overlapped the first, starts now at amplitude zero and can be reconstructed following the same procedure, repeating until no candidate events are left. For the details, please see the previous Section (and especially Fig. 3.18).

Criteria for inclusion In selecting the cPSCs to include for analysis, we aimed at having a broad representation and not just those which are kind to us in some sense. Settling on a principled way to extricate mediocre reconstructions from unkind events is difficult, and so we decided to discard only those cPSC whose residual standard deviation (after peeling) is larger than that of the original cPSC (32 of 1,085 and 14 out of 862 for the excitatory and inhibitory datasets respectively). Since the criterion is conservative, we expect the results below to be a lower bound, i.e. to contain some noise from not-so-well fitted traces that we refused to exclude from the analysis.

3.7.1 Onset Timing in Comparison to Timing of Steep Slopes

Figures 3.22 and 3.23 below summarize onset times and amplitudes of reconstructed PSCs (in black) and compare them with steep slope times and magnitudes (in red).

Involvement in ripples is cell-dependent That some cells are better attuned to the LFP rhythm was visible already in our locking studies with steep onsets. The present analysis confirms and expands that conclusion, without relying on any comparison with the ongoing LFP. PSCs at cells 1, 2, 27 and, to a lesser extent, 33 stand above the spontaneous rate, but are not clearly organized into ripple cycles (in some cases the amount of data is not conclusive). Not only participation in oscillations is cell-dependent, the variable timing of histogram peaks with respect to the SWR maximum indicates as well a certain phase dispersion across cells in the same condition. We will discuss phases in the next Section.

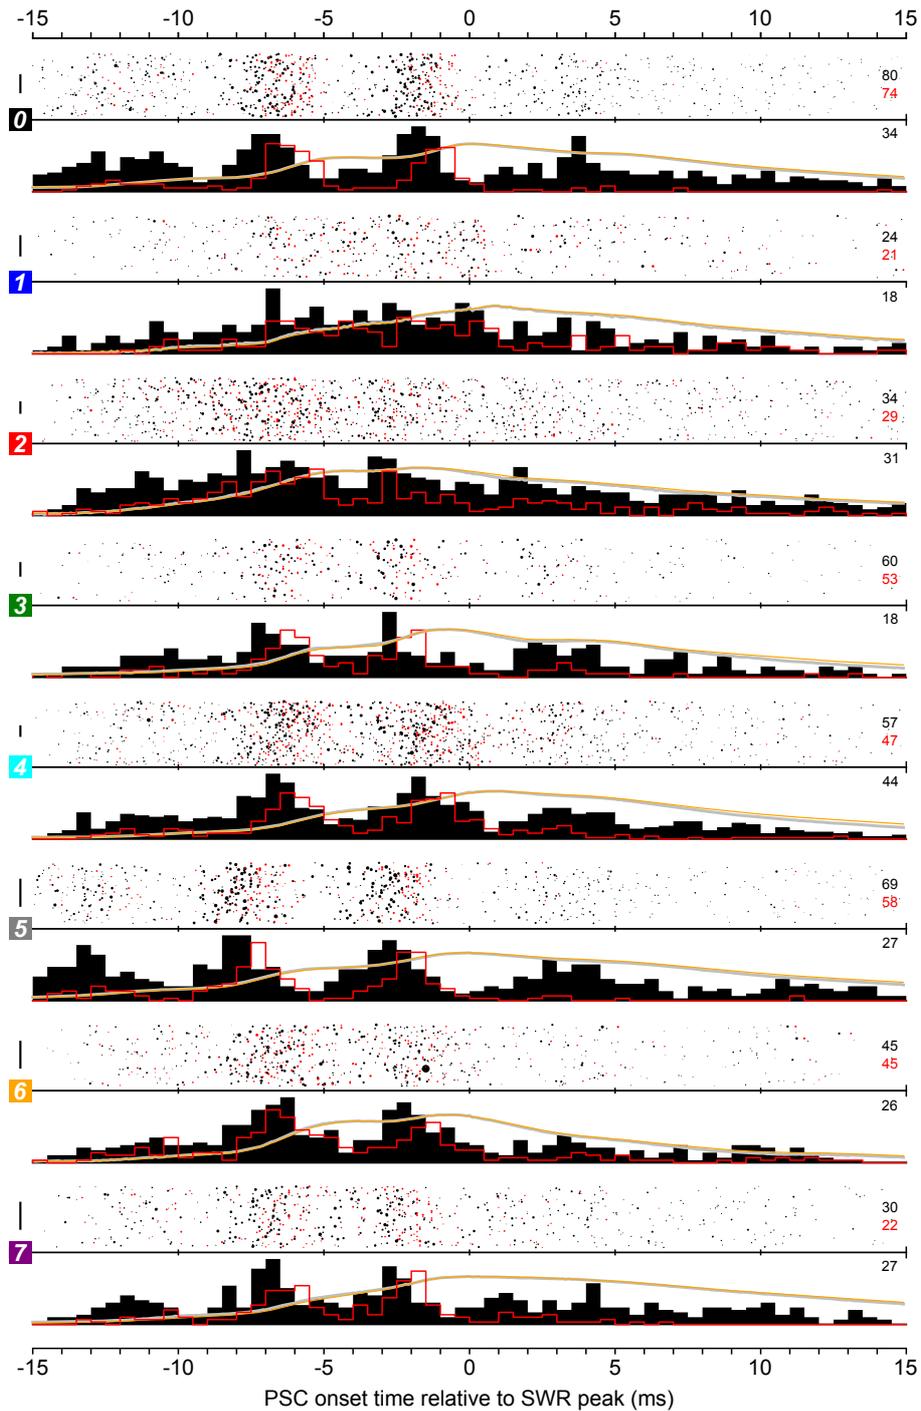


Figure 3.22. Summary of the Peeling Reconstruction Analysis (KGLu – 66). For each cell (identified by a colored square right of its data), a raster of presynaptic spikes is shown in the upper panel, whose in-ripple timing is summarized in the matching histogram below it. See caption to Fig. 3.20 for details. Reconstruction results (onsets and histogram) are in black; onset proxies derived from in-ripple 10% steepest slopes are in red. Inverted normalized cell-average cPSCs (grey) have been plotted on top of the inverted normalized average reconstructions (orange) to illustrate the goodness of fit at a population level. **Calibrations.** Dot area is proportional to reconstructed amplitude resp. steepness of slope for black and red dots, scaled by the cell’s median PSC amplitude resp. slope (printed in black and red rightmost, in pA and pA/ms respectively). Raster calibration bars over cell numbers span one minute of recording each; peak event counts of histograms are displayed at their upper right.

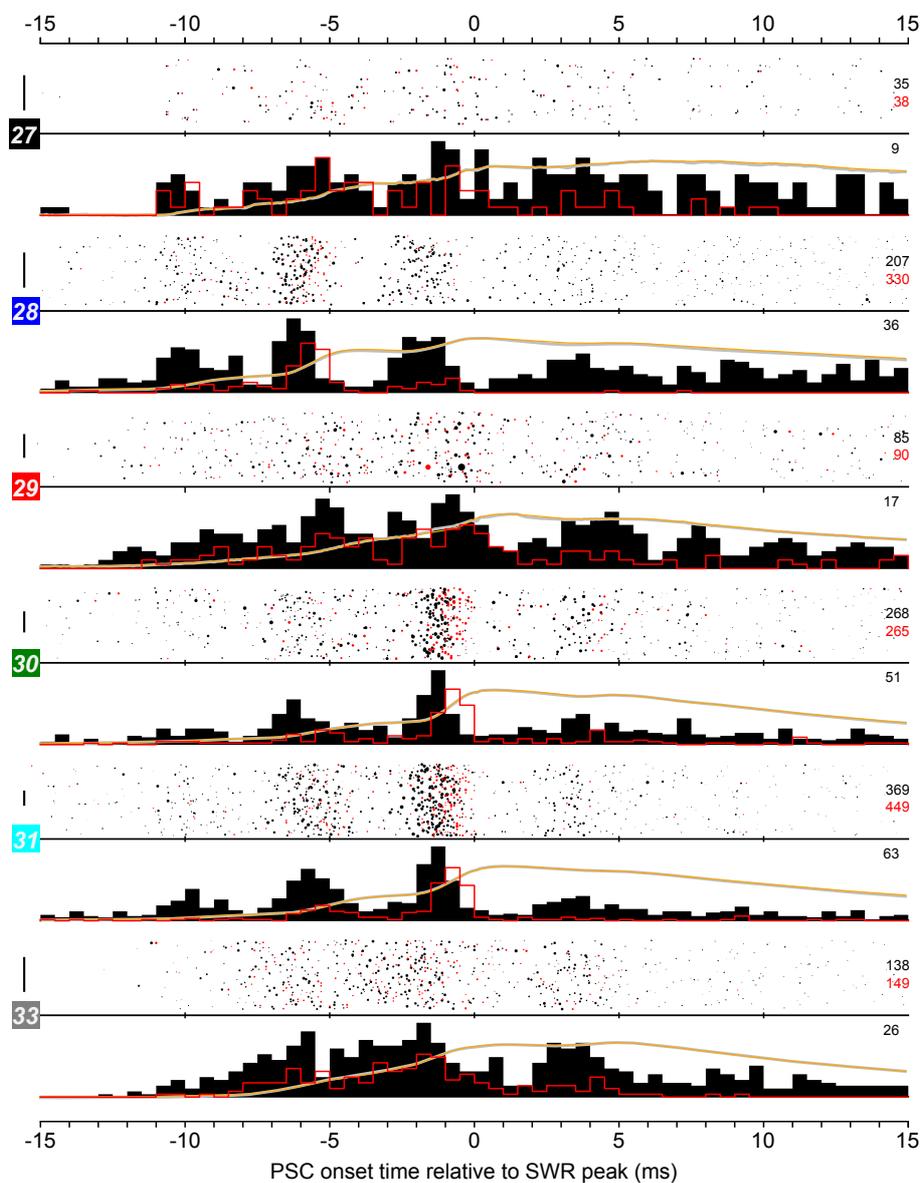


Figure 3.23. Summary of the Peeling Reconstruction Analysis (CsGlu —6). Rasters and histograms of detected events for the cells recorded at the reversal of excitation. Layout and calibrations as in Fig. 3.22.

Steep slopes are delayed and biased Steep slopes are simpler than reconstruction but suffer from two biases. One is ramp-up time. Slope-detected onsets are delayed by the time it takes the PSC slope to reach its maximum in the rise phase, where it is detected. This time is between 0.5 and 1 ms. Reconstruction data thus confirm a posteriori our assumption that this delay is short and rather homogeneous. Another bias is that the accumulation of events hinders detection of PSCs later in the ripple. Figure 3.24 depicts the problem. A first observation is that at 2% the central events (beats -1 and 0) are well detected, the initial ones (beats $-3, -2$) are slightly visible and the later ones (1) are completely neglected. Progressively increasing the threshold level (lighter red up to pink at 25%) makes late events visible, but at the expense of peak to valley contrast elsewhere. This is exemplified by the gap between beats -1 and 0 , where the contrast diminishes with threshold. By contrast, peeling reconstruction detects events at later phases of the ripple without compromising contrast elsewhere.

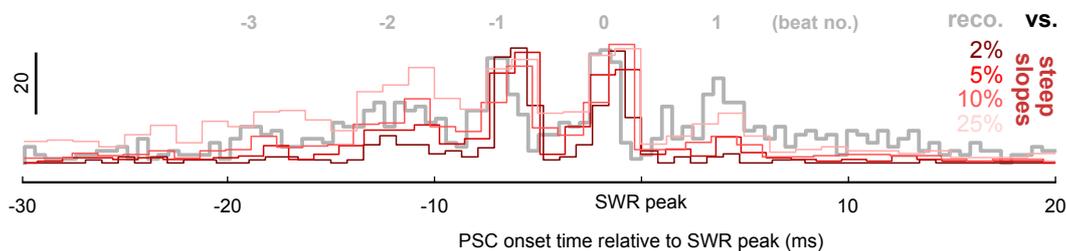


Figure 3.24. Steep Slopes vs. Peeling Reconstruction for Detection of PSCs. Histograms of EPSC counts during ripples for cell 0 comparing their detection by proxy of downward slopes (dark red to pink, 2% steepest to 25%) against full peeling reconstruction (grey). At top peaks are numbered (beat number) for reference in the text.

Locking statistics with reconstruction onsets We reanalyzed locking as in Section 3.5 but using onsets gained from reconstruction, without observing major differences (not shown). Slopes at 10% offer, as we just demonstrated, less performance, but they are biased in a robust way and involve less hypotheses. They also illustrate the evolution of our thought about how to unravel compound postsynaptic activity using time-domain techniques. For all these reasons we decided to maintain the discussion of locking within the framework of steep slopes.

In the next Section we study PSC parameters across cells, folding the time dimension.

3.7.2 In-Ripple PSC Rates, Amplitudes and Kinetics

We saw in Figure 3.21 that EPSCs during ripples matched the kinetics of spontaneous events, were larger than those (by up to a large factor), and arrived with frequency in the ripple range. That comparison was limited to one cell. We here ask whether PSCs are different across cells for cells in the excitatory (Fig. 3.25) and inhibitory dataset (Fig. 3.26), and what are the differences among EPSCs and IPSCs during ripples.

3.7.2.1 EPSCs During Ripples

We examine their frequency of arrival, their amplitude and their kinetics in turn.

Inter onset intervals The inter onset interval of EPSCs across cells is shown at top left in Figure 3.25. Distributions for different cells peak either at 2.5-3.0 ms or at 3.5-4.0 ms, suggesting a faster-than ripple succession of PSCs. The coincidence in all these distributions, some belonging to cells that show relatively mild ripple patterning (see Fig. 3.22, cells 1 and 2) is remarkable.

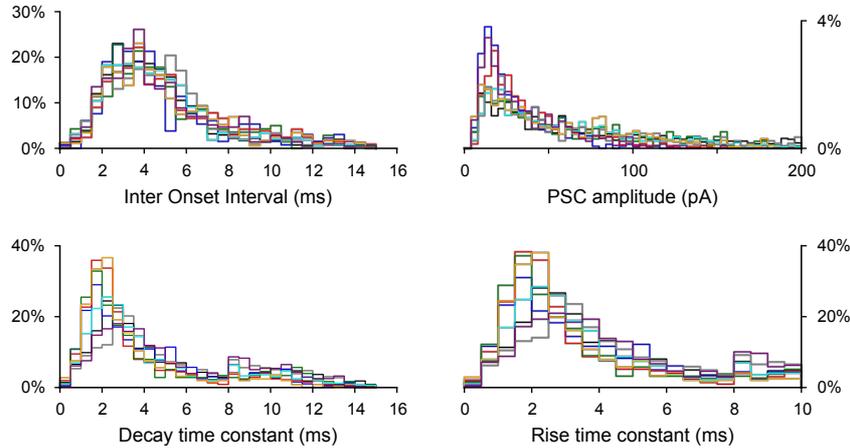


Figure 3.25. EPSC Event Rates, Amplitudes and Kinetics Across Cells. Histograms of inter-onset intervals and fit parameters for individual component EPSCs embedded in cPSCs.

Amplitudes PSC amplitude distributions differ in the relative importance of the low-amplitude peak, with less well ripple-connected cells showing smaller events, while agreeing on the general shape of the tail. No quantization structure is apparent in the distribution of amplitudes — Section 3.7.3 revisits this issue by breaking down PSCs in cycles according to their distance to the ripple peak.

Rise and decay time constants Kinetics are very similar across cells, but they differ in the proportion of fast decay events, which is $1.5 \times$ larger for some cells. A scatterplot of decay vs. rise time constants revealed that fast decay events are generally those where the least squares routine has settled for an alpha function, i.e. many points lay on the diagonal $\tau_d = \tau_r$. This superabundance at the diagonal suggests that least squares would reduce further τ_d so as to achieve a sharper turn of the exponential if it could (setting $\tau_d < \tau_r$ simply flips the sign of the PSC with model function Eq. 3.1). Faster events would however challenge the resolution of the filtered trace upon which reconstruction is based (400 Hz), and perhaps incur in overfitting. The secondary peaks at 10 or 8 ms reflect the boundary for the application of the penalty function for slow events. Refer to the previous Section for a discussion of the role of penalizing slow decays in making reconstruction work through long cPSCs.

3.7.2.2 Comparison of IPSCs vs. EPSCs During Ripples

IPSC feature distributions are presented in Fig. 3.26. Inter onset intervals are as in excitation for cell 33, mildly bimodal for cell 27 and somewhat larger (4.0-5.0 ms) for the rest. Only one cell (27) musters a PSC amplitude distribution with abundance of small events; the rest show a peak at 40 pA and tail off similar to EPSCs except that

their tail extends further to up to 400 pA at event counts similar to 200 pA (not shown in the Figure). While there are also very fast-decaying inhibitory events, a sizable proportion is in the long, fat tail that extends up to 15 ms. Likewise, inhibitory rises are considerably slower than EPSC rises. Heterodox cell 27, which stands out in every histogram, had only 43 SWR/cPSC events, of which 41 could be successfully fitted (the next cell by cPSC count had 118).

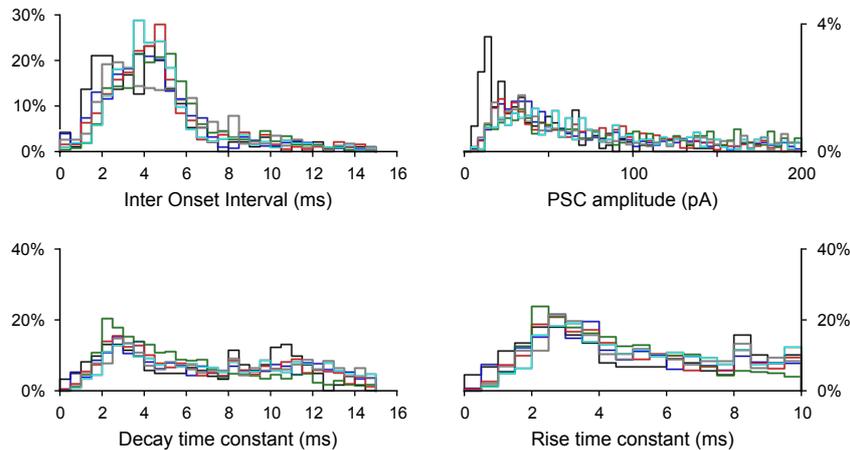


Figure 3.26. IPSC Event Rates, Amplitudes and Kinetics Across Cells. Histograms of inter-onset intervals and fit parameters for individual component IPSCs embedded in cPSCs.

In the next Section we drill down to the ripple cycles to investigate if several statistically distinct populations of PSCs are involved at different stages of the ripple.

3.7.3 Cycle-Wise PSC Features

How do the features of reconstructed PSCs depend on the ripple cycle? To answer this question, we adopted a criterion that assigns a ripple cycle number to each of the detected PSCs. Individual PSCs with onset time nearest to the SWR peak in each ripple were tagged as the *central* PSC (“0”). Proceeding outwards, we numbered previous PSCs as belonging to cycles $-1, -2, -3 \dots$ and succeeding ones $1, 2, 3 \dots$. The results obtained were qualitatively replicated when demanding that the central PSC onset time precedes the SWR peak as compared with just choosing the nearest.

Limitations of the analysis Another interesting view of the same data would be provided by binning the PSCs according to their phase in the SWR oscillation, which eliminates ambiguity with respect to the minority of cPSCs with skipped cycles.

3.7.3.1 Amplitudes: Towards In-Ripple Conductances

The statistics of ripple-cycle amplitudes are displayed in Figures 3.27 for excitatory PSCs, and 3.28 for inhibitory PSCs as stacked histograms, one stack per cell and one histogram panel per ripple cycle. Figure 3.28 also recapitulates the summary statistics for excitation next to those for inhibition in order to facilitate comparison. PSC amplitudes are normalized to the peak cPSC current to enhance comparability between cells.

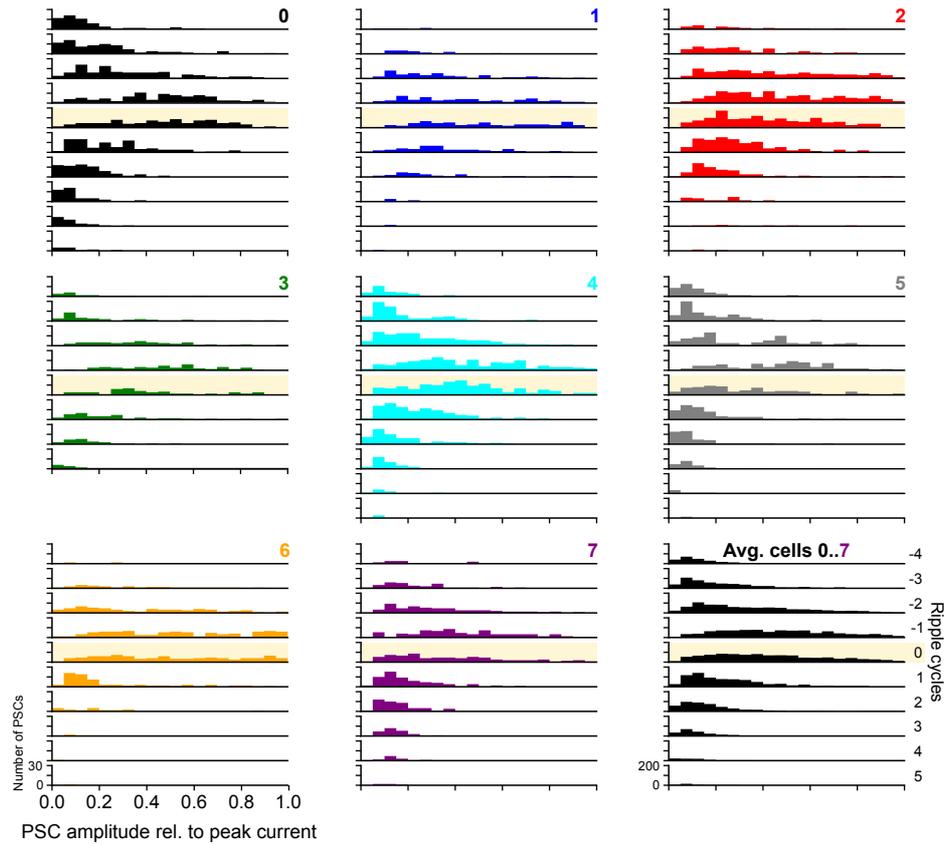


Figure 3.27. Amplitude Distribution of EPSCs Depends on Ripple Cycle (KGlū - 66). Each of the panels, corresponding to one cell, shows a stacked histogram of PSC amplitudes, one per ripple cycle. The cycle numbering was determined for each cPSC proceeding outwards from the closest PSC to the location of the SWR peak. That “central” PSC is marked as 0 and its corresponding histogram is highlighted. At bottom right, average of all cPSCs in the dataset at the KGlū -66 condition (*excitation*).

Looking first at Figure 3.27, the following aspects become apparent:

- Despite the normalization of amplitudes according to the peak current of each ripple, which has an average that is strongly cell-dependent (see Fig. 3.1), visible differences among cells remain. Results were similar when normalizing to the size of the largest individual PSC in each ripple.
- Amplitude distributions mid-ripple do not have a well-defined mode. In some cells the shape hints indeed at a multimodal character of the distribution, but these detailed features of individual cells are lost to a large extent upon averaging.
- Cells where the timing is well coupled to the ripple show a better defined shift of mean PSC amplitudes as well (0, 3, 5, for example).
- Pre- and post-ripple distributions (towards the outer cycles) are similar, thus improving our confidence in the reconstruction method.

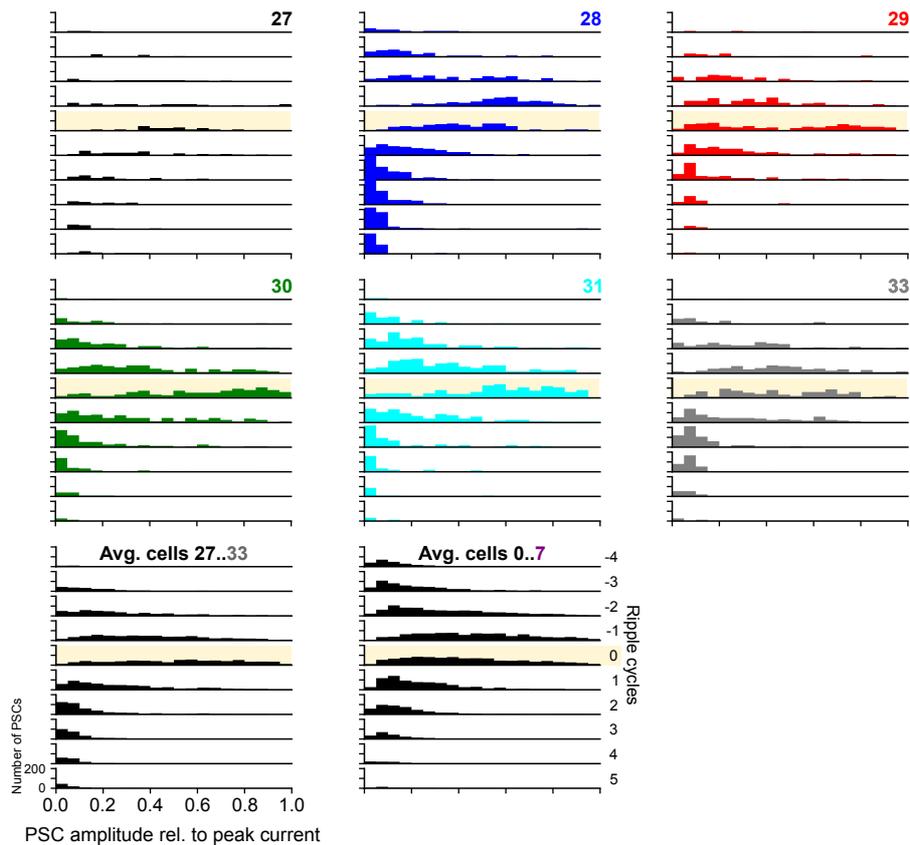


Figure 3.28. Amplitude Distribution of IPSCs Depends on Ripple Cycle (CsGlu – 6).

Each of the panels, corresponding to one cell, shows a stacked histogram of PSC amplitudes, one per ripple cycle. The cycle numbering was determined for each cPSC proceeding outwards from the closest PSC to the location of the SWR peak. That “central” PSC is marked as 0 and its corresponding histogram is highlighted. The bottom row shows the average of all cPSCs in the dataset at the KGlu –6 condition (*inhibition*) and reproduces the corresponding summary figure for KGlu –66 (*excitation*) from 3.27 to facilitate comparison.

Inhibitory vs. excitatory amplitudes The summary panels in Figure 3.28 reproduce the slight delay of peak inhibitory conductances observed already (e.g. in Fig. 3.3). Inhibitory currents start later and end later, and more slowly. While excitatory PSC amplitudes at cycles –1 and 0 are almost identically distributed, inhibitory events at the central cycle are larger than those just before. There seems to be a comparative near-absence of very low amplitude cEPSCs, a remarkable fact when one remembers that the larger overall magnitude of inhibitory currents is discounted via normalization. With regard to quantization, inhibitory histograms do show several peaks, hence encouraging further research to establish the atomic unit of synaptic contribution during ripples. Like for excitation, however, they are washed out upon averaging.

Spikes timed by ripples Figure 2.26 shows that in current clamp under block of inhibition, spikes tend to happen primarily during cycle 0 and cycle 1.

Estimation of conductances The data presented here allow to build a probabilistic model of synaptic current during ripples to be hooked into a spike generation model. Indeed, amplitudes of individual PSCs during ripples in voltage clamp are our best handle on the underlying conductances, and can be used to derive a probability distribution of conductances as a function of time, a route that we will pursue elsewhere.

3.7.4 PSC Kinetics Are Stable Over the Course of the Ripple

Figure 3.29 shows side by side the summary figures for PSCs recorded from excitatory cells 0-7 and inhibitory cells 27-33. No individual cell histogram stacks are shown because the variability among cells of the time constants was substantially less than in the case of amplitudes.

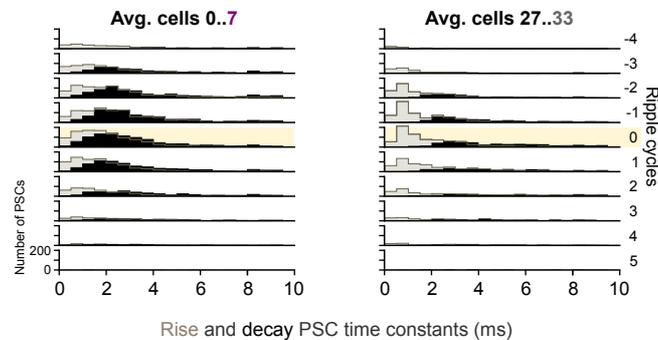


Figure 3.29. In-Ripple PSC Time Constants. Two stacked histograms compare the kinetics of PSCs for the KGlu -66 and KGlu -6 datasets, shown left and right respectively. Rise time constants τ_r are shown in grey-brown while the black bars show the decays τ_d . Ripple cycles were assigned as for amplitudes (Figs. 3.27 and 3.28).

Kinetics during ripples are rather stationary with the exception perhaps of the second post-peak cycle. Both excitation and inhibition display very fast characteristic rise times in the sub-2 ms range, with inhibition showing enhanced precision (note marked consecutive peaks of the rise constant at 1 ms during cycles $-2, \dots, 2$). Decay times for AMPA_R-mediated currents are slightly faster than for GABA_A_R ones, as seen in the consolidated histograms of the previous Section.

We can conclude from this negative result that EPSCs and IPSCs during ripples use each the same synaptic mechanisms over the course of the ripple.

Continuing in the spirit of dissecting the ripple into its cycles, in the next Section we present analysis that access subripple dynamics of onset timings.

3.7.5 Phase Drift Aligns Excitation and Inhibition During the Ripple

The preceding analyses of locking, reconstruction onset timings and in-ripple kinetics taught us that inhibitory currents do behave in delicately different ways to excitatory currents. First, the cPSC peak of inhibitory currents is reached later. Second, phases of inhibitory PSCs follow those of excitatory PSCs by 0.7 ms in average (54°). Third, their kinetics are less precise and slower. We wanted to know if timing discrepancies were constant during the ripples or there was a more sophisticated relation between the two types of current.

Figure 3.30 shows the dynamics of ripple-locked excitatory and inhibitory currents (from cell 0 and cell 28 respectively). This is just one example of the many pairs that can be built where inhibitory onsets start out in the ripple behind excitation and end up in alignment.

To establish this observation statistically, we display in Figure 3.31 the averaged fitted onset histograms for 8 and 6 cells, respectively. During the initial course of ripples, excitation is slightly phase-advanced, leading inhibition by ~ 1.5 ms. In later

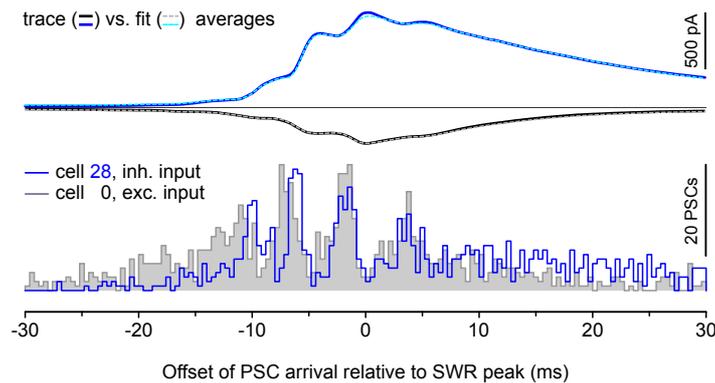


Figure 3.30. Relative Shifting of Excitatory vs. Inhibitory Phases during a Ripple. At bottom, histogram of onset times of currents incoming into cell 0 recorded with KGlu 0 at -66 mV vs. onsets for cell 28, recorded with CsGlu at -6 mV (both obtained via peeling reconstruction). At top, mean measured and mean reconstructed traces for both cells.

periods, the phases of the two components converge (phase difference plotted in orange in Figure 3.31). Such a phase delay is in qualitative agreement with the observation from vector strength analysis of steep slopes, that inhibition lags excitation in average by 0.7 ms over a ripple cycle.

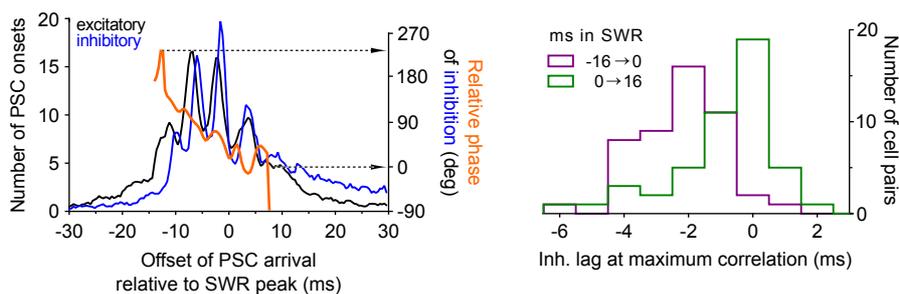


Figure 3.31. Phase Drift at Population Level. The progressive alignment of the inhibitory to the excitatory phase during the ripple can be verified for most of the 48 combinations of excitatory ($n=8$) and inhibitory ($n=6$) histograms. **Left.** Average excitatory (black) and inhibitory (blue) histograms of onsets. Their local Hilbert phase difference (orange) demonstrates a progressive reduction of the lag of inhibitory with respect to excitatory onsets. **Right.** Peak-correlation lag histograms derived from 48 excitatory/inhibitory onset histograms like Fig. 3.30 (each count represents a combination of one excitatory and one inhibitory cell). Inhibition follows excitation by ~ 1.5 ms early in the ripple (purple) but aligns later in the oscillation (green).

Note that in the Figure, the later start of inhibitory ripples, their delayed peak and their longer tail are all visible, in agreement with analyses conducted on raw currents.

The shrinking of the lag can be confirmed in another population analysis by comparing the lags of correlation peaks determined for all 48 excitatory-inhibitory cell pairs early vs. late in the ripple (Fig. 3.31, right panel). Cross-correlation peaks computed on the earlier period, between -16 and 0 ms relative to the SWR peak, clustered around -2 ms (median: -2.0 ms; purple histogram) whereas those computed between 0 and $+16$ ms clustered around 0 ms (median: 0 ms; green histogram). Collectively, these analyses reveal high precision of ripple-associated inputs and a progressive synchronization of excitation and inhibition over the course of ripples.

Chapter 4

On Sequence Memory

Leibold and Kempter introduced in 2006 a model for networks capable of storing sequences that is based on binary neurons, binary synapses, sparse connectivity and activity, the Willshawian clipped-Hebb plasticity rule (Willshaw et al., 1969) and a success metric for the readout, whereby items are not considered stored unless they can be retrieved at a given quality. In that work Leibold and Kempter calculate the capacity of sequential memory in a recurrent network by means of a probabilistic theory as well as a mean field approach. This foundational model is henceforth referred to as LK2006 (Leibold and Kempter, 2006).

In this Chapter we first define the technical terms of the previous paragraph. We describe the LK2006 model and extend it to account for inhibition. We demonstrate that inhibition can substantially increase the capacity and robustness of a sequence replay network by engaging into oscillatory feedback. The reader wishing to skip forward to the main results may consult Figure 4.18 and the last Section 4.4.

We start in **Section 4.1** with a description of the LK2006 *cellular model*, inspired by CA3 region in the hippocampus, and show with the help of simulations how it can perform replay of stored memories for certain well-chosen sets of its defining parameters.

The many neurons present in mammalian hippocampal tissue are expensive to simulate even when modelled quite simply. High neuron numbers are however of advantage to build a *mean-field model* of the cellular activity that allows easier study of the factors that determine replay, as we explain in **Section 4.2**. We concentrate particularly on the tradeoff between the neuronal firing threshold and the number of active neurons that represent a memory, and leave fixed the total number of neurons as well as their connectivity.

In **Section 4.3** we deploy the tools of dynamical systems theory in order to understand how to design networks that are capable of replaying a maximum number of stored memories. We also ask from the perspective of optimal detector theory what is the optimum threshold for memory replay, and obtain the answer that it depends on the network activity level.

Guided by the insights gathered along the way, in **Section 4.4** we set out to investigate the role of inhibition in sequence replay. A simple inhibitory feedback makes replay stable when memories are encoded with few active neurons, thereby enhancing the network's memory capacity for sequences. At the same time, inhibition makes replay more robust when neuronal thresholds are not homogeneous. Our results are obtained by manipulation of the neuronal threshold (*instantaneous inhibition*), confirmed by a mean-field theory incorporating a third dynamical equation (*dynamic inhibition*), and validated by cellular simulations.

4.1 Sequence Replay With Binary Neurons

Any small volume of neuronal tissue poses formidable challenges to our understanding: we confront a staggering array of neuronal types, in electrical communication via synaptic and non synaptic pathways whose configuration and transfer properties change over time in an experience-dependent manner. Obtaining analytical insight about the factors that influence the function of the network definitely requires simplified models. These models must neglect all complexities that have no bearing on the particular function under study, that affect separate (too long or too short) timescales or that would challenge their prediction power.

We start out in Section 4.1.1 by presenting such a simplified model of the hippocampal CA3 region as a cellular replay network whose state is updated in discrete time. Before any replay can take place, the network must be loaded with memories. We discuss this memorization, or imprinting process in Section 4.1.2. We next define formally the capacity of the network as a measure of its functional performance (Section 4.1.3). The theoretical capacity limit can be realized only for some network configurations: those that enable replay in sequences of specified length. We discuss this dynamical criterion in Section 4.1.4. Finally, we describe in Section 4.1.6 the architecture of the high-performance code that simulates replay networks (by Axel Kammerer) and discuss representative samples of its output in Section 4.1.7.

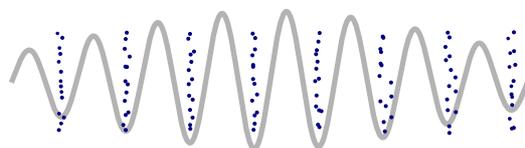
4.1.1 Assumptions of the LK2006 Model

Here we discuss the main abstractions involved in the model: discrete time to represent oscillation cycles as units of computation, coding by principal neurons only, binary character of neurons and synapses, sparseness of neuronal activation, sparseness of connectivity, and the definition of a combinatorial code based on the activation of subsets of neurons. We finally discuss how to evaluate the performance of a replay network.

4.1.1.1 Structural Assumptions

These are the constitutive hypotheses of the model.

Oscillations and discrete-time dynamics A central problem in neuroscience is the notion of simultaneity of neural activity. An activity pattern (combination of neurons that fire, and when) is synchronous in a practical sense if it fits within the integration time of the relevant readout structures. Here we assume that single cycles in neural oscillations provide a rhythmical basis for such an integration. Thus, one theta cycle (~ 200 ms) would subserve the encoding of a pattern in the sequence and one ripple cycle (~ 5 ms) its replay. On first approximation, it is enough to account for those network state changes across cycles, and to consider activity within a cycle to be effectively simultaneous. These assumptions allow to dispense with continuous time and write a discrete map for the dynamics of the variables of interest.



Coding by principal neurons Experiments in the hippocampus have concentrated so far mainly in the relation between activity of principal neurons and external stimuli. It seems that the firing of pyramidal neurons correlates better with the behavioral experience than that of inhibitory interneurons, and projection patterns of interneurons in cortex are highly non-specific (Fino and Yuste, 2011). The original model thus does away with inhibitory neurons; in our extension (Section 4.1.5) we do introduce an inhibitory population. Following the observation above, however, this inhibitory population has a dynamical, not coding function. In particular the synapses of inhibitory neurons are established at random, independently of the memories loaded on the network.

Neuronal state Neurons, or *units*, are assumed to have binary state. Firing is the state “1” and silence the state “0”. Units integrate their input linearly and compare it to a threshold θ ; for simplicity all units have the same threshold. A presynaptic spike is translated into a postsynaptic membrane potential increase of one. The linear integration performed by individual units extends back just one step (i.e. takes into account activity of afferents in the previous cycle only). We shall use a vector with N components, denoted $\mathbf{x}(t)$, to represent the binary firing state of the network at cycle number t i.e. $x_j(t) \in \{0, 1\}, j = 1, \dots, N$. In the rest of the Section we seek to write out explicitly a discrete update rule for the network state during replay with this form:

$$\mathbf{x}(t+1) = f(\mathbf{x}(t); \mathbf{p}_x) \quad (4.1)$$

where \mathbf{p} are the parameters that configure the network and neither f nor \mathbf{p} are time-dependent. We then explore its behavior in cellular simulations of replay. When including inhibition we will track an additional population of K neurons and write the update rule in terms of $\mathbf{u}(t)$ with $u_j(t) \in \{0, 1\}, j = 1, \dots, N + K$:

$$\mathbf{u}(t+1) = f(\mathbf{u}(t); \mathbf{p}_u). \quad (4.2)$$

Memory traces A *memory trace* is a pattern of activation of the network, hence also called *memory pattern*. Here, we assume that all such patterns involve the firing of same number of randomly chosen neurons, M (*pattern size*) and the silence of the remaining $F \equiv N - M$. Thus, patterns are represented by binary arrays ξ of length N with $\sum_{i=1}^N \xi_i = M$. We assume that patterns in the hippocampus represent an episodic memory built from the confluence of perceptions from several sensory modalities (Fig. 4.1).

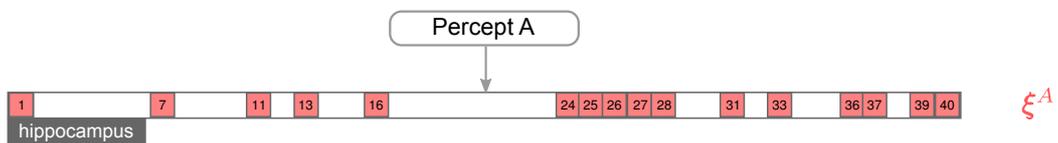


Figure 4.1. Combinatorial Neural Code for Sensory Percepts. A multi-sensory percept is encoded as a binary pattern of activity of N excitatory units, M of them active and $F = N - M$ inactive.

Associations An association is an ordered pair consisting of a cue pattern ξ^A and a target pattern, ξ^B . An association is considered to be stored if approximate activation of the cue leads to the network activating approximately the target in the next time step (symbolized $\xi^A \rightarrow \xi^B$).

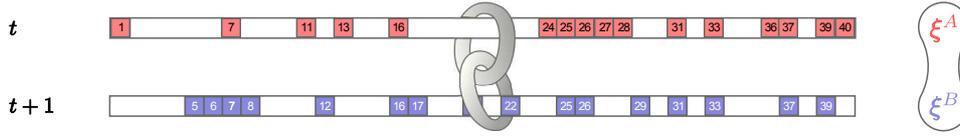


Figure 4.2. Association of Memory Traces. An association of a cue pattern ξ^A (active neurons in red) representative of a multi-sensory percept A with a target pattern ξ^B (active neurons in blue) representative of a percept B . An association can be seen as a link in a chain, or sequence, of memories.

Associations embedded in sequences A sequence of length Q is a chain of $Q - 1$ associations (“ \rightarrow ”),

$$[\xi^1 \rightarrow \xi^2 \rightarrow \xi^3 \rightarrow \dots \rightarrow \xi^{Q-1} \rightarrow \xi^Q].$$

Its successful storage requires that approximate initialization with the initial cue ξ^1 leads to the network autonomously replaying the rest of the sequence with high fidelity. Precisely, an appealing property of associative memories that becomes crucial in the context of sequences, is that they can retrieve a target pattern from a damaged cue pattern. There is obviously a tradeoff between the allowable deterioration of the patterns (the network’s resilience) and the overall capacity of the system. By feeding the target pattern back as the next cue pattern, here we challenge the capability of the network to *repeatedly* recall target patterns from distorted cue patterns.

Encoding and retrieval The storage of memories (encoding) and its retrieval are assumed to happen during different network operating regimes that are mutually exclusive. There is experimental support for the existence of an *encoding* and a *retrieval* state in the hippocampus (Chrobak et al., 2000), each tied to one dominating hippocampal rhythm. Briefly, the standard view on the switching between the two modes is that the retrieval state, associated with sharp waves, kicks in when a global inhibitory brake is released and thus fails to prevent CA3 bursts that propagate to CA1 (Chrobak and Buzsáki, 1996). Encoding, assigned to the theta oscillation, is by external sequential presentation of patterns $\rho = 1, \dots, P$ causing changes in the states of the synapses. Retrieval, or replay is the output of the spontaneous dynamics upon presentation of an initial cue, one memory per oscillation cycle, $t = 1, \dots, P$.

Morphological synapses Units are connected by a *morphological synapse* if there exists a physical substrate between them susceptible of acquiring functional transmission properties. The morphological connectivity is randomized with probability c_m , i.e. neuron i feeds into neuron j with probability

$$\Pr(w_{ij} = 1) = c_m. \quad (4.3)$$

However, morphological synapses are initially all *silent*. How do they become active?

Functional synapses and storage of memories The process of storing memories in the network consists in the activation or strengthening of particular synapses to a *potentiated* state. The state of potentiation is represented with the binary variable s , which equals one if the synapse transmits presynaptic spikes into contributions to the postsynaptic membrane potential, and zero otherwise. The probability that an existing synapse feeding from neuron i to neuron j is in the potentiated state grows during memorization. Accordingly, it is denoted

$$\Pr(s_{ij} = 1) = q_\rho, \quad (4.4)$$

where the subindex ρ counts how many memories have been stored so far. We abbreviate $q = q_P$ at the end of the storage process: it is the (final) *probability of potentiation*. The resulting effective connectivity is called c , corresponding to the fraction of synapses that is operational at the end of storage out of the total number N^2 of combinatorially possible synapses, i.e. $c = q c_m$. The difference between morphological and functional connectivity is the silent connectivity $c_s \equiv c_m(1 - q)$. The ratio $q \equiv c/c_m$ is the complementary fraction of morphological synapses that remain available for encoding of even further memories. In order to compare with the literature it is convenient to express those *resources for plasticity* rather in terms of the ratio of silent to potentiated synapses, $r \equiv c_s/c = (1 - q)/q$, which is more readily estimated in experiment (anatomically in CA1; Megas et al., 2001 or physiologically, in CA3; Montgomery et al., 2001). In the model presented here, a synapse that has been potentiated once cannot be silenced (a possible mechanism for *forgetting*) and cannot be further potentiated either.

4.1.1.2 Quantitative Assumptions

The model can gain biological relevance by the judicious choice of its parameters.

Sparse activity A prominent feature of neural activity, and in particular of hippocampal pyramidal neurons, is that firing is infrequent or *sparse*. We define the fraction of neurons whose activity could represent a pattern as the *coding ratio* $f \equiv M/N$ of the system. The network is expected to replay the memories at high quality, i.e. with a “firing rate” (computed over one time bin, or cycle) close to this number. The assumption of sparseness is then

$$\text{sparse activity: } M \ll N. \quad (4.5)$$

In order for the mean-field model (Section 4.2) to furnish a good approximation, both M and N have to be substantially large, i.e. $1 \ll M \ll N$. As we shall see below, the biologically reasonable values of $N = 10^5$ and $f \simeq 1\%$ do well enough^{4.1}. These numbers are supported by anatomical studies (West et al., 1991 give $N \simeq 2.5 \times 10^5$ for *regio inferior* in rat hippocampus) and electrophysiological recordings (Csicsvari et al., 2000; f ranges from 1 to 10% as the window goes from one ripple cycle to the whole oscillation, a rate of 2 Hz is equivalent to 0.01 spikes per ripple cycle assuming that all of the active neurons in a pattern, and only them, fire in a cycle).

Sparse recurrent connectivity The hallmark of the CA3 region of the hippocampus is the abundance of recurrent connections between pyramidal cells. Recurrence being an essential ingredient of an associative network, models have often been inspired by CA3. The reference value of effective connectivity of about 5% is large (neocortex: 1 in 70 to 100, CA3 1 in 60, CA1 1 in 100; Deuchars and Thomson, 1996), but still far from full connectivity as e.g. assumed by Willshaw et al. (1969). Consequently, it is important to transpose models of memory based on densely connected networks to the more biologically realistic case of sparse connectivity. This is done here by the introduction of the morphological connectivity $c_m < 1$ and the assumption of sparseness.

$$\text{sparse connectivity: } c_m \ll 1. \quad (4.6)$$

^{4.1} A coding ratio of 1% represents a firing rate of about 2 Hz assuming that the M neurons spike in a replay interval with the duration of a ripple cycle (about 5 ms), see *Oscillations and discrete-time dynamics* in Section 4.1.1.

Synaptic weights The weights of synapses represent the translation factor between a presynaptic spike and a postsynaptic effect in the form of a change of the target neuron's membrane potential. The synaptic weights of recurrent excitatory synapses is set to one, i.e. $w_{EE} = 1$ because any freedom in this factor would be compensated by a scaling in the firing threshold θ .

4.1.1.3 Full Specification of the Cellular Model

The binary firing state $x_i(t+1)$ of neuron i at iteration $t+1$ can be related to the firing state of the whole network at time t , $x_j(t)$, $j = 1, \dots, N$ via the binary connectivities w_{ij} and the binary synaptic states s_{ij} using the following deterministic update rule:

$$x_i(t+1) = \begin{cases} 1 & \text{if } \sum_{j=1}^N w_{ij} s_{ij} x_j(t) > \theta \\ 0 & \text{otherwise.} \end{cases} \quad (4.7)$$

In vector form, it reads

$$\mathbf{x}(t+1) = \Theta(E^E \mathbf{x}(t) - \boldsymbol{\theta}), \quad (4.8)$$

where E^E is the effective recurrent connectivity matrix with elements $e_{ij} \equiv w_{ij} s_{ij}$, $\boldsymbol{\theta}$ is a vector of length N with the identical thresholds θ as entries everywhere and $\Theta(x)$ is the elementwise Heaviside step function with value one for $x > 0$, and zero elsewhere.

4.1.1.4 Assessment of the Replay Functionality

How to evaluate the performance of a replay network?

Hit and false alarm neurons For every pattern, the M neurons which should be active are called for short *On neurons* or *hit neurons* (light red in Fig. 4.1). The remaining $F = N - M$ —in white—are called *Off neurons* or *false alarm neurons*, since activating one of them at t is ringing the bell at the wrong time. We shall occasionally use the subindices 1 and 2 to indicate respectively On and Off neurons.

$$\begin{aligned} \text{On}_t \text{ neurons} &\equiv \{j(t): \xi_j^t = 1\}, \\ \text{Off}_t \text{ neurons} &\equiv \{j(t): \xi_j^t = 0\}. \end{aligned} \quad (4.9)$$

Activation variables In order to characterize the quality of replay it suffices to consider how many of the On, or hit neurons, and how many of the Off, or false alarm neurons are active at any one time step. We call these activation variables^{4.2} respectively *hits* $m_t \in 0, \dots, M$ and *false alarms* $n_t \in 0, \dots, N - M$.

$$\begin{aligned} \text{hits} & \quad m_t \equiv \mathbf{x}(t) \cdot \boldsymbol{\xi}^t \\ \text{false alarms} & \quad n_t \equiv \mathbf{x}(t) \cdot (\mathbf{1} - \boldsymbol{\xi}^t). \end{aligned} \quad (4.10)$$

4.2. It may make more sense to the reader at this point to rather consider *success variables*: hits and true negatives $N - n_t \in 0, \dots, N - M$, since then our goal of good replay corresponds to the concurrent maximization of both quantities. The equations that we will derive below are more symmetric, however, if we stick to either activation (or inactivation variables) as opposed to success or failure variables.

From the point of view of a readout, hits represent the signal and false alarms the background, or noise, that jeopardises the identification of a memory item.

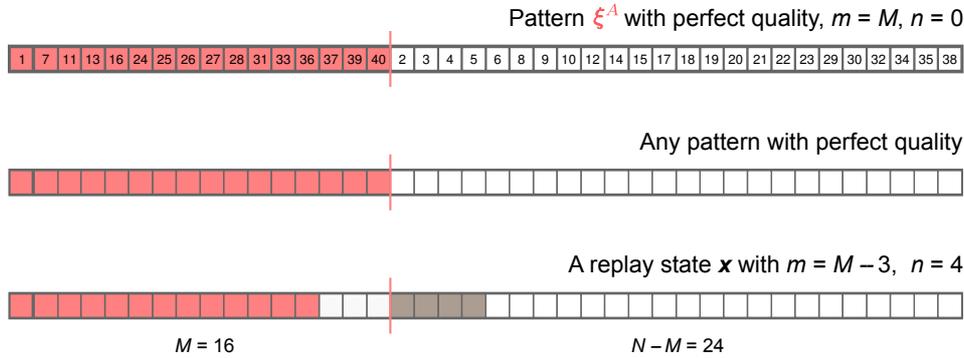


Figure 4.3. Activation Variables. At top, pattern ξ^A from previous examples, now with all firing neurons laid out contiguously. Pattern ξ^A here is reproduced perfectly, i.e. $m = M = 16$, $n = N - M = 24$. The transition to activation variables is sketched in the middle: in order to assess quality, pattern identity is not necessary. At bottom, a realization of the pattern during replay with three false dismissals $M - m = 3$ and four false alarms, $n = 4$; the $\binom{M}{m}$ ways to obtain m hits in M On neurons, as well as the $\binom{n}{N-M}$ ways to arrange n false alarms in N Off neurons are statistically equivalent when deciding which neurons should fire next.

Quality Both false dismissals and false alarms in a cue pattern challenge the correct recall of the target pattern. Deciding on a quality measure requires knowledge of the readout structure, and in particular of the relative importance of false dismissals and false alarms. A naive choice is

$$\Gamma_t \equiv m_t/M - n_t/F. \quad (4.11)$$

This definition emphasizes the false dismissals in the regime of interest of low coding ratios, $f = M/N \ll 1$. Note that $-1 \leq \Gamma_t \leq 1$. A predominance of negative values would induce us to switch our perspective and observe silences as responsible for coding, instead of discharges.

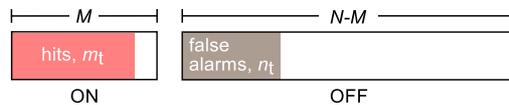


Figure 4.4. Populations of On and Off Neurons. During successful replay the majority of the On neurons should stay active (hits) and the majority of the Off neurons should be silent (correct dismissals). The inactive (white) neurons in the ON group are false dismissals (g_t). The inactive (white) Off neurons are true negatives.

4.1.2 Imprinting

Here we discuss an adapted Willshaw synaptic plasticity rule for sparsely connected networks and the noise that random morphological connectivity induces in pattern recall.

Biological motivation According to paired recordings in the hippocampal CA3 region of the rat, only about 10% of all possible connections have a synaptic morphological correlate. Importantly, despite the many morphological contacts between CA3 neurons, there exist “all-silent” pairs for which none of the connecting synapses is active. It has been estimated that, of all morphological connections, at most one half is silent in a functional, loaded hippocampal network (Montgomery et al., 2001; Nusser et al., 1998).

Here, the network is loaded with associations as described by Willshaw et al. (1969). As part of our program to understand factors that influence the capacity of biologically relevant networks, and following LK2006, we assume however *sparse* connectivity under the same imprinting rule.

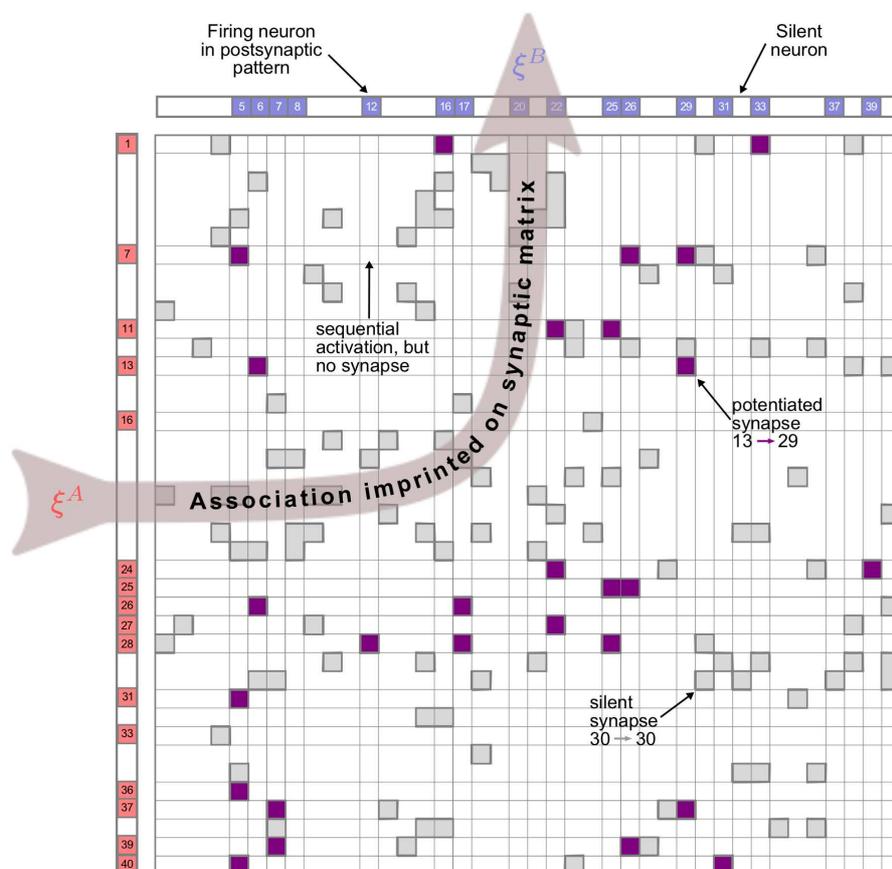


Figure 4.5. Storage of an Association in the Synaptic Matrix. Synaptic matrix after a sparsely connected network is loaded once using the Willshaw rule with the association shown in Fig. 4.2. Synapses are represented as intersections in the grid. White squares stand for non-connected cells. Grey squares signal those of the morphological connections that remain silent. Purple squares are the memory that binds the traces ξ^A and ξ^B , and correspond to morphological synapses potentiated by the sequential firing of a unit active in the cue pattern (red) and one active in the target pattern (blue; see also Fig. 4.6.) Note that the same cell may participate in several memory traces (e.g. number 7), whether associated or not. In this illustration, total neuron number is $N = 40$, pattern size $M = 16$, morphological connectivity $c_m = 20\%$ and one-step loading $c_1 = 27/1,600 \simeq 1.7\%$ (purple). This network would have a maximum capacity of $P \simeq 4$ associations (Eq. 4.12 below), for a value of remaining plasticity resources $r = 1$ (down from $r_1 = 10.9$ in the Figure, with just one association).

In detail, associations between patterns presented in sequence are stored by potentiating all of the morphological synapses available between sequentially active neurons. Already potentiated synapses are not further strengthened, thereby justifying the name *clipped Hebb rule* (Nadal, 1991). The loading of the network with a sequence consists in the serial presentation of the memory associations $\rho = 1, \dots, P$ for their step-wise imprinting on the plastic synapses (Fig. 4.5). As the loading of sequences proceeds, so grows the average fraction of morphological synapses in the potentiated state^{4.3}, $E[q_\rho] = c_\rho/c_m$. To investigate the maximum number of associations whose replay can be supported by a network thus imprinted, it is convenient to express the pattern loading P in terms of the effective connectivity reached at the end of the process, $c_{\rho=P}$, which we abbreviate as c . By choosing c while we require replay to be executed at a given quality, we effectively constrain other parameters such as e.g. pattern size M and firing threshold θ .

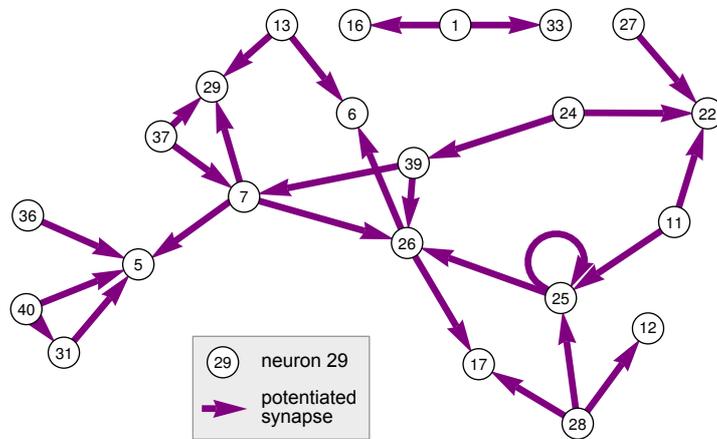


Figure 4.6. Effective Synaptic Network Induced by an Association. The association established in Fig. 4.5. An arrow $i \rightarrow j$ represents a morphological synapse that has been potentiated, corresponding with entry ij of the effective synaptic matrix E^E (Eq. 4.7).

Noise and variability There is no *dynamical* randomness in this model, i.e. the prescription for the network state update is entirely deterministic without noise. However, the random morphological connectivity substrate (grey squares in Fig. 4.5) causes some associations to be more faithfully represented than others and thus a corresponding variability in the activations. The loss in quality indirectly brought about by the random connectivity challenges the resilience of the system in a manner that is representative of that induced by a dynamical source of noise.

4.1.3 Capacity

In this Section we address the problem of quantifying how many memories can one *load* in a network. We postpone to the next Section the question as to whether the network can effectively *replay* them. Indeed, networks can be loaded with too many memories for effective replay much as paper grocery bags have a practical weight limit below their theoretical volume capacity.

^{4.3} An average comes into play because of the randomized state of the morphological connectivity makes uncertain just how many synapses will be potentiated at each step.

Orthogonal patterns The number of patterns stored at a fixed loading level corresponding to potentiation of cN^2 synapses can be roughly estimated under the assumption that patterns are *orthogonal*, i.e. no neuron participates in two different patterns. In that case, the number of patterns is the ratio between potentiated and morphological synapses

$$P = cN^2 / c_m M^2 = c / (c_m f^2).$$

Assuming orthogonality of patterns is warranted if the coding ratio $f = M/N$ is very low; then the P patterns of size M , which are composed of a random selection of cells, are unlikely to collide.

Nonorthogonal patterns A synapse is potentiated if the presynaptic partner *and* the postsynaptic partner are both activated in sequence. The active members of a pattern are chosen at random and independently of those of any other pattern, with probability $f = M/N$. The probability that a single association potentiates any one synapse is thus the product f^2 . Due to the clipping nature of the plasticity rule, one cannot distinguish if a given synapse was potentiated one or several times. Hence we rather consider the complementary situation, *viz.* the silent state, which occurs for one synapse with probability $1 - f^2$. Since patterns are independent, a synapse remains silent after imprinting P associations with probability simply given by the product $(1 - f^2)^P$. Equivalently, a synapse will have been potentiated with probability $1 - (1 - f^2)^P$. Our criterion for halting the imprinting procedure was that the network would reach an effective connectivity of c , i.e. that cN^2 synapses out of $c_m N^2$ would be in the potentiated state. At the end of the imprinting, thus, $c/c_m = 1 - (1 - f^2)^P$, whence we obtain the number of associations as

$$P(c, c_m, f) = \frac{\log(1 - c/c_m)}{\log(1 - f^2)}. \quad (4.12)$$

The approximation for non overlapping patterns which started the Section is recovered for low $f \ll 1$ by expanding $(1 - f^2)^P \simeq 1 - Pf^2$.

Capacity The customary definition of capacity normalizes the number of pattern by dividing by the number of synapses available at each neuron, $c_m N$, since this number is independent of the network size:

$$\alpha \equiv \frac{P}{c_m N}. \quad (4.13)$$

This definition facilitates comparison with single-neuron based measures of capacity (e.g. for the perceptron).

4.1.4 Successful Replay as a Quality Criterion

In the previous Section we introduced the number of patterns P that can be stored in the network depending on the total loading level c reached during the imprinting process. The argument establishes a structural link between these two quantities based in the increment in c that is brought about by every imprinting step ρ , but it does

not tell whether those patterns can effectively be *retrieved* at the synaptic utilization level entailed by c . Retrieval of a target pattern is triggered by presentation of the corresponding cue, but, clearly, if too many patterns have been stored, the retrieval will not be error-free. Let us recall the original Willshaw criterion. Willshaw et al. assumed full connectivity, $c_m = 1$, set the threshold at M , and derived an optimal loading $c = 0.5$ (their p) to exhaust the capacity under the demand that the target pattern be reproduced with total fidelity (not a single spuriously firing neuron in the target pattern).

The approach to optimizing capacity adopted by Leibold and Kempter adds a dynamical element: it is required that all associations in a sequence are reproduced one after the other with a fidelity above a minimum quality level. A dynamical rule that allows to calculate how the quality evolves as the sequence is replayed is necessary. There exist only two kinds of error in reproducing a memory pattern: inactive neurons that should have been active, i.e. *false dismissals* g , and active neurons that should have remained silent, i.e. *false alarms* n . Hence it is sufficient to track just these two *error variables* as replay proceeds along the sequence (and it is superfluous to know the actual firing state of the neurons themselves). Note that from one iteration to the next, the identities of the hit and false alarm neurons will change, yet the error variables could remain at the same level.

Neuronal activation is calculated by comparing synaptic input to thresholds, and for this reason we will discuss replay not in terms of error variables but instead in terms of *activation variables*: number of hits $m = M - g$ and number of false alarms n at each time step.

4.1.5 Inhibitory Population

An inhibitory population can be incorporated into the original cellular model of Leibold and Kempter by adding a number K of binary units with threshold η that upon firing do not increase but *decrease* the postsynaptic membrane potential.

Synaptic weights Synapses between inhibitory and excitatory neurons can have strengths different from the excitatory recurrent synapses. We introduce w_{EI} and w_{IE} for the synaptic weight of excitatory to inhibitory synapses and inhibitory to excitatory synapses respectively. These synapses are still binary, but when potentiated instead of acquiring the normalization value $w_{ij} = w_{EE} \equiv 1$ they take up the value w_{EI} or w_{IE} respectively.

Unstructured connectivity For the inhibitory population there is no plasticity and thus no silent connectivity (see *Coding by principal neurons* in Section 4.1.1): all the wiring involving inhibitory neurons is indifferent to the storage of memories and subserves, thus, just a dynamical role.

Network state update rule To write the new update rule, we expand the N -long activity vector $\mathbf{x}(t)$ to include K additional elements $\mathbf{r}(t)$ representative of the state of the inhibitory population (see Eqs. 4.1 and 4.2),

$$\mathbf{u}(t) \equiv (\mathbf{x}(t), \mathbf{r}(t))$$

and grow blockwise the new effective synaptic matrix from the original excitatory one E^E as

$$E \equiv \begin{pmatrix} E^E & w^{EI} E^{EI} \\ w^{IE} E^{IE} & w^{II} E^I \end{pmatrix}. \quad (4.14)$$

The vector of thresholds now holds two subvectors for the two populations, $\lambda \equiv (\boldsymbol{\theta}, \boldsymbol{\eta})$ and the cellular update rule is, with the new definitions,

$$\mathbf{u}(t+1) = \Theta(E \mathbf{u}(t) - \lambda). \quad (4.15)$$

Assumptions about the connectivity of inhibitory neurons The connectivity matrix from inhibitory to excitatory neurons E^{IE} is assumed throughout to have all elements equal to one (recent imaging research in cortex makes this hypothesis plausible, see Fino and Yuste, 2011). The random effective connectivity matrix E^{EI} from excitatory to inhibitory neurons is of dimension $K \times N$. Its entries are set to one with probability $c_{EI} < 1$, which provides a source of variability in spite of the fixed inhibitory-to-excitatory connectivity. We assume no recurrent inhibitory connectivity for now, i.e. $E^I = 0$.

4.1.6 Cellular Simulations

In this Section we describe the algorithm for the simulation of the replay network presented above and an efficient computer implementation in C++. The adjective *cellular* refers to the fact that the state of cells, or units, is computed at each time step in order to assess the replay functionality of the network, and stands in opposition to *mean-field* simulations where only a minimum of replay-related variables is evolved in time for the same purpose. The code for these cellular simulations was written by Axel Kammerer; all simulation results in this work have been kindly provided by him.

The algorithm has two stages: creation of the synaptic matrix from randomized patterns (*imprinting*) and testing of the recall performance (*replay*). We describe them in turn.

4.1.6.1 Imprinting

The imprinting procedure needs a function to generate a collection of P random patterns and a function to produce the effective connectivity matrix E from the morphological connectivity and the potentiation induced by the Willshaw rule.

Generation of patterns Given c , c_m , and f , the number of patterns P to be generated and loaded is calculated using Equation 4.12. After imprinting of exactly P patterns as described below, the number of activated synapses is calculated in order to confirm that the desired effective connectivity c has been reached.

For each pattern $\boldsymbol{\xi} \in \{\boldsymbol{\xi}^1, \dots, \boldsymbol{\xi}^P\}$ we use the following randomization procedure. An integer j is randomly drawn from the set $\{0, \dots, N-1\}$. If ξ_j is zero (that neuron has not been chosen for the pattern so far), then we activate it, $\xi_j = 1$. The count of active neurons in pattern $\boldsymbol{\xi}$ increases by one. To first order, the expected number of collisions is $\sim f^2 M$. Hence the total amount of random numbers that we need to generate is of the order $M + f^2 M$, i.e. well below the N random numbers required by a naive implementation. The process stops when we have $\sum \xi_i = M$, i.e. the desired M active neurons per pattern.

pattern generation pseudocode

```

active_neurons =  $\emptyset$ 
while count(active_neurons) <  $M$ :
     $j$  = choose integer in range( $N$ ):
    if not  $j$  in active_neurons:
        add  $j$  to the set active_neurons

```

Patterns are stored as a binary vector of length N to facilitate comparison with the binary activity vectors generated by the network during the replay phase, rather than using a memory-efficient sparse representation.

Morphological and functional connectivity (imprinting) For each synapse $i \rightarrow j$, the morphological connectivity w_{ij} is decided by drawing a random number $y \in [0, 1)$ and setting $w_{ij} = \Theta(c_m - y)$. The synapse is potentiated from its initial state $s_{ij} = 0$ if it is found to link patterns activated in sequence upon scanning over the pairs of patterns (associations) generated as described above (i.e. there exists $t \in [1, \dots, P]$ such that $\xi_i^t = 1 = \xi_j^{t+1}$). The product of the two, $e_{ij} = w_{ij} s_{ij}$, is stored as the row-sparse effective connectivity matrix.

Unstructured inhibitory connectivities In networks with inhibition, the creation of inhibitory synapses benefits from the simplified connectivity scheme described in Section 4.1.5 (*Assumptions about the connectivity of inhibitory neurons*). Indeed, simulations assume that each inhibitory unit projects to all of the excitatory excitatory ones, i.e. $c_{IE} = 1$. The random, nonsaturated connectivity matrix E^{EI} from excitatory to inhibitory neurons is generated row-wise as follows. Given the probability of finding an activated synapse, c_{EI} , the distances d between activated synapses in a row (entries with value one) are drawn from the distribution $p(d) = c_{EI}(1 - c_{EI})^d$.

4.1.6.2 Replay

Replay is triggered by activating a set of neurons and applying the update rule Eq. (4.7) with the connectivity matrix E from the imprinting process above.

Initialization A copy of the first stored pattern ξ^1 is used to decide which neurons are active at step 1, $\mathbf{x}(t=1) = \xi^1$.

Step forward and multithreading The update step $\mathbf{x}(t+1) = \Theta(E^E \mathbf{x}(t) - \boldsymbol{\theta})$ involves N scalar products of the rows e_i of the effective connectivity matrix (stored as lists of indexes of active units) with the activity vector $\mathbf{x}(t)$ from the previous time step, of length N . Clearly, state update is independent across units, which suggests broadcasting the computation to different computational units (processor cores). Thus, the update step is parallelized in several threads (usually about eight), each handling $N/\text{numthreads}$ dot products. The state of the network after each such multiplication must be synchronized across the threads.

Assessment of the quality of replay The numbers of hits m and false alarms n are extracted from Eqs. 4.10 by calculating the dot products between activity vector $\mathbf{x}(t)$ and the t -th pattern ξ^t or its complement $(\mathbf{1} - \xi^t)$, respectively. The activity vector is discarded as soon as the next iteration has been computed, and only the values of m , n are committed to disk.

Parameter studies At each pattern size M several thresholds θ are tried, guided by the insights from the mean-field model (Section 4.2 below). For every M, θ pair, a few (<10) simulation runs are performed. In this way, the stochasticity induced by the realizations of the random morphological connectivity can be assessed. In each simulation a total of $Q = 100$ iterations is computed. The final state of the network is considered to be representative of the fixed point at $Q \rightarrow \infty$.

Inhibitory state In networks with inhibition it is necessary to update at each time step also the state of the inhibitory neurons. We follow here the notation introduced above in Section 4.1.5. The assumptions on the connectivity (non-recurrency) allow to write the inhibitory state $r(t)$ in terms only of the state of the excitatory neurons $x(t)$, following the rule

$$r(t+1) = \Theta(E^{\text{IE}} x(t) - \eta).$$

The update uses the same multithreaded code for sparse dot products as described above for the excitatory population. By virtue of the all-to-all inhibitory to excitatory connectivity, the update rule for the excitatory populations is simplified with respect to Eq. 4.15 —it only depends on a global inhibition bias $\sum_{k=1}^K r_k(t)$

$$x(t+1) = \Theta(E^{\text{E}}(x(t) - \sum_{k=1}^K r_k(t)) - \theta).$$

4.1.6.3 Computational Cost

The simulations described above need handling of very large connectivity matrices whose size grows quadratically in the total number of neurons N . A simulation run of the algorithm described above implemented in C++ takes 2-4 hours on a 2011 amd64 system with 24 6-core processors and 128 Gb of RAM. But precisely in the large numbers of neurons present in the mammalian hippocampus lies the hope for a simple model. Much as thermodynamics disregards positions and speeds of individual gas molecules and describes instead the state of a gas in terms of macroscopic variables such as pressure, volume and temperature, the theory we describe in Section 4.2 takes advantage of the large network and pattern sizes to directly estimate the evolution of the figures of merit for the replay process: hits and false alarms.

Memory requirements A total of P patterns are stored as length N arrays of `char`, thus spending one byte^{4.4} per neuron. The excitatory connectivity matrix E is stored as N sparse rows, each of which is represented, in average, by Nc integers of size 4 bytes (C++ `int` type). The L matrix is similarly encoded, with an average number of $NK c_{\text{EI}}$ integers indicating the position of active excitatory-to-inhibitory synapses in each of the K rows. The connectivity from inhibitory to excitatory neurons is total, $c_{\text{IE}} = 1$, and thus demands no additional storage. In addition to these major memory requirements, the algorithm keeps copies of the current state of inhibitory and excitatory neurons, and a few other scalar variables, which we neglect in comparison to the main data structures. While these memory demands, summarized in Table 4.1 are exceeded by the specifications of standard personal computers at the time of writing, they induce heavy loads on the memory-processor bus (Alted, 2010) that considerably slow down the computation to dozens of hours.

4.4. A possible space optimization would be to resort to the STL's `vector<bool>` or `bitset` data types.

data structure	size (bytes)	example size
loaded patterns	$PN \text{ sizeof(char)}$	2.64 Gb
E^E matrix	$cN^2 \text{ sizeof(int)}$	1.90 Gb
E^{EI} matrix	$c_{EI}KN \text{ sizeof(int)}$	0.08 Gb
Total	$N \left(\frac{\log(1-c/c_m)}{\log(1-M^2/N^2)} + 4cN + 4c_{EI}K \right)$	4.63 Gb

Table 4.1. Resources Required by Cellular Simulations. Main memory requirements of the example for the costliest simulation considered while preparing this work, with $N = 10^5$, $M = 500$, $K = 2 \cdot 10^4$, $c = 5\%$, $c_m = 10\%$, $c_{EI} = 0.01$.

4.1.7 Dynamical Regimes

What are the possible trajectories of the activation variables during replay? Cellular simulations of the excitatory-only network produce three dynamical regimes (Fig. 4.7), depending on the five parameters c , c_m , N , M and θ . At constant c , c_m and N the balance between the pattern size M and the common threshold θ decides which will be the operating regime of the network:

- i. *all silent* — the number M of active neurons in a pattern is not enough to keep On neurons active, or the threshold θ is too high.
- ii. *all active* — the pattern size is too high, or the threshold is too low, to prevent spilling of the initial activity of the On neurons to the Off neurons.
- iii. *sequence retrieval* — the combination of M and θ permits ongoing activation of the On neurons while most Off neurons remain silent. The network fulfills its purpose.

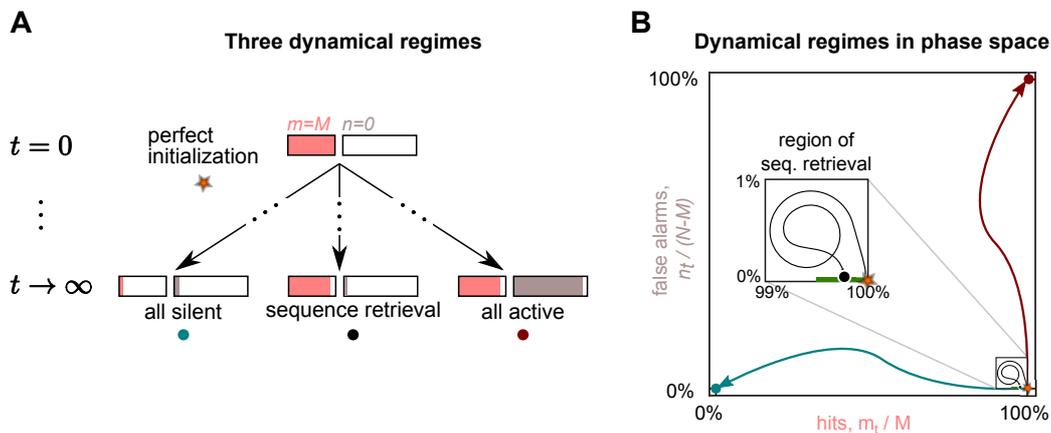


Figure 4.7. Dynamical Regimes in Phase Space. **A.** Perfect initialization can lead to several states: a silent state with no active neurons (left, *all silent* state), an “epileptic” state where all neurons are active regardless of their On/Off character (right, *all active* state), or the desired state of high activation of On neurons and low activation of Off neurons (center, *sequence retrieval* state). **B.** Phase space spanned by the activation variables m , n . The three trajectories shown correspond to the dynamical regimes described in A. The inset magnifies the retrieval region, where a spiral-like line symbolizes a trajectory for which every pattern is recalled with more than 99% of hits and less than 1% of false alarms. The green stripe at low left of the inset indicates the span of more than 95% of hits and less than 0.04% of false alarms used for phase-space plots below.

Figure 4.8 shows how these final states come about in a network where all parameters are fixed except the neuronal threshold θ . There exist a band of thresholds compatible with good replay separating the too-high thresholds that silence the network from the too-low thresholds that cause indiscriminate firing. A second feature of note is the transient replay for the transitional θ values between regimes: the network is capable of recalling adequately a few memory items before falling silent or becoming epileptic.

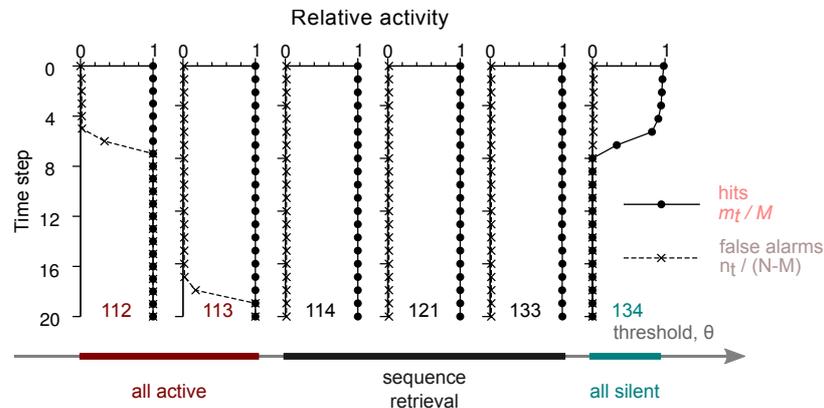


Figure 4.8. Three Asymptotic Dynamical Regimes in Cellular Simulations. Trajectories of the activity variables obtained for different θ and fixed M converge towards either the all-active regime (too low θ ; dark red), the all-silent regime (too high θ ; green) or the regime of sequence retrieval (black). Retrieval takes place transiently for too low or too high θ . Modified from Leibold and Kempner, 2006 (Fig. 1, parameters $M = 1,600$, $N = 10^5$, $c = 0.05$, $c_m = 0.1$).

Networks with inhibition will have to be characterized by the asymptotic values of the activation variables as well as the level of inhibition, provided it remains stable. For the moment, let us observe (Fig. 4.9) that the states of no replay are characterized as in the excitatory case by continued full activation of the excitatory subnetwork or its complete silence, and the transition from those regimes to retrieval is likewise smooth, i.e. through increasingly long epochs of retrieval. The fact that these states may be accompanied of different inhibition levels, sometimes oscillatory, will be discussed below in Section 4.4.2.

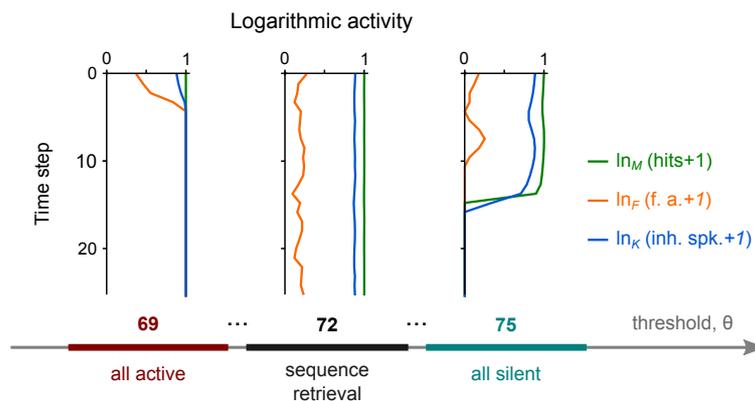


Figure 4.9. Three Dynamical Regimes in Cellular Simulations With Inhibition. Example trajectories from simulations of a retrieval network with inhibition. The threshold levels $\theta = 69$ (leftmost panel) and $\theta = 75$ (rightmost panel) are at the boundary of the region of sequence retrieval, represented in the central panel. Parameters were $M = 1,150$, $N = 10^5$, $K = 5,000$, $w_{IE} = 0.012$, $c_{IE} = 1$, $w_{EI} = 1$, $c = 0.05$, $c_m = 0.1$, $c_{EI} = 0.01$, $c_{IE} = 1$, and $\eta = 13$.

4.2 Construction of a Mean-Field Model

Despite the great simplifications introduced in the previous Section 4.1 with respect to the complexity inherent in biological neuronal networks, simulations remain computationally expensive. This is due to the large amounts of memory needed to store the neuronal and synaptic state and the costly randomization of patterns for the imprinting process. After lengthy simulations, we usually inquire whether replay was successful, folding the final state $\mathbf{x}(T)$, with N variables, into just two activation variables, m_T and n_T . In this Section we derive an analytical mean-field model that keeps track of those two variables only, and takes advantage precisely of the high neuron numbers characteristic of mammalian replay networks. The practical gain of a successful mean-field model is one of speed; additionally, one can systematically study the scaling laws for memory from its analytic formulas. This program was exhaustively conducted by Leibold and Kempster for the case of networks without inhibition, see Leibold and Kempster (2006). The additions necessary to build inhibition into the model are explained here in parallel to the description of their original model.

We outline the foundations of Leibold and Kempster's Markovian probabilistic theory in Section 4.2.1. Next, we assume the limit of infinite sequences and present the resulting mean-field model (Section 4.2.2), which builds upon the estimation of the statistics of synaptic input (Section 4.2.4). The model provides an update rule for the activation variables, and thus the two types of error associated with memory replay. In the following Sections, we explore the stability properties of the mean-field map. We do so both in phase space and in a phase diagram that shows the inherent tradeoff between neuronal firing threshold and the number of active neurons involved in representing memories (Section 4.3).

4.2.1 Markovian Dynamics and Binomial Statistics

Here we briefly present the building blocks of the Markovian model in Leibold and Kempster (2006): transition matrix, reduced connectivities and the averaging of the activation variables.

Transition matrix for the activation state In the LK2006 model for sequence storage and replay, neurons integrate input only from the previous time step. This Markovian character of the model allows to write a transition matrix

$$T(m_{t+1}, n_{t+1} | m_t, n_t) \quad (4.16)$$

that represents the probability for the state m_t, n_t to become m_{t+1}, n_{t+1} after one time step. Iterated application of such map from the m_0, n_0 of the initial cue, taking into account all possible trajectories through intermediate steps, allows to produce a distribution of expected m and n for later iterations.

Reduced connectivities The transition matrix approach needs the definition of effective *potentiated* connectivities between neuron groups On and Off. This makes the success of the approach dependent on sufficient neuron numbers ($N \gg M \gg 1$) so that the average connectivity among the groups can be taken as a probability. The imprinting process ensures that the average effective connectivity from On_t to On_{t+1} neurons is $c_{\text{On} \rightarrow \text{On}} = c_m$, i.e. saturates the available morphological connectivity. This is the fraction of potentiated synapses expected in the association depicted in Figure 4.5 (purple-colored synapses). The reduced connectivities $c_{\text{On} \rightarrow \text{Off}} = \text{prob}(\text{On}_{t-1} \rightarrow \text{Off}_t)$

and $c_{\text{Off} \rightarrow \text{On}} = \text{prob}(\text{Off}_{t-1} \rightarrow \text{On}_t)$ as well as $c_{\text{Off} \rightarrow \text{Off}} = \text{prob}(\text{Off}_{t-1} \rightarrow \text{Off}_t)$ are obtained (Leibold and Kempner, 2006) from normalization conditions^{4.5}:

$$\begin{pmatrix} c_{\text{On} \rightarrow \text{On}} & c_{\text{On} \rightarrow \text{Off}} \\ c_{\text{Off} \rightarrow \text{On}} & c_{\text{Off} \rightarrow \text{Off}} \end{pmatrix} = \begin{pmatrix} c_m & c(1 - rM/F) \\ c(1 - rM/F) & c(1 + rM^2/F^2) \end{pmatrix}.$$

We shall assume the sparse coding regime $f = M/N \ll 1$, where reduced connectivities simplify (Fig. 4.10):

$$\begin{pmatrix} c_{\text{On} \rightarrow \text{On}} & c_{\text{On} \rightarrow \text{Off}} \\ c_{\text{Off} \rightarrow \text{On}} & c_{\text{Off} \rightarrow \text{Off}} \end{pmatrix} \xrightarrow{M \ll F} \begin{pmatrix} c_m & c \\ c & c \end{pmatrix}. \quad (4.17)$$

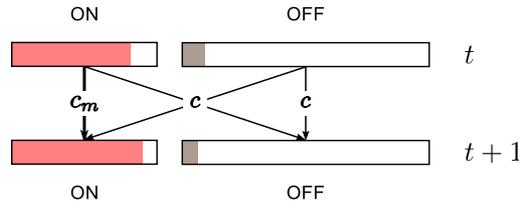


Figure 4.10. Reduced Connectivity of the Excitatory Network. The connectivity of the excitatory network can be summarized for $1 \ll M \ll N$ in terms of the probabilities of connections among neuron groups On and Off (see Eq. (4.17)).

So far, we have considered two source and target populations: the On neurons (group 1, activation variable m , total M) and the Off neurons (group 2, activation variable n , total F), both of excitatory character. We now add to consideration a pool of inhibitory neurons (group 3, activation variable k , total K ; see Section 4.1.5). As a source population, they project to and receive input from the excitatory groups with random probabilities of connection c_{IE} and c_{EI} respectively. The reduced connectivity matrix is as follows (sources in rows, targets in columns; compare Fig. 4.10 and Eq. 4.17):

$$\begin{pmatrix} c_{\text{On} \rightarrow \text{On}} & c_{\text{On} \rightarrow \text{Off}} & c_{\text{On} \rightarrow \text{Inh}} \\ c_{\text{Off} \rightarrow \text{On}} & c_{\text{Off} \rightarrow \text{Off}} & c_{\text{Off} \rightarrow \text{Inh}} \\ c_{\text{Inh} \rightarrow \text{On}} & c_{\text{Inh} \rightarrow \text{Off}} & c_{\text{Inh} \rightarrow \text{Inh}} \end{pmatrix} = \begin{pmatrix} c_m & c & c_{\text{EI}} \\ c & c & c_{\text{EI}} \\ c_{\text{IE}} & c_{\text{IE}} & c_{\text{II}} \end{pmatrix}. \quad (4.18)$$

Averages of activation variables In order to compare the distributions of activation variables with realizations of the dynamics in cellular simulations (Section 4.1.6) Leibold and Kempner take averages of final states $\langle m_t \rangle$, $\langle n_t \rangle$ with respect to the distributions given by the transition matrix T . They then verify that the Markov model captures the main features of the cellular simulations, in particular the existence of the three phases all-silent, all-active and replay and the approximate thresholds for the transitions between them at fixed M (compare Fig. 3 with Fig. 1 there). We direct the reader to their account for further details on the reduced connectivities, the transition matrix, and the resulting distributions for m_t , n_t (Leibold and Kempner, 2006, Section 4.2.2). We shall henceforth omit the averaging brackets and assume that m_t , n_t always refer to the respective averages over realizations of the random morphological connectivity matrix.

4.5. Note that c_{11} there is $c_{\text{On} \rightarrow \text{On}}$ here, c_{00} there is $c_{\text{Off} \rightarrow \text{Off}}$ here, and so on.

Infinite sequences as a proxy for long sequences Upon simulating the Markovian dynamics for increasing lengths of stored sequences, $Q = 1, \dots, 8$, Leibold and Kempster realized that the capacity $\alpha = P/cN$, which is measured in terms of total number of associations in replayable stored sequences, progressively decreases towards a saturation value as Q increases (see their Section 7.1, Leibold and Kempster, 2006). It seemed justified to take advantage of this weak dependence of capacity on sequence length for long sequences ($Q \gtrsim 8$) and formulate a simplified dynamics for the means of the activation variables in the limit $Q \rightarrow \infty$.

The next Section presents a map that updates the expected activation values m_t, n_t from one time step to the next in dependence of the configuration of the network.

4.2.2 Mean-Field Model for Activation Averages in the Case $Q \rightarrow \infty$

For infinite sequences we can use the fixed-point distributions of hits and false alarms given by the Markovian map (Eq. 4.16), or rather the fixed-point distributions of their mean values, given that the distributions are unimodal and the network size N is sufficiently large. We seek thus a map

$$\begin{aligned} m_{t+1} &= T_{\text{On}}(m_t, n_t, k_t; \mathbf{p}) \\ n_{t+1} &= T_{\text{Off}}(m_t, n_t, k_t; \mathbf{p}) \\ k_{t+1} &= T_{\text{Inh}}(m_t, n_t, k_t; \mathbf{p}), \end{aligned} \quad (4.19)$$

where \mathbf{p} stands for the collection of parameters that characterize the map, $\mathbf{p} = \{M, N, K, c, c_m, c_{\text{El}}, c_{\text{IE}}, c_{\text{II}}, \theta, \eta\}$.

In vector notation, $\mathbf{y}_{t+1} = \mathbf{T}(\mathbf{y}_t; \mathbf{p})$ (compare Eqs. 4.1 and 4.8 specifying the *cellular* update rule). Again, m, n and k stand for mean values at the iteration indicated by the respective subindices. In order to characterize their fixed-point distributions we investigate next the distributions of inputs to active neurons.

4.2.3 Synaptic Input as a Random Process and the Optimal Detector

In the mean-field framework, the synaptic input to a neuron is assumed to be a random process statistically independent of the input to any other neuron. Each *excitatory* target neuron confronts a signal detection problem (Helstrom, 1994; Poor, 1994), in particular, a binary hypothesis test: given the incoming level of synaptic input h , should the detector neuron infer that it is an On neuron (fire, activation hypothesis A_{On}) or an Off neuron (remain still, activation hypothesis A_{Off})^{4.6?}

$$\begin{aligned} A_{\text{On}}: \quad h &= \overbrace{h(\text{from all}; c)}^{\text{noise}} + \overbrace{h(\text{from On}; c_m) - h(\text{from On}; c)}^{\text{signal}} \\ \textit{versus} \\ A_{\text{Off}}: \quad h &= \overbrace{h(\text{from all}; c)}^{\text{noise}}. \end{aligned} \quad (4.20)$$

^{4.6.} The notation $A_{\text{On}}, A_{\text{Off}}$ is used for the hypotheses of firing and silence instead of the customary H_1 and H_0 for the alternative and null hypotheses respectively.

The signal is the excess potential incoming into the cell from the associated On presynaptic pattern due to their enhanced mean afference c_m (see Fig. 4.10). Reasonable priors for these alternative hypotheses in the course of replay are $p(A_{\text{On}}) = M/N = f$ and $p(A_{\text{Off}}) = 1 - f$.

The firing decision can be right, in which case costs $\varepsilon_{\text{On};\text{On}}$ and $\varepsilon_{\text{Off};\text{Off}}$ will be respectively incurred, or it may have missed the true On or Off nature of the neuron at that time step as established by the memorization procedure, thereby incurring costs $\varepsilon_{\text{On};\text{Off}}$ if it mistakenly fired or $\varepsilon_{\text{Off};\text{On}}$ if it remained still while it was supposed to fire. The decision, which can be thought of as the task of detecting the presynaptic signal to fire, can be implemented by means of a randomized strategy: selecting one or the other firing state with probability dependent on the synaptic input level. The neuronal firing threshold is the parameter for our detector implementing such a decision strategy.

The probabilities of a positive firing decision for neurons from each of the post-synaptic populations, $\text{Pr}_{\text{On}}(1|h)$ (hit) and $\text{Pr}_{\text{Off}}(1|h)$ (false alarm) can be turned into an update rule for the corresponding average activation variable m, n of the On or Off group (cf. Sec. 4.2.1) just by multiplying by the number of neurons—thus extracting a statistical frequency from the probability of each independent detection experiment:

$$\begin{aligned} m_{t+1} &= M \text{Pr}_{\text{On}}(x(t+1) = 1|h(t); \theta) \\ n_{t+1} &= F \text{Pr}_{\text{Off}}(x(t+1) = 1|h(t); \theta). \end{aligned} \quad (4.21)$$

The principle of the mean-field rule is thus to identify the probability that a post-synaptic neuron fires, or, equivalently, the active fraction of the target population it belongs in, with the fraction of the probability mass function (p.m.f) of its synaptic input H that is over threshold, i.e. the cumulative distribution function $p(x=1) = F_H(h=\theta)$. Hence, a complete characterization of the contributions to h from the source populations is in order.

4.2.4 Synaptic Input Distributions From the Afferent Populations

We recall at this point the cellular rules Eq. 4.7 and 4.15 that prescribe how the firing state of a neuron is to be updated based on the connectivity and the network firing state of the afferents in the previous time step:

$$u(t+1) = \begin{cases} 1 & \text{if } \sum_j w_j s_j u_j(t) > \lambda \\ 0 & \text{otherwise.} \end{cases} \quad j \in \text{On}(t) \cup \text{Off}(t) \cup \text{Inh}(t) \quad (4.22)$$

The total input to a cell depends on the wiring from, and the state of, its presynaptic partners. For each potential presynaptic partner j we can ask if it is morphologically connected i.e. if $w_j = 1$, if it should have been potentiated (i.e. if it was presynaptic in an association with the considered cell, $s_j = 1$) and if it is active at this point in time, i.e. if $u_j(t) = 1$. These three random binary processes are essentially independent, and hence their probabilities of success can be multiplied to obtain the probability of the presynaptic cell effecting one unit of synaptic input, again a Bernoulli process $\text{Pr}(h_j = 1) = \text{Pr}(w_j s_j u_j(t) = 1)$. The total synaptic input is a random variable (r.v.)

consisting of the sum of Bernoulli processes over all potential presynaptic partners. When each of N elements follows a Bernoulli process with the same probability of success p , the random variable constructed as the sum of the Bernoullis is binomially distributed with probability $p = \Pr(h_j = 1)$. We denote the probability of obtaining $h \in \mathbb{N}$ heads (ones) out of $N \in \mathbb{N}$ independent identically distributed binary trials, each with Bernoulli probability p by

$$B_{N,p}(h) \equiv \binom{N}{h} p^h (1-p)^{N-h}. \quad (4.23)$$

Since the Bernoulli distribution is the binomial for just one trial, i.e. $\text{Ber}_p = B_{1,p}$, the statement above is a particular case of the result about the sum of independent binomial random variables of the same Bernoulli p —their convolution (\star) is again a binomial, i.e. if $Y \sim \sum_i^n X_i$ with $X_i \sim B_{N_i,p}(x_i)$, then this *reproductive property* of the binomial distribution (Johnson et al., 2005) is

$$Y \sim X_1 \star X_2 \star \dots \star X_n \sim B_{\sum_i N_i,p}(y)$$

We use \sim to signify “is distributed like”.

Exc	On	m_t/M
	Off	n_t/F
Inh		k_t/K
at t (pre.)		$\Pr(u=1)$

Table 4.2. Firing Probabilities of Presynaptic Neurons According to Group.

We can take advantage of the reproductive property to obtain the combined synaptic input from a group of cells to a target cell, as long as the probability of synaptic contribution $\Pr(h_j)$ observes the same distribution. The firing probability $\Pr(u_j = 1)$ defines three candidate source groups: On, Off and Inh (Table 4.2). The morphological connection probability (Table 4.3) to a neuron in any of the three target groups On, Off and Inh is also constant within any of those three source groups.

		Exc		Inh	at $t+1$ (post.)
		On	Off		
Exc	On	c_m	c_{EI}		
	Off				
Inh		c_{IE}	c_{II}		
at t (pre.)					$\Pr(w=1)$

Table 4.3. Morphological Connection Probabilities. Probability of existence of a synapse between two neurons according to the group of the presynaptic (rows) and that of the postsynaptic cell (columns).

The potentiation probability is however only homogeneous for the Inh source group (last row in Table 4.4) to all target groups, and for the On(t) group when considering input to On($t + 1$) neurons. The π subindex breaks up *target* groups in subgroups according to how many associations their neurons have been postsynaptic in. We will come back to it after discussing the simpler cases.

		Exc		Inh	at $t + 1$ (post.)
		On $_{\pi}$	Off $_{\pi}$		
Exc	On	1	q_{π}	1	
	Off	q_{π}			
Inh		1		1	
at t (pre.)					Pr($s = 1$)

Table 4.4. Potentiation Probabilities. Probability of potentiation of a synapse between two neurons according to the group of the presynaptic (rows) and that of the postsynaptic cell (columns).

Distribution of synaptic inputs from a homogeneous population The inhibitory neurons $j \in \text{Inh}$ all fire with probability $\Pr(u_j(t) = 1) = k_t/K$. They have a constant probability of morphological afference, $\Pr_{\text{Exc}}(w_j = 1) = c_{\text{IE}}$ if we are considering an excitatory postsynaptic neuron and $\Pr_{\text{Inh}}(w_j = 1) = c_{\text{II}}$ if we are rather counting inputs to a postsynaptic inhibitory neuron. Since they are not subject to plasticity, their synapses both to excitatory and inhibitory neurons are always potentiated, $\Pr_{\text{All}}(s_j = 1) = 1$. Hence, the summed input from the inhibitory population follows a binomial law in every case, e.g. for excitatory targets

$$H_{\text{Exc}}^{\leftarrow \text{Inh}} \sim B_{K, c_{\text{IE}}k_t/K}(h) = \binom{K}{h} (c_{\text{IE}}k_t/K)^h (1 - c_{\text{IE}}k_t/K)^{K-h}. \quad (4.24)$$

The same is true about excitatory inputs of any kind to the inhibitory cells. The summed input of On(t) cells to On($t + 1$) cells follows a binomial as well, because the synapses linking them are potentiated with certainty by the imprinting procedure, $\Pr_{\text{On}}(s_j = 1) = 1$ for all $j \in \text{On}(t)$.

All other combinations need a more involved analysis. For Off(t) feeding onto On($t + 1$) neurons as well as for all excitatory neurons feeding onto Off($t + 1$) it is necessary to distinguish in how many patterns π the postsynaptic neuron is involved. Inputs will only be binomial with a homogeneous probability of potentiation q_{π} when considering target neurons of the more restricted group On $_{\pi}$ or Off $_{\pi}$. To assess quality of replay, we aim however at obtaining equations for hits and false alarms, i.e. probability of activation of postsynaptic On or Off neurons respectively, without distinction of how many associations they are postsynaptic in. For that, an averaging procedure over the subgroups of unequal participation π will be needed that accounts for their relative frequency.

To complete the characterization of input distributions, we calculate the probability of potentiation of postsynaptic On $_{\pi}$ and Off $_{\pi}$ neurons, q_{π} , and the weight with which this probability will be incorporated to calculate the distribution of inputs to a generic On or Off neuron, $\Pr(\pi)$.

Probability of potentiation q_π The probability of potentiation $\Pr_{\text{Exc}\pi}(s = 1) = q_\pi$ depends on how many associations $0 \leq \pi \leq P$ the specific postsynaptic neuron participates in. If the postsynaptic neuron was never associated as a target, none of its synapses will have been activated by the Willshaw rule and thus $q_0 = 0$. Because the Willshaw rule does not further strengthen already activated synapses, it is again advantageous to calculate the probability that a synapse remains unpotentiated (cf. *Nonorthogonal patterns* in Section 4.1.3). That probability, for a single association, is simply the probability of the neuron not being active, $1 - f$. Because of independence, it is $(1 - f)^\pi$ for π associations and thus

$$q_\pi = 1 - (1 - f)^\pi. \quad (4.25)$$

Probability of being π times postsynaptic partner The probability that any excitatory neuron is involved in π associations as a target is binomial:

$$\Pr(\pi) = B_{P,f}(\pi) = \binom{P}{\pi} f^\pi (1 - f)^{P - \pi}. \quad (4.26)$$

In a network with $M = 1,000$, $N = 10^5$, $c = 5\%$, $c_m = 10\%$ at its maximal loading of $P = 6,731$ and with firing sparseness $f \sim 0.01$ an excitatory neuron is in average 67 times postsynaptic.

Summary of synaptic input distributions Table 4.5 puts together the factors from the previous Tables 4.2-4.4 and summarizes the probability of a synaptic contribution from a source to a target neuron. Each entry there is to be understood as the binomial probability p in $B_{A,p}$ that characterizes the respective input distribution $H_{\text{target}}^{\text{source}}$; *target* here refers to the target class of *one* representative neuron whereas *source* is the *group* that gives rise to that input. The group sizes are $A = M, F, K$ for On, Off and Inh neurons respectively.

		Exc		Inh	at $t + 1$ (post.)
		On $_\pi$	Off $_\pi$		
Exc	On	$c_m m_t / M$	$c_m q_\pi m_t / M$	$c_{\text{EI}} m_t / M$	Pr($h = 1$)
	Off	$c_m q_\pi n_t / F$		$c_{\text{EI}} n_t / F$	
Inh	$c_{\text{IE}} k_t / K$		$c_{\text{II}} k_t / K$		
at t (pre.)					

Table 4.5. Probabilities of Spike Transmission. Binomial probability of synaptic contribution of from a presynaptic neuron whose group is indicated in the rows to a postsynaptic neuron whose group is in columns.

Input sum from several source populations Each of the postsynaptic groups, On, Off and Inh receives inputs from each of the presynaptic groups, i.e. these inputs are independent and combine additively, with coefficients w_1, w_2, w_3 to allow for synaptic weights:

$$H_{\text{target}} = w_1 H_{\text{target}}^{\text{On}} + w_2 H_{\text{target}}^{\text{Off}} + w_3 H_{\text{target}}^{\text{Inh}} \quad (\text{target} \in \{\text{On}_\pi, \text{Off}_\pi, \text{Inh}\}; \pi \in 0, \dots, P).$$

The corresponding random variable, *weighted total synaptic input* is thus a convolution of the contributions from the different source populations:

$$H_{\text{target}} \sim w_1 H_{\text{target}}^{\text{On}} \star w_2 H_{\text{target}}^{\text{Off}} \star w_3 H_{\text{target}}^{\text{Inh}}. \quad (4.27)$$

Writing out the explicit form H_{target} requires the convolution of three different binomials, a rather cumbersome calculation offering little insight. We may note that each of the simple binomials may be substituted with its corresponding Gaussian limit distribution with good approximation, which allows to profit from the fact that the convolution of Gaussians is again a Gaussian.

Gaussian approximations A generally admitted rule of thumb that enables to approximate the binomial $B_{A,p}$ by the normal $\mathcal{N}(Ap, Ap(1-p))$ requires that $Ap(1-p) > 9$ (Johnson et al., 2005). We thus ask: what are the values of the activation variables that fulfill the rule for a given source-target pair, i.e. for a fixed number of source neurons A , firing probability a_i/A and effective connectivity $c_a q$? The quadratic equation that results from the criterion above leads to the following brackets for the activity levels:

$$a_-/A = \frac{1 - \sqrt{1 - 36/A}}{2c_a q} < a_i/A < \frac{1 + \sqrt{1 - 36/A}}{2c_a q} = a_+/A.$$

We summarise in Table 4.6 the figures for a number of source-target combinations and exemplary values of N, M, θ, c, c_m .

Source population	Target neuron class	activity at t	brackets (a_-/A to a_+/A)	values during replay
Inh \rightarrow	On or Off	$\frac{k}{K} \in$	0.002 ... 0.998	$0 \leq \frac{k}{K} \leq 1$
Off \rightarrow	Inh	$\frac{n}{F} \in$	0.009 ... 1	$0 \lesssim \frac{n}{F}$
	$\frac{(\text{On/Off})_{\pi=1}}{(\text{On/Off})_{\pi=10}}$		0.091 ... 1	
	$\frac{(\text{On/Off})_{\pi=10}}{(\text{On/Off})_{\pi=6,920}}$		0.010 ... 1	
On \rightarrow	$\frac{\text{On}}{\text{Inh}}$	$\frac{m}{M} \in$	0.091 ... 1	$\frac{m}{M} \lesssim 1$
	$\frac{\text{Off}_{\pi=1}}{\text{Off}_{\pi=11}}$		0.908 ... 1	
	$\frac{\text{Off}_{\pi=11}}{\text{Off}_{\pi=6,920}}$		— ... 1	
	$\frac{\text{Off}_{\pi=6,920}}{\text{Off}_{\pi=6,920}}$		0.868 ... 1	
			0.091 ... 1	

Table 4.6. Accuracy of Gaussian Approximations to the Binomials. Brackets for the activity levels that warrant an approximation of the binomial distribution of single-source synaptic input by a Gaussian when taking exemplary values for the population size A of $N = 10^5$, $M = 10^3$, $K = 5 \times 10^3$, for the morphological connectivity of $c_m = 0.1$, $c_{\text{IE}} = 1$, $c_{\text{EI}} = 0.01$ and for the probability of potentiation $q = 1$ or $q_\pi = 1 - (1 - f)^\pi$ as appropriate (see Tables 4.2-4.4). With these numbers, $f = 0.01$ and $P = 6,931$. Notice same brackets for Exc \rightarrow On $_{\pi=6,920}$ and Exc \rightarrow Inh; this is due to $q_{6,920} \simeq 1$. Given that $\pi \in (0, 6,931)$, the chosen values $\pi = 10$ and $\pi = 6,920$ are representative of extreme levels of participation.

The Gaussian approximation works well in the activity regimes characteristic of replay, with exception of input from On to Off neurons involved in a low number of associations $\pi \lesssim 11$ (m_-/M evaluates to 9.082). The fraction of such neurons is, by Eq. 4.26, $\Pr(\pi < 11) = 5 \times 10^{-19}$, i.e. rarely will one find a single one in populations of $F \sim 10^5$.

In our regime of operation and for a wide range of parameters, in conclusion, the contributions to the synaptic input seen by a postsynaptic cell can be each well approximated by a Gaussian (and where they can not, there are next to none such postsynaptic neurons).

The sum of Gaussian random variates is distributed like a Gaussian if they are independent. The mean of the sum is the sum of the means, and the variance of the sum is the sum of the variances if the source distributions are uncorrelated. Hence we need only know the individual mean and variances of the synaptic inputs from source populations in order to fully determine the distribution of inputs seen by a neuron in each of the postsynaptic classes.

Note that by employing the Gaussian approximation we adopt a probability distribution function instead of a discrete probability mass function, i.e. synaptic input is not anymore defined over \mathbb{N} but over \mathbb{R} (including unphysical negative values, that are extremely unlikely in any scenario of interest).

Mean and variance of the synaptic input distributions A binomial distributed like $B_{A,p}$ has mean Ap and variance $Ap(1-p)$. These are also the mean and variance of the approximating Gaussian. With this, we can immediately build the input distribution to inhibitory cells. The input from On cells has binomial probability $c_{\text{EI}}m/M$. Hence its mean and variance are, respectively $c_{\text{EI}}m$ and $c_{\text{EI}}m(1 - c_{\text{EI}}m/M)$. The same expressions with n instead of m and F instead of M hold for input from the Off population. Recurrent input from inhibitory cells follows a binomial with probability $c_{\text{II}}k/K$ and has average and variance $c_{\text{II}}k$, $c_{\text{II}}k(1 - c_{\text{II}}k/K)$ respectively. We have to take into account the respective synaptic efficacies as well, w_{IE} and w_{II} . The moments that define $H_{\text{Inh}} \sim \mathcal{N}(\mu_{\text{Inh}}, \sigma_{\text{Inh}}^2)$ are, then,

$$\mu_{\text{Inh}} = w_{\text{EI}}c_{\text{EI}}(n+m) + w_{\text{II}}c_{\text{II}}k \quad (4.28)$$

$$\sigma_{\text{Inh}}^2 = w_{\text{IE}}^2c_{\text{EI}}\{m(1 - c_{\text{EI}}m/M) + n(1 - c_{\text{EI}}n/F)\} + w_{\text{II}}^2c_{\text{II}}k(1 - c_{\text{II}}k/K). \quad (4.29)$$

Synaptic input averaged for On and Off cell classes Similar expressions can be obtained for each of the On_π and Off_π populations. For these excitatory target populations, however, we are instead interested in constructing H_{On} and H_{Off} , i.e. a distribution of synaptic inputs for the ‘‘average’’ On or Off target neuron, respectively (the counts of hits and false alarms that we need to assess the quality of sequence retrieval throughout the dynamics do not care about the precise number of associations that a neuron has been target in).

The inhibitory contribution to the means and variances of excitatory populations is a simple binomial $H_{\text{Exc}}^{\text{Inh}} \sim B_{K, c_{\text{IE}}k/K}$ that does not require averaging because the potentiation probability is 1, i.e. not dependent on postsynaptic participation in associations:

$$\mathbb{E}[H_{\text{Exc}}^{\text{Inh}}] = -w_{\text{IE}}c_{\text{IE}}k \quad (4.30)$$

$$\text{Var}[H_{\text{Exc}}^{\text{Inh}}] = w_{\text{IE}}^2c_{\text{IE}}k(1 - c_{\text{IE}}k/K). \quad (4.31)$$

To construct a synaptic input distribution to e.g. the Off population from the On population, however, we need weight the input to a neuron in the Off_π class by the probability to find such a neuron in the target population ($\text{Pr}(\pi)$, Eq. 4.26). We thus obtain the input to an “average” neuron (by participation in associations as postsynaptic). To make the discussion independent of the source population and target neuron class, we take the total number of source neurons to be A , their firing probability to be $\text{Pr}(u=1)=a/A$, the probability that they are morphologically connected to the target neuron as $\text{Pr}(w=1)=c_a$ and the probability that this connection is potentiated as $\text{Pr}(s=1)=q_\pi$, given by Eq. 4.25. For brevity we shall denote H_π the random variable of the synaptic input to neurons that are target in π associations, with pmf $\text{Pr}_\pi(h)$, and just H its average over associations, with pmf $\text{Pr}(h)$.

Input is then distributed $H_\pi \sim B_{A, c_a q_\pi a/A}$ with mean and variance

$$\begin{aligned} \text{E}[H_\pi] &= c_a q_\pi a \\ \text{Var}[H_\pi] &= c_a q_\pi a (1 - c_a q_\pi a/A). \end{aligned}$$

This distribution is just of input to neurons involved in π associations. For the average neuron we average over associations

$$\text{Pr}(h) = \sum_{\pi=0}^P \text{Pr}(\pi) \text{Pr}_\pi(h). \quad (4.32)$$

Note that the distribution of inputs for an average neuron is now given by a weighted sum over many unimodal pdfs. Figure 4.11 shows the distributions of inputs to neurons participating in $\pi=1, \dots, P$ associations as they are averaged to form a single input distribution to the “mean neuron” (thick black curve). The other colored lines correspond to Gaussian approximations that will be explained below.

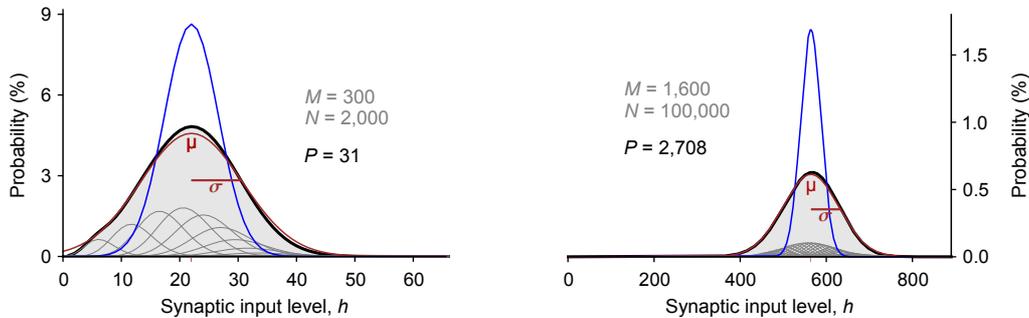


Figure 4.11. Synaptic Input and Involvement in Associations. The synaptic input to an excitatory cell as a function of the number of associations it participates in (grey curves for $\pi=0, \dots, P$) can be averaged to a Gaussian-looking distribution (in black). The accuracy of the Gaussian approximation (in dark red) greatly improves as the network is made sparser to accommodate more patterns P (compare left and right panel). Neglecting the correlations induced by the Willshaw imprinting rule leads to a substantial underestimation of the variance (blue curve). Parameters were $c=0.05$, $c_m=0.1$; M, N as indicated on each panel in grey.

The mean of Eq. 4.32 is simply the weighted average of the component $\text{E}[H_\pi]$, i.e.

$$\text{E}[H] = \sum_{\pi=0}^P \text{Pr}(\pi) \text{E}(H_\pi) = \sum_{\pi=0}^P \text{Pr}(\pi) A c_a q_\pi a/A = c_a a \text{E}_\pi[q]. \quad (4.33)$$

Synaptic input is just being weighted by the mean potentiation strength across participation as postsynaptic in patterns, $E_\pi[q] = c/c_m$ (see Appendix B for the derivation). The variance $\text{Var}[H] = E[H^2] - E[H]^2 = E[H^2] - c_a^2 a^2 E_\pi^2[q]$ requires the expectation of the squared input,

$$\begin{aligned}
E[H^2] &= \sum_{\pi=0}^P \Pr(\pi) \underbrace{\sum_{h=0}^A h^2 \Pr_\pi(h)}_{E[H_\pi^2]} \\
&= \sum_{\pi=0}^P \Pr(\pi) \left\{ \underbrace{c_a q_\pi a (1 - c_a q_\pi a / A)}_{\text{Var}[H_\pi]} + \underbrace{c_a^2 q_\pi^2 a^2}_{E[H_\pi]^2} \right\} \\
&= c_a a \sum_{\pi=0}^P \Pr(\pi) \{ q_\pi (1 - c_a q_\pi a / A) + c_a q_\pi^2 a \} \\
&= c_a a \sum_{\pi=0}^P \Pr(\pi) \{ q_\pi + (1 - 1/A) c_a a q_\pi^2 \} \\
&= c_a a \{ E_\pi[q] + (1 - 1/A) c_a a E_\pi[q^2] \}.
\end{aligned}$$

Subtracting the square of the mean (Eq. 4.33) we have

$$\begin{aligned}
\text{Var}[H] &= c_a a \{ E_\pi[q] + (1 - 1/A) c_a a E_\pi[q^2] \} - c_a^2 a^2 E_\pi^2[q] \\
&= c_a a + c_a^2 a^2 E_\pi[q^2] (1 - 1/A) - c_a^2 a^2 E_\pi^2[q] \\
&= c_a a + c^2 a^2 (\text{CV}_\pi^2[q] + 1) (1 - 1/A) - c^2 a^2 \\
&= c_a a + c^2 a^2 \text{CV}_\pi^2[q] (1 - 1/A) - c^2 a^2 / A \\
&= c_a a (1 - c_a a / A) + c^2 a^2 \text{CV}_\pi^2[q] (1 - 1/A).
\end{aligned}$$

Here we introduced the squared coefficient of variation $\text{CV}_\pi^2[q]$, whose explicit expression is given in the Appendix B.

Summary of averages and variances With the general expressions just computed and the parameters in Tables 4.2-4.4 in place of the generic A, a, c_a we can now provide expressions for the mean total inputs μ to neurons in the On, Off and Inh target classes and their variances σ^2 :

$$\begin{aligned}
\mu_{\text{On}} &= c_m m + c n - w_{\text{IE}} c_{\text{IE}} k \\
\sigma_{\text{On}}^2 &= c_m m (1 - c_m m / M) + c n (1 - c n / F) + c^2 n^2 \text{CV}_\pi^2[q] (1 - 1/F) + \\
&\quad + w_{\text{IE}}^2 c_{\text{IE}} k (1 - c_{\text{IE}} k / K)
\end{aligned}$$

$$\begin{aligned}
\mu_{\text{Off}} &= c m + c n - w_{\text{IE}} c_{\text{IE}} k \\
\sigma_{\text{Off}}^2 &= c m (1 - c m / M) + c n (1 - c n / F) + \\
&\quad + c^2 m^2 \text{CV}_\pi^2[q] (1 - 1/M) + c^2 n^2 \text{CV}_\pi^2[q] (1 - 1/F) + \\
&\quad + w_{\text{IE}}^2 c_{\text{IE}} k (1 - c_{\text{IE}} k / K)
\end{aligned}$$

$$\begin{aligned}
\mu_{\text{Inh}} &= w_{\text{EI}} c_{\text{EI}} (m + n) - w_{\text{II}} c_{\text{II}} k \\
\sigma_{\text{Inh}}^2 &= w_{\text{IE}}^2 c_{\text{EI}} \{ m (1 - c_{\text{EI}} m / M) + n (1 - c_{\text{EI}} n / F) \} + w_{\text{II}}^2 c_{\text{II}} k (1 - c_{\text{II}} k / K).
\end{aligned}$$

Simplified expressions for averages and variances The coefficient of variation tends to zero linearly with the coding ratio f (Appendix B) and therefore for extremely sparsely activated networks it is possible to ignore all variance terms where it is a factor, which amounts to decreeing that excitatory populations are not wired in a structured fashion but at random, just like the inhibitory population. The corresponding variance formulas do not have pattern-participation weighting, i.e. are like that for σ_{Inh}^2 . It is the implicit assumption taken in the theory and simulations of Leibold and Kempter, 2006 (p. 927; to compare with formulæ there, remove also all terms to and from inhibitory populations). This approximation can substantially underestimate the variances already for coding ratios $f \simeq 0.01$. In a network with $M = 1,000$, $N = 10^5$, $c = 5\%$ and $c_m = 10\%$ operating at a regime close to replay with $m/M = 0.9$ and $n/F = 0.1$, the standard deviation σ_{Off} goes to about 48 up from 23 when taking into account imprinting correlations as introduced by a non vanishing $\text{CV}_\pi^2[q]$; this effect is strong due to the high values of n compatible with good quality under the error weighting given by Eq. 4.11 albeit not so prominent for the lower n generally observed during retrieval.

Another, more conservative simplification, consists in disregarding the spread of firing probabilities over the population of A source neurons. Instead of introducing $\text{Pr}(u_j = 1) = a/A$, $j = 1, \dots, A$ as one of the factors the binomial probability of synaptic contribution $\text{Pr}(h = 1)$ we can restrict the binomial to the already firing neurons, in number a , and set $\text{Pr}(u_j = 1) = 1$, $j = 1, \dots, a$. Synaptic contributions become somewhat simpler upon this substitution of $H \sim B_{A, c_a q \pi a/A}$ by $H \sim B_{a, c_a q \pi}$. This is equivalent to substituting the binomial over firing neurons by just its average $A a/A = a$ and it is in this sense that we speak of disregarding its spread. The effect is making the non-correlational terms of the variance linear.^{4.7} This approximation should be particularly good at very low or very high activations, where the variance of the binomial is lower. Retrieval entails precisely those activation levels for the excitatory neurons, although it does not require a high or low inhibitory activation. For our standard network operating in the retrieval regime, σ_{Off} goes to about 51 up from 48. We collect the resulting means and variances, which will be the ones exclusively used henceforth, in Table 4.7.

$\mu_{\text{On}} = c_m m + c n - w_{\text{IE}} c_{\text{IE}} k$
$\sigma_{\text{On}}^2 = c_m m (1 - c_m) + c n (1 - c) + c^2 n \text{CV}_\pi^2[q] (n - 1) + w_{\text{IE}}^2 c_{\text{IE}} k (1 - c_{\text{IE}})$
$\mu_{\text{Off}} = c (m + n) - w_{\text{IE}} c_{\text{IE}} k$
$\sigma_{\text{Off}}^2 = c (m + n) (1 - c) + c^2 (m + n) \text{CV}_\pi^2[q] (m + n - 1) + w_{\text{IE}}^2 c_{\text{IE}} k (1 - c_{\text{IE}})$
$\mu_{\text{Inh}} = w_{\text{EI}} c_{\text{EI}} (m + n) - w_{\text{II}} c_{\text{II}} k$
$\sigma_{\text{Inh}}^2 = w_{\text{IE}}^2 c_{\text{EI}} (m + n) (1 - c_{\text{EI}}) + w_{\text{II}}^2 c_{\text{II}} k (1 - c_{\text{II}})$

Table 4.7. Means and Variances of Synaptic Input to On, Off and Inh Populations. These expressions are the basis for all further mean-field calculations in this document.

4.7. Observe that the simplified variance contribution $c_a a(1 - c_a)$ interpolates linearly the previous formula $c_a a(1 - c_a a/A)$ between the values $a=0$ and $a=A$. This is not the best local linear approximation to the nonlinear variance in the vicinity of those points. The respective series expansions yield, instead,

$$\begin{array}{ll} \text{at } a \simeq 0, & c_a a, \\ \text{at } a \simeq A, & c_a (1 - c_a) A + c_a (1 - 2c_a) (a - A). \end{array}$$

Note that means and variances in LK2006 assume *both* random imprinting and no variability of firing neurons, whereas with exception of illustrative Figure 4.8 here, we limit ourselves to the second, safer approximation (compare blue and dark red curves in Fig. 4.11). Something can still be done to render the formulæ more friendly to analytical work: in Appendix B we present an approximation to the correlational terms in the variance that stems from expanding $CV_{\pi}^2[q]$ in a power series on f around zero. This approximation is very accurate for the low coding ratio regimes of interest and makes those terms less unwieldy.

4.3 Phase-Space Tools to Analyze Replay

In this Section attention is focused on the two-dimensional, i.e. excitatory-only, mean-field model of sequence replay from LK2006 (with updated variances). After explicitly writing out the mean-field rule, we proceed to the analysis of replay in phase space and in the phase diagram, two important techniques that shall see continued application in the last Section, when inhibition is incorporated.

4.3.1 The Mean-Field Rule Fully Specified

Recall from Eq. 4.20 that the principle of the mean field rule is to identify the fraction of activated neurons of the On/Off population in the next time step with the fraction of the input distribution to an On/Off neuron that was over its threshold θ at the last time step (Fig. 4.12).

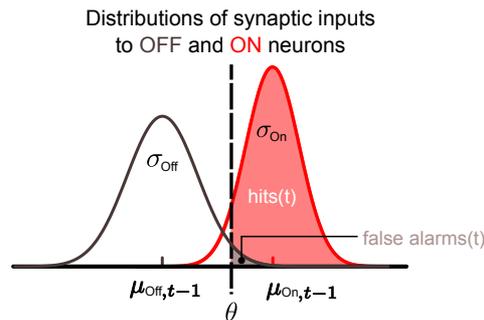


Figure 4.12. To Fire or Not to Fire?. Unimodal pdfs for synaptic inputs to Off (brown outline) and inputs to On neurons (red outline). Filled areas (overthreshold) represent probabilities of activation of Off and On postsynaptic neurons, respectively. The tail in the distribution of hits leftward of the threshold represents the false negatives (silent On neurons), whereas the bulk of the distribution of synaptic inputs to Off neurons leftward of the threshold is the amount of true negatives.

This amounts to integrating probability distribution functions that were assumed to be unimodal, whence the resulting cumulative distribution function will be a sigmoid

function, i.e. both T_{On} and T_{Off} in Eq. 4.19 will be of sigmoidal shape:

$$\begin{aligned} m_{t+1}/M &= \text{sigmoid}(\mu_{\text{On}}(m_t, n_t), \sigma_{\text{On}}(m_t, n_t); \theta) && \text{hits} \\ \underbrace{n_{t+1}/F}_{\% \text{ active next}} &= \underbrace{\text{sigmoid}(\mu_{\text{Off}}(m_t, n_t), \sigma_{\text{Off}}(m_t, n_t); \theta)}_{\% \text{ inputs over threshold}} && \text{false alarms.} \end{aligned}$$

The preceding Section furnished us with Gaussian approximations to both pdfs of synaptic input that are accurate for sufficiently large pattern size P . The mean field equations result from integrating the Gaussian φ centered on μ with width given by σ ,

$$\frac{1}{\sigma} \varphi\left(\frac{h-\mu}{\sigma}\right) \equiv \frac{1}{\sigma \sqrt{2\pi}} \exp\left(-\left(\frac{h-\mu}{\sqrt{2}\sigma}\right)^2\right).$$

to the right of the threshold θ

$$\begin{aligned} \text{sigmoid}(\mu, \sigma; \theta) &= \frac{1}{\sigma} \int_{\theta}^{\infty} \varphi\left(\frac{h-\mu}{\sigma}\right) dh && h' \equiv h - \theta \\ &= \frac{1}{\sigma} \int_0^{\infty} \varphi\left(\frac{h'+\theta-\mu}{\sigma}\right) dh' && z \equiv (-h' + \mu - \theta)/\sigma \\ &= \int_{\frac{\mu-\theta}{\sigma}}^{-\infty} \varphi(-z) dz \\ &= \Phi(z). \end{aligned}$$

Here the cdf of the Gaussian is

$$\Phi(z) = [1 + \text{erf}(z/\sqrt{2})]/2,$$

where $\text{erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. Its argument z depends on the target population considered ($\mu_{\text{On/Off}}, \sigma_{\text{On/Off}}$ given in Table 4.7):

$$z_{\text{On/Off}} = \frac{\mu_{\text{On/Off}} - \theta}{\sigma_{\text{On/Off}}}.$$

The mean-field equations for the excitatory-only system are thus, in detail,

$$\begin{aligned} \frac{m_{t+1}}{M} &= \frac{1}{2} \left[1 + \text{erf}\left(\frac{c_m m_t + c n_t - \theta}{2^{1/2} \sqrt{c_m m_t (1-c_m) + c n_t (1-c) + c^2 n_t \text{CV}_{\pi}^2[q](n_t-1)}}}\right) \right], \\ \frac{n_{t+1}}{F} &= \frac{1}{2} \left[1 + \text{erf}\left(\frac{c(m_t + n_t) - \theta}{2^{1/2} \sqrt{c(m_t + n_t)(1-c) + c^2(m_t + n_t) \text{CV}_{\pi}^2[q](m_t + n_t - 1)}}}\right) \right]. \end{aligned} \quad (4.34)$$

4.3.2 Receiver-Operating Characteristics

Receiver operating characteristics (or ROC curves) are the customary graphical tool to analyze binary classifiers, such as are in our case On and Off neurons. ROC curves

represent true positives (our hits; correctly activated On neurons, m) vs. false negatives (our false alarms; incorrectly firing Off neurons, n). The binary classifier has a degree of freedom to set the desired tradeoff between the two kinds of errors, false dismissals ($g = M - m$) and false alarms^{4,8}, n . In our detector system the ROC curve is parameterized by the neuronal firing threshold, θ .

In Figure 4.13 we represent a number of ROC curves for our excitatory network capturing the essential ingredient of time development in our system (we swap the axes and plot n vs m instead of m vs n to preserve comparability with LK2006 and later parts of this work). Concentrate, to start with, on the ROC curve whose cusp reaches closest to perfect replay ($m = M$, $n = 0$ at lower right). The equations for the errors Eq. 4.34 do depend on θ but also on the value of m , n at the previous iteration. In the Figure, the successive ROC curves as one progresses towards the upper left correspond to the updated values of m_t , n_t as we apply the rule always with the initial fixed threshold. The red dots show how the initial threshold tends to escape from the optimal region and so replay risks instability (errors too large to trigger a sufficiently good next pattern). In the particular cases illustrated, the fixed point of the map is reached after a few iterations during which the threshold becomes progressively too low (left) or too high (right panel). With a less fine tuning of the initial threshold (farther away from θ_{opt} at the cusp), the map does not converge before the threshold becomes clearly inadequate and replay bursts in the epileptic, all-active state or stalls in silence.

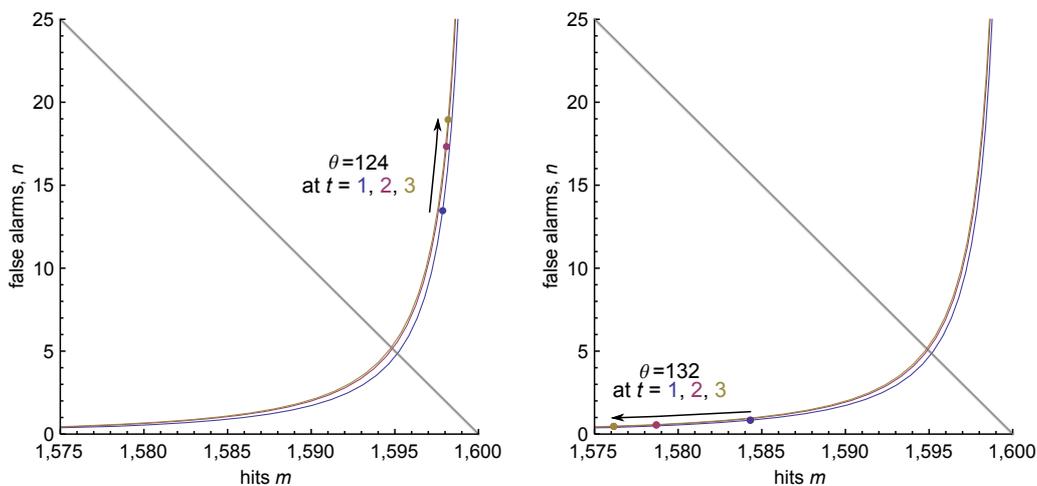


Figure 4.13. Receiver Operating Characteristic (ROC Curves) for the Neuron Detector.

In each of the two panels, replay is started with a perfect pattern $m_0 = M$, $n_0 = 0$ and a set threshold ($\theta = 124$ at left, $\theta = 132$ at right). Different ROC curves with decaying best performance (higher distance from cusp to lower right corner) correspond to the successive states of the mean-field dynamics, and are drawn until the system converges. As a result of the evolution of the system, the threshold (disc) becomes progressively maladapted for the decision task (black arrow indicates sense of the dynamics). Parameters are $N = 10^5$, $M = 1,600$, $c_m = 0.1$, and $c = 0.05$.

4.8. The error tradeoff will be affected by the relative weights assigned to each kind of error (see Section 4.2.3). In the discussion about ROC curves here for illustration we do not incorporate the error weighting scheme induced by the quality formula Eq. 4.11, which is heavily skewed and penalizes false dismissals much above false alarms.

On the methodological front, this observation prompts us to seek tools adapted to the analysis of the dynamical evolution of the system, because only in that scenario can the requirement of sustained replay be quantified. From the point of view of the performance of the memory network, the present analysis spawns interest in possible mechanisms that adaptively correct the threshold during the dynamics, so as to hold it close to the optimal static threshold for each iteration.

4.3.3 Mean Field Dynamics in the Phase Plane

To understand the qualitative dynamics in absence of an analytical solution we turn to the tools of dynamical systems theory (Strogatz, 1994). We first take a look at trajectories of the discrete map in (m, n) phase space. The trajectory bundles are best understood by examining the nullclines of the system, which are the loci of points where one further iteration does not change at least one of the coordinate e.g. $m_t - T_{\text{On}}(m_t, n_t) = 0$. We finally summarise the region of pattern sizes M and thresholds θ where replay is dynamically successful in a *phase diagram*, a $\theta(M)$ plot generalizing in the mean-field the classification of dynamical regimes first seen in Figure 4.8. The phase diagram allows to explore the factors that affect the minimum pattern size M_{opt} and thus the capacity of the system for sequence memory (see Eqs. 4.12 and 4.13).

4.3.3.1 Trajectories in the Phase Plane

In the forthcoming sections we shall be discussing the attractors of the dynamics of the *mean-field map*. The underlying cellular dynamics does not, however, have a fixed point and does not thus fall in the category of attractor memory networks (Hopfield, 1982). For the representation to be informative, we have chosen to display phase space in scaled coordinates (n/F vs. m/M), and only its lower-right corner close to the point of perfect replay (see green box in Fig. 4.7). A nonlinear scaling would distort the trajectories, and once trajectories escape the lower-right corner of phase-space they go straight to the all-silent fixed point at lower left or to the all-active one at upper right. Finally, in this representation, the displacements in both directions give rise to equivalent changes in the replay quality (Eq. 4.11).

4.3.3.2 Nullclines and Fixed Points

The nullclines of a phase portrait are the points where the flow is exclusively horizontal (n -nullcline, given in implicit form by $n_t - T_{\text{Off}}(m_t, n_t) = 0$) or vertical (m -nullcline, $m_t - T_{\text{On}}(m_t, n_t) = 0$). They break up phase space in regions according to the sign of the change in m and n . Nullclines do not tell directly about the stability of the regions they demarcate because they can be crossed by the flow, but their crossings are the fixed point of the system, where the flow comes to a halt ($m_{t+1} - m_t = 0; n_{t+1} - n_t = 0$).

Figure 4.14 shows the broader shape of the nullclines. The corners of the m and

n nullcline for high n are shown separately in two insets. The lower-right corner is zoomed-in (yellow) to facilitate analysis of the fixed point of retrieval.

Fixed points The crossings of the nullclines reveal three stable fixed points close to $(m/M, n/F) = (0, 0)$ (all-silent), $(1, 1)$ (all active) and $(1, 0)$ (retrieval), in this example with $M = 1,000$, $\theta = 79$. These fixed points are indicated by discs in Figure 4.14. The remaining fixed points are saddle points (repulsive in one direction and attractive in the other), and are indicated by squares.

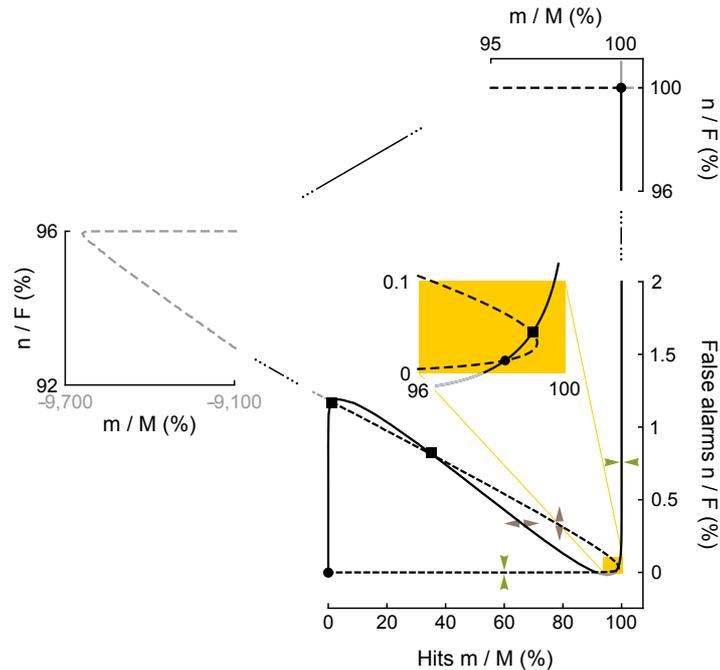


Figure 4.14. Overview of Nullclines and Fixed Points of the Mean-Field Model. The n -nullcline is shown dashed (mirror-Z-shaped); the m -nullcline full (N-shaped). Nonphysical sections of the nullclines with m/M or $n/F \notin [0, 1]$ are greyed out. Arrowheads indicate the repulsive (brown) or attractive (green) character of each of the nullcline branches. The main plot (bottom panel) shows the whole range of $m/M \in [0, 100\%]$ and the range of $n/F < 2\%$ that is of interest for retrieval dynamics. The turning point of the m -nullcline is shown at left; the all-active fixed point is shown in the inset at top. The yellow region zooms into the area of replay. It exhibits a saddle fixed point (square) and the stable fixed point (disc) of high-quality retrieval. Parameters were $M = 1,000$, $\theta = 79$, $c = 0.05$, $c_m = 0.1$, and $N = 10^5$.

Shape of the nullclines The shape of the nullclines can be understood by referring to the sigmoid function $(1 + \text{erf}(z))/2$ (depicted in Fig. 4.15). The sigmoid is centered on $z = 0$. It is mostly flat and close to zero for $z < -2$, mostly linear for $-2 < z < 2$ and mostly flat and close to 1 for larger z (this derives from the fact that two standard deviations about the mean contain about 95% of the area under a Gaussian curve). In our case, the scaling of z is provided by the standard deviation of the synaptic input to the respective population, $\sigma_{\text{On/Off}}$, and its origin $z = 0$ by the difference of the mean input $\mu_{\text{On/Off}}$ with the threshold θ . Let us discuss for example the m nullcline (black line

in 4.14) and disregard for the moment the inhomogeneity in m, n of the scaling σ_{On} :

$$m_t = [1 + \text{erf}((c_m m_t + c n_t - \theta)/\sigma_{\text{On}})]/2.$$

For $m_t \simeq 0$ (m nullcline) there will be a range of n_t such that $(c n_t - \theta)/\sigma_{\text{On}} < -1$. In the Figure that range stretches until about $n_t/F = 0.012$ or $n_t \simeq 1,200$. In view of $c = 0.05$ and $\theta = 80$ we have $(60 - 80) \lesssim -\sigma_{\text{On}}$. Now σ_{On} in the region of phase space $m \simeq 0, n = 1,200$ is about 9 (Table 4.7). We see that the qualitative condition $z \lesssim -2$ is fulfilled. Once we reach that point the linear section of the sigmoid kicks in (Fig. 4.15). A tradeoff between activation coming from hits and false alarms governs this linear section of the nullcline. There $m_t \propto (c_m m_t + c n_t - \theta)/\sigma_{\text{On}}$. The negative slope corresponds to the increasing compensation of false alarms by hits as the number of the latter grows. Eventually, the region where $m_t \simeq M$ is reached. If the threshold had been tuned to support replay, it should be slightly lower than $c_m M$ so that a majority of On neurons fire, and the nullcline will reach almost $m_t = M, n_t = 0$. From there on, further activation of Off neurons at maximum hit activity is already so much above threshold ($z_{\text{On}} > 2, \sigma_{\text{On}}(M, 0) \simeq 10$) that the m -nullcline overlaps completely with the $m_t = M$ vertical line.

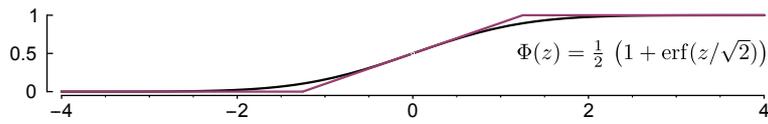


Figure 4.15. The Sigmoid Function Has Three Roughly Linear Regimes. The cumulative distribution function of the normal is a shifted and scaled error function. It is the prototypical threshold nonlinearity (sigmoid), with three regions that can, each, be approximated linearly. The leftmost region $z \in (-\infty, -2)$ is the under threshold region. Regardless of z the value of the function is very nearly zero. The central region $z \in [-2, 2]$ is approximately linear with slope $1/\sqrt{2\pi}$ (21.75°). The rightmost region is the saturation region. Regardless of $z \in (2, \infty)$ the value of the function is very nearly 1. The region limits $-2, 2$ have been set arbitrarily for illustration; they contain 95% of the area of the Gaussian. The intersections of the central straight line with 0 and 1 are at $\pm\sqrt{\pi}/2 \simeq 1.25$.

A similar reasoning can be followed for the n -nullcline although the quantitative details are different due to the sparsity of the network ($M \ll F$) and the different connectivity ($c_m = 2c$). Importantly, the connectivities influences the slopes of the diagonal section of the nullclines, thus controlling whether they cross to generate the fixed point of stable replay.

Effect of threshold on the phase plane The nullclines of the mean-field dynamical system are shown in Figure 4.16 B for a threshold setting that, similar to similar to Figure 4.14, allows for stable retrieval: there exists an asymptotically stable fixed point

$$(m_\infty, n_\infty) = [T_{\text{On}}(m_\infty, n_\infty), T_{\text{Off}}(m_\infty, n_\infty)]$$

with many hits $m_\infty \simeq M$ and few false alarms $n_\infty \ll N - M$.

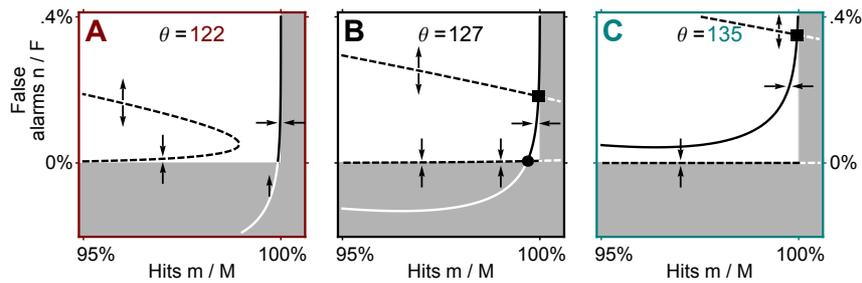


Figure 4.16. Effect of the Firing Threshold on the Nullclines and Fixed Points. B. The m nullcline (solid line) intersects twice with the n nullcline (dashed), producing stable (disc) and unstable (square) fixed points. Arrows indicate attractive or repulsive character of the nullcline; gray areas correspond to unphysical values $n < 0$ or $m > M$. **A.** Threshold too low; no fixed points. The trajectories starting out nearby a perfect pattern escape to the all-active fixed point at $(m/M, n/F) = (1, 1)$. **C.** Threshold too high. Only the unstable fixed point remains; trajectories escape to the all-silent fixed point at $(0, 0)$.

If the firing threshold is too low or the pattern size is too large (Fig. 4.16 A), the nullclines do not cross in a retrieval regime: After initialization at the condition of perfect retrieval $(m_0, n_0) = (M, 0)$, all neurons immediately start to fire and the network falls into an all-active state, $(m, n) \simeq (M, F)$. If the firing threshold is too high or the pattern size is too low (Fig. 4.16 C), only an unstable fixed point exists in the retrieval region. After initialization at perfect retrieval, the network immediately falls into an all-silent state $(m, n) \simeq (0, 0)$.

Note that in Figure 4.16 B as compared to Figure 4.14 (yellow inset) the m nullcline intersects the horizontal branch of the n nullcline at $n = 0$. This fact allows to apply the implicit function theorem to the nullclines so as to extract proper functions $m(n)$ and find out in closed form whether they intersect at $n = 0$. These formulas will work unfortunately only for the case where the loss of stability is, as depicted in C, by too high a threshold, because the curvature of the branch of the n nullcline $n(m)$ in the vicinity of $m = M$ is not negligible. The original idea of applying the implicit function theorem is due to Axel Kammerer, and we have sketched it here as a possible starting point for further analytical work, since it can provide a good approximation of the upper edge of the stability wedge (cf. next Section) for many networks of interest.

4.3.3.3 Phase Diagram: the Wedge of Stability

In this Section we investigate systematically which combinations of M, θ enable high-quality, sustained replay. For that, we employ a *phase diagram* that displays the asymptotic stability of a range of M, θ combinations by a color code (choropleth map).

This phase diagram (Fig. 4.17) reveals the three phases of a sequence memory network that we have come to expect from cellular simulations: all silent, all active, and retrieval. The region in which retrieval is possible is wedge-shaped with a thin tip at low pattern sizes M . We tested the region of retrieval obtained from the mean-field equations with extensive computer simulations of the corresponding networks of binary neurons (white discs in Fig. 4.17; see Section 4.1.6). As expected, owing to the finite size of the simulated network, the region of retrieval is overestimated by mean-field theory, yet the deviations are relatively small. According to Eq. 4.12, the number P of stored associations increases with decreasing coding ratio $f = M/N$, and thus the network displays the highest memory capacity at the wedge tip $M = M_{\text{opt}}$.

There the stability of the fixed point is particularly sensitive to noise and thus the high capacity is not accessible unless the dynamics can be stabilized. This could be done by a mechanism that senses the escape of trajectories to the non-replay fixed points and adapts the threshold accordingly. We therefore investigated whether inhibition can provide such stabilization by effectively adapting the threshold.

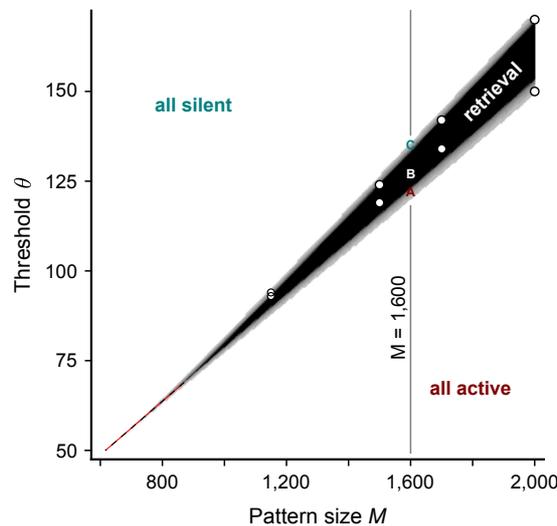


Figure 4.17. Phase Diagram. M - θ space consists of three areas: *All silent*; the sequence dies out. *All active*; all neurons fire at maximum rate. *Sequence retrieval* (black); the fraction of hits is much larger than the fraction of false alarms for infinitely many time steps. The diagonal dashed line at bottom left separates the all-silent and all-active phases for M -values at which no retrieval phase exists. Areas in light gray correspond to transient retrieval of at least 3 time steps. White discs mark the boundary of the retrieval region as obtained from simulations of N binary neurons for exemplary values of M . Note colored labels A, B, C marking the M , θ combinations of the respective phase-space panels in Figure 4.16. Parameters are $N = 10^5$, $M = 1,600$, $c_m = 0.1$, and $c = 0.05$.

4.3.4 Bayes-Optimal Threshold

In this Section we ask what is the threshold function of the the current activity levels m , n that improves capacity the most.

Telling distributions apart from one observation We established in Section 4.2.3 the dichotomic decision problem that is to be solved by each excitatory neuron. There, we laid out two hypotheses (Eq. 4.6): the detector neuron is supposed to fire (belongs in the On population of the current time step; activation hypothesis A_{On}), or it is not (Off population; A_{Off}). The respective priors $p(A_{\text{On}}) = f$ and $p(A_{\text{Off}}) = 1 - f$ are given by the coding ratio f , which stipulates how many active neurons code for a pattern at each time step. The substantive data for the decision is a one-shot sample from the distribution of synaptic input levels. Following the reasoning of previous sections, we may simplify such distribution as Gaussians whose location (mean) and spread (variance) depend on whether the detector neuron is target in a pattern or not at the current cycle of the replay process. The problem is analogous to deciding whether a person is from Holland or from Spain based exclusively on her stature and the respective country averages and standard deviations.

Bayes detector A Bayesian strategy to solve this problem optimally seeks to maximize the *probability of success* S , i.e. the probability of taking the right decision:

$$S = \Pr(\text{spike}|A_{\text{On}}) \Pr(A_{\text{On}}) + \Pr(\text{silence}|A_{\text{Off}}) \Pr(A_{\text{Off}}).$$

Given the spike generation model ($\text{spike} \Leftrightarrow h \geq \theta$), the conditional probabilities of spike or silence above correspond to integrals of the respective probability densities over regions of synaptic input separated by the boundary θ :

$$S(\theta) = \Pr(h \geq \theta | A_{\text{On}}) \Pr(A_{\text{On}}) + \Pr(h < \theta | A_{\text{Off}}) \Pr(A_{\text{Off}}).$$

The forward model that computes the synaptic input distribution given the neuron type is already available from the development of the mean-field rule in Section 4.3.1 above:

$$S(\theta) = \frac{m}{M} \Pr(A_{\text{On}}) + \left(1 - \frac{n}{F}\right) \Pr(A_{\text{Off}}).$$

(here and elsewhere in this Section we have elided the cycle index t). The equations leading to Eq. 4.34 allow to rewrite the success probability in terms of Gaussian cdfs:

$$S(\theta) = f \Phi\left(\frac{\mu_{\text{On}} - \theta}{\sigma_{\text{On}}}\right) + (1 - f) \left(1 - \Phi\left(\frac{\mu_{\text{Off}} - \theta}{\sigma_{\text{Off}}}\right)\right).$$

Optimal threshold The threshold θ_{opt} that maximizes the success probability can be readily obtained by demanding $dS(\theta)/d\theta = 0$. Since $\partial\Phi(z)/\partial\theta = -e^{-z^2/2}/(\sqrt{2\pi}\sigma)$ we have

$$\frac{dS}{d\theta} = \frac{1}{\sqrt{2\pi}} \left(f \frac{e^{-z_{\text{On}}^2/2}}{\sigma_{\text{On}}} - (1 - f) \frac{e^{-z_{\text{Off}}^2/2}}{\sigma_{\text{Off}}} \right)$$

i.e. the optimal threshold is at the crossing point of the firing probability-weighted Gaussians. The resulting equation is quadratic in θ :

$$z_{\text{Off}}^2 - z_{\text{On}}^2 = 2 \log \frac{1 - f}{f} \frac{\sigma_{\text{On}}}{\sigma_{\text{Off}}},$$

and has roots

$$\theta_{\pm} = \frac{c(m+n)\sigma_{\text{On}}^2 - (c_m m + c_n)\sigma_{\text{Off}}^2 \pm \sigma_{\text{On}}\sigma_{\text{Off}} \sqrt{(c - c_m)^2 m^2 + (\sigma_{\text{Off}}^2 - \sigma_{\text{On}}^2) \log \left[\frac{M^2}{F^2} \frac{\sigma_{\text{Off}}^2}{\sigma_{\text{On}}^2} \right]}}{\sigma_{\text{On}}^2 - \sigma_{\text{Off}}^2}.$$

Generally one of the thresholds is positive and the other negative when $m \simeq M$, $n \ll F$, which enables heuristical identification of the sign leading to maximization of $S(\theta)$ (evaluating the second derivative is cumbersome).

Dependence of the optimal threshold on the activity level Foreshadowing our interest in the adaptive regulation of the threshold, we ask how the threshold should change with the excitatory activities m and n . Figure 4.18 shows that the level curves

of the function $\theta_{\text{opt}}(m, n)$ in the phase-space region of retrieval (grey) can be very well approximated by a Taylor expansion (dashed; $\partial_x f \equiv \partial f / \partial x$):

$$\theta_{\text{opt}}(m, n) = \theta_{\text{opt}}(M, 0) + \partial_m \theta_{\text{opt}}(M, 0)(m - M) + \partial_n \theta_{\text{opt}}(M, 0)n.$$

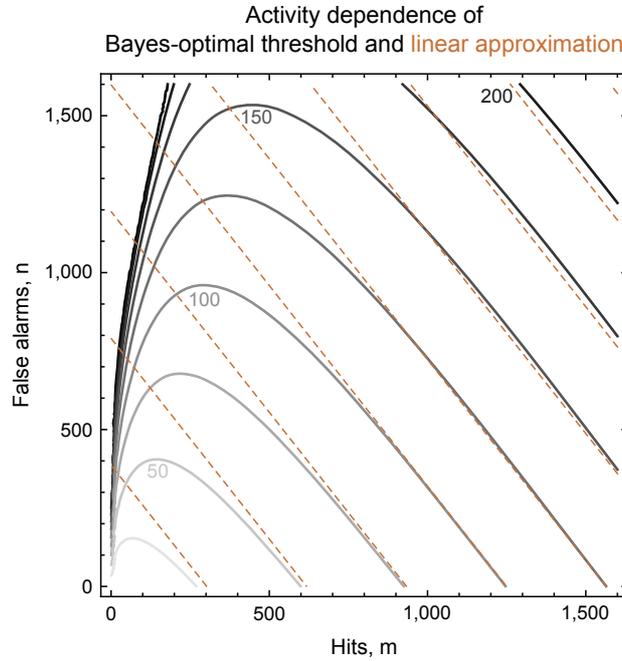


Figure 4.18. The Optimal Detector Threshold is Linear in the Network Activity.

Level curves (grey) of the optimal threshold as a function of the numbers of firing On (abscissa) and Off neurons (ordinate). In this exemplary network with $c = 0.05$, $c_m = 0.1$, $N = 10^5$ and with patterns of $M = 1,600$ neurons, the linear approximation (dashed brown; $\theta(m, n) \simeq 1.118 + 0.079m + 0.062n$) is very accurate in the retrieval region of phase space.

As intuitively expected, when the activity grows the optimal threshold increases. Hits are weighted slightly above false alarms, because they can elicit more activity owing to their denser connectivity. The line $\theta(M, 0) = \partial_m \theta_{\text{opt}}(M, 0)(m - M)$ (perfect replay) cleanly bisects the stability wedge of Figure 4.17: the Bayes-optimal threshold finds the core of the numerically calculated mean-field wedge.

Note that Figure 4.18 demonstrates the optimality of a threshold linearly increasing with activity levels (m, n or any linear combination thereof), but also that the particular coefficients obtained in the vicinity of $M, 0$ (lower right) provide a good quantitative approximation. We take advantage of this fact to study how the optimal threshold depends on the pattern size M . In order to do that, one must decide for a trajectory along a particular phase space direction. The most relevant dangers for the stability of our system are, as was shown in the ROC curves (Fig. 4.13) and later confirmed by analysis of the nullclines in phase space (Fig. 4.16), those that run away from the stable retrieval points along the axes of m (direction decreasing m) or of n (direction increasing n). The respective magnitudes of the coupling coefficients $\partial_m \theta_{\text{opt}}$ and $\partial_n \theta_{\text{opt}}$ are shown in units of c in Figure 4.19. The average network connectivity c provides a good unit because, as will be shown below, coupling coefficients between threshold and activity behave like global connectivity offsets.

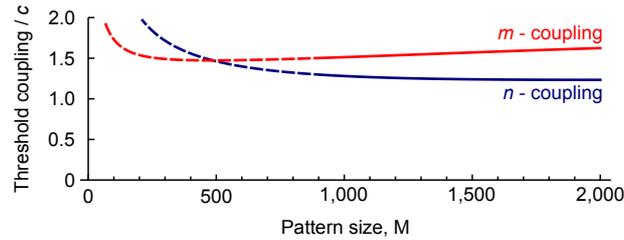


Figure 4.19. Threshold Coupling to Activity Along the Escape Trajectories. Coefficients that couple the optimal detector threshold θ_{opt} and excitatory population activity of On and Off cells (red and dark blue, respectively; dashed where $M < M_{\text{opt}} = 880$).

The coupling coefficients are rather M -independent (for all M s for which the θ_{opt} is indeed stable, $M > 880$, see Fig. 4.17; recall also that the linear approximation breaks below about that value, Fig. 4.18). They range between c and the $c_m = 2c$ of the example network, and the larger value of the coupling to m reflects well the need for the network to hold back more strongly those better-connected cells that belong to an association currently in replay.

Finally, we observe that trajectories with $m - n = \text{const.}$, such as occur close to equilibrium (Fig. 4.13) are roughly perpendicular to the isolines in Figure 4.18 and show a coupling coefficient of about $2c$ for the allowable pattern sizes M (not shown).

4.4 Inhibition Increases Capacity for Sequence Replay

The relation between Bayes-optimal threshold and network activity elucidated above, together with the insight about the need of an adaptive threshold form the motivation of the present Section. Here we investigate whether an inhibitory feedback can help the network replay at lower pattern sizes and thus at higher capacity.

4.4.1 Instantaneous Inhibition

We first analyze the role of inhibition assuming an instantaneous negative feedback proportional to the total number $m + n$ of active neurons. This feedback effectively increases the threshold such that the dynamics is derived from Eqs. 4.34 by substituting $\theta \rightarrow \theta + b(m_i + n_i)$, where the positive b is a feedback gain. This feedback gain is an effective connectivity (see below), and it is advantageous to express it as a multiple $b = \beta c$ of the effective average connectivity c , introducing the dimensionless inhibitory coupling β .

The main effect of inhibition is as follows. When the threshold θ is too low (as in Fig. 4.16 A), inhibition moves the n -nullcline rightward; when θ is too high (as in Fig. 4.16 C), inhibition moves the m -nullcline downward; and when θ lies along the midline of the uninhibited wedge, inhibition moves both nullclines—as depicted in Fig. 4.20 A for a pattern size M below the M_{opt} of the purely excitatory model. Thus, inhibition restores the stable fixed point and therefore effectively enlarges the retrieval phase (Fig. 4.20 B). In particular, inhibition lowers the optimal pattern size

M_{opt} . While $M = 800$ is not stable in the thin tip of the uninhibited system attending to our criterion (that requires $m/M > 0.9$ and $n/F < 0.1$), it becomes stable somewhere between $b = 0.5 c$ and $b = 0.8 c$. Memory capacity, which scales as $\alpha \propto M_{\text{opt}}^{-2}$, is thereby enhanced by the inhibitory feedback (by a factor of about 2 in the example of Fig. 4.20 C). We verified that the scaling of memory capacity observes a logarithmic dependence $M_{\text{opt}} \propto \ln N$ regardless of the presence of inhibition (Fig. 4.20 D). The capacity, thus, grows with network size as $\alpha \propto N / (\ln N)^2$.

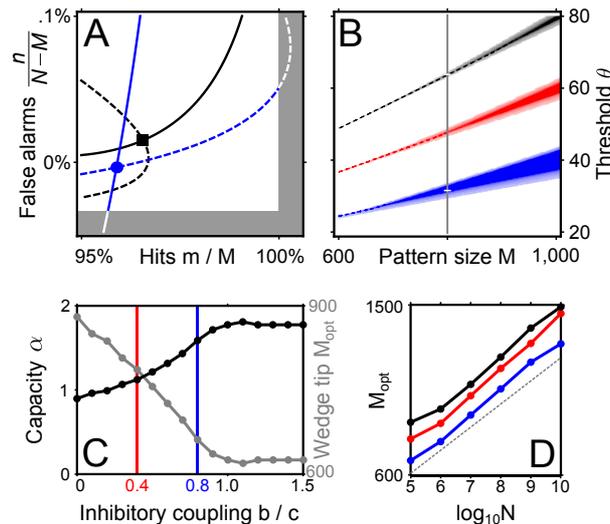


Figure 4.20. Instantaneous Inhibition. Increasing inhibitory feedback from $b = 0$ (black) through $b = 0.4 c$ (red) all the way to $b = 0.8 c$ (blue) stabilizes previously unstable replay. **A.** Nullclines at $M = 800$ and $\theta = 63.7$ (red omitted for clarity). Firing threshold for the blue nullclines is offset to $\theta - b M = 31.7$ to account for lower mean input h (see text). **B.** Phase diagram. Grey vertical line indicates M value used in A. **C.** Capacity (black) and minimum pattern size M_{opt} (gray) as a function of b . **D.** The minimum pattern size grows sublinearly with N so that the capacity shows an overall increase. The dashed grey line indicates a logarithmic dependence $M \propto \ln N$. Parameters were $N = 10^5$, $c = 0.05$, $c_m = 0.1$.

Threshold offset induced by inhibition With the introduction of an inhibitory feedback the mean level of synaptic activity during retrieval is reduced in a predictable manner and the threshold has to adjust correspondingly. For the area of retrieval in phase space as well as for all phase diagram points within the wedge, where $m \simeq M$, $n \simeq 0$, the threshold offset can be approximated as $-b M$. Applying this offset to the respective thresholds renders networks with different levels of inhibitory feedback comparable both in the phase plane and in the phase diagram. The correction in the phase plane was adopted in Figure 4.20 A, because for the uncorrected threshold there would be no crossing of the nullclines under inhibition. This can be seen in the uncorrected phase diagram of panel B, where $\theta = 63.7$, as used for the uninhibited (black) nullclines, lies far above the blue stability wedge. Figure 4.21 illustrates the removal of the offset both in the phase plane (inset) and in the phase diagram (main plot). There, a strong inhibitory feedback ($b = 0.8 c$; blue wedge) rescues from instability the network with $M = 880$, $\theta = 69$ (yellow disc), which has a threshold too low for replay without inhibition.

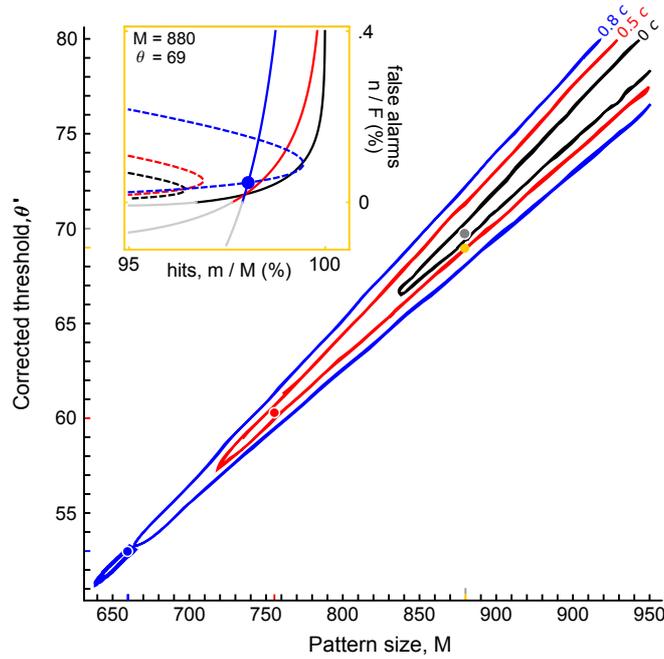


Figure 4.21. Replay-Level Threshold Normalization. A network is rescued from falling in the all-active regime by adding an inhibitory coupling strong enough to broaden the region of retrieval to its operating point ($M = 880$, $\theta = 69$; yellow dot). Maximum capacity networks are marked as black, red and blue dots ($b = 0, 0.4c, 0.8c$ respectively) on the phase diagram; They do not appear exactly at the wedge of the tip because the plotting of the wedge uses an adaptive mesh, while the capacity optimization routine scans a (coarser) homogeneous mesh.

Inhibition works both to cool down and heat up the network It is intuitive that inhibition, by supplying a negative feedback on the overall activity, should prevent the runaway excitation phenomenon that menaces the stability of recurrent excitatory networks. But how does it rescue replay from above the midline in Figure 4.21? Inhibition has an effect on the z -score of the synaptic input to On and Off cells. The average synaptic inputs to On and Off populations were, respectively,

$$\begin{aligned}\mu_{\text{On}} &= c_m m + c n \\ \mu_{\text{Off}} &= c(m + n).\end{aligned}$$

From those, inhibitory coupling subtracts without distinction

$$\theta_\beta = \beta c(m + n) = \beta \mu_{\text{Off}}.$$

Still there is one free parameter in a settable θ' , so that the mapping with instantaneous inhibition reads

$$\begin{aligned}m &= M \Phi\left(\frac{\mu_{\text{On}} - \beta \mu_{\text{Off}} - \theta'}{\sigma_{\text{On}}}\right) \\ n &= F \Phi\left(\frac{(1 - \beta) \mu_{\text{Off}} - \theta'}{\sigma_{\text{Off}}}\right).\end{aligned}$$

This θ' will be set as in the uninhibited case by demanding stable retrieval. If for a given M the threshold θ was of stable retrieval in the uninhibited map, the inhibited map will choose a θ' that is correspondingly lower, and that by about $b m = \beta c M$, as discussed above. We thus obtain the relationship of the parameter θ' with inhibition to the parameter θ without inhibition, during retrieval:

$$\theta' \simeq \theta - \beta c M$$

This allows to recast the formulas of the inhibited system

$$\begin{aligned} m &= M \Phi \left(\frac{\mu_{\text{On}} + \beta (c M - \mu_{\text{Off}}) - \theta}{\sigma_{\text{On}}} \right) \\ n &= F \Phi \left(\frac{\mu_{\text{Off}} + \beta (c M - \mu_{\text{Off}}) - \theta}{\sigma_{\text{Off}}} \right), \end{aligned}$$

We can identify $\Delta\theta = -\beta c [(M - m) - n] = -\beta c (g - n)$ as a supplementary and dynamic correction to the threshold that is proportional to the difference between false dismissals and false alarms. Inhibition, not caring about whether it is triggered by On or Off neurons, will remain absent as long as the errors compensate in number (diagonal of phase space). However, it does activate when the trajectory escapes towards full activation ($n \rightarrow F$) or silence ($g \rightarrow M$). Its meaning can only be understood in a dynamical scenario, since for the static optimal detector it just contributes a shift in the mean synaptic input (a rigid displacement to the right of the bumps in Fig. 4.12) and thus brings no advantage. When the network is going *fast* to full activation (Fig. 4.13, left), $\Delta\theta > 0$: inhibition depresses both hits and false alarms. That it does so in equal measure does not diminish its practical effect of bringing the system to cross the diagonal region of convergence (Fig. 4.14). The opposite effect sets in when the state goes to silence well parallel to the m axis (Fig. 4.13, right). The adaptive contribution becomes $\Delta\theta < 0$ and the effect is of effective net excitation, again directing the system towards the diagonal, where it will be collected by the normal dynamics. This is the scenario of recovery of stability if the system was initialized close enough to perfect replay.

Note that this global inhibition, however, has no power to single out On neurons at the origin of false dismissals and weigh them preferentially, as a strategy in line with the asymmetric definition of quality by Eq 4.11 would demand.

This line of argument explains the broadening of the wedge along the θ axis in the phase diagram: previously unstable threshold settings are stabilized by inhibition. One can argue considering m, n fixed and M variable. Variations in M at fixed m will lead to increased false dismissals, which can be compensated for by inhibition. Indeed, increasing M will result in that whenever there are false dismissals, they are more in absolute terms at the same level m/M (for it is this ratio that is used to decide stability), hence causing a stronger negative inhibition (net excitation). Decreasing M will instead result in an imbalance $n > g$, with the opposite correcting behaviour (positive inhibition). Hence, a range of $M + \delta M, M - \delta M$ for every stable M at given θ becomes accessible by the use of feedback inhibition.

Inhibition acts similar to resources for plasticity One can observe that the action of inhibition is mediated by a reduction by equal amounts of the effective connectivities c and c_m to $c - \beta c$ and $c_m - \beta c$ respectively. It has been shown (Leibold and Kempter, 2006) that free resources for plasticity, i.e. $r \equiv c_m/c - 1$, which depend on the quotient of both connectivities, greatly improve the region of stability. By switching on inhibition those become, effectively,

$$r_{\text{inh}} = \frac{c_m - \beta c}{c - \beta c} - 1 = \frac{c_m/c - \beta}{1 - \beta} - 1 = \frac{r + 1 - \beta}{1 - \beta} - 1 = \frac{r}{1 - \beta}.$$

In this approximation, increasing the value of β in the range $[0, 1)$ has an “equivalent” *nonlinear* boosting effect on the resources for plasticity. As the quotation marks indicate, this aspect of inhibition has to be cautiously handled, for it disregards the effects that addressing r has on the variances and that are not procured by inhibition.

Defective cue patterns Consider now the scenario of a defective initialization of the network, or an in-replay state of a perfectly initialized network somewhat away from optimal performance at $m = M, n = 0$, for example $m = M, n = \varepsilon \ll F$. The map reads

$$\begin{aligned} m &= M \Phi \left(\frac{c_m M + c \varepsilon - \beta c (M + \varepsilon) - \theta'}{\sigma_{\text{On}}} \right) \\ n &= F \Phi \left(\frac{c M + c \varepsilon - \beta c (M + \varepsilon) - \theta'}{\sigma_{\text{Off}}} \right). \end{aligned}$$

We established above that the new M term $-\beta c M$ is immaterial because it is absorbed by θ' , which is free to choose for replay. See here that both cdfs will now be evaluated at lower argument if the inhibition is switched on, $\beta > 0$. A sufficiently high β (e.g. $\beta = 1$) can compensate the effect of the deviation ε on the false alarms, at the expense of also reducing the number of hits. How relevant proportionally this reduction is depends on the quantity $\beta c \varepsilon / \sigma_{\text{On/Off}}$, i.e. it will be more relevant for the more localized distribution, which is usually σ_{Off} .

Upon division by M in the argument throughout

$$\begin{aligned} m &= M \Phi \left(\frac{c_m + c(1 - \beta)\varepsilon/M - [\beta c + \theta'/M]}{\sigma_{\text{On}}/M} \right) \\ n &= F \Phi \left(\frac{c + c(1 - \beta)\varepsilon/M - [\beta c + \theta'/M]}{\sigma_{\text{Off}}/M} \right), \end{aligned}$$

the relative importance of the perturbation is seen to depend on $(1 - \beta)/M$. The perturbation could be minimized by increasing M or by increasing the inhibitory coupling. This is in line with the observation that unstable states can become stable in a less capable network (with larger pattern size M) or in a more inhibited one (with a broader wedge).

If the initial cue is defective rather because it contains false dismissals $\gamma > 0$, i.e. $m = M - \gamma, n = 0$ we have

$$\begin{aligned} m &= M \Phi \left(\frac{c_m(M - \gamma) - \beta c(M - \gamma) - \theta'}{\sigma_{\text{On}}} \right) \\ n &= F \Phi \left(\frac{c(M - \gamma) - \beta c(M - \gamma) - \theta'}{\sigma_{\text{Off}}} \right), \end{aligned}$$

or

$$\begin{aligned} m &= M \Phi \left(\frac{c_m M - (c_m - \beta c) \gamma - [\theta' + \beta c M]}{\sigma_{\text{On}}} \right) \\ n &= F \Phi \left(\frac{c(M - (1 - \beta) \gamma) - [\theta' + \beta c M]}{\sigma_{\text{Off}}} \right). \end{aligned}$$

Here inhibition will only manage to hide the perturbation for the hit neurons if it is coupled with $\beta = c_m/c$, which corresponds to $\beta = 2$ for the exemplary values of c, c_m used in this work. Every bit above $\beta = 1$ will come at the expense of augmenting the false alarms in return. Hence $\beta = 1$ strikes a balance, by taking advantage of the higher localization of the distribution of synaptic inputs to Off neurons for cues contaminated with false alarms and minimizing the effect of false dismissals too without side effects.

4.4.2 Dynamic Inhibition

We next asked whether the effects observed with instantaneous inhibition can be reproduced with dynamic inhibition, i.e. in a scenario where inhibitory activity has its own dynamics, and whether additional features would arise. To this end, we extended the discrete map by including a third dynamical variable k_t that accounts for the number of spikes in an inhibitory pool following Section 4.1.5, Eq. 4.19, and the distributions of synaptic input of Table 4.7. In brief, the additional equation is

$$T_{\text{Inh}}(m, n) = \Phi[(\mu_{\text{Inh}} - \eta)/\sigma_{\text{Inh}}], \quad (4.35)$$

or, expanding the moments

$$\frac{k_{t+1}}{K} = \frac{1}{2} \left[1 + \text{erf} \left(\frac{c_{\text{EI}} w_{\text{EI}} (m_t + n_t) - \eta}{2^{1/2} \sqrt{w_{\text{EI}}^2 c_{\text{EI}} (1 - c_{\text{EI}}) (m_t + n_t)}} \right) \right]. \quad (4.36)$$

The inhibitory action on the sequence-related variables m and n is implemented by replacing the thresholds in Eqs. 4.34 by $\theta \rightarrow \theta + w_{\text{IE}} c_{\text{IE}} k_t$, and the variances by

$$\sigma_{\text{On/Off}}^2(m, n) \rightarrow \sigma_{\text{On/Off}}^2(m, n) + k w_{\text{IE}}^2 c_{\text{IE}} (1 - c_{\text{IE}}).$$

Initialization To test for sequence retrieval, the map is initialized with a perfect pattern and matching inhibition, $(m_0, n_0, k_0) = [M, 0, K T_{\text{Inh}}(M, 0)]$. The resulting phase diagram reveals again regions of stable and transient retrieval (Fig. 4.22).

In agreement with the instantaneous linear inhibition model, the retrieval region in the phase diagram extends to lower pattern sizes M (higher capacities). However, the non-linearity of the sigmoidal Gaussian cdf in Eq. 4.35 introduces a shearing of this region, which can be explained as follows: The Gaussian cdf is roughly linear in the vicinity of the inhibitory threshold η and virtually flat elsewhere (see Fig. 4.15). Hence, as an approximation, inhibition has no effect at low total activities $m + n$, it adds a constant $w_{\text{IE}} c_{\text{IE}} K$ to the threshold θ at high total activities, and it establishes a nearly linear coupling for intermediate regimes, similar to the instantaneous inhibition model from Figure 4.20 and the demand of Bayes optimality (Section 4.3.4). During sequence retrieval, total activity is approximately constant, $m + n \simeq M$, and therefore the retrieval region of the dynamical inhibition model can be understood as a combination of three retrieval regions of the instantaneous inhibition model for different feedback gains b and thresholds θ .

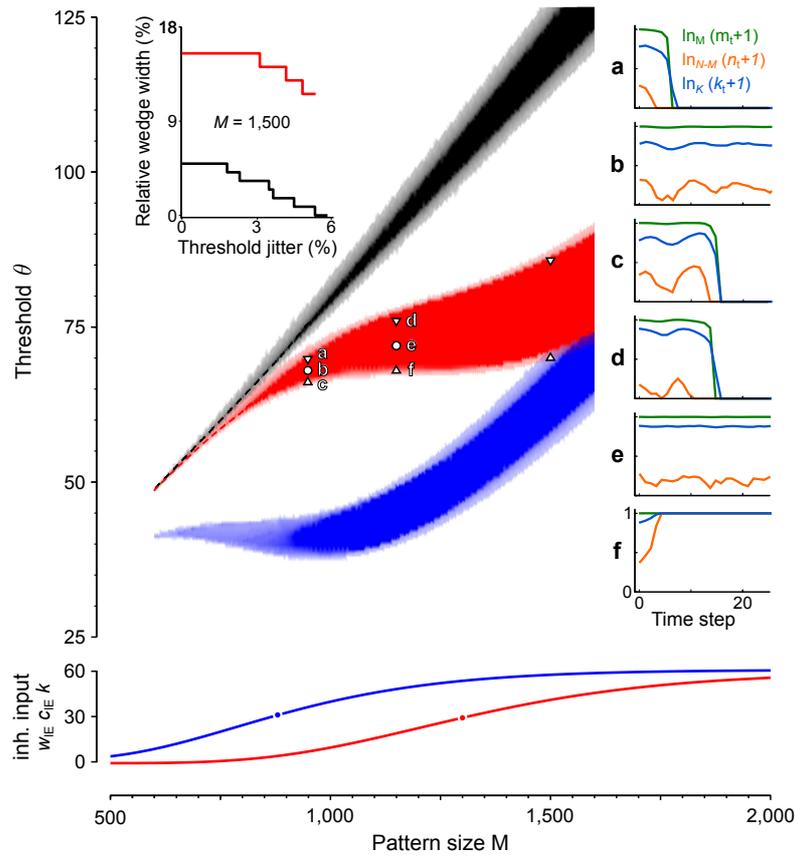


Figure 4.22. Dynamical Feedback Inhibition. **Main panel.** Retrieval regions in M - θ space for the 2-dimensional model without inhibition (black) vs. the 3-dimensional model with increasing dynamical feedback inhibition (red, blue). Triangles on the red region mark the first values of θ (in integer steps) for which sequence retrieval is unstable in simulations. Parameters are $K = 5,000$, $w_{IE} = 0.012$, $c_{IE} = 1$, $w_{EI} = 1$, $c_{EI} = 0.01$, and $\eta = 13$ for the red region; $\eta = 8.8$ for the blue region. **a-f.** Example trajectories from network simulations for M - θ pairs as indicated by white markers a-c at $M = 950$ and d-f at $M = 1,150$ with the activities m_t in green, n_t in orange, and k_t in blue. **Bottom panel.** Inhibitory synaptic input at perfect replay ($m = M$, $n = 0$) provides maximum feedback coupling at different pattern sizes: $M = 880$ for the blue wedge vs. $M = 1,300$ for the red wedge (discs). **Inset.** Robustness of sequence retrieval against threshold jitter with (red) and without inhibition (black) at $M = 1,500$ for simulated networks with threshold noise.

The broadening in θ of the region of retrieval (for constant M) with both instantaneous and dynamic inhibition suggests that sequence memory becomes more robust. Mean field theories, however, generally overestimate the regions of stable sequence retrieval. To assess the predictive power of our mean field results, we ran further simulations (see Section 4.1.6) where neuronal thresholds θ were jittered according to a Gaussian process. The results show that the increase of the relative range of thresholds by inhibition indeed withstands threshold noise (Fig. 4.22, Inset).

For maximum robustness at high capacities, the area of retrieval should be broadest at minimum $M = M_{\text{opt}}$. We suggest two heuristic criteria for the parameters of dynamic inhibition. First, to achieve maximum sensitivity of inhibition to excitation, the linear region of $T_{\text{Inh}}(m, n)$ should be centered at the average total input $m + n \simeq M$ during retrieval. This requirement is granted by setting the inhibitory

threshold to $\eta = \mu_3(M, 0)$. Second, the slope at this average total input should yield maximum capacity according to the instantaneous inhibition model (Fig. 4.20 C and Section 4.3.4), i.e., it should take a value of at least c . This requirement can be met by appropriately choosing the coupling factor $w_{IE} K$. The blue region in Fig. 4.22 illustrates the outcome of such an optimization at $M = 880$ with an effective slope of $1.62 c$, against $M = 1,300$ and slope of $1.33 c$ of the red region. The blue region of stable retrieval is almost flat in M - θ space, thus also increasing the range of pattern sizes for which retrieval is possible at fixed θ . To the left of the region of stable retrieval, we observe a substantial region of transient sequences. Such large regions of transient retrieval only occur for slopes larger than c (not shown). The minimum pattern size M_{opt} of stable retrieval, however, does not decrease further for slopes above c (as in Fig. 4.20 C).

Inhibitory neurons need not sample well the excitatory population (c_{EI} can be small because w_{EI} can be large); if there are enough presynaptic partners, because of $N \gg 1$, interneurons will obtain anyways a good average. Regarding c_{IE} , it really acts only in combination with K : it is possible either to have only one inhibitory neuron projecting to every excitatory or more of them that target random groups, as long as excitatory neurons receive as a population the inhibition brake.

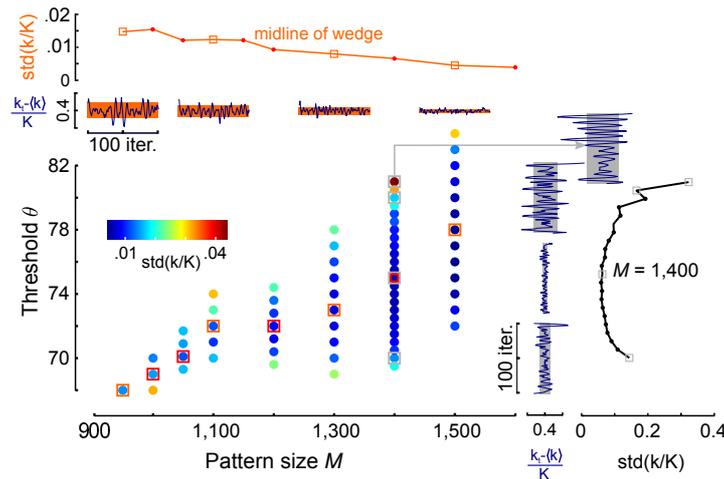


Figure 4.23. Inhibition is Tuned Up at the Edge of Stability. Transition to instability is marked by increase in amplitude of gamma-like oscillations. At lower left colored discs mark combinations of M, θ for which numerical simulations revealed stable retrieval. The standard deviation over time of the inhibitory activity k_t (normalized to K) is represented by the color code as indicated. At top, oscillation amplitudes (measured as $\text{std}(k/K)$) are shown for M, θ along the midline of the wedge. Examples of k_t/K are overlaid on orange bars of height $\text{std}(k/K)$. At right, oscillation amplitudes $\text{std}(k/K)$ are shown in dependence of the threshold only at fixed $M = 1,400$. Examples of k_t/K are overlaid on grey bars of width $\text{std}(k/K)$.

Simulations confirm the shape of the fundamental regimes all active, all silent, and retrieval predicted by the three-dimensional mean-field model. Figure 4.22 (a-f) displays simulation results (m_t, n_t, k_t) for typical situations. Interestingly, all-silent states can also sometimes be observed for low threshold values, where inhibition overcompensates the false alarms and transiently allows for sequence retrieval before the network falls into silence (Fig. 4.22 c). In the retrieval phase, the network typically exhibits oscillatory behavior with periods of about 5 to 10 time steps (Fig. 4.23).

Inhibitory activation oscillates with an amplitude, as quantified by $\text{std}(k_t/K)$, that depends on the location in the phase diagram; this reflects the corrections to the excitatory variables induced by their fast variation when in trajectories far from a fixed point, i.e. when the system is configured with M , θ values at the edge of stability (borders and tip of the wedge).

4.4.3 Recurrent Inhibition

As a coda for our investigation of inhibitory feedback we briefly and unsystematically explore the impact of recurrently coupling the inhibitory pool. Anatomical and electrophysiological investigations in the hippocampus show that inhibitory interneurons project to dendritic shafts of other interneurons. An inhibitory population that is coupled to itself is governed by the dynamical equation (compare Eq. 4.35):

$$\frac{k_{t+1}}{K} = \frac{1}{2} \left[1 + \text{erf} \left(\frac{c_{\text{EI}} w_{\text{EI}} (m_t + n_t) - c_{\text{II}} w_{\text{II}} k_t - \eta}{2^{1/2} \sqrt{w_{\text{EI}}^2 c_{\text{EI}} (1 - c_{\text{EI}}) (m_t + n_t) + w_{\text{II}}^2 c_{\text{II}} (1 - c_{\text{II}}) k_t}} \right) \right]. \quad (4.37)$$

r, c	color	$K/10^3$	$w_{\text{IE}}/10^{-3}$	c_{EI}	w_{EI}	c_{II}	w_{II}	η	note
1,1	blue	5	12	0.01	1	—	—	8.8	Fig. 4.22; cyan: $k_0 = K/2$.
1,2	green	20	4	0.0085	1	0.0003	1	4	light green: $k_0 = K/2$; $M_{\text{opt}} \gtrsim 700$.
1,3	brown	58	0.8	0.5	1	0.5	0.01	50	
2,1	black	58	0.8	0.5	1	0.65	0.007	45	$M_{\text{opt}} \simeq 700$.
2,2	pale blue	20	4	0.001	1	0.0005	1	4	$k_0 = K/2$; $M_{\text{opt}} \gtrsim 700$.
2,3	orange	50	0.8	0.5	1.3	0.5	0.008	600	broken wedge.
3,1	red	80	0.5	0.5	1	0.5	0.006	180	
3,2	purple	64	0.86	0.52	1	0.94	0.0048	44	$M_{\text{opt}} \gtrsim 700$.
3,3	navy	28	1.1	0.4	0.4	0.45	0.002	120	broken wedge.

Table 4.8. Parameters Used for the Exploration of Recurrent Inhibition. Position in the grid of Fig. 4.24 is indicated as row, column (r,c). All networks have in common $N = 10^5$, $c = 0.05$, $c_m = 0.1$, and $c_{\text{IE}} = 1$, and were initialized with $k_0 = T_{\text{Inh}}(M, 0, 0)$ unless noted.

Note that finding a fully self-consistent initialization is harder to find because of the direct dependency of k_{t+1} on k_t . We decided to compare two approaches in Figure 4.24, where we show the regions of stability for different combinations of parameters. One, as in the previous Section, assumes that the starting activity of the inhibitory pool is consistent with the activation of M excitatory hit neurons only, as corresponds to pattern retrieval, and an initially silent inhibitory population, i.e. $k_0/K = T_{\text{Inh}}(M, 0, 0)$. The other consists in setting $k_0 = K/2$, and without recurrent inhibition often leads to disintegration of the whole wedge. However, recurrent inhibition is able to sustain an asymptotic state of replay in an island *at low pattern sizes* (see green and pale blue wedges in Fig. 4.24). Remarkably, there exist parameter combinations for which a network with recurrent inhibition slightly improves on the capacity achieved without it, but finding them requires patient exploration. Armed with the intuitions of Eq. 4.4.3 about the relationships between the three weights w_{EI} , w_{IE} , w_{II} , the three corresponding connectivities, the inhibitory threshold η , the number of inhibitory neurons K and the initialization k_0 , it is possible to reduce M_{opt} by about

10%, as shown by the saturated-color wedge tips and verified by calculating individual trajectories e.g. on the light green and light blue wedges. Table 4.8 shows the parameter values used for the Figure. Based on our sampling of parameter space, which is no substitute for a more systematic optimization approach or proper analytical bounds, it seems that recurrent inhibition can help robustness at the tip and push the stable region to slightly lower M_{opt} . It is however unlikely that it can operate higher gains than e.g. the $M_{\text{opt}} \gtrsim 700$ achieved in some of the examples below.

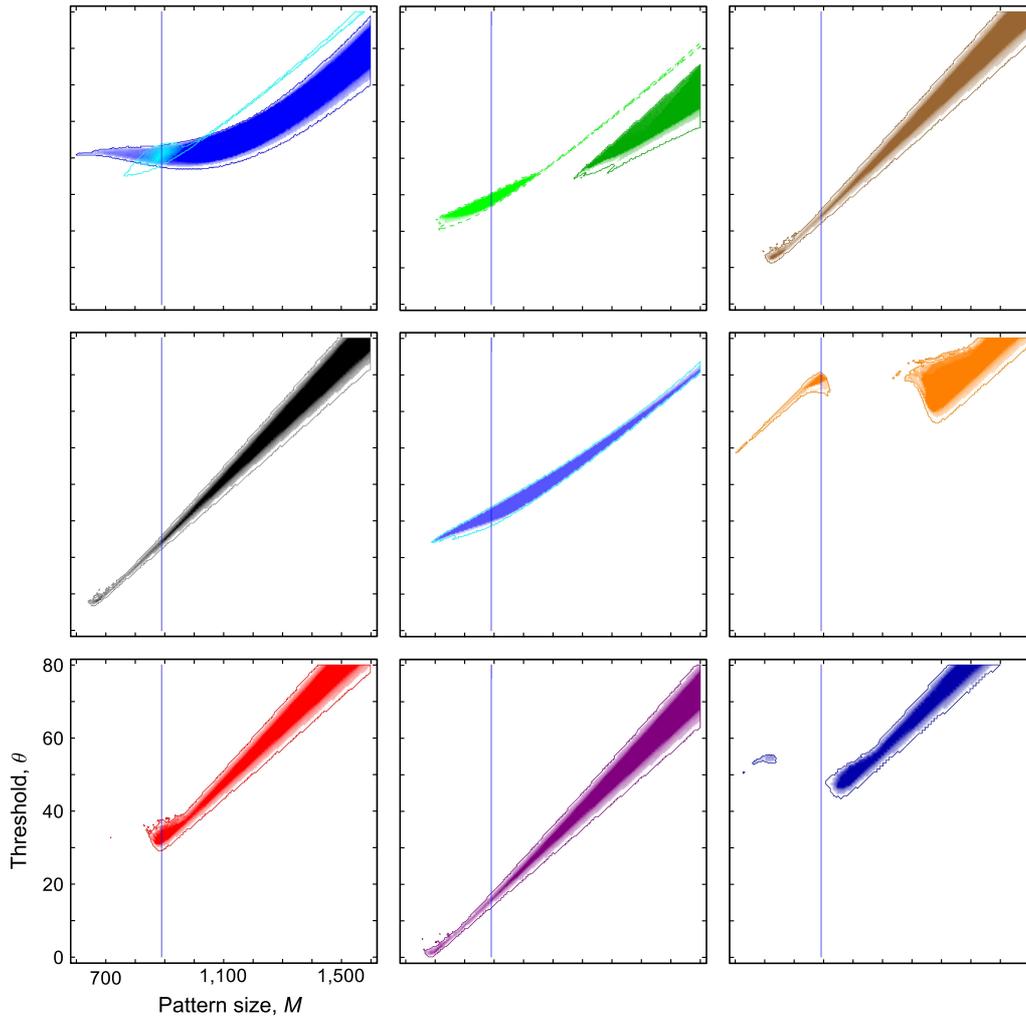


Figure 4.24. Phase Diagram Examples in Presence of Recurrent Inhibition. Nine networks with recurrent inhibition (Table 4.8) are compared to a network where $c_{II}=0$ (blue, top left, identical to blue in Fig. 4.22). Initialization with $k_0=K/2$ is demonstrated in lighter color for the first two networks and in the central example for both colors; all others are initialized with $k_0=KT_{\text{Inh}}(M, 0, 0)$. A blue vertical line marks pattern size $M=890$, the approximate tip of the stable region without recurrent inhibition in the top-left panel.

Part IV

Discussion

Chapter 5

Ripples, Assemblies, and Sequences

In this final Chapter we collect our results and discuss them in the light of the present state of knowledge. We also take the liberty to suggest future lines of work. Figure 5.1 presents the unifying view of the two strands of this Thesis: the descriptive of signal analysis of noisy electrophysiological data, and the prescriptive of a simplified mathematical network model. These are bound together by the concept of neuronal assemblies, that make up the memory traces and the heartbeat of the ripple.

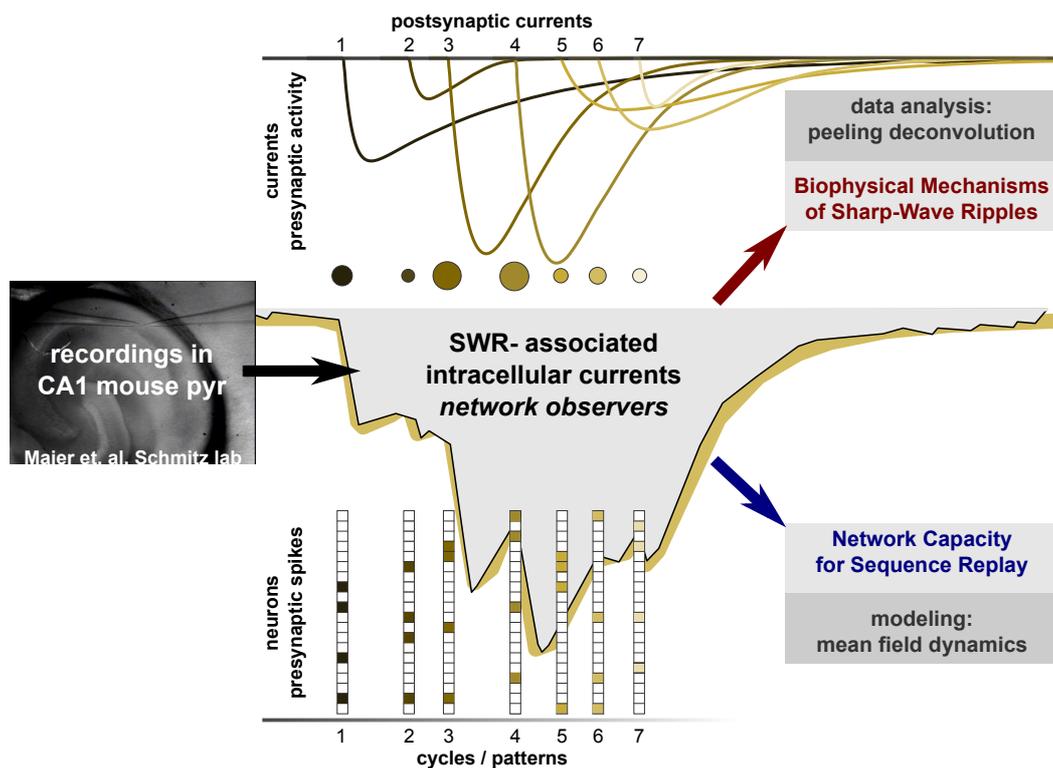


Figure 5.1. Summary of Research Presented in this Thesis. We used as a starting point currents measured in voltage clamp at CA1 pyramidal neurons (center, black line enclosing grey area) concurrent to sharp-wave ripples in the LFP (not shown) as recorded from a mouse slice model (IRDIC image at left, cell 5 is patched). Analysis of the composition of the currents (top) revealed a rhythmic succession of EPSCs (colored curves; discs below sized by amplitude) and IPSCs (not shown). The notable modulation of pyramidal cell activity in a 5 ms rhythm together with literature reports of replay of behavioral sequences during ripples inspired a theoretical model of capacity for sequence replay based on coactivating assemblies at each ripple cycle (bottom, cycle-discrete schematic raster plot).

In the first part (upper panel; dark red) we studied currents incoming to CA1 pyramidal cells during SWR *in vitro*. We used a novel peeling reconstruction approach in combination with more traditional signal analysis techniques to obtain a quantitative characterization of the contributing excitatory and inhibitory PSCs, including their rhythmicity, cycle-dependent amplitude and relative timing.

In the second part (lower panel; dark blue), inspired by the above results on ripple locking of currents and numerous reports that SWR support the retrieval of behavioural spatial sequences, we built a mean-field model of a hippocampal recurrent network with inhibition. With it we could obtain the optimal number of neurons active in an assembly and what their threshold should be so as to concurrently optimize capacity and robustness for the replay of sequences.

5.1 Biophysical Mechanisms of Sharp-Wave Ripples

Hippocampal SWR have been researched predominantly in rodents, *in vivo* (Buzsáki, 1986; Buzsáki et al., 1992; Ylinen et al., 1995; Chrobak and Buzsáki, 1996; Csicsvari et al., 2000; Buzsáki et al., 2003; Klausberger et al., 2003), *in vitro* (Maier et al., 2003; Nimmrich et al., 2005; Wu et al., 2005a; Behrens et al., 2005; Both et al., 2008; Ellender et al., 2010; Böhner et al., 2011), and *in silico* (Traub et al., 1999; Memmesheimer, 2010; Taxidis et al., 2011; Stacey et al., 2011). Yet, their generation and extinction mechanisms remain unclear.

Here, we combined intra and extracellular observations from *in vivo* head-fixed mice (Crochet and Petersen, 2006; Margrie et al., 2002; Poulet and Petersen, 2008) with an in-depth analysis of intra and extracellular recordings from an *in vitro* mouse model (Maier et al., 2009) to investigate synaptic input onto CA1 pyramidal cells during hippocampal ripples. We then studied the functional constraints that befall on networks using SWR as a possible mechanism for the retrieval of stored sequences of memories (see next Section, 5.2, for discussion of our theoretical findings).

5.1.1 Executive Summary

We present here in short form the methods that we have employed, the results that we attained with them and some of the many larger questions they inspire and, as yet, leave unsolved.

Techniques Studies addressing the cellular physiology of ripples have focused on multiunit recordings in freely moving animals (Ylinen et al., 1995, Csicsvari et al., 2000, Buzsáki et al., 2003) and juxtacellular staining in anesthetized rats (Klausberger et al., 2003). However, in contrast to extracellularly measured activity, synaptic input to participating neurons is revealing of the *effective action* of the network on those cells, i.e. it allows evaluation of what fraction of the observed extracellular spiking activity arrives to which cell, and when. Despite this fact, synaptic inputs during ripples have not been studied systematically.

Enough temporal precision for the quantitative analysis of fast synaptic inputs can only be achieved by recording synaptic *currents*; the decay of the potential is too slow for that purpose. The use of the voltage clamp configuration, however, entails limitations that are particularly severe for recordings *in vivo*. Hence, after establishing the existence of SWR-associated fast currents with a strong, ripple-modulated excitatory component *in vivo*, we proceeded to a quantitative evaluation from *in vitro* data. Dual LFP-LFP, LFP-whole-cell, and LFP-dual whole-cell recordings were used along with pharmacological tools to characterize synaptic inputs to CA1 principal cells. Electrophysiological time series were addressed with an array of classical techniques for the analysis of oscillations (event-triggered averages, vector strength, spectral power) added to custom time-domain techniques to unweave the individual PSCs from the SWR-associated compound post-synaptic currents.

Results We found that cPSCs are phase-locked to ripples and coherent among principal neurons. These currents contained modulated excitatory components: First, they could be observed at a membrane voltage with low driving force for Cl^- , i.e., a recording condition for which GABA_A receptor-mediated synaptic inhibition is negligible. Further, at that holding potential, their kinetics were consistent with those of non-SWR-associated EPSCs. Phasic PSCs in ripples persisted during ripples after intracellular block of Cl^- -mediated inhibition and maintained coherence with the LFP oscillation. The ripple-locked excitatory current contribution was shown to be coherent as well spatially, over almost the whole CA1. Extending the approach of pharmacologically blocking inhibition at the single-cell level to current clamp recordings further revealed that ripple-locked excitation can regulate spike timing. Experiments in CA1 minislices devoid of the CA3 and subicular subfields demonstrated that, *in vitro*, ripple-coherent excitatory cPSCs can emerge locally within the CA1 network. Incoming PSCs exhibited a strongly cycle-dependent amplitude over the course of the ripple, but rather cycle-independent kinetics comparable to spontaneous events. Finally, inhibitory currents synchronized to excitatory currents, recovering over the course of the ripple a lag of about 1.5 ms to end in perfect phase alignment.

Open questions Despite the above results, the major questions surrounding the physiology and function of SWR remain open. Among them: How are SWR triggered? Do they end spontaneously or are they actively terminated? Do all ripple cycles have a similar purpose and capability for processing or is there an *early* and a *late ripple*? What is the functional role of excitation vs. inhibition during ripples? Are there stable assemblies of coactive neurons? If so, are different assemblies expressed in the successive ripples? or are they expressed even with the temporal precision of single ripple cycles? What can coordinate firing of pyramidal cells with ripple cycle precision? Is it synaptic? What is the minimal amount of hippocampal tissue that can self-organize to produce a ripple? And from the point of view of methods: What are the specific strengths of LFP vs. intracellular current recordings for monitoring network activity, and how can they be combined? What is the real extent of the somatic clamp? What conclusions of *in vitro* electrophysiology are translatable to the *in vivo* situation? In particular, what is the transverse extent of a local circuit in CA1? How does the slicing trauma differentially affect excitatory and inhibitory synapses? The list goes on and on. In the following Sections we touch upon some of these issues in the light of our results and the current state of research.

5.1.2 Excitation and Inhibition

Here we discuss the evidence for the existence of both phasic inhibition and excitation and our results on their phase relationships.

Phasic inhibition Our results challenge the prevailing view that sharp wave-associated ripples are shaped by phasic synaptic inhibition *alone*. This view is based on several observations about ripples *in vivo*:

1. At the intracellular level, using sharp microelectrode recordings on CA1 pyramidal cells in anesthetized rats, Ylinen et al. (1995) varied the pipette Cl^- concentration and showed that ripple-associated postsynaptic potentials displayed phase shifts as expected from inhibitory PSPs.
2. In the same study, a current source density (CSD) profile was produced from an extracellular shank electrode that showed ripple modulation in the pyramidal layer. Since perisomatic afference to CA1 pyramidal cells is mostly inhibitory and the LFP patterns are widely considered to be governed by local synaptic activity, the observed extracellular ripple was suggested to be caused by the inhibitory inputs to the soma. Hindering dissection of the actual responsible, the observed current source could be indeed active, due to inhibitory-driven efflux, or passive, due to the compensation of a concurrent active sink at the dendrites caused by excitatory-driven ion influx there.
3. Third, extracellular recordings *in vivo* revealed that soma-targeting interneurons increase their discharge rate during ripples and fire rhythmically with the network oscillation (Csicsvari et al., 1999b). The increase in rate is still only about half of that for pyramidal neurons (factor 3.8 vs. 8.6), but a larger proportion of the interneuronal population is recruited during ripples (about 60% vs. 30%, Csicsvari et al., 2000). Interneurons start firing earlier and cease later in the ripple (Csicsvari et al., 1999b). More importantly, single putative interneurons fire ripple-locked at high frequencies, whereas the paucity of pyramidal firing makes it necessary to conduct any locking analysis at the population level, or on extremely long recordings. These studies also showed, however, that pyramidal rates couple better to ripple amplitude (in SD of baseline activity) than interneuronal rates do (see in particular Csicsvari et al., 1999a, Fig. 3).
4. Ripple-locked firing of soma-targeting basket cells has been confirmed more recently by juxtacellular recordings with *post hoc* morphological reconstructions in anaesthetized rats (Klausberger et al., 2003).

Phasic excitation Phasic excitation is to be expected from the abundant reports using *in vivo* multielectrode recordings, which demonstrate that pyramidal cells, as a population, fire phase-locked to ripples. Our study adds that phasic excitatory inputs at ~ 200 Hz are reliably expressed and prominent during ripples as seen from other pyramidal cells. Further, they are effective in regulating spike timing. This is demonstrated by the ripple-locked spiking of cells whose inhibitory inputs have been blocked pharmacologically with DNDS. Phasic involvement of the majority cell population in the hippocampus ($>90\%$; Buhl and Whittington, 2006) does not come as a surprise. Indeed, while oscillations are theoretically possible in an inhibitory-only population

when driven by a stationary excitatory forcing, it is likely that in a mixed population with reciprocal connections both types of cells entrain each other into the rhythm (Brunel, 2000; Isaacson and Scanziani, 2011).

Preferred phase of excitatory and inhibitory inputs It would be tempting to compare our results on preferred current phases with the *in vivo* reports of phase locking of individual pyramids (Buzsáki et al., 2003; Csicsvari et al., 1999b; O’Neill et al., 2006) and interneurons (Klausberger and Somogyi, 2008). Unfortunately, the times are incomparable without additional knowledge: different delay lines are involved in carrying an emitted spike to an extracellular electrode and to the soma of a voltage-clamped postsynaptic cell. Those delays include the axonal propagation delay, the synaptic transmission delay and the electrotonic ramp-up of the current inside the postsynaptic neuron all the way to the soma. Additionally, the phases of steep slopes used in our locking analyses are delayed by ca. 1 ms with respect to the fitted onsets. These considerations prevent direct comparisons. At most, phase differences of excitation and inhibition could be compared between our analysis and the extracellular electrode-based literature, but this would still require assuming that the delays are similar for both types of input.

Interplay of excitation and inhibition during the ripple We compared our results at the reversal of inhibition (-66 mV) to those at higher holding potentials (-45 mV), reaching up to the reversal of inhibition (-6 mV; Cs-based solution). We found that, as potentials were increased, the polarity of the currents flipped from inward to outward in a cell-dependent manner. Those cells that showed a partial polarity change at an intermediate holding potential systematically exhibited an inward-signed overall current first, followed by an outward current^{5.1} with the somewhat longer tail-off also seen in fully inhibitory traces. This is reminiscent of the gradual tilt towards inhibition observed in sensory cortices in response to impulse-like stimuli (Isaacson and Scanziani, 2011).

To refine this initial observation we improved on our relatively assumption-free slope-based technique to detect PSCs embedded in cPSCs by developing a time-domain iterative reconstruction method with two stages. Detection, first, is performed by deconvolution. Peeling, second, exploits the information garnered from extrapolating the uncorrupted section of the first incoming PSC to subtract it altogether from the remaining trace. Using this technique, we were able to decompose the vast majority of the monophasic cPSCs (excitatory or inhibitory) into individual PSCs, and collect statistics on their amplitude, kinetics, and timing by ripple cycle. The large amplitudes of PSCs observed around the SWR peak could be explained either by superlinear time integration (whereby single presynaptic partners see their input amplified when it occurs in fast sequence) or by the well-orchestrated coactivation of a presynaptic

5.1. This is somewhat similar to the “biphasic” events reported by Ellender et al. (2010) in their Figure 3 albeit with the caveat that we study mouse, CA1 and currents that follow a consistent pattern for each cell whereas these authors recorded in CA3 of the rat the membrane potential of one pyramidal cell expressing EPSPs, IPSPs or mixed PSPs in a single recording. They also observed a distance-dependent ratio of excitation to inhibition (their Fig. 3 H), with excitation predominant for sharp-waves originated close to the patched cell, which would be in line with our observations. In general results like those of Ellender et al. seem to beg for a study of the synaptic composition of sharp-wave ripples that takes into account the coupling between location in the slice and time along the ripple that is induced by the propagation of the ripple (Csicsvari et al., 2000; our own observations indicating lags of at most one ripple cycle between distant locations in mouse CA1).

assembly. According to the assembly hypothesis, virtually every pyramidal cell would be capable of sensing the activation of any assembly, but only selected ones belonging in the sequential schedule of activation would be driven strongly enough to surmount the shunt inherent to the SWR-induced high-conductance, and fire.

The kinetics of the in-ripple and spontaneous PSCs were largely compatible and homogeneous across cells, which supports the view that cPSCs are built from PSCs drawn from the network-wide pool of afference.

Finally, ripple-associated excitatory and inhibitory currents express an exquisite temporal precision and converge in phase, as inhibitory phasic input reduces its initial lag with respect to excitation over the course of ripples. This may be interpreted as a quenching mechanism. Indeed, with recurrent excitation, synaptic or otherwise, there exists the risk that activity avalanches out of control into an epileptic-like state. Interneurons could switch from emphasizing the oscillation at the beginning by firing in phase opposition, to bridling it by discharging in phase alignment towards the end. This phenomenon could help explain the partial polarity reversal noted above for some cells at an intermediate holding potential. The diversity of interneuronal ripple phase-tuning according to cell phenotype (Klausberger and Somogyi, 2008) could be instrumental in achieving this staggeringly precise phase drift of about a quarter of a cycle over a few cycles. The dampening of the ripple has been so far thought to occur passively as the depolarization from CA3 wears off. Since our data shows that ripples are also transient in the minislice, i.e., in absence of CA3 input altogether, the active interplay and phase alignment of excitation and inhibition appears as an alternative or complement to explain their limited duration.

Note that the short delays between excitation and inhibition are in principle compatible with both feedforward (excitatory afference incoming to both pyramids and interneurons) and feedback inhibitory circuits (local interneurons fed only by local pyramidal cells). However, in a very simplistic model where axonal delays in the local circuit would be of the same order for excitatory and inhibitory cells one would expect feedforward inhibition to be faster by one delay with respect to feedback inhibition. This runs counter to the observed initial delay of inhibition. Notwithstanding and since the real situation is much more complex, it cannot be excluded that the network architecture provides yet another route to explain the observed synchronization.

5.1.3 Origin of Ripple-Coherent Spikes

Two possible, nonexclusive sources of ripple-coherent EPSCs are conceivable.

CA3 coupling First, they could represent input from synaptically coupled CA3 pyramidal neurons. Indeed, phase-coupling of CA3 pyramidal cell spikes with CA1 field ripples has been demonstrated in vitro (Both et al., 2008), although in vivo locking is limited to slow ripples (<140 Hz; Csicsvari et al., 1999a Fig. 4; Sullivan et al., 2011, Fig. 8). We found ripple-modulated and locked EPSCs also in CA1 minislices, thus ruling out this possibility as the *only* origin of ripple-locked EPSCs. Although observation of SWR in minislices is not new (Maier et al., 2003; Nimrich et al., 2005) our analysis shows that the timing and polarity of ripple-associated currents in minislices resemble those in full slices, thus supporting the idea that minislice SWRs display the full phenotypic range of the original phenomenon.

A fine-grained approach to test the influence of CA3 output on CA1 ripples was pioneered by Nakashiba et al. (2008). They reported a tetanus-toxin-based triple transgenic mouse that allows reversible inhibition of CA3 pyramidal cells by a two week withdrawal of doxycycline from the diet. Although comparably slow, this is more targeted than pharmacological inhibition or tissue lesions. Importantly, it affects the whole cell and not just one receptor subtype as was the case with deletion of the NMDAR gene in earlier work from the same lab (Tsien et al., 1996). Following up, Nakashiba et al. (2009) showed that CA1 ripples persisted after silencing CA3 input. These surviving “mutant” ripples displayed lower oscillation frequency on average (~ 120 Hz vs. 148 Hz), but during those of the mutant ripples that fell into the faster ripple frequency band, the experience-associated correlation of CA1 cell pair firings was maintained (their Fig. 3F). It is not clear, though, which signal was the immediate trigger for the CA1 network to generate mutant ripples. Feasible inputs might arise from entorhinal cortex or thalamus (Nakashiba et al., 2009). Collectively, though Schaffer collateral input onto CA1 may be obligatory for the transfer of information involved in memory consolidation, transmission from CA3 to CA1 does not seem to be required for the occurrence of ripple oscillations in CA1 (see also Buzsáki et al., 1992).

Local CA1 origin A second framework to explain the origin of ripples posits their purely local emergence, as a result of recurrent synaptic input alone. Through paired recordings, Deuchars and Thomson (1996) found that recurrent excitatory connectivity among CA1 pyramidal cells is mediated by synapses between the axon and the basal dendrite. Their measurements in rat slices point at a frequency of one in 100 random pairs, considerably lower than 1 in 16 for the CA3-CA1 pathway but within a factor two of the CA3 recurrence (1 in 50-100 after Miles and Wong, 1986; guinea pig slices). Complementary anatomy work on one pair showed that in spite of the cells being connected at only two sites, both on third-order basal dendrites (with all the implied synaptic filtering), the postsynaptic EPSP was large (1.5 mV amplitude). Of note, mouse tissue has a larger neuronal packing density, which is perhaps a reason why it seems more difficult to establish reliable models of SWR in vitro in the rat: Wu et al. (2005a) observe SWRs in only 40% of their thick CA3 slices and propagation towards CA1 is rarer (Liset Menéndez de la Prida, pers. comm.). The low firing rates of pyramidal cells result in insufficient statistics for correlational analysis with present tetrode technology, even in vivo (where rates are higher). This makes it difficult to test the functional utilization of the recurrent connections unveiled by ultrastructural anatomy combined with in vitro paired recordings. Our own results, especially our observation of ripple-coherent EPSCs in CA1 minislices are consistent with recurrent CA1 connectivity.

Lack of concomitant CA3 input could underlie the lower SWR occurrence in minislices. According to this idea the local CA1 network would be lifted less often to a transient self-sustaining oscillation (Stacey et al., 2011) from the synaptic noise due to CA1 alone. The reduction in amplitude could be connected with decreased CA3-induced inhibitory activity, but regrettably we lacked the experimental data (CA1 cells in the minislice held at the reversal potential of excitation) to test this hypothesis here.

Why are spikes so elusive? Perhaps it should not be surprising that there are so few spikes during ripples. SWR are, after a fashion, the high-conductance states of the hippocampus (Destexhe et al., 2003), featuring 150% of the baseline conductance (Böhner et al., 2011). The major effect of inhibitory inputs may be indeed of a shunting nature, making it difficult for excitation to elevate the membrane potential closer to threshold.

Our cells held in voltage clamp reveal regular presynaptic spiking activity as PSCs. Yet the overall network activity is very sparse, both from excitation and inhibition. This is especially visible when in current clamp, where induced depolarization is needed to observe even a few spikes in a fraction of the cells during ripples. This is not exclusive of the mouse slice; for example Ellender et al. (2010) observe rates of 0.3 Hz across participating cells in rat slices (13% of all cells). We conceive two ways to account for this fact. One is that the code is very sparse and thus the probability of choosing an assembly-participating cell is correspondingly small. This would also result in scarce postsynaptic currents. Proponents of the axonal coupling hypothesis (see below) would perhaps argue that spikes are generated in the axon, and do not necessarily propagate back to the soma, especially if the axon branches off a dendrite and not directly from the soma, or if the spike is generated in a thin axon collateral (Böhner et al., 2011).

5.1.4 Emergence of the SWR Phenomenon

We enquire what phenomena local to the slice may trigger the irregularly occurring SWR, a spatially extended pattern that combines a sharp wave LFP depolarization prominent in stratum radiatum and a ripple principally visible in stratum pyramidale.

5.1.4.1 Emergence of Sharp Waves

Cellular processes involved in triggering sharp waves are still subject to investigations. Recently, it has been proposed that sharp waves in CA3 may be induced by rebound depolarization following strong inhibitory activity (Ellender et al., 2010). Regardless of the mechanism by which the depolarizing wave is generated in CA3, sharp wave-associated excitation arriving from CA3 is commonly believed to trigger sharp waves in CA1 (Buzsáki, 1986), which secondarily give rise to ripples (see Taxidis et al., 2011 for a model based on this assumption).

5.1.4.2 Emergence of Ripples

The mechanisms responsible for the generation and maintenance of ripples are also a matter of debate. We have observed in several measures of locking that inhibitory and excitatory current inputs are similarly well coordinated. It is difficult to estimate from the available data to what extent the rhythm is set up predominantly from CA3 tonic excitation engaging a highly disciplined oscillation of a CA1 interneuronal network, or whether the oscillations lean rather on the local inhibition-excitation-inhibition feedback loop (Brunel, 2000). Computational models of gamma oscillations that synchronize based on inhibitory interactions have trouble recreating oscillations beyond the 100 Hz barrier (but see Stacey et al., 2011 for an attempt based on increasing the noise). That is why modelers have turned to two other possible mechanistic pictures.

Synaptic coupling First, ripples may reflect the synchronous discharge of pyramidal cells during the replay of memory sequences. This is the idea implicit in Leibold and Kempter (2006) and that we have further developed in this Thesis to incorporate the effect of inhibition. As outlined above and in the description of our theoretical model, an active assembly becomes visible to every cell postsynaptic to an assembly member, but only a subset of those cells will be excited by presynaptic partners enough so as to fire. Interestingly, in the epilepsy model provided by injection of kainic acid in the rat hippocampus, pyramidal cells alone are capable of establishing a fast rhythm (>200 Hz) due to in-phase firing at the maximal single-cell burst frequency or a very fast rhythm (>400 Hz; “emergent”) that results from out-of-phase cellular bursting with dephasing in the range of one oscillation cycle (Ibarz et al., 2010). This study combining modelling and juxtacellular recordings *in vivo* exemplifies how a fast network rhythm can emerge from desynchronized unit bursts at frequencies considerably below that of the resulting LFP oscillation.

Electrical coupling by axonal gap junctions Alternatively, or additionally to recurrent excitatory synaptic coupling alone, electrical coupling of CA1 principal cell axons may generate oscillations in the ripple frequency range (Traub et al., 1999; Traub and Bibbig, 2000). Indeed, gap-uncoupling agents disrupt the ripple *in vivo* (halothane; Ylinen et al., 1995) as well as *in vitro* (octanol, carbenoxolone Maier et al., 2003), while sparing the sharp-wave. However, their specificity is often questioned.

Gap junctions in the hippocampus proper have been ultrastructurally confirmed so far only in dendrodendritic appositions, indirectly between calretinin immunopositive (Gulyás et al., 1996), and directly between parvalbumin positive interneurons (Fukuda and Kosaka, 2000) and other GABAergic cells. Dendritic coupling suffers from low-pass filtering, which can hinder fast rhythms (Bennett and Zukin, 2004). The evidence for pyramidal axo-axonal coupling is based on pharmacology, physiology and dye coupling (Schmitz et al., 2001), but to our knowledge only one study found so far pyramidal-pyramidal connections in electron micrographs, and they were not between axons (Mercer et al., 2007). Anatomical evidence is hard to find, perhaps because, as modelling work shows (Traub et al., 1999), hypothetical gap junctions would be efficient in producing fast rhythms already at a very low density of 1 to 3 per axon.

The physiological evidence for axonal coupling is mainly supplied by the transfer of spikes onto spikelets, which is believed to be a consequence of electrical coupling between axons of cortical pyramidal neurons (Draguhn et al., 1998; Schmitz et al., 2001; Mercer et al., 2007; Wang et al., 2010). Along this line, recent work has demonstrated bursts of spikelets in the hippocampus during behavior *in vivo* (Harvey et al., 2009) and even at ripple frequency (Epsztein et al., 2010). The observed ripple-locked EPSCs could thus correspond to rhythmic output of a gap junction-coupled network of CA1 principal cells. This promising indirect evidence from *in vivo* studies contrasts with the results of an *in vitro* study (Both et al., 2008). The authors reported that antidromic rhythmic or tonic stimulation of CA1 pyramidal cell axons did not elicit SWRs and that SWRs survived the removal of distal axons in the alveus.

What role do gap junctions play in connection to the plasticity phenomena that are presumed to underlie the consolidation properties of SWR?. Interestingly, modulation of gap junctional function has been reported at different timescales, dependent on voltage, pH, and Ca^{2+} intracellular concentration (Neyton and Trautmann, 1986). In the brain, junctional conductance between inhibitory neurons of the rat's thalamic reticular nucleus is significantly reduced for hours upon stimulation of their metabotropic glutamate receptors (Landisman and Connors, 2005). Also pertinent to plasticity in electrical synapses, Barnes et al. (1987) compared young rats with senescent rats and found an increase in the electrotonic coupling of pyramidal cells in CA1 as measured from dye coupling (dendritic, somatic or dendrosomatic). Besides identifying a possible homeostatic role of gap junctions in keeping excitability in the face of age-induced loss of afferents, Barnes et al. speculatively suggest that the establishment of gap junctions may reflect a Hebbian mechanism for associative storage, contingent upon coactivation. The increase of electrical coupling over life would then reflect acquired memory. Joining the authors in their speculation, we may add that indeed a symmetric (O'Neill et al., 2006) and extremely fast connection would constitute an excellent substrate for well-timed, reliable assemblies. A suggestive question in this framework is whether individual cells or systemic processes could disable the gap-junctional networks opportunistically by modulatory mechanisms to allow for the encoding of new memories.

To make progress in the falsification of the axo-axonal network hypothesis, it would be necessary to improve on the specificity of gap junction blockers (our own analysis with carbenoxolone is inconclusive because of its poor specificity, Maier et al., 2011) or gain experimental access to the axon-specific membrane dynamics, perhaps with fast voltage-sensitive dyes (Popovic et al., 2011; Canepari et al., 2011). A third vector of advance may be provided by intracellular manipulations in vitro that alter the junctional conductance (Neyton and Trautmann, 1986); Barnes et al. (1987), for example, used Ca^{2+} loading successfully to block dendritic coupling.

In summary, a network of pyramidal cells organized by axo-axonal coupling via gap junctions, a recurrent synaptic excitatory network, or a closely knit inhibitory subnetwork connected both through chemical and electrical synapses (Amitai et al., 2002) all possess the ability under the right conditions to generate fast oscillations of variable spatial recruitment capability (Stacey et al., 2011).

5.1.5 Discussion of Methods

Current research in neuroscience devotes extraordinary quantitative attention to methodological questions. Because the brain is so complex and multifarious, almost invariably investigations must flank the question under study from several methodological angles: anatomical, histological, electrophysiological and computational to name but a few. As research is conditioned by the scope and validity of the methods employed; we enumerate below the main methodological issues connected to our investigations.

Reliability of voltage clamp and dendritic levels of afference Voltage clamp of electrotonically extended neurons over their whole dendritic tree is very uncertain (Williams and Mitchell, 2008). Our experiments are nevertheless relevant because voltage clamp is likely to be good at controlling the perisomatic inputs. It is precisely those inputs that have been suggested to be responsible for the ripple in the pyramidal cell layer (Ylinen et al., 1995); our holding the cell soma at the reversal of inhibition is expected to make the driving force for them negligible. Thus the inward currents we observed were either truly excitatory or, rather unlikely, due to distal inhibitory afference made visible by an insufficient reach of the clamp into the farther membrane.

The fast kinetics of the observed excitatory currents are somewhat surprising in view of the electrotonic filtering imposed by propagation from far-out locations. It may well be that these events emerge from contacts of Schaffer collaterals or other CA1 axons on the basal dendrites (Deuchars and Thomson, 1996; Memmesheimer, 2010).

In vivo vs. in vitro network and slicing collateral damage Slicing causes trauma to the retained tissue: growth factors, potassium and neurotransmitters are released into the extracellular milieu and the background activity of the network ceases. Kirov et al. (1999) studied via electron microscopy the synapse density of *rat* hippocampal slices vs. perfusion-fixed tissue representative of the *in vivo* situation. Slices showed about 50% more spines two hours after the procedure and the number remained stable for at least 11 hours more. The stabilization of this excess seems coupled to return of synaptic activity to the slice one hour post-trauma. Not all spine types are affected by this process equally. The time span of slice recordings is too short for reactive synaptogenesis or sprouting to play a role. In view of these facts, we hypothesize that the network wiring that produces the observed SWR is almost certainly very different from the naive network before slicing. Yet, the new synaptic connections that are established do follow the chemical and geometrical constraints present (that laminar boundaries are respected even by sprouting; see Frotscher et al., 1997), so as to be able to reconstitute a network that expresses oscillations very similar to SWR as observed *in vivo*. We thus see the network *in vitro* as statistically representative of the *in vivo* network (excluding the severed connections running through the section boundary) but do not expect it to conserve the identity of the existing synapses or their strength. To put it simply — no memories of the living animal are to be caught in the slice. As for the situation in the minislice, we ignore what are the possible additional side effects of sectioning, in particular, Schaffer collaterals. Modeling has shown that oscillations can emerge in a recurrent network from a spontaneously firing subpopulation. If sectioning of Schaffer collaterals steadily liberates glutamate, SWRs would be facilitated by the additional excitability even in absence of CA3 synaptic input. If the widely held hypothesis that the CA3 excitatory tonic volley is necessary for the generation of CA1 ripples is to be made compatible with the minislice SWR a better understanding of such potential effects is needed. Another source of worry is whether the slice is thick enough to sufficiently represent the original network. The common features of the *in vitro* SWR (regional initiation and propagation, pharmacological sensitivity and intracellular correlates) are retained. In addition, the more compact structure of the mouse tissue with respect to the rat, where extra thick slices (1 mm vs. typical 0.4 mm) were deemed of advantage for studies of oscillations depen-

dent on CA3 recurrent axons (Wu et al., 2005a) further legitimizes the slice model employed in this work. Schaffer collaterals are rather expansive across slice sectioning planes, i.e. along the longitudinal axis in CA (Johnston and Amaral, 2004). In this light, it is perhaps not surprising that the phenomenon survives deafferentation of CA1 — the input from CA3 could have been not needed for its appearance in the first place (if it had been, there are chances that SWR would not be visible in the slice at all).

Probing the network with patched cells vs. LFP electrodes This work has relied on combined LFP and intracellular signals to characterize collective network activity. Both are averages of neuronal activity, and both are biased and ill-specified, in different ways. In other words, they (mostly) represent weighted averages of synaptic activity, with different sets of weights that are unknown and vary in space and with time. Of note, weights of contributions to both LFP and currents can be positive and negative, i.e. contributions can mutually cancel. This renders the inverse problem of obtaining the synaptic activity from the measured LFP or patch signal an essentially ill-posed one, with multiple formal solutions. The only way forward is to introduce additional conditions, for example demanding a laminar structure for the study of LFPs (Makarova, 2011) or establishing a voltage clamp that minimizes currents of specific sign, which was our approach.

We assigned a somewhat preponderant role to the LFP as the reference for the SWR oscillation, instead of using our quite regular, but more variably shaped, cPSCs. Besides the historical reasons (hippocampal rhythms are defined by virtue of their LFP electrographic signature), there is the fact that patched cells show cell class-specific afference patterns and even within the same class the measured currents can differ widely: thus their sampling of the network is neither universal nor predictable.

In spite of its canonical role in establishing the network oscillatory state, LFPs present a number of challenges for interpretation. CA1 being a laminated structure, the LFP varies spatially in a more predictable fashion than currents across cells or even across single cell compartments (Spruston, 2008), but its range depends on the extent of synaptic correlations (Lindén et al., 2011) and what frequency band is examined. The LFP also varies across layers along the axis from dendrite to axon and depends on possible anisotropies of the extracellular space as well as on ionic concentrations that may vary with behavioral state. Finally, it may capture far-field effects and it is certainly influenced by overlapping spikes; precisely the ones that are customarily referenced to it in unit-to-network oscillatory synchronization studies. Notwithstanding these defects of LFPs as a tool to characterize network activity, they have been successfully employed *in vitro* to identify coactive cell groups during SWRs: Reichinnek et al. (2010) clustered SWRs employing a machine learning approach and established stable categories with well-defined associated sets of coactivated cells. Currents provide a less geometrical and rather more topological portrait of network activity, i.e. they emphasize connectivity onto the patched cell vs. spatial adjacency to the measuring electrode. Weighted averages of postsynaptic currents at least fix the postsynaptic partner, which offers a considerable interpretive advantage: spikes are all converging onto a postsynaptic cell of known phenotype, although they cannot at present be individuated according to source through a “PSC sorting” procedure in the spirit of spike sorting. It is indeed highly attractive to sample the network as

it is seen by the cells that have to decide their firing on the presynaptic activity. In comparison, extracellular measurements and sophisticated analysis allow to ascribe spikes to source units, but cannot tell about their targets. Current measurements at the soma have limitations due to the poor spatial range of the voltage clamp technique for extended neurons (Williams and Mitchell, 2008). This downside currently renders them unusable for quantitative work on *in vivo* data and conditions their utility *in vitro* to the availability of knowledge about the spatial afference pattern of presynaptic neurons, a condition that is probably most likely to be met in hippocampal CA1 than anywhere else (Andersen et al., 2007).

Simultaneous intracellular voltage clamp recordings with extracellular multielectrodes may eventually bridge the gap between these two representations of network activity by e.g. enabling spike-triggered averages of currents, shedding light on the number and type of presynaptic partners of a given cell. At the same time, this may provide insight as to exactly how much information can be gleaned about sites of afference and, eventually, presynaptic delays from the shape of the currents as measured postsynaptically. Our peeling reconstruction algorithm may help realize these more sophisticated analyses, by disentangling signals that overlap considerably in time. This overlap is precisely the electrographic signature of assemblies in action, the concept that has guided our investigation in this Thesis. Peeling reconstruction may also be generally applied to other overlapping signals, such as voltage measurements, or Ca^{2+} fluorescence imaging. The reconstruction approach, however, requires a considerable set of assumptions that must be carefully checked in each particular domain of application. In general, we have verified that time-domain methods (steep slopes, reconstruction) have advantages over spectral methods in addressing the characteristics of these high SNR signals, especially when events have to be weighted equally regardless of their amplitude.

5.1.6 Assemblies

Assemblies in action are the implementation of memory traces. They represent also the connection between our work on the mechanistic aspects of SWR as a substrate for sequence memory and our computational investigation of memory capacity.

Atomic span for the expression of assemblies Units fire with exquisite timing during ripples and they provide sequence-like content as evinced by the *in vivo* literature on replay. A particularly pressing question in this context is what is the temporal span for the expression of assemblies: one assembly per whole ripple oscillation, two (as in an association), or one per ripple beat? In this Thesis, both while analyzing data and when modeling, we worked under the hypothesis that SWR, though transient and fast-oscillating, can indeed host one assembly per cycle. Reports of ripple reactivation representing whole trajectories lend credence to our viewpoint, while the most convincing experiments to date linking SWR activity to memory performance (Dupret et al., 2010) would rather hint at a single assembly being expressed per ripple (their Fig. 5b). A clustering study of SWR *in vitro* (Reichinnek et al., 2010) found

multicellular activity patterns to be selective for certain SWR clusters, but lacked resolution and cell numbers to make a statement about single-cycle-locked assemblies, in addition to the inherent limitations common to *in vitro* studies discussed above. Substantiating this debate is much beyond the traction of the data we have presented here, but there is one noteworthy aspect of our results that may become significant in future investigations. The slight lag between excitation and inhibition provides an extremely precise window of opportunity for integration of afferent input, and thus, expression of assemblies. The window is however not uniform, but closes progressively as the ripple advances beyond peak amplitude. Our data seems to indicate that there are a few periods, perhaps just one or two, just before the peak of the ripple, that are particularly suited to intense, well-timed excitation. A pictorial metaphor of this is that of two haircombs sliding one over the other. For a very brief instant a thin window between pins appears before pin matches to gap and vision of the background is occluded. In this picture, the whole ripple would be a sophisticated oscillatory scenario for one or two assemblies to be expressed just before the peak with exquisite precision and considerable power.

How many cells are in an assembly? There are many confounding factors that hinder the quantification of assembly size, but our cycle-dependent amplitude data provides a hint that can orient further investigation: the maximal PSC amplitudes observed are about ten times the average of spontaneous events. Still, since we do not know if the voltage-clamped cells would have spiked (but we expect them not to in a majority of cases, in view of the rarity of spikes in current clamp) this size may be a lower bound on assembly size, if it is to be related to the number of PSCs a listener, non-participating cell sees during the ripple.

One has to discard the possibility that bursts with inter spike intervals below 10 ms from single pyramidal cells have a definitory role in the ripple oscillation, i.e. that it is not a population rhythm. First, bursts as observed in theta phase precession experiments are limited to very few spikes (two or three). Second, in this preparation we observe very few spikes (see above), which makes bursts all the more unlikely. Third, the low input resistance during SWR is detrimental to bursting. And finally, the large PSC amplitudes observed during ripples would be hard to explain only with one presynaptic, bursting partner.

5.1.7 Summary

We demonstrated coherent excitatory currents in CA1 pyramidal neurons during ripples. These oscillation-locked inputs represent synchronous activity of assemblies of principal cells as has been proposed in several *in vivo* and modeling studies (Wilson and McNaughton, 1994; Harris et al., 2003; Leibold and Kempter, 2006; O'Neill et al., 2008; Dragoi and Tonegawa, 2010). Inhibition incoming to the same cells started ripples in phase opposition and ended in phase alignment. This was shown using a novel peeling deconvolution method for the analysis of fast currents that can be applied to other signals of neurophysiological interest. These results may shed light on the mechanistic aspects of SWR generation and concurrent expression of assemblies.

5.2 Theoretical Model of Sequence Replay

Neural network implementations of associative memory have been profusely studied, first only with excitation and dense connectivity (Willshaw et al., 1969), then introducing sparse connectivity (Gibson and Robinson, 1991) and eventually adding global feedback inhibition, e.g. Golomb et al. (1990), Hirase and Recce (1996). However, in real neural networks interneurons do much more than instantaneously adjust the gain by controlling the threshold of excitatory neurons. They also display abundant recurrent connectivity (Amitai et al., 2002; Yuste, 2011), so that it is useful to consider them as a population and study their behaviour in time beyond threshold scaling.

Here, combining cellular simulations with a mean-field model, we studied large, sparse associative memory networks inspired in hippocampus in order to assess sequence replay in the presence of inhibition. We first addressed global instantaneous inhibition and asked what would be the feedback coupling function that optimizes capacity. We next introduced a full inhibitory population, on an equal footing with the excitatory population, and studied how its activity evolved during the replay phenomenon. We finally considered self-coupling in the inhibitory population and explored its functional impact on the capacity and robustness of sequence replay.

5.2.1 Executive Summary

Here are in short form the methods that we used to address the question of optimizing memory capacity for sequences, the results we obtained and some of the further questions stimulated by our investigation.

Techniques We extended an existing mean-field model that derives activation ratios from the probability distributions of synaptic inputs (Gibson and Robinson, 1991; Leibold and Kempster, 2006). We studied how replay evolved with time in mean-field phase space and how it depended on network parameters in the phase diagram from the theory of discrete dynamical systems. Optimal detection theory with the Bayes criterion was used to find the best threshold for excitatory neurons. Cellular simulations were employed throughout as ground truth to check the validity of mean-field findings.

Results Inhibition added to our neural network model for sequence memory indirectly boosted capacity by stabilizing replay modes that were unstable in the excitatory-only network. We showed analytically in the mean field approximation that the best global instantaneous inhibition is given by a threshold that grows linearly with excitatory activity. A population of inhibitory neurons that facilitates replay close to the capacity limit displays oscillatory activity, with increasing amplitudes betraying the onset of instability at high loads. Recurrence in the inhibitory population was found to marginally enhance capacity and robustness.

Open questions Some of following open questions exceed the scope of the present model, as it is difficult to obtain additional understanding of memory networks staying within its assumptions. We wonder how to systematically investigate the effect of the high degree of inhibitory recurrence that seems to be present in biological networks. If replay at high capacity is inherently a transient phenomenon, how are represented extended experiences spanning several replay containers (Davidson et al., 2009)? Is

the frequency relation between excitatory and feedback inhibitory activity dependent on the time-discrete nature of the model? The quantitative parameters of both CA1 and CA3 fit within our general modeling framework, but what so far uncaptured assumption makes biological sequences tend to be expressed rather in CA1?

5.2.2 Replay Benefits From an Adaptive Threshold

Our replay model contains variability derived from the random choice of the patterns and the random morphological connectivity of the neural network. The variability implies that, given a cue assembly, there will be in the target pattern both active neurons that should be inactive and silent ones that should be active. When the number of active neurons required to encode memory is pushed down so as to increase capacity, every neuron is excited close to its threshold: slightly below if it is a nonparticipating neuron and slightly above if it belongs in the target memory. If, due to the randomness, the average activity increases at any given iteration, it will be difficult when the pattern size is ambitiously driven down for extra capacity, to prevent a positive feedback loop from taking over and activating the whole network (conversely for the opposite case leading to a silent network). A threshold proportional to overall activity provides the necessary feedback to keep the network stable at high capacity. But this threshold can only be tuned in a sequence setting, where the previous iteration provides information about the likely level of global activity in the present one. Hence, the temporal continuity of cue-target associations in a sequence enables the choice of a useful threshold, as long as the network is not destabilized in a single step.

Considering each neuron as a binary classifier confronted with a decision task, we employed ROC curves to visualize the need for an adaptive threshold. They showed how an initially near-optimal threshold became progressively unsuitable as replay proceeded.

5.2.3 Static Inhibition and the Optimal Threshold

Instantaneous feedback inhibition (as provided by a hypothetical fast, globally connected interneuron) can stabilize the retrieval of memory sequences and thereby increase both memory capacity and robustness. The optimal instantaneous inhibitory feedback is a roughly linear function of the total network activity. This was found out numerically by Hirase and Recce (1996). Here the same conclusion was reached semianalytically using a probabilistic approach in the framework of optimal detection theory. Additionally, we could identify the desirable coupling coefficient to excitation as a function of pattern size. Static inhibition admits in a first approximation an interpretation as a reduction of the effective excitatory connectivity, which results in a net increase of the resources for plasticity.

5.2.4 Dynamic Inhibition and Oscillations in the Gamma Range

Extending the model to dynamic inhibition, we find that, at the edges of stable replay, inhibition induces strong oscillations. At a rate of one each 5 to 10 ripple oscillations, they roughly correspond to the gamma range (20-40Hz).

Gamma oscillations are ubiquitous in the brain and their synchronization is generally attributed to local inhibitory networks (Wang and Buzsáki, 1996). Several cognitive functions have been related to increased gamma power and coherence, such as sensory integration, attention, and memory (Jutras and Buffalo, 2010). Specifically, gamma coherence between subregions in the hippocampal formation and prefrontal cortex has been shown to correlate with involvement in a short-term memory task (Sigurdsson et al., 2010). This finding fits well into the general view of gamma rhythms as a mechanism that facilitates communication between brain areas (Colgin et al., 2009).

In our model, gamma occurs as a side effect of feedback stabilization during replay, and by construction, the rhythm does not entrain the excitatory population. It is a suggestive speculation, based on the reports above, that the enhanced power of a brain rhythm can be a dynamic signature of a functional state, in this case that gamma-frequency oscillations be the signature of memory overloading. If gamma, additionally, does indeed correlate with information transfer to other brain areas, then the implication would be that the critically loaded hippocampus periodically rids itself of memories during periods of enhanced gamma (but the hippocampus may still be instrumental for recall, see Goshen et al., 2011). If the establishment of the transfer rhythm is conditional upon overloading at the source, then memory networks would need to be critically loaded for transfer to be initiated and consolidation would be lazy, i.e. activated on an as-needed basis.

5.2.5 Transient Retrieval

Our model identifies parameter regions in which transient retrieval occurs that lasts for only a few time steps. These regions of transient retrieval extend far into low pattern sizes for strong inhibitory feedback, and thus correspond there to the regimes of largest memory capacity. Neuronal networks exhibiting activity sequences hence operate with optimal memory performance in a regime of hyperexcitability that is stabilized by delayed inhibition. This transient retrieval regime is consistent with the dynamic features of sequence replay during sharp wave ripple complexes in the hippocampus. These last for 5 to 10 cycles of an approximately 200 Hz oscillation. Our data analysis work on an in-vitro model of SWR (see above) shows that they are accompanied by delayed inhibitory feedback, though the frequency is roughly similar to that of excitatory activation, and the phase detail cannot be compared with a theoretical model that only has one time step per ripple cycle.

A way out of the capacity vs. stability dilemma In large environments, sequence replay in-vivo can span several ripple episodes (Davidson et al., 2009), allowing long sequences to be constructed by concatenating multiple fragments riding on transient, time-limited ripples. An interpretation of such fragmentation in the context of our theoretical findings is that the dilemma between stability and capacity is solved by chaining several inherently unstable replay epochs, each tuned to maximum capacity. An open question is what is the optimal tradeoff between sequence length and number of successive ripples. At the mechanistical root, it is also unclear how information can be transmitted between successive ripples.

5.2.6 Capacity Bounds Depend on Model Assumptions

Throughout this work, we considered the connectivity parameters c and c_m as constants, based on the assumption that the morphological connectivity c_m is mainly determined by geometrical constraints such as the size of the cell surface, or the volume requirement of wiring. A degree c of functional connectivity results from the specific learning rules that ensure that the network always remains plastic—in order to store new memories a large fraction of synapses has to be able to change its state. In the parameter regime used for our analysis, this requirement is fulfilled by fixing $c/c_m=0.5$. Moreover, the network is sparse, based upon the experimental evidence that indicates that hippocampal networks are sparsely connected (Miles and Wong, 1986).

A specific assumption underlying our analysis is that of a constant and low pattern size M . While the firing sparsity of hippocampal principal neurons is well established, remaining even during ripples below 10 Hz (Csicsvari et al., 1999b), the pattern size itself is likely to be variable in reality. This entails a potential loss of capacity.

Another significant simplification is the discreteness in time. Dynamical interactions of synaptic currents and membrane processes during sharp-wave ripples may also reduce capacities. In this sense the capacity values that we have derived here can only be considered as upper bounds useful for determining scaling behavior. Extending the model to more realistic dynamics is necessary to investigate how tight is this bound for a real spiking network.

5.2.7 More Sophisticated Models Needed to Formulate Predictions

Our time-discrete model of replay ignores the possibility of desynchronization of assemblies that exists in continuous time, as we have just advanced. But it also limits the conclusions that can be reached about the neurophysiological substrate. First, it makes it impossible to examine the dynamics on time scales shorter than one iteration step. In order to be able to accommodate our experimental findings on the phase shift of inhibition with respect to excitation, we need a new model where these fine-grained interactions can be taken into account. Second, not only temporal information gained from experiment is disregarded, but also the cycle-dependent amplitude of ripple currents.

A framework is needed that can relate the observed currents to coding, i.e. neuronal spiking. The mediator between these two objects is, naturally, the neuronal membrane potential. In order to estimate it, an integration equation is needed that, given the conductances, can yield the time-dependent membrane potential. At first, a simple equation with few parameters, like the integrate-and-fire model, can be employed. Conductances being notoriously voltage-dependent, it may be necessary to estimate them from a range of holding potentials in voltage clamp experiments, using for instance peeling reconstruction wherever the polarity allows. The resulting integrated membrane potential can be validated by comparison to current clamp experiments.

A probabilistic model based upon ripple-cycle-dependent conductance distributions can tell us how likely is it for neurons to spike at a given cycle in the ripple. The spiking can then be built into a network model that can be explored for self-consistency, i.e. the sustained spiking output of the network should lead to the observed transient SWR bursts in the LFP.

Even recent models of ripples, whether based on synaptic, directional (Memmesheimer, 2010; Taxidis et al., 2011) or nonsynaptic, symmetric communication (Vladimirov et al., 2012) do not address the functional aspect of sequence replay. We thus identify an acute need to bridge the observed functional role of SWR (replay of behavioral storylines encoded by spike sequences) and its mechanistic constitution as a short-lived LFP burst of intertwined excitatory and inhibitory activity. The model we suggest would provide a stepping stone towards the incorporation of further biological constraints from other studies of SWR, while keeping the replay capability as a functional constraint. The critical test for the integration of both is the ability of assemblies to remain well-differentiated and propagate their activities across ripple cycles.

5.3 General Conclusion and Outlook

In this Thesis we have pursued the elusive signature of neuronal assemblies in action. We have searched them in the loudest but also briefest of all neuronal rhythms: sharp-wave ripples in the hippocampus. To find them, we listened to the network both outside the and inside the cells. We found a data analysis scalpel in peeling reconstruction, allowing to single out the voice of separate assemblies. But to understand why memory needs assemblies, how many cells make up a voice that can be heard amidst the permanent chatter of the brain, and how inhibitory cells, that silence their neighbors, help increase the number of such voices that rise from the choir of memories, we needed to go further and build a mathematical model.

From the detailed analysis of the measurements we learned that excitation and inhibition are both critically involved in shaping the ripple, but also that their accord goes beyond just walking in step. The playing catch of inhibition with excitation remains mysterious: what its function is emerges as a central question derived from the results in this Thesis.

From the theoretical model we learned that even a much less time-specific inhibition than what we observed in the measurements can help the stable replay of memories expressed by excitatory activity, and so parsimoniously allow the network to hold more of them. To find how our relatively simple model in discrete time translates to a continuous setting is the next challenge. Such a model would rejoin the threads in this Thesis, requiring first peeling reconstruction on the recordings of ripple currents to estimate the conductances and then a spiking network to deliver a model for sharp-wave ripples that is biologically constrained, and, at the same time, can make predictions about both function and mechanism.

Appendix A

Electrophysiology Methods

We present here the details of the electrophysiology *in vivo* and *in vitro*, with special attention to the intracellular solutions employed.

A.1 Animal Surgery and Electrophysiology In Vivo

In vivo electrophysiology was carried out in the laboratory of James F.A. Poulet; detection of SWR on *in-vivo* data was carried out at Dietmar Schmitz's laboratory, both in Berlin.

Surgery All *in vivo* experimental procedures followed previously described methods (Crochet and Petersen, 2006; Poulet and Petersen, 2008). Male C57Bl/6 mice, between 3 and 6 week old, were anesthetized and implanted with a lightweight metal head holder. After surgery, animals were allowed to recover for at least one day before starting habituation to head-restraint. Habituation continued for several days until the animal sat calmly for a period of 1-2 h. On the day of the experiment two small craniotomies, for LFP and whole-cell recordings, were made under isoflurane anesthesia (1.5%). Animals were then allowed to recover for at least 2 hours before recordings. Coordinates for craniotomies were determined stereotactically on the left hemisphere: 2 or 3 mm posterior of bregma, and 2 mm lateral of the midline.

Recordings For LFP recordings glass pipettes (5-7 M Ω) filled with Ringer's solution were used. To determine the recording depth of CA1 *stratum pyramidale*, LFP electrodes in both craniotomies were lowered slowly until clear ripple activity was detected, usually at about 1,200-1,300 μ m depth. Then one pipette was retracted and replaced by a patch pipette. Whole-cell recordings were made with 5-7 M Ω glass electrodes filled with intracellular solution KGlu 0 of Table A.1. The liquid junction potential was accounted for by subtracting 7 mV from all recorded voltages (Lee et al., 2009). All *in vivo* signals were amplified 100 times with a Multiclamp 700B (Axon Instruments, Union City, USA), filtered at 10 kHz, digitized at 20 kHz (ITC-18; HEKA Elektronik, Lambrecht, Germany), and stored (IgorPro; WaveMetrics, Lake Oswego, USA).

Anatomical reconstruction (Figs. 2.3 A and 2.6 A) For *in vivo* experiments, mice were deeply anesthetized with urethane immediately after the experiment and perfused with 4% paraformaldehyde (PFA). After overnight fixation, brains were cut into 100 μm thick coronal slices. Biocytin-filled cells were subsequently visualized with 3,3'-diaminobenzidine tetrahydrochloride (0.015%) using a standard ABC kit (Vectorlabs, Burlingame, CA, USA) and reconstructed on a light microscope at 40 \times with a NeuroLucida 3D system (MicroBrightField, Williston, VT, USA).

SWR detection SWRs *in vivo* were detected with an algorithm similar to Csicsvari et al. (1999b), as recounted in detail in Fig. S1 of Maier et al. (2011). Band-pass-filtered (120–300 Hz) extracellular data were rectified and smoothed with a moving average of 10 ms window. Then events exceeded a threshold of $6 \times \text{SD}$ of noise power were selected. Finally, all too-short events, i.e. of duration below 12 ms at a threshold level of $2 \times \text{SD}$ of noise were rejected. The time of the maximum of each LFP ripple was used as a reference. Centered on it, stretches of 400 ms duration of extra and intracellular traces were cut out and stored for analysis.

A.2 In Vitro Experimental Procedures

In vitro experiments were carried out at the laboratory of Dietmar Schmitz in Berlin.

Ethics statement Animal maintenance and experiments were in accordance with the respective guidelines of local authorities (Berlin state government, T0100/03 and G188-09) and followed the German animal welfare act and the European Council Directive 86/609/EEC.

Slice preparation Horizontal slices (400 μm) were prepared from ventral to mid-hippocampus of C57Bl/6 mice 4 to 8 weeks old, and maintained at the surface of oxygenated artificial cerebrospinal fluid (ACSF) at $\sim 35^\circ\text{C}$. ACSF contained (in mM): 119 NaCl, 2.5 KCl, 1.3 MgSO_4 , 2.5 CaCl_2 , 10 glucose, 1 NaH_2PO_4 , 26 NaHCO_3 . Osmolarity of ACSF was routinely checked (290–310 mosmol/l). Slices were incubated for 1 to 4 hours before their transfer to a submerged chamber for recordings at $\sim 32^\circ\text{C}$.

Minislice preparation The minislices were always cut during the slicing procedure and placed in the interface chamber together with all other slices.

Electrophysiological recordings in vitro Data were recorded in Igor Pro using a Multiclamp 700A amplifier. Extracellular LFPs were recorded with ACSF-filled glass electrodes (resistance: 0.2–0.3 $\text{M}\Omega$). Signals were amplified 1,000 times, low-pass filtered at 2 kHz or 4 kHz and digitized at 5 kHz or 10 kHz. Whole-cell recordings were performed with borosilicate glass electrodes (2–5 $\text{M}\Omega$) filled with the intracellular solutions of Table A.1. Which condition uses what solution is indicated in Table 2.1.

Sol.	Modes	Base	Blocker	KCl	EGTA	MgATP	MgSO ₄	Na ₃ GTP	Na ₂ PCr	Other
0	vc/cc	135 KGlu	—	4	—	4	—	0.3	10	—
1	vc/cc	120 KGlu	—	10	5	3	2	—	—	1 GTP
2	vc	120 CsGlu	—	10	5	3	2	1	—	—
3	vc	120 CsF	1 DIDS	10	5	—	—	—	—	—
4_±	vc/cc	70 KGlu	+0.5 DNDS	45	—	4	—	0.4	5	5 CaCl ₂

Table A.1. Composition of Intracellular Solutions. First column numbers the pipette solutions (0 for in vivo, 1-4 for in vitro). Second column indicates the modes in which the solution was used, voltage clamp —vc— or current clamp —cc—. Third column indicates the main anion and/or cation of the solution. Mnemonic names are highlighted in bold and used with the solution numbers and holding potentials to identify solutions throughout the document. Successive columns indicate concentration of different molecules, in mM. The buffering agent HEPES (4-(2-hydroxyethyl)-1-piperazineethanesulfonic acid) was present in all solutions at 10 mM, as was biocytin (2 mg/ml in vivo, 0.3-0.5% all others, only applied to some cells). Intracellular pH was adjusted to ~ 7.4 with KOH (solutions 1, 3 and 4; 0 to 7.3) or CsOH (solution 2, 1 M). *DIDS*: 4,4'-diisothiocyanatostilbene-2,2'-disulfonic acid. *DNDS*: for 4,4'-dinitrostilbene-2,2'-disulfonic acid, disodium salt. *PCr*: phosphocreatine; *KGlu* and *CsGlu*: potassium and cesium gluconate, respectively. *GTP*: guanosine triphosphate. *ATP*: adenosine triphosphate. *EGTA*: ethylene glycol tetraacetic acid (a Ca²⁺ buffer).

Solutions 0 and 1 (**KGlu 0** and **KGlu 1** respectively) are intended, when in voltage clamp, for the electrotonic isolation of excitatory currents by measuring at the reversal potential of chloride. The Nernst reversal potential for Cl⁻ with solution KGlu 1 is -67.3 mV (the unknown ionic composition of the extracellular milieu prevents a similar estimation for KGlu 0, used in vivo). Solution 2 (**CsGlu**) is intended for the electrotonic isolation of inhibitory currents by measuring at the reversal potential of glutamatergic currents, estimated from the Goldman-Hodgkin-Katz equation to be -4 mV. Using cesium instead of potassium as the main cation in solution 2 is needed to block potassium channels and prevent very large potassium outward currents. Solution 3 (**DIDS**) is intended for the pharmacological isolation of excitatory currents. Use of fluoride as the main anion in solution 3 reduces calcium currents and improves duration (the reader can find references that substantiate this fact in Isokawa et al. (1997)). In addition, fluoride, together with lack of ATP, also helps minimize the GABA_A-mediated Cl⁻ conductances, whereas intracellular Cs blocks GABA_B-mediated potassium conductances (references in Nelson et al., 1994, note 8).

De- and hyperpolarizing current steps (0.2-1 s) were applied in the whole-cell current-clamp configuration to characterize the intrinsic properties of each cell; only cells that showed typical spiking characteristics of principal cells were retained for analysis. Series resistance was monitored continuously throughout recordings; cells were rejected if it exceeded 20 M Ω or varied $> 30\%$ during recordings. No compensation was used. Voltages are expressed throughout after liquid junction potential correction; the liquid junction potential was experimentally determined following Neher and Rudy (1992).

GABA uncaging Caged GABA (20 ml at 100 μ M) was reperused at 2.5–3.0 ml/min. Uncaging was done using an UV pulsed laser (Rapp OptoElectronic, Wedel, Germany) attached with a 200 μ m optical fiber coupled into the epifluorescence port of the micro-

scope with an OSI-BX adapter (Rapp OptoElectronic) and focused on the specimen by the objective lens. This yielded an illuminated circle of 20-50 μm . Laser flash duration was 5 ms. Laser power under the objective corresponding to the stimulus intensity levels used was monitored with a photodiode array-based photodetector (PDA-K-60, Rapp OptoElectronic) and did not change over time. GABA was uncaged over the cell soma in the presence of 10 μM NBQX and 50 μM APV.

Anatomical reconstruction After recording, slices were transferred to a fixative solution containing 4% paraformaldehyde (PFA) and 0.2% saturated picric acid in 0.1 M phosphate buffer. Slices were re-sectioned into 70 μm thin sections. Biocytin-filled cells were subsequently visualized with 3,3'-diaminobenzidine tetrahydrochloride (0.015%) using a standard ABC kit (Vectorlabs, Burlingame, CA, USA) and reconstructed on a light microscope at 40 \times with a Neurolucida 3D system (MicroBrightField, Williston, VT, USA).

Applied drugs 2,3-dioxo-6-nitro-1,2,3,4-tetrahydrobenzo(f)quinoxaline-7-sulfonamide (NBQX), DIDS, and caesium fluoride were purchased from Sigma Aldrich, Germany. 6-Imino-3-(4-methoxyphenyl)-1(6H)-pyridazinebutanoic acid hydrobromide (gabazine) and D-(-)-2-Amino-5-phosphonopentanoic acid (D-APV) were obtained from Biotrend, Cologne, Germany. γ -aminobutyric acid, α -carboxy-2-nitrobenzyl ester, trifluoroacetic acid salt (O-(CNB-caged) GABA) was purchased from Molecular Probes, Eugene, OR, USA. DNDS was kindly provided by Dr. Robert J. Bridges, Rosalind Franklin University, Chicago, US.

Slow Currents: NMDA and GABA_B Receptors

NMDA and GABA_B currents are not relevant in the experimental scenario defined above under the voltage clamp conditions described in Chapter 2.

NMDA currents NMDA currents do not conform to Ohm's law by virtue of the Mg^{2+} block-induced nonlinearity in the conductance: $g = g(V)$ is zero unless the cell is strongly depolarized, then it becomes roughly linear (Koch, 2005, Fig. 4.8). But the reversal potential of NMDA receptors is close to 0 mV. Thus, at membrane potentials close to the reversal of Cl^- , the conductance of NMDARs is negligible, while close to the reversal of AMPAR-mediated currents the driving force is very low. Hence at both isolating potentials of choice, slow excitatory currents are not expected to be present. At the intermediate potential of -45 mV we did not appreciate substantially slower spontaneous EPSCs or longer sharp waves. The effect of NMDA on the extracellular waveform has been investigated by Colgin et al. (2005) in CA3. The finding was that bath application of NMDA blockers substantially increased ($\sim 2 \times$) the size both of sharp waves and of ripples, probably due to an indirect effect on the Ca^{2+} flow that is concomitant for NMDA conductances to Na^+ influx. The amount of Mg^{2+} present in the extracellular milieu is critical for the expression of NMDA currents in the physiological situation; the situation has been characterized in vitro in stimulation experiments (Herron et al., 1986), but it is not clear whether short high-frequency trains of EPSPs can lift the voltage-dependent block, and our voltage-clamp experiments do not help resolve this question for the case of ripples.

GABA_B currents Contrary to glutamate receptors agonized by AMPA and NMDA, GABA_A and GABA_B receptors are not colocalized (Koch, 2005). GABA_B receptors are metabotropic receptors that cause, mediated by a G protein, the opening of K⁺ channels with a very low reversal potential (below -90 mV) and slow kinetics (decay in the order of 100 ms). We do not observe such slow currents in any of our recordings; a possible explanation is that the driving force for them at -66 mV is comparatively small, while the presence of cesium in the pipette solution when recording at -6 mV blocks those conductances from the inside of the cell (see above).

Appendix B

Participation in Associations

The probability of potentiation of a synapse whose target has participated in $\pi \in [0, P]$ associations depends only on the firing sparsity of the network, i.e. $q_\pi = 1 - (1 - f)^\pi$. In the process of constructing the synaptic input to an average target neuron in the On or Off population, the moments of the distribution of q_π appear naturally. Here we calculate the mean, variance and coefficient of variation of q according to the distribution of participations,

$$\Pr(\pi) = B_{P,f}(\pi) = \binom{P}{\pi} f^\pi (1 - f)^{P-\pi}. \quad (\text{B.1})$$

B.1 Average and Variance of the Potentiation State

The average potentiation evaluates to

$$\begin{aligned} E_\pi[q] &= \sum_{\pi=0}^P \Pr(\pi) q_\pi \\ &= \sum_{\pi=0}^P \binom{P}{\pi} f^\pi (1 - f)^{P-\pi} (1 - (1 - f)^\pi) \\ &= 1 - \sum_{\pi=0}^P \binom{P}{\pi} f^\pi (1 - f)^{P-\pi} (1 - f)^\pi \\ &= 1 - (1 - f)^P \sum_{\pi=0}^P \binom{P}{\pi} f^\pi \\ &= 1 - (1 - f)^P (1 + f)^P \\ &= 1 - (1 - f^2)^P \\ &= c/c_m. \end{aligned} \quad (\text{B.2})$$

In the last step we have recast the result in terms of the connectivities of the network by employing the expression for the maximal capacity $(1 - f^2)^P = 1 - c/c_m$ from Section 4.1.3. The second moment about zero is calculated similarly:

$$\begin{aligned} E_\pi[q^2] &= \sum_{\pi=0}^P \Pr(\pi) q_\pi^2 \\ &= \sum_{\pi=0}^P \binom{P}{\pi} f^\pi (1 - f)^{P-\pi} (1 - (1 - f)^\pi)^2 \\ &= \sum_{\pi=0}^P \binom{P}{\pi} f^\pi (1 - f)^{P-\pi} (1 - 2(1 - f)^\pi + (1 - f)^{2\pi}) \\ &= 1 - 2(1 - f^2)^P + \sum_{\pi=0}^P \binom{P}{\pi} f^\pi (1 - f)^{P-\pi} (1 - f)^{2\pi} \quad \text{cf. Eq. B.2} \\ &= 1 - 2(1 - f^2)^P + (1 - f)^P \sum_{\pi=0}^P \binom{P}{\pi} (f(1 - f))^\pi \\ &= 1 - 2(1 - f^2)^P + (1 - f)^P (1 + f(1 - f))^P \\ &= 1 - 2(1 - f^2)^P + (1 - f^2)^P (1 + f(1 - f)/1 + f)^P \\ &= 1 + (1 - f^2)^P \left((1 - f^2/1 + f)^P - 2 \right). \end{aligned}$$

The coefficient of variation, defined as the ratio of the standard deviation of a distribution to its mean, provides a normalized estimate of the dispersion of the distribution. We need rather its square for the computation of the variances of synaptic inputs, $CV_{\pi}^2[q] \equiv \text{Var}_{\pi}[q]/E_{\pi}^2[q]$. Its explicit expression is, using the calculations above,

$$CV_{\pi}^2[q] = \frac{E_{\pi}[q^2]}{E_{\pi}[q]^2} - 1 = \frac{c_m^2}{c^2} \left[1 + (1 - f^2)^P \left((1 - f^2/_{1+f})^P - 2 \right) \right] - 1.$$

B.2 Correlations of Potentiation are Linear in Coding Ratio

The correlations in the potentiation state induced by the imprinting rule are reflected as additional terms in the variances of the activations. Those terms contain as a factor the square coefficient of variation just calculated. It turns out that this squared coefficient of variation is linear in f in the regime of low coding ratios $f = M/N \ll 1$, which is anyways a precondition for the validity of our mean field model:

$$\begin{aligned} CV_{\pi}^2[q] &= -\left(1 - \frac{c_m}{c}\right)^2 \log\left(1 - \frac{c}{c_m}\right) f + O(f^2) \\ &= r^2 \log(1 + r^{-1}) f + O(f^2) \end{aligned}$$

In the second line we have introduced the resources for plasticity, $r \equiv c_m/c - 1$. For our CA3-inspired standard network, $r \sim 1$ and thus $CV_{\pi}^2[q] \simeq f \log 2$. The linear approximation is excellent for low f and it will never lead to underestimate the variance for realistic f ($\lesssim 0.9$, Fig. B.1).

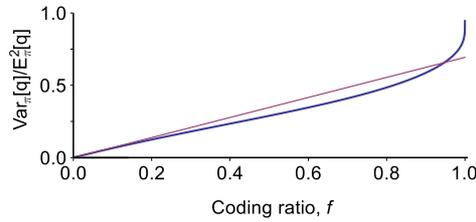


Figure B.1. The Squared Coefficient of Variation is Linear in f Under Sparse Coding.

The squared coefficient of variation of the synaptic states over pattern participation, $CV_{\pi}^2[q] = \text{Var}_{\pi}[q]/E_{\pi}^2[q]$ as a function of the coding ratio $f \equiv M/N$ is shown in blue vs. its linear approximation in red. The range of coding ratios explored in this work is highlighted in the axis of abscissæ.

Care should be exercised in the regions of phase space of high n , where the $\sim n^2$ factor in the variances that multiplies $CV_{\pi}^2[q]$ will amplify and make significant the small overestimation in the linear approximation for about $f > 0.1$. In practice this is never a problem for our purposes, for we concentrate in the regions of phase space at or around retrieval, where n is still small. Thoughtful consideration of the relative importance of the variance terms under situations of interest, such as the aforementioned regime of sequence retrieval, allows to disregard e.g. the correlational contributions to σ_{0n}^2 , which simplifies further analytical work.

Bibliography

- [1] Alted, F. (2010). Why modern CPUs are starving and what can be done about it. *Computing in Science & Engineering*, **12**(2), 68–71.
- [2] Amaral, D. G. and Lavenex, P. (2007). Hippocampal neuroanatomy. In *The Hippocampus Book*. Oxford University Press, USA, 1st edition.
- [3] Amitai, Y., Gibson, J. R., Beierlein, M., Patrick, S. L., Ho, A. M., Connors, B. W., and Golomb, D. (2002). The spatial dimensions of electrically coupled networks of interneurons in the neocortex. *Journal of Neuroscience*, **22**, 4142–4152.
- [4] Andersen, P., Morris, R., Amaral, D., Bliss, T., and O’Keefe, J., editors (2007). *The Hippocampus Book*. Oxford University Press, USA, 1st edition.
- [5] Bähner, F., Weiss, E. K., Birke, G., Maier, N., Schmitz, D., Rudolph, U., Frotscher, M., Traub, R. D., Both, M., and Draguhn, A. (2011). Cellular correlate of assembly formation in oscillating hippocampal networks in vitro. *Proceedings of the National Academy of Sciences of the United States of America*, **108**(35), E607–616.
- [6] Barnes, C. A., Rao, G., and McNaughton, B. L. (1987). Increased electrotonic coupling in aged rat hippocampus: A possible mechanism for cellular excitability changes. *The Journal of Comparative Neurology*, **259**(4), 549–558.
- [7] Battaglia, F. P., Benchenane, K., Sirota, A., Pennartz, C. M. A., and Wiener, S. I. (2011). The hippocampus: hub of brain network communication for memory. *Trends in Cognitive Sciences*, **15**(7), 310–318.
- [8] Behrens, C. J., Boom, L. P. v. d., Hoz, L. d., Friedman, A., and Heinemann, U. (2005). Induction of sharp wave-ripple complexes in vitro and reorganization of hippocampal networks. *Nature Neuroscience*, **8**(11), 1560–1567.
- [9] Bennett, M. V. L. and Zukin, R. S. (2004). Electrical coupling and neuronal synchronization in the mammalian brain. *Neuron*, **41**(4), 495–511.
- [10] Boashash, B. (1992). Estimating and interpreting the instantaneous frequency of a signal. i. fundamentals. *Proceedings of the IEEE*, **80**(4), 520–538.
- [11] Both, M., Bähner, F., von Bohlen und Halbach, O., and Draguhn, A. (2008). Propagation of specific network patterns through the mouse hippocampus. *Hippocampus*, **18**(9), 899–908.
- [12] Brunel, N. (2000). Dynamics of sparsely connected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, **8**(3), 183–208.
- [13] Buhl, E. and Whittington, M. (2006). Local circuits. In *The Hippocampus Book*, pages 297–319.
- [14] Buhry, L., Azizi, A. H., and Cheng, S. (2011). Reactivation, replay, and preplay: How it might all fit together. *Neural Plasticity*, **2011**, 1–11.
- [15] Buzsáki, G. (1986). Hippocampal sharp waves: Their origin and significance. *Brain Research*, **398**(2), 242–252.
- [16] Buzsáki, G. (1989). Two-stage model of memory trace formation: A role for "noisy" brain states. *Neuroscience*, **31**(3), 551–570.
- [17] Buzsáki, G. (1996). The hippocampo-neocortical dialogue. *Cerebral Cortex*, **6**(2), 81–92.

- [18] Buzsáki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, **304**(5679), 1926–1929.
- [19] Buzsáki, G., Lai-Wo S., L., and Vanderwolf, C. H. (1983). Cellular bases of hippocampal EEG in the behaving rat. *Brain Research Reviews*, **6**(2), 139–171.
- [20] Buzsáki, G., Horvath, Z., Urioste, R., Hetke, J., and Wise, K. (1992). High-frequency network oscillation in the hippocampus. *Science*, **256**(5059), 1025–1027.
- [21] Buzsáki, G., Buhl, D. L., Harris, K. D., Csicsvari, J., Czeh, B., and Morozov, A. (2003). Hippocampal network patterns of activity in the mouse. *Neuroscience*, **116**(1), 201–211.
- [22] Canepari, M., Popovic, M., Vogt, K., Holthoff, K., Konnerth, A., Salzberg, B. M., Grinvald, A., Antic, S. D., Zecevic, D., Canepari, M., and Zecevic, D. (2011). Membrane potential imaging in the nervous system. pages 25–41. Springer New York.
- [23] Carr, M. F., Jadhav, S. P., and Frank, L. M. (2011). Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval. *Nature Neuroscience*, **14**(2), 147–153.
- [24] Cash, S. and Yuste, R. (1999). Linear summation of excitatory inputs by CA1 pyramidal neurons. *Neuron*, **22**(2), 383–394.
- [25] Chrobak, J. J. and Buzsáki, G. (1994). Selective activation of deep layer (V-VI) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat. *The Journal of Neuroscience*, **14**(10), 6160–6170.
- [26] Chrobak, J. J. and Buzsáki, G. (1996). High-frequency oscillations in the output networks of the Hippocampal–Entorhinal axis of the freely behaving rat. *The Journal of Neuroscience*, **16**(9), 3056–3066.
- [27] Chrobak, J. J., Lörincz, A., and Buzsáki, G. (2000). Physiological patterns in the hippocampo-entorhinal cortex system. *Hippocampus*, **10**(4), 457–465.
- [28] Colgin, L. L., Jia, Y., Sabatier, J.-M., and Lynch, G. (2005). Blockade of NMDA receptors enhances spontaneous sharp waves in rat hippocampal slices. *Neuroscience Letters*, **385**(1), 46–51.
- [29] Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., and Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, **462**(7271), 353–357.
- [30] Compte, A., Reig, R., Descalzo, V. F., Harvey, M. A., Puccini, G. D., and Sanchez-Vives, M. V. (2008). Spontaneous high-frequency (10–80 Hz) oscillations during up states in the cerebral cortex in vitro. *The Journal of Neuroscience*, **28**(51), 13828–13844.
- [31] Craver, C. F. (2007). *Explaining the brain: mechanisms and the mosaic unity of neuroscience*. Oxford University Press.
- [32] Crochet, S. and Petersen, C. C. H. (2006). Correlating whisker behavior with membrane potential in barrel cortex of awake mice. *Nature Neuroscience*, **9**(5), 608–610.
- [33] Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999a). Fast network oscillations in the hippocampal CA1 region of the behaving rat. *Journal of Neuroscience*, **19**(16), RC20.
- [34] Csicsvari, J., Hirase, H., Czurkó, A., Mamiya, A., and Buzsáki, G. (1999b). Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *The Journal of Neuroscience*, **19**(1), 274–287.
- [35] Csicsvari, J., Hirase, H., Mamiya, A., and Buzsáki, G. (2000). Ensemble patterns of hippocampal CA3–CA1 neurons during sharp wave-associated population events. *Neuron*, **28**(2), 585–594.
- [36] Davidson, T. J., Kloosterman, F., and Wilson, M. A. (2009). Hippocampal replay of extended experience. *Neuron*, **63**(4), 497–507.
- [37] Destexhe, A., Mainen, Z. F., and Sejnowski, T. J. (1994). Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *Journal of Computational Neuroscience*, **1**(3), 195–230.

- [38] Destexhe, A., Rudolph, M., and Paré, D. (2003). The high-conductance state of neocortical neurons in vivo. *Nature Reviews Neuroscience*, **4**(9), 739–751.
- [39] Deuchars, J. and Thomson, A. (1996). CA1 pyramid-pyramid connections in rat hippocampus in vitro: Dual intracellular recordings with biocytin filling. *Neuroscience*, **74**(4), 1009–1018.
- [40] Diba, K. and Buzsáki, G. (2007). Forward and reverse hippocampal place-cell sequences during ripples. *Nature Neuroscience*, **10**(10).
- [41] Diekelmann, S. and Born, J. (2010). The memory function of sleep. *Nature Reviews Neuroscience*, **11**(2), 114–126.
- [42] Dragoi, G. and Tonegawa, S. (2010). Preplay of future place cell sequences by hippocampal cellular assemblies. *Nature*, **469**(7330), 397–401.
- [43] Draguhn, A., Traub, R. D., Schmitz, D., and Jefferys, J. G. R. (1998). Electrical coupling underlies high-frequency oscillations in the hippocampus in vitro. *Nature*, **394**(6689), 189–192.
- [44] Dudek, S. M. and Friedlander, M. J. (1996). Intracellular blockade of inhibitory synaptic responses in visual cortical layer IV neurons. *Journal of Neurophysiology*, **75**(5), 2167–2173.
- [45] Dupret, D., O’Neill, J., Pleydell-Bouverie, B., and Csicsvari, J. (2010). The reorganization and reactivation of hippocampal maps predict spatial memory performance. *Nature Neuroscience*, **13**(8), 995–1002.
- [46] Ego-Stengel, V. and Wilson, M. A. (2010). Disruption of ripple-associated hippocampal activity during rest impairs spatial learning in the rat. *Hippocampus*, **20**(1), 1–10.
- [47] Ellender, T. J. and Paulsen, O. (2010). The many tunes of perisomatic targeting interneurons in the hippocampal network. *Frontiers in Cellular Neuroscience*, **4**.
- [48] Ellender, T. J., Nissen, W., Colgin, L. L., Mann, E. O., and Paulsen, O. (2010). Priming of hippocampal population bursts by individual perisomatic-targeting interneurons. *The Journal of Neuroscience*, **30**(17), 5979–5991.
- [49] Epsztein, J., Lee, A. K., Chorev, E., and Brecht, M. (2010). Impact of spikelets on hippocampal CA1 pyramidal cell activity during spatial exploration. *Science*, **327**(5964), 474–477.
- [50] Fino, E. and Yuste, R. (2011). Dense inhibitory connectivity in neocortex. *Neuron*, **69**(6).
- [51] Frotscher, M., Heimrich, B., and Deller, T. (1997). Sprouting in the hippocampus is layer-specific. *Trends in Neurosciences*, **20**(5), 218–223.
- [52] Fukuda, T. and Kosaka, T. (2000). Gap junctions linking the dendritic network of GABAergic interneurons in the hippocampus. *The Journal of Neuroscience*, **20**(4), 1519–1528.
- [53] Gerstner, W. and Kistler, W. M. (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity*. Cambridge University Press, 1st edition.
- [54] Ghahramani, Z. and Jordan, M. I. (1997). Factorial hidden markov models. *Machine Learning*, **29**(2), 245–273.
- [55] Gibson, W. G. and Robinson, J. (1991). Statistical analysis of the dynamics of a sparse associative memory. *Neural Networks*, **5**(4), 645–661.
- [56] Girardeau, G., Benchenane, K., Wiener, S. I., Buzsáki, G., and Zugaro, M. B. (2009). Selective suppression of hippocampal ripples impairs spatial memory. *Nature Neuroscience*, **12**(10), 1222–1223.
- [57] Golomb, D., Rubin, N., and Sompolinsky, H. (1990). Willshaw model: Associative memory with sparse coding and low firing rates. *Physical Review A*, **41**(4), 1843–1854.
- [58] Goshen, I., Brodsky, M., Prakash, R., Wallace, J., Gradinaru, V., Ramakrishnan, C., and Deisseroth, K. (2011). Dynamics of retrieval strategies for remote memories. *Cell*, **147**(3), 678–689.
- [59] Grewe, B. F., Langer, D., Kasper, H., Kampa, B. M., and Helmchen, F. (2010). High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nature Methods*, **7**(5), 399–405.

- [60] Gulyás, A. I., Hájos, N., and Freund, T. F. (1996). Interneurons containing calretinin are specialized to control other interneurons in the rat hippocampus. *The Journal of Neuroscience*, **16**(10), 3397–3411.
- [61] Handy, T. C. (2009). *Brain signal analysis: advances in neuroelectric and neuromagnetic methods*. MIT Press.
- [62] Harris, K. D. (2005). Neural signatures of cell assembly organization. *Nature Reviews Neuroscience*, **6**(5), 399–407.
- [63] Harris, K. D., Csicsvari, J., Hirase, H., Dragoi, G., and Buzsáki, G. (2003). Organization of cell assemblies in the hippocampus. *Nature*, **424**(6948), 552–556.
- [64] Harvey, C. D., Collman, F., Dombeck, D. A., and Tank, D. W. (2009). Intracellular dynamics of hippocampal place cells during virtual navigation. *Nature*, **461**(7266), 941–946.
- [65] Hasselmo, M. E. (1999). Neuromodulation: acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, **3**(9), 351–359.
- [66] Hebb, D. O. (1949). *The organization of behavior: a neurophysiological theory*. Wiley and Sons, New York.
- [67] Heerden, P. J. v. (1963). A new optical method of storing and retrieving information. *Applied Optics*, **2**(4), 387–392.
- [68] Helstrom, C. W. (1994). *Elements of Signal Detection and Estimation*. Prentice Hall.
- [69] Herron, C. E., Lester, R. A. J., Coan, E. J., and Collingridge, G. L. (1986). Frequency-dependent involvement of NMDA receptors in the hippocampus: a novel synaptic mechanism. *Nature*, **322**(6076), 265–268.
- [70] Hille, B. (2001). *Ion channels of excitable membranes*. Sinauer.
- [71] Hirase, H. and Recce, M. (1996). A search for the optimal thresholding sequence in an associative memory. *Network: Computation in Neural Systems*, **7**(4), 741–756.
- [72] Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, **79**(8), 2554.
- [73] Ibarz, J. M., Foffani, G., Cid, E., Inostroza, M., and Menendez de la Prida, L. (2010). Emergent dynamics of fast ripples in the epileptic hippocampus. *The Journal of Neuroscience*, **30**(48), 16249–16261.
- [74] Insausti, R. (1993). Comparative anatomy of the entorhinal cortex and hippocampus in mammals. *Hippocampus*, **3s**, 19–26.
- [75] Isaacson, J. S. and Scanziani, M. (2011). How inhibition shapes cortical activity. *Neuron*, **72**(2), 231–243.
- [76] Isokawa, M., Levesque, M., Fried, I., and Engel, J. (1997). Glutamate currents in morphologically identified human dentate granule cells in temporal lobe epilepsy. *Journal of Neurophysiology*, **77**(6), 3355–3369.
- [77] Johnson, N. L., Kemp, A. W., and Kotz, S. (2005). *Univariate discrete distributions*. John Wiley and Sons.
- [78] Johnston, D. and Amaral, D. G. (2004). Hippocampus. In *The Synaptic Organization of the Brain*. Oxford University Press.
- [79] Johnston, D. and Wu, S. M.-S. (1994). *Foundations of Cellular Neurophysiology*. The MIT Press, 1st edition.
- [80] Jones, M. W. and McHugh, T. J. (2011). Updating hippocampal representations: CA2 joins the circuit. *Trends in Neurosciences*, **34**(10), 526–535.
- [81] Jutras, M. J. and Buffalo, E. A. (2010). Synchronous neural activity and memory formation. *Current Opinion in Neurobiology*, **20**(2), 150–155.

- [82] Kammerer, A., Tejero-Cantero, Á., and Leibold, C. (2012). Inhibition enhances memory capacity: optimal feedback, transient replay and oscillations. *Journal of Computational Neuroscience*, pages 1–12.
- [83] Kirov, S. A., Sorra, K. E., and Harris, K. M. (1999). Slices have more synapses than perfusion-fixed hippocampus from both young and mature rats. *The Journal of Neuroscience*, **19**(8), 2876–2886.
- [84] Klausberger, T. and Somogyi, P. (2008). Neuronal diversity and temporal dynamics: The unity of hippocampal circuit operations. *Science*, **321**(5885), 53–57.
- [85] Klausberger, T., Magill, P. J., Márton, L. F., Roberts, J. D. B., Cobden, P. M., Buzsáki, G., and Somogyi, P. (2003). Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature*, **421**(6925), 844–848.
- [86] Koch, C. (2005). *Biophysics of computation: information processing in single neurons*. Oxford University Press.
- [87] Koniaris, E., Drimala, P., Sotiriou, E., and Papatheodoropoulos, C. (2011). Different effects of zolpidem and diazepam on hippocampal sharp wave—ripple activity in vitro. *Neuroscience*, **175**(0), 224–234.
- [88] Landisman, C. E. and Connors, B. W. (2005). Long-term modulation of electrical synapses in the mammalian thalamus. *Science*, **310**(5755), 1809–1813.
- [89] Le Van Quyen, M., Foucher, J., Lachaux, J.-P., Rodriguez, E., Lutz, A., Martinerie, J., and Varela, F. J. (2001). Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *Journal of Neuroscience Methods*, **111**(2), 83–98.
- [90] Lee, A. K., Manns, I. D., Sakmann, B., and Brecht, M. (2006). Whole-cell recordings in freely moving rats. *Neuron*, **51**(4), 399–407.
- [91] Lee, A. K., Epsztein, J., and Brecht, M. (2009). Head-anchored whole-cell recordings in freely moving rats. *Nature Protocols*, **4**(3), 385–392.
- [92] Leibold, C. (2011). What autoassociative network models may tell us about the neurobiology of memory. *Network: Computation in Neural Systems*, **22**(1-4), 231–250.
- [93] Leibold, C. and Kempster, R. (2006). Memory capacity for sequences in a recurrent network with biological constraints. *Neural Computation*, **18**(4), 904–941.
- [94] Lindén, H., Tetzlaff, T., Potjans, T. C., Pettersen, K. H., Grün, S., Diesmann, M., and Einevoll, G. T. (2011). Modeling the spatial reach of the LFP. *Neuron*, **72**(5), 859–872.
- [95] Longstaff, A. (2011). *Bios Instant Notes Neuroscience*. Taylor & Francis.
- [96] Lu, L., Airey, D. C., and Williams, R. W. (2001). Complex trait analysis of the hippocampus: Mapping and biometric analysis of two novel gene loci with specific effects on hippocampal structure in mice. *The Journal of Neuroscience*, **21**(10), 3503–3514.
- [97] Magee, J. C. (1999). Dendritic ih normalizes temporal summation in hippocampal CA1 neurons. *Nature Neuroscience*, **2**(6), 508–514.
- [98] Magee, J. C. and Cook, E. P. (2000). Somatic EPSP amplitude is independent of synapse location in hippocampal pyramidal neurons. *Nature Neuroscience*, **3**(9), 895–903.
- [99] Maier, N., Nimrich, V., and Draguhn, A. (2003). Cellular and network mechanisms underlying spontaneous sharp wave-ripple complexes in mouse hippocampal slices. *The Journal of Physiology*, **550**(3), 873–887.
- [100] Maier, N., Morris, G., Jochenning, F. W., and Schmitz, D. (2009). An approach for reliably investigating hippocampal sharp wave-ripples in vitro. *PLoS ONE*, **4**(9), e6925.
- [101] Maier, N., Tejero-Cantero, Á., Dorn, A. L., Winterer, J., Beed, P. S., Morris, G., Kempster, R., Poulet, J. F., Leibold, C., and Schmitz, D. (2011). Coherent phasic excitation during hippocampal ripples. *Neuron*, **72**(1), 137–152.

- [102] Makarova, J. (2011). Parallel readout of pathway-specific inputs to laminated brain structures. *Frontiers in Systems Neuroscience*, **5**.
- [103] Mallat, S. and Zhang, Z. (1993). Matching pursuits with time-frequency dictionaries. *Signal Processing, IEEE Transactions on*, **41**(12), 3397–3415.
- [104] Margrie, T., Brecht, M., and Sakmann, B. (2002). In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. *Pflügers Archiv European Journal of Physiology*, **444**, 491–498.
- [105] Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions B*, **262**(841), 23–81.
- [106] McClelland, J. L., McNaughton, B. L., and O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, **102**(3), 419–457.
- [107] Megías, M., Emri, Z., Freund, T., and Gulyás, A. (2001). Total number and distribution of inhibitory and excitatory synapses on hippocampal CA1 pyramidal cells. *Neuroscience*, **102**(3), 527–540.
- [108] Memmesheimer, R.-M. (2010). Quantitative prediction of intermittent high-frequency oscillations in neural networks with supralinear dendritic interactions. *Proceedings of the National Academy of Sciences*, **107**(24), 11092–11097.
- [109] Mercer, A., Bannister, A. P., and Thomson, A. M. (2007). Electrical coupling between pyramidal cells in adult cortical regions. *Brain Cell Biology*, **35**(1), 13–27.
- [110] Mikula, S., Trotts, I., Stone, J. M., and Jones, E. G. (2007). Internet-enabled high-resolution brain mapping and virtual microscopy. *NeuroImage*, **35**(1), 9–15.
- [111] Miles, R. and Wong, R. K. (1986). Excitatory synaptic interactions between CA3 neurones in the guinea-pig hippocampus. *The Journal of Physiology*, **373**.
- [112] Montgomery, J. M., Pavlidis, P., and Madison, D. V. (2001). Pair recordings reveal all-silent synaptic connections and the postsynaptic expression of long-term potentiation. *Neuron*, **29**(3), 691–701.
- [113] Nadal, J. P. (1991). Associative memory: on the (puzzling) sparse coding limit. *Journal of Physics A: Mathematical and General*, **24**, 1093.
- [114] Nakashiba, T., Young, J. Z., McHugh, T. J., Buhl, D. L., and Tonegawa, S. (2008). Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. *Science*, **319**(5867), 1260–1264.
- [115] Nakashiba, T., Buhl, D. L., McHugh, T. J., and Tonegawa, S. (2009). Hippocampal CA3 output is crucial for ripple-associated reactivation and consolidation of memory. *Neuron*, **62**(6), 781–787.
- [116] Neher, E. and Rudy, B. (1992). Correction for liquid junction potentials in patch clamp experiments (Chapter 6). In *Ion Channels*, volume Volume 207, pages 123–131. Academic Press.
- [117] Nelson, S., Toth, L., Sheth, B., and Sur, M. (1994). Orientation selectivity of cortical neurons during intracellular blockade of inhibition. *Science*, **265**(5173), 774–777.
- [118] Neyton, J. and Trautmann, A. (1986). Physiological modulation of gap junction permeability. *Journal of Experimental Biology*, **124**(1), 93–114.
- [119] Ng, A. (2011). Machine learning (Stanford CS 229 - lecture notes).
- [120] Nimmrich, V., Maier, N., Schmitz, D., and Draguhn, A. (2005). Induced sharp wave-ripple complexes in the absence of synaptic inhibition in mouse hippocampal slices. *The Journal of Physiology*, **563**(3), 663–670.
- [121] Nusser, Z., Lujan, R., Laube, G., Roberts, J. B., Molnar, E., and Somogyi, P. (1998). Cell type and pathway dependence of synaptic AMPA receptor number and variability in the hippocampus. *Neuron*, **21**(3), 545–559.

- [122] O'Keefe, J. (2007). Hippocampal neurophysiology in the behaving animal. In *The Hippocampus Book*. Oxford University Press, USA, 1st edition.
- [123] O'Keefe, J. and Nadel, L. (1978). *The Hippocampus as a Cognitive Map*. Oxford University Press, USA.
- [124] O'Neill, J., Senior, T., and Csicsvari, J. (2006). Place-selective firing of CA1 pyramidal cells during sharp Wave/Ripple network patterns in exploratory behavior. *Neuron*, **49**(1), 143–155.
- [125] O'Neill, J., Senior, T. J., Allen, K., Huxter, J. R., and Csicsvari, J. (2008). Reactivation of experience-dependent cell assembly patterns in the hippocampus. *Nature Neuroscience*, **11**(2), 209–215.
- [126] Poor, H. V. (1994). *An introduction to signal detection and estimation*. Birkhäuser.
- [127] Popovic, M. A., Foust, A. J., McCormick, D. A., and Zecevic, D. (2011). The spatio-temporal characteristics of action potential initiation in layer 5 pyramidal neurons: a voltage imaging study. *The Journal of Physiology*, **589**(17), 4167–4187.
- [128] Poulet, J. F. A. and Petersen, C. C. H. (2008). Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature*, **454**(7206), 881–885.
- [129] Rabinovich, M., Volkovskii, A., Lecanda, P., Huerta, R., Abarbanel, H. D. I., and Laurent, G. (2001). Dynamical encoding by networks of competing neuron groups: Winnerless competition. *Physical Review Letters*, **87**(6), 068102.
- [130] Rall, W. (1967). Distinguishing theoretical synaptic potentials computed for different somadendritic distributions of synaptic input. *Journal of Neurophysiology*, **30**(5), 1138–1168.
- [131] Reichert, H. (1993). *Introduction to Neurobiology*. Oxford University Press, USA, 1st edition.
- [132] Reichinnek, S., Künsting, T., Draguhn, A., and Both, M. (2010). Field potential signature of distinct multicellular activity patterns in the mouse hippocampus. *Journal of Neuroscience*, **30**(46), 15441–15449.
- [133] Richardson, M. J. E. and Silberberg, G. (2008). Measurement and analysis of postsynaptic potentials using a novel voltage-deconvolution method. *Journal of Neurophysiology*, **99**(2), 1020–1031.
- [134] Roux, S. G., Cenier, T., Garcia, S., Litaudon, P., and Buonviso, N. (2007). A wavelet-based method for local phase extraction from a multi-frequency oscillatory signal. *Journal of Neuroscience Methods*, **160**(1), 135–143.
- [135] Schmitz, D., Schuchmann, S., Fisahn, A., Draguhn, A., Buhl, E. H., Petrasch-Parwez, E., Dermietzel, R., Heinemann, U., and Traub, R. D. (2001). Axo-axonal coupling: A novel mechanism for ultrafast neuronal communication. *Neuron*, **31**(5), 831–840.
- [136] Shiavi, R. (2007). *Introduction to applied statistical signal analysis: guide to biomedical and electrical engineering applications*. Academic Press.
- [137] Sigurdsson, T., Stark, K. L., Karayiorgou, M., Gogos, J. A., and Gordon, J. A. (2010). Impaired hippocampal-prefrontal synchrony in a genetic mouse model of schizophrenia. *Nature*, **464**(7289), 763–767.
- [138] Spruston, N. (2008). Pyramidal neurons: dendritic structure and synaptic integration. *Nature Reviews Neuroscience*, **9**(3), 206–221.
- [139] Stacey, W. C., Krieger, A., and Litt, B. (2011). Network recruitment to coherent oscillations in a hippocampal computer model. *Journal of Neurophysiology*, **105**(4), 1464–1481.
- [140] Strien, N. M. v., Cappaert, N. L. M., and Witter, M. P. (2009). The anatomy of memory: an interactive overview of the parahippocampal-hippocampal network. *Nature Reviews Neuroscience*, **10**(4), 272–282.
- [141] Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Westview Press.

- [142] Sullivan, D., Csicsvari, J., Mizuseki, K., Montgomery, S., Diba, K., and Buzsáki, G. (2011). Relationships between hippocampal sharp waves, ripples, and fast gamma oscillation: Influence of dentate and entorhinal cortical activity. *The Journal of Neuroscience*, **31**(23), 8605–8616.
- [143] Taxis, J., Coombes, S., Mason, R., and Owen, M. R. (2011). Modeling sharp wave-ripple complexes through a CA3-CA1 network model with chemical synapses. *Hippocampus*.
- [144] Traub, R., Schmitz, D., Jefferys, J., and Draguhn, A. (1999). High-frequency population oscillations are predicted to occur in hippocampal pyramidal neuronal networks interconnected by axo-axonal gap junctions. *Neuroscience*, **92**(2), 407–426.
- [145] Traub, R. D. and Bibbig, A. (2000). A model of high-frequency ripples in the hippocampus based on synaptic coupling plus Axon–Axon gap junctions between pyramidal neurons. *The Journal of Neuroscience*, **20**(6), 2086–2093.
- [146] Tsien, J. Z., Huerta, P. T., and Tonegawa, S. (1996). The essential role of hippocampal CA1 NMDA Receptor–Dependent synaptic plasticity in spatial memory. *Cell*, **87**(7), 1327–1338.
- [147] Tufte, E. R. (1997). *Visual Explanations: Images and Quantities, Evidence and Narrative*. Graphics Press.
- [148] Tufte, E. R. (2001). *The Visual Display of Quantitative Information*. Graphics Press, 2nd edition.
- [149] Vladimirov, N., Tu, Y., and Traub, R. D. (2012). Shortest loops are pacemakers in random networks of electrically coupled axons. *Frontiers in Computational Neuroscience*, **6**.
- [150] Vogelstein, J. T., Packer, A. M., Machado, T. A., Sippy, T., Babadi, B., Yuste, R., and Paninski, L. (2010). Fast nonnegative deconvolution for spike train inference from population calcium imaging. *Journal of Neurophysiology*, **104**(6), 3691–3704.
- [151] Wang, X.-J. and Buzsáki, G. (1996). Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *The Journal of Neuroscience*, **16**(20), 6402–6413.
- [152] Wang, Y., Barakat, A., and Zhou, H. (2010). Electrotonic coupling between pyramidal neurons in the neocortex. *PLoS ONE*, **5**(4), e10253.
- [153] Wehr, M. and Laurent, G. (1996). Odour encoding by temporal sequences of firing in oscillating neural assemblies. *Nature*, **384**(6605), 162–166.
- [154] West, M. J., Slomianka, L., and 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*, **231**(4), 482–497.
- [155] Williams, S. R. and Mitchell, S. J. (2008). Direct measurement of somatic voltage clamp errors in central neurons. *Nature Neuroscience*, **11**(7), 790–798.
- [156] Willshaw, D. J., Buneman, O. P., and Longuet-Higgins, H. C. (1969). Non-holographic associative memory. *Nature*, **222**(5197), 960–962.
- [157] Wilson, M. A. and McNaughton, B. L. (1994). Reactivation of hippocampal ensemble memories during sleep. *Science*, **265**(5172), 676–679.
- [158] Wu, C., Asl, M. N., Gillis, J., Skinner, F. K., and Zhang, L. (2005a). An in vitro model of hippocampal sharp waves: Regional initiation and intracellular correlates. *Journal of Neurophysiology*, **94**(1), 741–753.
- [159] Wu, C., Luk, W. P., Gillis, J., Skinner, F., and Zhang, L. (2005b). Size does matter: Generation of intrinsic network rhythms in thick mouse hippocampal slices. *Journal of Neurophysiology*, **93**(4), 2302–2317.
- [160] Ylinen, A., Bragin, A., Nadasdy, Z., Jando, G., Szabo, I., Sik, A., and Buzsáki, G. (1995). Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *The Journal of Neuroscience*, **15**(1), 30–46.
- [161] Yuste, R. (2011). Dendritic spines and distributed circuits. *Neuron*, **71**(5), 772–781.

Tools

Data analysis for Chapter 2 was performed by collaborators using Matlab (except for coherences). Automatic symbolics, numerics and plotting for Chapter 4 (*On Sequence Memory*) were implemented in Mathematica 8.0 and, occasionally, in Python.

The effort of many individuals involved in Free Software projects has provided me with considerable freedom to elaborate and express my ideas by the use of a computer. I wish to restate here my debt of gratitude with each of them. A non limitative list of software I used follows. Version numbers are provided for reproducibility and to help track down potential issues.

Processing — Python 2.6.5 and its modules Numpy 1.5.0, Scipy 0.8.0, ipython 0.10.1 and rpy 2.1.9 for access to the Rayleigh test in the R 2.10.1 package circular 0.3-8.

Storage of electrophysiological data — h5py 1.3.0 to access time series stored in HDF5 binary files, SQLite 3.7.4 via Python's sqlite3 module, version 2.4.1 for the metadata, which was stored using a custom schema.

Source code edition and management — GNU Emacs, Mercurial and Git for source code edition and management.

Bibliography — Zotero and Bib_TE_X.

Writing this document — GNU T_EX_{MACS}, a free multiplatform interactive typesetter, tree-based programmable document editor, computer algebra frontend, and more.

Knowledge management — tiddlyweb/tiddlyspace, tiddlywiki and org-mode.

Graphics — Matplotlib 1.0.0. Inkscape 0.48. The GIMP.

General tools — The Firefox and Chrome browsers, Ubuntu / Debian, the GNU toolset and the Linux kernel.

This work was funded with public resources. All results are available from the author upon request at alvaro@minin.es.

Acknowledgements

How did this all start? A few years ago, motivated by the problem of the network of airports at the DLR and a readings with Nacho and Marta, I grew very interested in complex networks. I wanted to return to Academia to learn more about graphs and how to combine their discrete topology with the continuous manifolds of our spatial experience into *graphifolds*. One good day, talking to Philipp about it, he said—look, why don't you apply to our Graduate School of Systemic Neurosciences? The deadline is *in two weeks*. Intimidated, I replied: Philipp, I don't know if I *like* Neuroscience. Philipp gave the mythical, and true answer: you will *first* apply and *then* become interested. And so it was—over the next few months the lengthy selection process left time enough to become acquainted with Neuroscience and realize that the most complex graph I could ever wish for was embedded right into my head! And not only that—it was a biological structure, and I knew *nothing, absolutely nothing* about biology. The challenge was set.

Thus my earliest gratitude goes to good friends Marta, Nacho, and Philipp. To José Manuel Udías, who supported my search for a field where I could unleash all my passion. And from beginning to end, to Christian Leibold, who taught by example how to say much with few words, and how to generally *be* good. For scientific advice and advice about Science, I am also grateful to Stefan Glasauer, Christian Machens, Gonzalo de Polavieja, Nikolaus Maier and Liset Menéndez de la Prida.

That the time in Munich became an unreservedly happy part of my life owes much to the framework in which it all happened. Adding to the charms of the city, Bayern and the Alps, a truly special environment: the hard work of dedicated men and women. Thank you Benedikt, and all the team of the GSN. It takes as well a great effort to foster an atmosphere of openness and intellectual stimulation as Andreas Herz has successfully done around the BCCN. Christian appears again in this setting, for he has no minor role in both enterprises, though in his characteristic manner, this may pass unnoticed to the casual observer. For my financial sustenance I am indebted to the GSN and to the DFG (LE-2250/2-1,2). Ich möchte an dieser Stelle meine Dankbarkeit auch an Deutschland kundtun: ich bin nur die Summe aller Möglichkeiten, die mir angeboten wurden, und dieses großzügige und offene Land bot mir viele. Möge dieses schöne Land, das ich liebe, ein Leuchtturm dieser Werte bleiben.

Commuting to Martinsried by bike or U-Bahn was always a pleasure, knowing that I could spend the day with Hannes and Philipp, *Breiman colegas* of Eisbach and Westpark in our first lab on the third floor; friends who went with me *durch dick und dünn*. At that time I also met Marc, from whom I learned, in and out of science. Later came Alex Mathis, in whom we trust, the delightful collaboration with Axel, the joyous curmudgeoning with Gonzalo, the excitement with Hannes and Anna around informational syntropies, and of course Alex Loebel, whose image cheers me up even before he starts telling any of his jokes. Saurabh shared a precious jewel that Sebastian

helped cultivate. Back home it would be all fun, and cooking and cycling to collect Holunderblüte with Hugo and Bruno, and like it is the case with Valentina, and Fabien, and Orian, and Andreas and Miguel and Neil, there is simply no way to capture the joy of our friendship in a few lines, nor there is need for it here. Life seems to be just all that time that passes before seeing you again.

Dans la vallée du haut Var habitent un homme libre, et une petite fille qui craint un renard méchant mais absolument rien d'autre. J'ai appris tant de vous Alfred, Théa!

Less frequently than I wished, I enjoyed in Munich the company of Armin, Onur, Cristina, André, Julie, Sophie, Christian Porres, Eric, Eva-Maria, Javi, Frederik, Yuri, Carlos, Amory, Giulia, André, Heidi and Peter, Andi and María, and Wolfgang.

Il y a eu des moments où je me savais proche de tout ce qu'importe et que n'est jamais ce qui est grand ou qui brille. Cette joie des petits moments était bien souvent des côtés des Bleus de Munich—des amis, un ballon, l'herbe et les arbres majestueux de l'Englischer Garten, une bière bien gagnée après le match, des fêtes conviviales. Vous me manquez trop; je vais seulement nommer ici, pour tous, notre Capitain Thomas, celui aussi qui m'ouvrit la porte de votre belle société, et très cher ami qui personnifie à merveille les valeurs d'une groupe accueillante, noble... et gagnante!

No olvido de donde vengo: vengo del páramo de Castilla, vengo de los mares dorados de trigo mesetario, vengo de las montañas festoneadas de blanco en lontananza, vigilantes del curso de sus ríos, vengo del chopo, la viña y la encina. Vengo de mis abuelos minín y conce, de José María y Pilar que sembraron un país áspero y pobre de generosidad y esfuerzo. Que luego cultivaron mis padres, regándolo de esperanza, de concordia y de futuro. Ese país, que es el mío, lo dignifican día a día, Pablo, y Marc, y Marta, y Dani, y Pampa, y Helio, y Arantxa y Julián, y Luis, que hacen de él un lugar al que merece la pena volver, para mí y para mis exiliados queridos: Alicia, Samuel, Javi, África, Nacho y Almudena. La última curva la dulcificaron Miguel, Irene, Carlota, y Marta y Hugo y Ricardo, conspiradores de grandes alegrías. Llevo esos últimos meses y días en Munich en el recuerdo más íntimo, una felicidad simple y despreocupada, infantil como la de Pablo, obra de arte de Pilar y Joachim en cuya casa siempre me sentí en casa. Closer to the finish line, in Oxford, I was fortunate to meet Jean-Michel, with whom to explore the meaning of ultimate focus on the small, important things, and a cycling metaphor for almost any aspect of life (not to forget some epic culinary experiences). I treasure also the absolutely nonrandom kindness of Megan, to whom I thank as well, on behalf of *the* reader, the shortening of quite a few unbearably long sentences. The friendly support of David and Natalia meant much to me and was essential amidst inevitable times of distress.

Acabando estas líneas, una mirada de complicidad para Marta, una piedra, cómo no, para Ruffa, guapa, juguetona y leal, y, para mis padres, un abrazo de los que aprietan.

Eidesstattliche Erklärung

Hiermit versichere ich, dass ich die vorliegende Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Álvaro Tejero-Cantero
Oxford, 9 Mai 2012