

Modelling and Analysis of Cortico-Hippocampal Interactions and Dynamics During Sleep and Anaesthesia

Ioannis Taxidis, MSc

Thesis submitted to The University of Nottingham
for the degree of Doctor of Philosophy

July 2011

Abstract

The standard memory consolidation model assumes that new memories are temporarily stored in the hippocampus and later transferred to the neocortex, during deep sleep, for long-term storage, signifying the importance of studying functional and structural cortico-hippocampal interactions. Our work offers a thorough analysis on such interactions between neocortex and hippocampus, along with a detailed study of their intrinsic dynamics, from two complementary perspectives: statistical data analysis and computational modelling.

The first part of this study reviews mathematical tools for assessing directional interactions in multivariate time series. We focus on the notion of Granger Causality and the related measure of generalised Partial Directed Coherence (gPDC) which we then apply, through a custom built numerical package, to electrophysiological data from the medial prefrontal cortex (mPFC) and hippocampus of anaesthetized rats. Our gPDC analysis reveals a clear lateral-to-medial hippocampus connectivity and suggests a reciprocal information flow between mPFC and hippocampus, altered during cortical activity.

The second part deals with modelling sleep-related intrinsic rhythmic dynamics of the two areas, and examining their coupling. We first reproduce a computational model of the cortical slow oscillation, a periodic alteration between activated (UP) states and neuronal silence. We then develop a new spiking network model of hippocampal areas CA3 and CA1, reproducing many of their intrinsic dynamics and exhibiting sharp wave-ripple complexes, suggesting a novel mechanism for their generation based on CA1 interneuronal activity and recurrent inhibition. We finally couple the two models to study interactions between the slow oscillation and hippocampal activity. Our simulations propose a dependence of the correlation between UP states and hippocampal spiking on the excitation-to-inhibition ratio induced by the mossy fibre input to CA3 and by a combination of the Schaffer collateral and temporoammonic input to CA1. These inputs are shown to affect reported correlations between UP states and ripples.

Acknowledgements

Firstly, I need to express my deepest gratitude to my (practically three) supervisors for giving me the opportunity to become involved in such an exciting scientific field. This dissertation would not have been possible without their perpetual enthusiasm and immense knowledge. In specific, a big thank you to Markus Owen for always offering his deep mathematical and programming skills to my aid but most importantly for trusting me and allowing me to develop my research plans, while always providing a valuable guidance and solid supervision. A huge thanks to Rob Mason for introducing me to the fascinating world of experimental neuroscience and for always being there for advice and guidance on biological and many other issues. A special thanks to Steve Coombes who has been like a third supervisor to me and a close collaborator. His deep knowledge on the field of mathematical neuroscience and his constant overall guidance and support from the initial to the final level have deeply shaped this work. I am further grateful to all three above for their suggestions for participation in so many conferences, summer schools and workshops, thus introducing me to the vast neuroscientific community and its fascinating research subjects.

I also owe my deepest gratitude to my early collaborator Ben Coomber for his invaluable help during the first stages of this work and to Theodore Kypraios who, apart from a friend, has also been a great teacher of statistics to me. Furthermore, I would like to thank my internal and external examiners during my final examination, for their insightful comments and kind feedback, and all the scientists whom I have contacted in person or by email and have given me precious advice.

I am also very grateful to the European Commission's Framework Programme 6, for funding my research through a Marie Curie Early Stage Training fellowship, and the administrative and IT support personnel (particularly the indispensable Hilary Lonsdale and Dave Parkin) in the Department of Mathematical Sciences for readily providing help and support in many different ways.

These past three and a half years, apart from being scientifically stimulating, have also been deeply enjoyable, thanks to the presence of so many friends in Nottingham and abroad. In particular, I would like to thank my colleagues Helmut Schmidt, Jonathan Laudanski, Kyle Wedgewood and Jaroslav Hlinka for the stimulating discussions about science and life, and for being such amazing friends and officemates. A big thank you to all my friends in the mathematics neuro-group and all my colleagues abroad, who opened my horizons in neuroscience and made conferences such great experiences. To Marios Garganourakis and Nikos Fanidakis and all my friends in Melton Hall for their general support and amazing friendship. Moreover, I have to specially thank Daniel Michelbrink, Elena Elia, Eleni Verykouki, Vicky Panagiotopoulou and Sotiris Prokopiou for making every day in the office so enjoyable and fun. The friendship of all these people is truly precious.

A very special thanks to Athena, who through her love and presence in my life (but mostly in my mind) has influenced me during this work in one way or another.

Last but not least, I need to thank my parents for their caring love and unlimited support in so many different ways. This dissertation is dedicated to them, along with my eternal love.

Contents

List of Abbreviations	xiii
1 Introduction	1
I Assessing Cortico-Hippocampal Directed Interactions	6
2 Causal Interactions in Multivariate Autoregressive Processes	7
2.1 Introduction	7
2.2 Multivariate Time Series	8
2.2.1 Basic Measures	8
2.2.2 Stationarity	9
2.3 Multivariate Autoregressive Models	10
2.3.1 Autoregressive Processes	11
2.3.2 Stability Criterion	12
2.3.3 Choosing the Optimal VAR Model order	14
2.4 Spectral Analysis of Time Series	16
2.4.1 Power Spectrum and Nyquist Theorem	16
2.4.2 Spectral Estimation of a VAR Model	17
2.5 Structural Analysis of Multivariate Time Series	18
2.5.1 Correlation	18
2.5.2 Coherence	19

2.5.3	Directed Coherence	19
2.5.4	Directed Transfer Function	20
2.5.5	Granger Causality	22
2.5.6	Granger Causality Index	23
2.5.7	Partial Directed Coherence	24
2.5.8	Comparing Coherence, PDC, DTF and dDTF	27
2.5.9	Generalised Partial Directed Coherence	29
2.6	Discussion	31
3	Assessing Cortico-Hippocampal Connectivity During Anaesthesia	34
3.1	Introduction	34
3.2	Methods of Spike Train Preprocessing	36
3.2.1	Firing Rates	36
3.2.2	Bin Length Effect	37
3.2.3	Overlap Effect	39
3.2.4	Spike Times	40
3.3	Data Processing and Choice of VAR Model Order	42
3.4	VAR Modelling and Structural Analysis with FunCAT	44
3.5	Results of gPDC Analysis	47
3.6	Discussion	52
II	Modelling Cortical and Hippocampal Dynamics and Their Interactions	58
4	Cortical Sleep Rhythms and a Model of the Slow Oscillation	59
4.1	Introduction	59
4.2	Cortical Sleep Rhythms	60
4.2.1	Spindles	60
4.2.2	Delta Waves	61

4.3	Slow Oscillation	62
4.3.1	General Characteristics of the Slow Oscillation	62
4.3.2	Grouping of Spindles and Delta Waves by the Slow Oscillation	64
4.4	Computational Models of the Slow Oscillation	65
4.4.1	The Bazhenov et al. Model	66
4.4.2	The Compte et al. Model	67
4.4.3	The Hill and Tononi Model	69
4.5	Implementation of the Compte et al. Slow Oscillation Model in Brian	71
4.5.1	Single Cell Models	71
4.5.2	Modelling Synapses	73
4.5.3	Modelling the Network Connectivity	76
4.5.4	Numerical Methods	77
4.5.5	Modelling Slow Oscillations	77
4.6	Discussion	81
5	Network Models of the Hippocampal CA3 and CA1 Areas	85
5.1	Introduction	85
5.2	Hippocampal Morphology	86
5.3	Models of Hippocampal Cells	88
5.3.1	The Traub Model	88
5.3.2	The Pinsky-Rinzel Model	90
5.3.3	A Modification of the Pinsky-Rinzel Model	91
5.3.4	The Wang-Buzsáki Model	92
5.4	Description of a CA3 and a CA1 Network Model	93
5.4.1	General Model Architecture and Single Cell Models	93
5.4.2	Synaptic Interactions	94
5.4.3	General Connectivity	95
5.4.4	CA3 Connectivity	96

5.4.5	CA1 Connectivity	99
5.4.6	Numerical Methods	100
5.5	Simulation Results	100
5.5.1	Validation of the CA3 Connectivity	100
5.5.2	Spread of Activity Over the CA3 Array	101
5.5.3	Rhythmic Population Bursts in the CA3 Model	102
5.5.4	Interneuronal Gamma Oscillations in the CA1 Model	104
5.5.5	Gamma Oscillations in the Full CA1 Model	106
5.6	Discussion	106
5.6.1	The CA3 Model	107
5.6.2	The CA1 Model	109
6	A Computational Model of Sharp Wave-Ripple Complexes	112
6.1	Introduction	112
6.2	The Sharp Wave-Ripple Complex	112
6.2.1	Sharp Waves	112
6.2.2	Ripples	114
6.2.3	Possible Mechanisms for Ripple Generation	116
6.3	A Model for Ripples Based on Axo-Axonal Gap Junctions	117
6.4	Connecting the CA3 and CA1 Models and Detecting Ripples	119
6.4.1	Schaffer Collateral Connectivity	119
6.4.2	Simulations of Extracellular Recordings	121
6.4.3	Ripple Detection	122
6.5	Simulation Results	123
6.5.1	Sharp Wave-Ripple Complexes in the Full CA3-CA1 Model . . .	123
6.5.2	The Role of Inhibition and Excitation in Ripples	132
6.6	Discussion	133
6.6.1	Model Architecture	134

6.6.2	SWRs in the Full CA3-CA1 Model	134
6.6.3	Pharmacological Effects on Ripples	137
6.6.4	SWRs and Memory Consolidation	138
6.6.5	Gap Junctions	140
6.6.6	A Proposed Mechanism for SWRs	141
7	A Model of Cortico-Hippocampal Interactions	143
7.1	Introduction	143
7.2	Slow Oscillation-Ripple Correlations in Electrophysiological Recordings	144
7.3	Data Analysis on Slow Oscillation-Ripples Correlations	148
7.4	A Model of Cortico-Hippocampal Connectivity	152
7.4.1	General Connectivity and PSPs	152
7.4.2	PFC-to-CA3: Dentate Gyrus Mossy Fiber Input	155
7.4.3	CA1 Response to the PFC-to-CA3 Model	159
7.4.4	PFC-to-CA1: Entorhinal Cortex Temporoammonic Input	162
7.4.5	CA1-to-PFC: Monosynaptic Projections	166
7.5	Discussion	170
7.5.1	Mossy Fibre Input	171
7.5.2	Temporoammonic Input	173
7.5.3	Feedback to the Cortex and Memory Consolidation	175
7.5.4	Future Steps	176
8	Summary - Future Directions	177
8.1	Summary of Results	177
8.2	Some Future Directions	180
	References	182

List of Figures

2.1	Effect of VAR model order on PDC	27
2.2	VAR modelling a set of coupled oscillators	29
2.3	Comparing coherence, PDC, DTF and dDTF	30
2.4	Comparing PDC with gPDC	32
3.1	mPFC-hippocampal recording sites and LFP segments	35
3.2	mPFC-hippocampus PDC over one segment of rat LFP recordings . . .	37
3.3	Firing-rate PDC dependence on bin length	38
3.4	Firing-rate PDC dependence on bin overlap	40
3.5	French-Holden algorithm applied on a spike train	42
3.6	mPFC-hippocampus gPDC for LFPs from rats 1 and 2	48
3.7	Changes in gPDC during mPFC activated states	50
3.8	mPFC-hippocampus gPDC for LFPs from rats 3 and 4	51
3.9	mPFC-hippocampus gPDC for FHA-processed spike trains from rat 4 .	53
3.10	Schematic of information flow from gPDC analysis	54
4.1	Sleep stages and EEG patterns	60
4.2	Spindles in the cat thalamus	61
4.3	Slow oscillation in EEG signals and intracellular recordings	64
4.4	Corticothalamic circuit for the SO grouping of spindles and delta waves	66
4.5	The Bazhenov et al. model	67
4.6	The Compte et al. model	68

4.7	The Hill and Tononi model	70
4.8	Firing pattern of the pyramidal cell model in Compte et al. (2003) . . .	72
4.9	Firing pattern of the Wang-Buzsáki interneuron	72
4.10	Shape of $f(V_{pre})$ in Compte et al. (2003)	74
4.11	Synaptic variables in Compte et al. (2003)	75
4.12	Connectivity matrices in Compte et al. (2003)	78
4.13	Slow oscillations in the cortical network model	79
4.14	Average firing rates in the cortical model	80
4.15	Various ionic conductances in the cortical model	81
4.16	The cortical model with synaptic types blocked	82
5.1	Basic hippocampal circuitry	87
5.2	Firing patterns of CA3 pyramidal cells	90
5.3	f-I curve of the modified Pinsky-Rinzel model	92
5.4	Schematic diagram of the full CA3-CA1 model	97
5.5	Connectivity matrix of the pyramidal-to-pyramidal connections in CA3	98
5.6	Spread of activity in the CA3 model	101
5.7	Population bursts in the CA3 model	103
5.8	Gamma oscillations in the CA1 interneuronal network	105
5.9	Gamma oscillations in the CA1 model	107
6.1	CA1 recordings of an SWR	113
6.2	Depth profile of an SWR	115
6.3	Block diagram of the Traub et al. (2000) model	119
6.4	Histogram of the synapses-per-Schaffer connection on CA1 pyramidal cells	121
6.5	Ripple detection	123
6.6	Activity of the full CA3-CA1 model	124
6.7	Spectral analysis of the synaptic activity in the CA3-CA1 model	125
6.8	Ripple correlation with unit firing	126

6.9	Cell participation in ripples	128
6.10	Postsynaptic currents and average firing rates in ripples	129
6.11	Synchrony of ripples in CA1	130
6.12	Ripple dependence on CA3 bursts	131
6.13	CA1 power spectra after synaptic modifications	132
7.1	Correlation between cortical and hippocampal patterns during SWS . .	145
7.2	Relationships between SO, hippocampal firing and SWRs	147
7.3	LFPs from a rat mPFC/hippocampus and SO/SWR detection	149
7.4	Correlations of spikes and ripples with the average SO cycle	151
7.5	Schematic diagram of the modelled cortico-hippocampal connectivity . .	152
7.6	Raster plots of the PFC-CA3 model	157
7.7	CA3 firing during SO in the PFC-CA3 model	158
7.8	SO-correlated CA3 firing rates in the PFC-CA3 model	159
7.9	CA1 firing during SO in the PFC-CA3 model	160
7.10	Ripples during SO in the PFC-CA3 model	161
7.11	Raster plot of the PFC-CA1 model	163
7.12	CA1 firing rates for various TA inputs in the PFC-CA3-CA1 model . . .	164
7.13	Firing properties in the PFC-CA3-CA1 model	165
7.14	Ripple activity in the PFC-CA3-CA1 model	167
7.15	SO-correlated CA1 firing rates in the PFC-CA3-CA1 model	168
7.16	PFC raster plot depending on the CA1 feedback	169
7.17	Raster plot and ripple activity in the full PFC-CA3-CA1 model	170

List of Tables

5.1	Synaptic parameter values in the CA3 and CA1 models	95
5.2	Connectivity parameter values in the CA3 and CA1 models	98
6.1	Synaptic and connectivity parameter values for Schaffer collaterals . . .	120
7.1	Synaptic parameter values in the PFC-CA3-CA1 model	153

List of Abbreviations

AHP	after-hyperpolarisation
AIC	Akaike information criterion
AR	autoregressive model
BIC	Bayesian information criterion
CA	cornu ammonis
CTFT	continuous time Fourier transform
DC	directed coherence
dDTF	direct DTF
DG	dentate gyrus
DTF	directed transfer function
DTFT	discrete time Fourier transform
EC3	entorhinal layer 3
EPSP	excitatory postsynaptic potential
ffDTF	full-frequency DTF
FFT	fast Fourier transform
FHA	French-Holden algorithm
FPE	final prediction error
GCI	Granger causality index

LIST OF TABLES

gPDC	generalised PDC
IN	interneurons
IPSP	inhibitory postsynaptic potential
KA	kainic acid
LFP	local field potential
LIA	large-amplitude irregular activity
MF	mossy fibres
mini	miniature EPSP/IPSP
mPFC	medial prefrontal cortex
PC	principal component
PDC	partial directed coherence
PFC	prefrontal cortex
PSD	power spectral density
PSP	postsynaptic potential
PY	pyramidal cells
RE	reticular neurons
RMS	root mean square
RSS	residual sum of squares
SD	standard deviation
SE	status epilepticus
SO	slow oscillation
SW	sharp wave
SWR	sharp wave-ripple complex
SWS	slow wave sleep

LIST OF TABLES

TA	temporoammonic pathway
TC	thalamocortical neurons
VAR	vector autoregressive model

Introduction

One of the major questions in neuroscience is where and how memories are created, stored and retrieved. The field of memory consolidation has attracted numerous studies, based on different approaches, paradigms and hypotheses. Two of the main brain areas that are currently generally accepted as being involved in the consolidation of memories, are the neocortex (part of the cerebral cortex, involved in higher functions) and hippocampal areas (parts of the limbic system, involved in emotion, behavior, memory formation and other functions) [66]. The standard model for memory consolidation assumes that new memories are created through processes based on synaptic plasticity (changes in strength of neuronal connections) and are stored temporarily in the hippocampus and in closely connected medial temporal lobe regions. Various cognitive and physiological studies indicate that remote memories are stored in the neocortex [33]. The transferring of memories from hippocampus to cortex for long-term storage is generally thought to take place during deep stages of sleep.

One unresolved issue, directly related with memory transfer, is how the hippocampus communicates with the neocortex and how information flows between the two structures, particularly during deep sleep and anaesthesia. An influential suggestion has been that information travels from cortex to hippocampus during wakefulness and in the opposite direction during sleep [35], but a series of recent observations (reviewed in [209]) imply more complex interactions. These observations focus on two intrinsic rhythmic oscillations found in the cortex and hippocampus, and in their temporal and functional relations. The two rhythms are the slow oscillation, a low-frequency (<1 Hz) alteration between activated and deactivated states in the cortex and the so-called sharp wave-ripple complex in the hippocampus [34, 241]. The latter is a transient episode of synchronous fast oscillatory activity (~ 200 Hz), seen in field recordings of the hippocam-

pus, engaging the majority of hippocampal neurons. Ripples have been suggested to act as vehicles for information/memory transfer to the cortex, supported by the observation that temporal sequences between correlated hippocampal cells, acquired during wakefulness, are preserved during sleep [189, 235], replayed at faster timescales within the short window of ripples [129]. Most studies on temporal correlations between the cortical slow oscillation and hippocampal ripples indicate that the occurrence of ripples, during sleep and anaesthesia, is mainly timed by the activated states of the slow oscillation [104, 149, 188]. This suggests a general flow of information from cortex to hippocampus instead of the opposite direction. On the other hand, other studies report opposite correlations, at the level of pairs of single neurons, showing that hippocampal cell firing during ripples precedes spiking activity in cortical targets [234]. Finally, memory transfer from hippocampus to cortex is also supported by the replay of learning-associated temporal spiking patterns in the cortex during ripples [111, 164].

The above observations, along with many more, often contradictory studies (reviewed throughout this work), reveal that the issue of the temporal and functional cross-talk between hippocampus and cortex during sleep and anaesthesia is still open. In any case, though, rhythmic patterns, such as the slow oscillation and ripples and their temporal couplings, seem to be key elements involved in this communication and possibly in the ensuing memory consolidation [33]. Motivated by this issue of cortico-hippocampal functional connectivity, we embark here on a detailed study of functional and structural interactions between the neocortex and the hippocampus, during sleep and anaesthesia, along with their corresponding rhythmic patterns and intrinsic dynamics. We approach this interaction from two different perspectives: data analysis and computational modelling. The first part of the thesis is dedicated to the statistical analysis of electrophysiological data acquired from the two brain areas in question. We review some of the available statistical tools and apply them in an assessment of directional information flow between prefrontal cortex and hippocampus in anaesthetised rats. In the second part we focus on the intrinsic rhythmic dynamics of the two areas. We implement an established computational model of the cortical slow oscillation and develop a new model of hippocampal areas CA3 and CA1, which reproduces many of their intrinsic dynamics and exhibits sharp wave-ripple complexes, based on a novel mechanism. We finally couple the two computational models in an effort to study the interactions between cortical activated states and hippocampal ripples.

This thesis is split into the two aforementioned parts, with chapters 2 and 3 constituting

the data-analysis part, and chapters 4 to 7 constituting the computational modelling one. An outline of each chapter follows below:

Chapter 2

This chapter describes some basic principles of multivariate time series analysis and autoregressive modelling of random processes. The main mathematical tools and statistical measures, parametric and nonparametric, for assessing functional connectivity and directional interactions between time series are reviewed. We focus on the notion of Granger Causality and the related measures of Partial Directed Coherence (PDC) and generalised PDC. A critical value for significance detection is introduced for the latter. We also compare the effectiveness of some of these measures by checking whether they yield correct connectivity profiles in a set of tests.

Chapter 3

In this chapter, the previously described measures are put into practice in a neuroscientific framework. We first present methods for spike train preprocessing and their conversion to continuous signals. The effects of such preprocessing and the relevant problems arising from it are discussed. We next describe the basic components of a custom built numerical package we developed in MATLAB, named FunCAT, which performs autoregressive modelling and functional connectivity analysis in multivariate time series. This toolbox is used on a set of electrophysiological data, recorded by Drs. Robert Mason and Benjamin Coomber (University of Nottingham, Medical School). Local field potentials and unit spike data were recorded from anaesthetised rats, which are later administered kainic acid to evoke an increase in neuronal excitability. We apply the measure of generalised PDC, using FunCAT, in order to assess directionality of information flow between medial prefrontal cortex and hippocampus, as well as intrahippocampally [202]. We also examine the effects of kainic acid administration on this connectivity. Based on our results, combined with the established neuroanatomy of the corresponding areas, we propose a general cortico-hippocampal functional connectivity scheme.

Chapter 4

Entering into the computational part of this work, we focus on intrinsic rhythms of the two areas in question, during deep sleep and anaesthesia. We begin with a short review of

the main cortical sleep rhythms, focusing on the slow oscillation, its characteristics, and its relation with the other rhythms. We then review some of the computational models of the slow oscillation that have been developed and we implement one of them [47] in the Python neural networks simulator “Brian”, after introducing some minor modifications. We show that our modified version accurately reproduces the main characteristics and properties of the original model and the corresponding experimental data.

Chapter 5

We next shift our focus to the hippocampus, and give a short description of its basic circuitry and a review of some of the established computational models for hippocampal pyramidal cells and interneurons. We then describe in full detail the two computational models that were developed for the rat hippocampal areas CA3 and CA1, providing biological justification for our modelling assumptions. The two models are explored separately to validate that they provide accurate representations of some of the basic intrinsic hippocampal dynamics. We show that they both successfully reproduce a range of anatomical characteristics and/or functional properties of the corresponding hippocampal areas. Most importantly, the CA3 model exhibits periodic quasi-synchronised population bursts, while the CA1 model oscillates in the gamma frequency range, reproducing gamma activity features, seen in neurophysiological studies. The intrinsic dynamics of the two models are discussed in detail.

Chapter 6

This chapter begins with a review of the main characteristics and properties of sharp wave-ripple complexes and with a short description of a previous computational model for fast ripple-like oscillations, based on axo-axonal gap junctions between pyramidal cells. We then couple our two models in a biologically realistic CA3-to-CA1 feedforward scheme. The CA3 population bursts evoke intense transient oscillatory (~ 150 - 200 Hz) responses in the synaptic activity of CA1. These responses are explored in detail and are shown to accurately reproduce many of the basic characteristics of the sharp wave-ripple complex seen in extracellular recordings in CA1. By exploring these characteristics through the model, we unravel the main components that give rise to them and we are eventually led to the formulation of a novel mechanism for the generation of ripple episodes in CA1, based on interneuronal activity and strong fast-decaying recurrent inhibition. Finally, this mechanism is discussed in detail, in relation to various electro-

physiological studies. We argue that the model can interpret a series of observations on ripples and provides a link with memory consolidation modelling.

Chapter 7

In the final chapter, we initially review the main studies on temporal correlations between the slow oscillation and hippocampal activity, focusing on some established relations between cortical activated states and ripples. In an effort to validate some of the main findings of these studies, we analyse a small set of local field potential and unit spike data, recorded by Drs. Robert Mason and Yvonne Mbaki (University of Nottingham, Medical School) from the medial prefrontal cortex and hippocampus of a naturally sleeping rat. Having reproduced a computational model for cortical slow oscillations, and developed a new model for hippocampal sharp wave-ripple complexes, we finally couple them in a closed feedback loop in order to study cortico-hippocampal interactions in relation with these intrinsic rhythmic dynamics. The functional/anatomical relevance of the implemented connectivity scheme between the cortical network, CA3 and CA1, is described in detail. We show how the spiking activity of CA3 and CA1, individually and combined, depends on the inhibition-excitation balance, induced by the cortical oscillatory output, and how this can affect the observed correlations between cortical slow oscillation and hippocampal spiking in general, and between cortical activated states and ripples in particular. Our findings are discussed in relation with the reviewed studies on cortico-hippocampal interactions during deep stages of sleep and some links to memory consolidation are provided.

Chapter 8

Since each chapter provides a thorough discussion on its subject, we end this work by a short summary of all the main results presented. Also described are some possible future directions on the issue of cortico-hippocampal interactions, based on plausible extensions of our computational model.

Part of the work described in Chapters 2 and 3 was published in *Biological Cybernetics* [202]. Chapters 5 and 6 were based on work published in *Hippocampus* [203]. A final paper, describing the work presented in Chapter 7, is in preparation.

Part I

Assessing Cortico-Hippocampal Directed Interactions

Causal Interactions in Multivariate Autoregressive Processes

2.1 Introduction

A time series $x(t)$ is a set of observations with each one recorded at a time t . Discrete time series are those in which the observation times form a discrete set (e.g. when the observations are made at fixed time intervals), whereas in continuous series the observations are recorded continuously over time [25].

When studying time series, one's purpose is usually to draw inferences from them. But in order to do that, one first needs to set up a hypothetical model to represent the series. Then it becomes possible to estimate parameters involved, check for goodness of fit to data and potentially use the fitted model to draw some conclusions on the mechanism that generated the time series. A model of an observed data set $x(t)$ is a specification of the joint distributions of a sequence of random variables $X(t)$, of which $x(t)$ is supposed to be a realization [25]. A complete probabilistic model of such a random variable sequence would specify all the joint distributions. But since it would contain far too many parameters to be estimated from the available data, one usually focuses only on the first- and second-order moments of the joint distributions (means and covariances) and on properties of the sequences that depend only on those. Although some information is usually lost through this approach, their use is further justified by the fact that the usual methods for minimum squared error linear forecasting depend

only on second-order properties [25]. Autoregressive processes are typical examples of such models of only first- and second-order moments.

We present here an elementary introduction into some fundamental notions of multivariate time series analysis and autoregressive modelling of random processes. We also review some basic related elements of signal spectral analysis and present the main statistical measures used to detect structural or functional connectivities between multivariate time series. Our main focus is on the definition of the notion of Granger Causality and its related measure of Partial Directed Coherence (PDC). Furthermore, we make a comparison of these statistical tools based on a simple test case with known connectivity. Finally we implement a critical value for the recently-introduced measure of generalised PDC and we show the advantage of this measure over the original PDC.

The analysis presented here will be applied in a neuroscientific framework in Chapter 3. There, the multivariate time series will consist of electrophysiological data in the form of local field potentials and unit spike data, simultaneously recorded from cortical and hippocampal sites.

2.2 Multivariate Time Series

Since, apart from serial dependence within a time series, there can also exist interdependencies between the various components of a set of series, one often needs to study not just each of the components separately, but the whole set simultaneously. Such multivariate time series are usually represented by a vector containing the values of all series at each time point.

2.2.1 Basic Measures

Consider a set of m time series: $x_i(t)$, $i = 1, \dots, m$, denoted by:

$$\mathbf{X}(t) = \begin{bmatrix} x_1(t) \\ \vdots \\ x_m(t) \end{bmatrix}, \quad t = 0, \pm 1, \pm 2, \dots \quad (2.2.1)$$

The first- and second-order moments are the mean vector and the covariance matrix respectively, given by:

$$\boldsymbol{\mu}(t) = E[\mathbf{X}(t)] = \begin{bmatrix} \mu_1(t) \\ \vdots \\ \mu_m(t) \end{bmatrix}, \quad \boldsymbol{\Gamma}(t, h) = \begin{bmatrix} \gamma_{11}(t, h) & \dots & \gamma_{1m}(t, h) \\ \vdots & \ddots & \vdots \\ \gamma_{m1}(t, h) & \dots & \gamma_{mm}(t, h) \end{bmatrix}, \quad (2.2.2)$$

where $\gamma_{ij}(t, h) = \text{Cov}(x_i(t+h), x_j(t)) = E[(x_i(t+h) - \mu_i(t+h))(x_j(t) - \mu_j(t))]$ is the covariance function between series $x_i(t)$ and $x_j(t)$, measuring how much the two series vary together at time t and time lag h . It is usually called cross-covariance when it involves two different time series and auto-covariance when the covariance of a time series with itself is computed. When $h = 0$, the auto-covariance coincides with the definition of variance, $\sigma_i^2(t)$, which is a measure of statistical dispersion, since it averages the squared distance of the time series values from its mean.

Finally the cross-correlation matrix is defined as:

$$\mathbf{R}(t, h) = \begin{bmatrix} \rho_{11}(t, h) & \dots & \rho_{1m}(t, h) \\ \vdots & \ddots & \vdots \\ \rho_{m1}(t, h) & \dots & \rho_{mm}(t, h) \end{bmatrix}, \quad \rho_{ij}(t, h) = \frac{\gamma_{ij}(t, h)}{\sqrt{\gamma_{ii}(t, 0)\gamma_{jj}(t, 0)}} = \frac{\gamma_{ij}(t, h)}{\sigma_i(t)\sigma_j(t)}, \quad (2.2.3)$$

where $\rho_{ij}(t, h)$ is the cross-correlation between series $x_i(t)$ and $x_j(t)$ at time t and lag h .

The auto-covariance and auto-correlation functions are very useful measures of the degree of dependence between the values of a time series at different times. Hence, they can be very useful in choosing an appropriate model or predicting future values according to the present and past values of the series.

2.2.2 Stationarity

A time series $x(t)$ is defined as (weakly) stationary if:

1. $\mu(t)$ is independent of t and
2. $\gamma(t+h, t)$ is independent of t for any lag h .

Strict stationarity is defined by the condition that any set of values of the time series, e.g. $(x(1), \dots, x(n))$ has the same joint distributions as the set $(x(1+h), \dots, x(n+h))$ for any n and h (but we will only be dealing with the weak definition from now on).

For stationary multivariate time series we can refer to $\boldsymbol{\mu}(t) = \boldsymbol{\mu}$ as their mean vector and to $\boldsymbol{\Gamma}(t+h, t) = \boldsymbol{\Gamma}(h)$ as their covariance matrix at lag h . Similarly, the cross-correlation function between series $x_i(t)$ and $x_j(t)$ becomes:

$$\rho_{ij}(h) = \frac{\gamma_{ij}(h)}{\sqrt{\gamma_{ii}(0)\gamma_{jj}(0)}} = \frac{\gamma_{ij}(h)}{\sigma_i\sigma_j}, \quad (2.2.4)$$

implying that $\rho_{ii}(0) = 1$.

A classical example of a stationary time series is white noise, i.e. random numbers with zero mean, σ^2 variance, and covariance given by:

$$\gamma(t+h, t) = \begin{cases} \sigma^2, & \text{if } h = 0; \\ 0, & \text{if } h \neq 0. \end{cases} \quad (2.2.5)$$

The qualitative meaning of stationarity is that the properties of any one segment of the data are much like those of any other segment, i.e. there is no systematic change in the mean or variance, and any purely periodic oscillations have been removed. Note that, in a strict sense, there is very rarely stationarity in real data sets, as this property is mainly defined for time series generated by models. However, real data are often considered stationary, implying that they exhibit characteristics suggesting that a stationary model can sensibly be fitted [39].

2.3 Multivariate Autoregressive Models

Two basic purposes of multivariate time series analysis is forecasting and detection of interrelationships. Forecasting has to do with being able to predict future values of a variable which produced a time series, based on its present and past values. One has to choose a predictor function that takes as input these values and has future values as output.

When dealing with multivariate time series, the future value of any of them can sometimes be forecasted not only through its own present and past, but through those of the other series as well. This holds when a common mechanism generates these series, or generally when there is some kind of interdependence between them. Apart from forecasting, one may be interested in studying such possible dynamical interdependencies between the variables that generated the multivariate series. Through which mechanism are these variables interrelated? How would a specific change in one variable affect the others? How can we trace such a change through the other variables? These are

questions that usually arise in such studies and can often be addressed by modelling the time series by an appropriate model which can be analysed.

2.3.1 Autoregressive Processes

Consider a given time series $x(t)$ of which we want to forecast/model the value at a time step t^* . If we assume that only a linear combination of a finite number p of past values of the series is needed for the forecasted/modelled value $\hat{x}(t^*)$, then:

$$\hat{x}(t^*) = v + a_1x(t^* - 1) + a_2x(t^* - 2) + \dots + a_px(t^* - p), \quad (2.3.1)$$

where v is a constant that allows for a nonzero mean. Since $\hat{x}(t^*)$ will usually not equal the real value $x(t^*)$, we denote the error, or residual, by $u(t^*) = x(t^*) - \hat{x}(t^*)$ and we can write:

$$\begin{aligned} x(t^*) &= \hat{x}(t^*) + u(t^*) \\ &= v + a_1x(t^* - 1) + a_2x(t^* - 2) + \dots + a_px(t^* - p) + u(t^*). \end{aligned} \quad (2.3.2)$$

Assuming that the time series $x(t)$ is a realization of a random variable $X(t)$ and that the same data generation mechanism holds for every time t , then the previous equation can be written in the form of a so-called autoregressive process of order p (AR_p):

$$X(t) = v + a_1X(t - 1) + a_2X(t - 2) + \dots + a_pX(t - p) + U(t), \quad (2.3.3)$$

where the quantities $X(t), X(t - 1), \dots, X(t - p)$ and $U(t)$ are random variables. Furthermore, we assume that the residuals $U(t)$ for different time points are uncorrelated, implying that all useful information in the p past values of $X(t)$ is used in the forecast/model, so that there are no systematic errors.

Now assume that we have a multivariate time series $\mathbf{X}(t)$ of size m , where each value of each series depends not only on the p past values of itself, but also on those of the other series. This vector autoregressive process of order p (VAR_p) can be written similarly to eq. (2.3.3):

$$\mathbf{X}(t) = \mathbf{V} + \mathbf{A}_1\mathbf{X}(t - 1) + \mathbf{A}_2\mathbf{X}(t - 2) + \dots + \mathbf{A}_p\mathbf{X}(t - p) + \mathbf{U}(t), \quad (2.3.4)$$

where \mathbf{V} is now a $m \times 1$ fixed vector and the \mathbf{A}_i 's are fixed $m \times m$ coefficient matrices:

$$\mathbf{A}_i = \begin{bmatrix} a_{11,i} & \dots & a_{1m,i} \\ \vdots & \ddots & \vdots \\ a_{m1,i} & \dots & a_{mm,i} \end{bmatrix}. \quad (2.3.5)$$

All the $\mathbf{X}(t)$'s are considered $m \times 1$ random vectors and the error vectors $\mathbf{U}(t) = (u_1(t), \dots, u_m(t))^T$ form a sequence of independent identically distributed random m -vectors with zero means, namely a white noise vector, with some fixed covariance matrix $\Sigma = E[\mathbf{U}(t)\mathbf{U}(t)^T]$ at time lag 0.

There are various standard mathematical techniques for computing, from a finite (observed) set of data points, estimates $\hat{\mathbf{A}}_i$ and $\hat{\Sigma}$ of the coefficients \mathbf{A}_i and the error covariance matrix Σ respectively of the VAR process that is assumed to have generated the data. The most widely used ones include the Yule-Walker estimators, the Maximum Likelihood Estimators, the ARFIT algorithm and others. Descriptions of these estimators can be found in [25, 39, 132, 138, 152].

2.3.2 Stability Criterion

Consider the VAR₁ model:

$$\mathbf{X}(t) = \mathbf{V} + \mathbf{A}_1\mathbf{X}(t-1) + \mathbf{U}(t). \quad (2.3.6)$$

If the generating mechanism of this process started at some initial time $t = 1$, we get:

$$\begin{aligned} \mathbf{X}(1) &= \mathbf{V} + \mathbf{A}_1\mathbf{X}(0) + \mathbf{U}(1), \\ \mathbf{X}(2) &= \mathbf{V} + \mathbf{A}_1\mathbf{X}(1) + \mathbf{U}(2) = \mathbf{V} + \mathbf{A}_1(\mathbf{V} + \mathbf{A}_1\mathbf{X}(0) + \mathbf{U}(1)) + \mathbf{U}(2) \\ &= (\mathbf{I} + \mathbf{A}_1)\mathbf{V} + \mathbf{A}_1^2\mathbf{X}(0) + \mathbf{A}_1\mathbf{U}(1) + \mathbf{U}(2) \\ &\vdots \\ \mathbf{X}(t) &= (\mathbf{I} + \mathbf{A}_1 + \dots + \mathbf{A}_1^{t-1})\mathbf{V} + \mathbf{A}_1^t\mathbf{X}(0) + \sum_{i=0}^{t-1} \mathbf{A}_1^i\mathbf{U}(t-i). \end{aligned} \quad (2.3.7)$$

Hence the vectors $\mathbf{X}(1), \dots, \mathbf{X}(t)$ are uniquely determined by the initial condition $\mathbf{X}(0)$ and the residuals $\mathbf{U}(1), \dots, \mathbf{U}(t)$.

Now assume that instead of an initial starting time, the process started in the infinite past, an assumption implicitly made in the definition (2.3.4) of VAR_p. If we rewrite eq. (2.3.7) in its general form:

$$\mathbf{X}(t) = (\mathbf{I} + \mathbf{A}_1 + \dots + \mathbf{A}_1^j)\mathbf{V} + \mathbf{A}_1^{j+1}\mathbf{X}(t-j-1) + \sum_{i=0}^j \mathbf{A}_1^i\mathbf{U}(t-i), \quad (2.3.8)$$

and assume that all eigenvalues of \mathbf{A}_1 have magnitude less than 1, then as j goes to infinity, the first right-hand side term will converge to:

$$\lim_{j \rightarrow \infty} (\mathbf{I} + \mathbf{A}_1 + \dots + \mathbf{A}_1^j)\mathbf{V} = (\mathbf{I} - \mathbf{A}_1)^{-1}\mathbf{V} \quad (2.3.9)$$

and the second term can be ignored in the same limit since \mathbf{A}_1^{j+1} will converge to zero rapidly. In the third term, the sequence \mathbf{A}_1^i will be absolutely summable and consequently the third term exists in mean square when $j \rightarrow \infty$ [132]. Hence, when the restriction for the eigenvalues of \mathbf{A}_1 holds, $\mathbf{X}(t)$ is a well-defined stochastic process which can be written in the form:

$$\mathbf{X}(t) = \boldsymbol{\mu} + \sum_{i=0}^{\infty} \mathbf{A}_1^i \mathbf{U}(t-i), \quad t = 0, \pm 1, \pm 2, \dots \quad (2.3.10)$$

where $\boldsymbol{\mu} \equiv (\mathbf{I} - \mathbf{A}_1)^{-1} \mathbf{V}$ is the mean value of the variable $\mathbf{X}(t)$. The distributions and joint distributions of $\mathbf{X}(t)$ at any time t^* are uniquely determined by the distributions of the stochastic variables $\{\mathbf{U}(t) : t \leq t^*\}$. The covariance matrix of $\mathbf{X}(t)$ can be shown to be given by [132]:

$$\boldsymbol{\Gamma}(h) = \sum_{i=0}^{\infty} \mathbf{A}_1^{h+i} \boldsymbol{\Sigma} \mathbf{A}_1^{i^T}. \quad (2.3.11)$$

The condition for the magnitude of the eigenvalues of \mathbf{A}_1 to be less than 1, defines a VAR_1 process as stable. Specifically, the stability criterion for a VAR_1 process is:

$$\det(\mathbf{I} - \mathbf{A}_1 z) \neq 0 \quad \text{for } |z| \leq 1 \quad (2.3.12)$$

This can easily be extended to a general VAR_p process with $p > 1$. This is because every VAR_p process $\mathbf{X}(t)$ can be written in a VAR_1 form, as:

$$\mathbf{Y}(t) = \tilde{\mathbf{V}} + \tilde{\mathbf{A}} \mathbf{Y}(t-1) + \tilde{\mathbf{U}}(t) \quad (2.3.13)$$

with:

$$\mathbf{Y}(t) = \begin{bmatrix} \mathbf{X}(t) \\ \mathbf{X}(t-1) \\ \vdots \\ \mathbf{X}(t-p+1) \end{bmatrix}, \quad \tilde{\mathbf{V}} = \begin{bmatrix} \mathbf{V} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \quad (2.3.14)$$

$$\tilde{\mathbf{A}} = \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ \mathbf{I} & 0 & \dots & 0 & 0 \\ 0 & \mathbf{I} & & 0 & 0 \\ \vdots & & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & \mathbf{I} & 0 \end{bmatrix}, \quad \tilde{\mathbf{U}}(t) = \begin{bmatrix} \mathbf{U}(t) \\ 0 \\ \vdots \\ 0 \end{bmatrix}. \quad (2.3.15)$$

So similarly to the VAR_1 case, a VAR_p model is considered stable if all the eigenvalues of the matrix $\tilde{\mathbf{A}}$ lie within the unit circle. Or equivalently if the general stability criterion

is satisfied:

$$\det(\mathbf{I} - \tilde{\mathbf{A}} z) = \det(\mathbf{I} - \mathbf{A}_1 z - \dots - \mathbf{A}_p z^p) \neq 0, \quad \text{for any } |z| \leq 1. \quad (2.3.16)$$

Then the VAR_p model can be written in the form (2.3.10), with:

$$\boldsymbol{\mu} = (\mathbf{I} - \tilde{\mathbf{A}})^{-1} \tilde{\mathbf{V}} \quad \text{and} \quad \boldsymbol{\Gamma}(h) = \sum_{i=0}^{\infty} \tilde{\mathbf{A}}^{h+i} \boldsymbol{\Sigma}_{\tilde{\mathbf{U}}} \tilde{\mathbf{A}}^{iT}, \quad (2.3.17)$$

where $\boldsymbol{\Sigma}_{\tilde{\mathbf{U}}}$ is the sparse $mp \times mp$ matrix:

$$\boldsymbol{\Sigma}_{\tilde{\mathbf{U}}} = E[\tilde{\mathbf{U}}(t) \tilde{\mathbf{U}}(t)^T] = \begin{bmatrix} \boldsymbol{\Sigma} & 0 & \dots & 0 \\ 0 & \ddots & & \vdots \\ \vdots & & \ddots & \vdots \\ 0 & \dots & \dots & 0 \end{bmatrix}. \quad (2.3.18)$$

Qualitatively when a VAR model is stable, it generates stationary time series, whereas unstable VAR processes can generate non-stationary time series (but not necessarily), that have trends or exhibit pronounced seasonal fluctuations.

2.3.3 Choosing the Optimal VAR Model order

One of the key issues in the creation of a VAR process that models a given time series as accurately as possible, is the choice of the model order p . If it is too low, the model will not capture the essential dynamics of the data set, whereas if it is too high, it will also capture an unwanted noisy component, leading to over-fitting and instability [161]. There exist mathematical information criteria that can detect the optimal model order, i.e. the order which will give a model that best fits the data with the least possible complexity. These criteria are functions that increase with the error covariance matrix of an estimated VAR model and the number of free model parameters. In other words they intend to strike a balance between the goodness of fit and the model complexity.

Amongst the most used criteria are the Akaike Information Criterion (AIC) [3], the Bayesian Information Criterion (BIC) [183], the Final Prediction Error (FPE) [2] and the Hannan-Quinn Criterion [95]. The former two are defined as:

$$\text{AIC} = \ln\left(\frac{\text{RSS}}{N}\right) + \frac{2k}{N}, \quad (2.3.19)$$

$$\text{BIC} = \ln\left(\frac{\text{RSS}}{N}\right) + \frac{k \ln(N)}{N}, \quad (2.3.20)$$

where RSS is the VAR model's residual sum of squares, k is the number of model parameters and N is the data length. Both AIC and BIC criteria have the general form: $\ln(\text{RSS}/N) + kf(N)/N$, with f some function of N , and they increase with RSS and k , penalizing both a bad fit and the redundant use of parameters. Hence they acquire their minimum value at the optimal model order. In a multivariate case of m time series $x_{1\dots m}(t)$, RSS is an $m \times m$ matrix, whose element $\{i, j\}$ is given by:

$$\text{RSS}_{ij} = \sum_{t=1}^N (x_i(t) - \hat{x}_i(t))(x_j(t) - \hat{x}_j(t)) = \sum_{t=1}^N u_i(t)u_j(t) \quad (2.3.21)$$

and can be approximated by $N\hat{\Sigma}_{ij}$ of the error covariance matrix of the VAR model, estimated by the sampled values of the time series. If the model order is p , the number of parameters will be: $k = m^2p$, since each of the p coefficients A_i has m^2 elements. Then the two criteria can be written as:

$$\text{AIC} = \ln(\det(\hat{\Sigma})) + \frac{2m^2p}{N}, \quad (2.3.22)$$

$$\text{BIC} = \ln(\det(\hat{\Sigma})) + \frac{m^2p \ln(N)}{N}, \quad (2.3.23)$$

where the determinant is taken to turn the criteria into scalars [39, 132]. Consequently, to estimate the optimal p_{opt} from a range of model orders, one needs to compute the error covariance matrix $\hat{\Sigma}$ for all orders and find that for which the chosen criterion is minimised.

This implies that a VAR model has to be constructed for each order before the application of the criterion, which is rather unpractical since the construction of VAR models for a large range of p is computationally demanding. Instead, it would be more practical, if $\hat{\Sigma}$ for a specific model order could be estimated, without having to construct the full model first. A technique that gives such an estimation is the downdating algorithm of $\hat{\Sigma}$ [152], through which $\hat{\Sigma}$ is computed only for the highest model order and estimated for all lower orders, without constructing the full VAR model.

Finally, the logarithm of the FPE criterion is defined as:

$$\text{FPE} = \ln(\det(\hat{\Sigma})) + m \ln\left(\frac{N + mp}{N - mp}\right) \quad (2.3.24)$$

and works similarly as the other two criteria [132, 152]. Both FPE and BIC criteria are applied, through the downdating algorithm, in the following analysis in this chapter and the next.

2.4 Spectral Analysis of Time Series

2.4.1 Power Spectrum and Nyquist Theorem

Assume $x(t)$ is a univariate stationary time series with zero mean and its value is measured at fixed time intervals T . Its Power Spectral Density (PSD) is defined as the Discrete Time Fourier Function (DTFT) of the signal's covariance function:

$$P(f) = \sum_{h=-\infty}^{\infty} \gamma(h) e^{-\frac{2\pi}{T} i f h}. \quad (2.4.1)$$

$P(f)$ is actually a power spectrum, as it represents the distribution of the (average) power of the signal over the frequency domain. In other words: $P(f)df$ is the infinitesimal power of the signal in the frequency band $\{f - df/2, f + df/2\}$. $P(f)$ is a periodic function with a period of $f^* = 1/T$: $P(f) = P(f + 1/T)$, and can thus be described in the frequency interval $f \in \{-1/2T, 1/2T\}$ symmetrically around the origin $f = 0$.

Often, a discrete time series is derived by sampling a continuous signal with some sampling rate f_S . One issue is whether the power spectrum of the sampled signal will be identical to that of its continuous origin. If a continuous signal $x(t)$ is bandlimited with maximum frequency f_0 , its Continuous Time Fourier Transform (CTFT) will be zero outside the frequency range: $\{-f_0, f_0\}$. The f_S sampling process, is expressed by the multiplication of the signal with a sampling function \mathcal{S} :

$$x_S(t) = x(t) \cdot \mathcal{S} = \frac{1}{f_S} \sum_{k=0}^{\infty} x\left(\frac{k}{f_S}\right) \delta\left(t - \frac{k}{f_S}\right), \quad (2.4.2)$$

where $x_S(t)$ is the sampled signal. The CTFT of $x_S(t)$ will be given by the convolution of the transforms of $x(t)$ and \mathcal{S} :

$$X_S(f) = X(f) \circ F(\mathcal{S}) = \sum_{k=0}^{\infty} X(f + k f_S). \quad (2.4.3)$$

This convolution will replicate the Fourier Transform $X(f)$ of $x(t)$ at frequency intervals equal to f_S , i.e. $X_S(f)$ will be an f_S -periodic replication of $X(f)$ [138]. It is straightforward that if the sampling rate is selected too low, such that: $f_S < 2f_0$, then the adjacent replicated transforms, will overlap with each other, a phenomenon called aliasing. The Nyquist Theorem states that to avoid aliasing, the sampling frequency must be higher than twice the signal's maximum frequency: $f_S > 2f_0$. Stated differently, if f_0 is not known but the signal is still sampled at some rate f_S , then the power spectrum of the sampled signal, will be alias-free up until the Nyquist frequency $f_S/2$.

2.4.2 Spectral Estimation of a VAR Model

The spectral estimation problem is how to find an estimate $\hat{P}(f)$ of the power spectrum $P(f)$ of a stationary stochastic variable, from a finite sample generated by that variable [25]. This can be quite difficult, mainly because usually the sample size N is small, either because of a prohibitive cost for obtaining large number of data or because the sample has to be small to ensure that it retains stationarity.

There are two general approaches, nonparametric methods, based on the PSD definition and its properties, and parametric methods, based on modelling the signal. Two most common nonparametric spectral estimators are the correlogram and the periodogram [138, 197]. Methods for computing the latter include the Fast Fourier Transform (FFT) algorithm [166], the Bartlett method [13] and the Welch method [229].

Parametric methods are those where one assumes the signal is derived from some generating model, a VAR model in this case, then estimates the parameters of this model and finally derives the power spectrum through these parameters. When the assumed model is a very good approximation of the real signal, parametric methods provide more accurate PSD estimations than nonparametric ones.

Assume that from an m -variate time series with zero mean, we have constructed a stable VAR $_p$ model of the form (2.3.4). It can be Fourier transformed, yielding:

$$\mathbf{X}(f) = \mathbf{A}(f)\mathbf{X}(f) + \mathbf{U}(f), \quad (2.4.4)$$

with $\mathbf{X}(f)$, $\mathbf{U}(f)$ and $\mathbf{A}(f)$ the DTFTs of the signal, the residuals and the VAR coefficients respectively. $\mathbf{A}(f)$ is given by:

$$\mathbf{A}(f) = \sum_{r=1}^p \mathbf{A}_r e^{-2\pi i f r}. \quad (2.4.5)$$

Since the model is stable, $\mathbf{A}(f)$ is invertible and thus we can rewrite the Fourier transformation as: $\mathbf{X}(f) = \mathbf{H}(f)\mathbf{U}(f)$, with:

$$\mathbf{H}(f) = \bar{\mathbf{A}}(f)^{-1} \equiv (\mathbf{I} - \mathbf{A}(f))^{-1} \quad (2.4.6)$$

the transfer function matrix, a matrix of linear filters in the frequency domain [114]. It can be shown [11, 138] that the power spectral matrix $\mathbf{S}(f)$ for the VAR model is given by:

$$\mathbf{S}(f) = \mathbf{H}(f)\mathbf{\Sigma}\mathbf{H}^H(f), \quad (2.4.7)$$

where the superscript H signifies the Hermitian transposition. $\mathbf{S}(f)$ is an $m \times m$ matrix whose diagonal elements are the power spectra estimates of the individual signals:

$S_{ii}(f) = P_i(f)$ and the off-diagonal ones are estimates of the cross-spectral densities between pairs of signals, defined as:

$$S_{ij}(f) = \sum_{h=-\infty}^{\infty} \gamma_{ij}(h) e^{-2\pi i f h}, \quad (2.4.8)$$

assuming $T=1$. This is the same definition as (2.4.1) for the power spectrum, with the auto-covariance substituted by the cross-covariance between the two signals, estimated through their sampled data.

2.5 Structural Analysis of Multivariate Time Series

When one wants to investigate whether two or more signals have some structural or causal relations one needs to analyse them as time series and use statistical tools. Some of the following tools have been used in various scientific fields and have recently proven to be very useful in neuroscience as well. Others have been developed specifically for the need to provide ways to detect and characterize functional relationships between neural structures.

2.5.1 Correlation

Correlation measures the linear relation as a function of a time lag h , within one signal (auto-correlation) or between a pair of signals (cross-correlation). The latter may reflect a causal relationship since it is intuitive that if one signal “causes” another, there should be a delay from the first signal to the second.

As shown in (2.2.3), correlation between time series $x_i(t)$ and $x_j(t)$ at lag h is defined as:

$$\rho_{ij}(h) = \frac{\gamma_{ij}(h)}{\sqrt{\gamma_{ii}(0)\gamma_{jj}(0)}} = \frac{\gamma_{ij}(h)}{\sigma_i \sigma_j}. \quad (2.5.1)$$

This function ranges from +1, showing complete direct correlation, to -1 showing complete anticorrelation, with 0 denoting lack of linear interdependence between the signals, at time lag h . Generally a positive correlation, indicates that both signals have a tendency to have similar values with the same sign, while a negative one shows a tendency to have similar values with opposite signs at this lag. If one implicitly assumes that the two signals are linearly interrelated, then the time lag for which ρ_{ij} is maximised, is an estimation of the time delay between them.

2.5.2 Coherence

Coherence expresses the degree of relative synchrony between two series as a function of frequency. The coherence between series $x_i(t)$ and $x_j(t)$ is their cross-spectral density function (2.4.8), normalised by their individual auto-spectral densities:

$$C_{ij}(f) = \frac{|S_{ij}(f)|^2}{|S_{ii}(f)||S_{jj}(f)|}, \quad (2.5.2)$$

where the spectral densities can be parametrically estimated through eq. (2.4.7) in a VAR model. Coherence is actually a Fourier transformed version of the correlation function and thus expresses linear associations. It ranges from 0, which indicates independence of the signals at a specific frequency, to 1, indicating complete synchrony. Note that this measure is sensitive to both phase and amplitude relationships between the signals. Therefore the relative importance of these two factors cannot be distinguished with this measure and other techniques need to be applied to focus on only one of them [93].

2.5.3 Directed Coherence

We now move to a class of exclusively parametric statistical tools that require a model to be constructed first so that they can be applied on the modelled data. A characteristic of the previous statistical tools is that, although they may reveal aspects like synchronous activity etc., they yield no information on the direction of the dependencies, i.e. the direction of information flow between the structures. In order to be able to draw conclusions on the directionality of functional relations, one needs to use other statistical techniques to analyse the time series of the signals. We review some of the most extensively applied ones here, but for further reading and comparative analysis we refer to [9, 85, 125, 161, 236] and references therein.

In order to replace simple coherence with a statistical measure that would uncover the directionality of the information flow apart from just the relative synchrony of pairs of structures, the measure of directed coherence (DC) was introduced, based on information theory arguments [172]. DC is based on a unique decomposition of the ordinary coherence, into two directional components, one representing the feedforward and the other the feedback aspect of the functional interaction between the two structures. DC is applied on a modified bivariate VAR_p model where a common noise source s has been

added to allow the two model residuals to be completely uncorrelated:

$$\begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} = \sum_{r=1}^p \mathbf{A}_r \begin{bmatrix} x_1(t-r) \\ x_2(t-r) \end{bmatrix} + \begin{bmatrix} b_{11} & b_{1s} & 0 \\ 0 & b_{2s} & b_{22} \end{bmatrix} \begin{bmatrix} u_1(t) \\ u_s(t) \\ u_2(t) \end{bmatrix}, \quad (2.5.3)$$

where b_{ij} are weight factors and $u_s(t)$ is the common noise source [85]. This yields an error covariance matrix $\mathbf{\Sigma}$ that is diagonal. The DTFT of this model is similar to the normal VAR_p DTFT:

$$\begin{bmatrix} X_1(f) \\ X_2(f) \end{bmatrix} = \mathbf{H}(f) \begin{bmatrix} b_{11} & b_{1s} & 0 \\ 0 & b_{2s} & b_{22} \end{bmatrix} \mathbf{U}(f) = \begin{bmatrix} \tilde{H}_{11}(f) & \tilde{H}_{1s}(f) & \tilde{H}_{12}(f) \\ \tilde{H}_{21}(f) & \tilde{H}_{2s}(f) & \tilde{H}_{22}(f) \end{bmatrix} \mathbf{U}(f), \quad (2.5.4)$$

with $\mathbf{H}(f)$ the transfer function matrix (2.4.6). The DC estimate of the linear interaction from x_1 to x_2 is defined on this model as:

$$DC_{2 \leftarrow 1} = \frac{\tilde{H}_{21}(f)}{\sqrt{|\tilde{H}_{22}(f)|^2 + |\tilde{H}_{2s}(f)|^2 + |\tilde{H}_{21}(f)|^2}}. \quad (2.5.5)$$

The generalization of this technique to multivariate signals ($m > 2$) rests on the DTFT of a VAR_p model to get the spectral density matrix as in (2.4.7). Through this representation of $\mathbf{S}(f)$, the definition of DC from any $x_j(t)$ to any $x_i(t)$ in the system, is given by [10]:

$$DC_{i \leftarrow j}(f) = \frac{\sigma_{jj} H_{ij}(f)}{\sqrt{S_{ii}(f)}}. \quad (2.5.6)$$

When the covariance matrix $\mathbf{\Sigma}$ of the VAR model is diagonal, $DC_{i \leftarrow j}(f)$ can be interpreted as the fraction of the power contributing to the total power in x_i that originates from x_j at some frequency f [11]. Thus DC ranks the interaction strength with respect to the interaction “target”. In this case, it can be shown that the ordinary coherence can be expressed as:

$$C_{ij}(f) = \beta_i^H(f) \beta_j(f), \quad (2.5.7)$$

where $\beta_i(f) = [DC_{i \leftarrow 1}(f) \dots DC_{i \leftarrow m}(f)]^T$, justifying the argument that DC decomposes ordinary coherence into an incoming and an outgoing directed-interaction component [11].

2.5.4 Directed Transfer Function

The directed transfer function (DTF) is a variation of DC, applicable to multivariate time series instead of just bivariate and surpassing the restriction of diagonal $\mathbf{\Sigma}$ [115].

It actually coincides with the magnitude of DC when all the diagonal elements of $\mathbf{\Sigma}$ are set to 1 and is defined as:

$$DTF_{i \leftarrow j}(f) = \frac{H_{ij}(f)}{\sqrt{\sum_{j=1}^m |H_{ij}(f)|^2}}. \quad (2.5.8)$$

Note that the sum in the denominator runs over all possible “sources” x_j implying that $DTF_{i \leftarrow j}(f)$ compares the inflow of information to signal x_i from signal x_j with the inflow of information from all other signals. So a high value of DTF, close to 1, would imply that most of the information inflow to x_i is from source x_j .

It has been shown, both theoretically and through numerical tests, that DTF succeeds in detecting direct interactions but can also pick up indirect ones [9, 11, 125, 236]. This means that in some cases $DTF_{i \leftarrow j}(f)$ may be nonzero even though x_j is sending information to x_i through some third signal x_k , with x_k incorporated in the VAR model as well. For this reason, a variation of DTF was introduced, the direct-DTF (dDTF) [123]. This measure is defined by multiplying a full-frequency DTF (ffDTF) given by:

$$\eta_{ij}^2(f) = \frac{|H_{ij}(f)|^2}{\sum_f \sum_{j=1}^m |H_{ij}(f)|^2}, \quad (2.5.9)$$

by a variation of coherence termed partial coherence:

$$\kappa_{ij}^2(f) = \frac{M_{ij}^2(f)}{M_{ii}(f)M_{jj}(f)}, \quad (2.5.10)$$

where $M_{ij}(f)$ is the minor of the spectral density matrix $\mathbf{S}(f)$, obtained by removing its i -th row and j -th column. So finally dDTF is given by:

$$dDTF_{i \leftarrow j} = \kappa_{ij}(f)\eta_{ij}(f). \quad (2.5.11)$$

Although it has been argued that dDTF describes only the direct interactions between multivariate time series [123, 125], this measure has not been studied in detail yet, nor extensively applied.

For the critical value for DTF (the value above which DTF is considered significantly non-zero), an empirical distribution technique, using surrogate data from each time series, can be applied [114]. An analytical approximation of the asymptotic statistics of DTF and a corresponding critical value were proposed in [67]. The steps followed there are practically identical to those for the derivation of a critical value for Partial Directed Coherence (see section 2.5.7) [178].

2.5.5 Granger Causality

The notion of Granger causality, introduced by Clive Granger in a seminal paper [86], was initially applied in economical studies but soon was adopted by other scientific fields and recently by neuroscience as well, since it gives new insights into the question of directed interaction between neurons, neural ensembles or even brain regions.

The essence of Granger Causality is based on the idea that the cause precedes its effect. So if a variable X_j “causes” another one X_i , then the former should help improve the forecast or modelling of the latter. Specifically Granger’s rationale can be summarised as: A variable $X_j(t)$ “Granger-causes” another one $X_i(t)$, if knowledge of $X_j(t)$ ’s past values significantly improves predictions of $X_i(t)$ ’s future values. For a more mathematically precise definition of Granger Causality, based on information theory, see [132]. A sampled time series $x_j(t)$ Granger-causes another $x_i(t)$ when the inclusion of past values of x_j : $\{x_j(t-1), x_j(t-2), \dots\}$ in the derivation of a forecast/model for the next value of x_i : $\hat{x}_i(t)$, makes this forecast/model more accurate. This relation is not reciprocal, meaning that $x_j(t)$ can Granger-cause $x_i(t)$, without $x_i(t)$ necessarily causing $x_j(t)$, and this lack of reciprocity is actually what permits gauging the direction of information flow between variables. One could extend the concept to the forecast of values further away in the future: $x_i(t+1), \dots, x_i(t+h)$ but we will not be dealing with this “multi-step causality”.

Granger causality can be assessed by modelling a multivariate time series by a VAR model. Assume we have, for example, a bivariate stationary time series: $\mathbf{X}(t) = (x_1(t), x_2(t))^T$, with zero mean, and its stable VAR_p model:

$$\begin{aligned} \begin{bmatrix} x_1(t) \\ x_2(t) \end{bmatrix} &= \begin{bmatrix} A_{11,1} & A_{12,1} \\ A_{21,1} & A_{22,1} \end{bmatrix} \begin{bmatrix} x_1(t-1) \\ x_2(t-1) \end{bmatrix} + \dots \\ &+ \begin{bmatrix} A_{11,p} & A_{12,p} \\ A_{21,p} & A_{22,p} \end{bmatrix} \begin{bmatrix} x_1(t-p) \\ x_2(t-p) \end{bmatrix} + \begin{bmatrix} u_1(t) \\ u_2(t) \end{bmatrix}. \end{aligned} \quad (2.5.12)$$

The effect of past values of one series on the forecast of the other is represented through the VAR coefficients. Hence a null hypothesis of no Granger causality from x_2 to x_1 holds if all coefficients $A_{12,i}$ are zero:

$$A_{12,r} = 0 \text{ for all } r = 1, 2, \dots, p. \quad (2.5.13)$$

If any entry $A_{12,r}$ is nonzero, e.g. only the entry $A_{12,1}$, then x_2 is causing x_1 since its value at time $t-1$, is used in improving the modelling of $x_1(t)$. The same analysis can

be extended to multivariate time series with higher dimension. If it can be shown that $A_{ij,r}$ is significantly nonzero for any $r = \{1, 2, \dots, p\}$ in a VAR_p model, then the null hypothesis of noncausality from x_j to x_i can be rejected. Hence all statistical Granger-causality tests are based on testing the null hypothesis of zero VAR coefficients (see [81] for a review on such tests).

A variation of Granger causality is the instantaneous causality where one also includes the effect of x_j at time t , which is instantaneous to the forecast. In a VAR model, the instantaneous effect of $x_j(t)$ on the modelling of $x_i(t)$ is represented through the cross-covariance of their residuals, Σ_{ij} . It can be shown that there is no instantaneous causality from x_j to x_i if and only if: $\Sigma_{ij} = 0$ [132]. Contrary to normal Granger causality, the instantaneous causality is reciprocal, since $\Sigma_{ij} = \Sigma_{ji}$, and hence cannot be related to any notion of directionality of interactions.

2.5.6 Granger Causality Index

The “strength” of Granger causality between two time series can be quantified through the Granger Causality Index (GCI). Assume we have two zero-mean time series $x_1(t)$ and $x_2(t)$. We can model each of them separately by an AR_p model in which case, only each series’s own past values are considered in its model. Alternatively we can model them in parallel through a VAR_p . The quality of each model can be assessed through the covariance of its residuals.

Let us denote the two error covariances of the AR models by $\Sigma_{1|1}$ and $\Sigma_{2|2}$ and those of the VAR model by $\Sigma_{1|1,2}$ and $\Sigma_{2|1,2}$, where the subscript $1|1$ implies that only past values of x_1 were considered for its modelling and $1|1,2$ implies that past values of x_2 were also considered. The linear GCI from x_2 to x_1 is defined as:

$$\text{GCI}_{1 \leftarrow 2} = \ln \left(\frac{\Sigma_{1|1}}{\Sigma_{1|1,2}} \right) \quad (2.5.14)$$

and is a time-domain, bivariate measure of structural interactions [80, 85, 98, 161]. If x_2 is causing x_1 then by including it in the autoregressive model of x_1 it will improve its modelling, which is translated into the reduction of the corresponding error covariance. In other words, $\Sigma_{1|1,2}$ should be smaller than $\Sigma_{1|1}$ and thus $\text{GCI}_{1 \leftarrow 2}$ should be large. A critical value for GCI is given by Fisher’s Test [85], which under the null-hypothesis of non-causality yields:

$$F_{1 \leftarrow 2}^{\text{GCI}} = \frac{\Sigma_{1|1} - \Sigma_{1|1,2}}{\Sigma_{1|1,2}}. \quad (2.5.15)$$

Extending this analysis to an m -variate case with $m > 2$, the GCI and its critical value are defined as:

$$\text{GCI}_{i \leftarrow j} = \ln \left(\frac{\Sigma_{i|m-j}}{\Sigma_{i|m}} \right), \quad \text{and} \quad F_{i \leftarrow j}^{\text{GCI}} = \frac{\Sigma_{i|m-j} - \Sigma_{i|m}}{\Sigma_{i|m}} \quad (2.5.16)$$

respectively, where $\Sigma_{i|m}$ is the covariance of $u_i(t)$ when all m series have been included in the VAR model, whereas $\Sigma_{i|m-j}$ is the error covariance in a VAR model where x_j was omitted [80, 85, 98, 161].

Finally, a frequency domain equivalent of the bivariate GCI has been proposed [80]:

$$I_{1 \leftarrow 2} = -\ln \left(1 - \frac{(\Sigma_{22} - \frac{\Sigma_{12}^2}{\Sigma_{11}})|H_{12}|^2}{S_{11}} \right), \quad (2.5.17)$$

and has been applied in neurophysiological studies [26, 55].

2.5.7 Partial Directed Coherence

Partial Directed Coherence (PDC) was introduced by Baccalá and Sameshima [11, 174] and, like DC and DTF, it is applied on the frequency representation of the VAR model of a multivariate set of time series.

Assume that we have a stable VAR_p model for a set of m signals and we Fourier Transform the VAR coefficients \mathbf{A}_r to the frequency domain as in eq. (2.4.5), to get the matrix $\bar{\mathbf{A}}(f)$ (2.4.6). The definition of partial directed coherence from x_j to x_i is given by:

$$\pi_{i \leftarrow j}(f) = \frac{|\bar{A}_{ij}(f)|}{\sqrt{\sum_k |\bar{A}_{kj}(f)|^2}}. \quad (2.5.18)$$

Since testing for Granger causality in the time domain implies assessing whether $A_{ij,r} = 0$ for all $r \in \{0, p\}$, or in frequency domain whether $\bar{A}_{ij}(f) = 0$ for all frequencies f , PDC is directly related to Granger causality. Indeed PDC from j to i is zero if and only if this Granger non-causality condition holds.

Note that the stability criterion (2.3.16) guarantees that the denominator of (2.5.18) is strictly positive and hence the PDC coefficients are well defined. This denominator is a normalization that bounds PDC to values from 0 to 1. Summing over all possible “targets” of x_j , implies that PDC compares the outflow of information from signal x_j to signal x_i with the total outflow of information from x_j to all signals. Thus, contrary to DTF, PDC ranks the interaction strength with respect to the origin of the information flow so that a high value of PDC, close to 1, would imply that most of the information

outflow from x_j goes to x_i . Also note that when $i = j$, PDC represents the fraction of time evolution of x_i that can be represented by exclusive consideration of x_i 's past on itself.

One advantage of PDC is that it involves matrix $\bar{\mathbf{A}}(f)$ directly and not the transfer function matrix $\mathbf{H}(f)$, as in DTF, which would require inversion of $\bar{\mathbf{A}}(f)$ for each frequency. This yields better computational efficiency and precision.

PDC is an extension of the measure of partial coherence which describes the linear interaction between x_i and x_j on the frequency domain after discounting the linear influence of all other time-series [21]. A definition of the non-parametric version of partial coherence was given in (2.5.10). When it is calculated based on a VAR model, partial coherence can be expressed as:

$$\kappa_{ij}(f) = \frac{\bar{a}_i^H(f) \Sigma^{-1} \bar{a}_j^H(f)}{\sqrt{(\bar{a}_i^H(f) \Sigma^{-1} \bar{a}_i^H(f)) (\bar{a}_j^H(f) \Sigma^{-1} \bar{a}_j^H(f))}}, \quad (2.5.19)$$

where $\bar{a}_i(f)$ is the i -th column of $\bar{\mathbf{A}}(f)$: $\bar{a}_i(f) = [\bar{A}_{1i}(f) \ \bar{A}_{2i}(f) \ \dots \ \bar{A}_{mi}(f)]^T$ [11]. PDC is essentially an extension of $\kappa_{ij}(f)$, in the same fashion as DC is to ordinary coherence, since it includes the notion of directionality of the interaction in its definition and one can factorize $\kappa_{ij}(f)$ with PDC [11].

Since PDC is a rather recently introduced measure, most of its statistical properties such as an analytical critical value, have not been thoroughly studied yet. PDC was shown to coincide with DTF in bivariate time series [11] which could allow the use of an empirical “spectral causality criterion” [182], according to which connection from x_j to x_i is significant if $|\pi_{i \leftarrow j}| > 0.1$ for some f [174]. It has been proposed that the same criterion can be applied as a critical value for higher dimensional time series with $m > 2$ [11].

One analytical approach to the critical value problem was presented by Schelter et al. [178]. They argue that since PDC is calculated using estimates $\hat{A}_{ij,r}$ of the real coefficients $A_{ij,r}$ of a VAR _{p} process, estimated through the sample data, the statistical properties of the estimates of PDC can be derived from the properties of the estimates $\hat{A}_{ij,r}$. Commonly used estimates of the VAR coefficients, such as least squares and Yule-Walker estimates, are asymptotically normally distributed with mean the real $A_{ij,r}$ and covariances that tend to:

$$\lim_{N \rightarrow \infty} N \text{cov}(\hat{A}_{ij,r}, \hat{A}_{ij,r+h}) = \hat{\Sigma}_{ii} \hat{\Gamma}_{jj}^{-1}(h) \quad (2.5.20)$$

as the signal length N goes to infinity. Here $\hat{\Sigma}_{ii}$ is the i -th diagonal element of the

estimated error covariance matrix $\hat{\Sigma}$ and $\hat{\Gamma}_{jj}^{-1}(h)$ is the j -th diagonal element of the h -lag corresponding submatrix of the inverse of the VAR process's estimated covariance matrix $\hat{\Gamma}_X$ given by:

$$\hat{\Gamma}_X = \begin{bmatrix} \hat{\Gamma}(0) & \hat{\Gamma}(1) & \dots & \hat{\Gamma}(p-1) \\ \hat{\Gamma}(-1) & \hat{\Gamma}(0) & \dots & \hat{\Gamma}(p-2) \\ \vdots & \vdots & \ddots & \vdots \\ \hat{\Gamma}(-p+1) & \hat{\Gamma}(-p+2) & \dots & \hat{\Gamma}(0) \end{bmatrix}. \quad (2.5.21)$$

Consequently, the real and imaginary part of the estimation $\hat{A}_{ij}(f)$ are also asymptotically jointly normally distributed. They further show that under the null hypothesis of non-causality: $|\bar{A}_{ij}(f)|^2 = 0$, the asymptotic distribution of the expression:

$$\frac{N}{\hat{C}_{ij}(f)} |\hat{A}_{ij}(f)|^2 \quad (2.5.22)$$

is a weighted average of two independent χ^2 -distributions with one degree of freedom, with:

$$\hat{C}_{ij}(f) = \hat{\Sigma}_{ii} \left(\sum_{h=1-p}^{p-1} \hat{\Gamma}_{jj}^{-1}(h) \cos(2\pi h f) \right). \quad (2.5.23)$$

The same asymptotic distribution has been shown for the maximum likelihood estimators of $\bar{A}_{ij}(f)$ as well [201]. Keeping in mind that:

$$|\hat{\pi}_{i \leftarrow j}(f)|^2 = \frac{|\hat{A}_{ij}(f)|^2}{\sum_k |\hat{A}_{kj}(f)|^2}, \quad (2.5.24)$$

after some mathematical manipulation [178], an analytical approximation of the PDC critical value at the $100 \cdot \alpha\%$ significance level is given by:

$$\sqrt{\frac{\hat{C}_{ij}(f) \chi_{1,1-\alpha}^2}{N \sum_k |\hat{A}_{kj}(f)|^2}}, \quad (2.5.25)$$

with $\chi_{1,1-\alpha}^2$ the $1 - \alpha$ quantile of the χ^2 -distribution with one degree of freedom. Note that this critical value depends on the frequency f and also that since it is a pointwise level, it can be exceeded for some (few) frequencies even though the values of PDC could actually be insignificant for these frequencies.

To test the effect of the VAR model order p in PDC calculation and the effectiveness of the critical value, we follow [178] in implementing a system of two uncoupled relaxators and damped oscillators given by the following expressions:

$$\begin{aligned} x_1(t) &= 2 \cos(2\pi/15) \exp(-1/5) x_1(t-1) - \exp(-2/5) x_1(t-2) + u_1(t), \\ x_2(t) &= 2 \cos(2\pi/15) \exp(-1/5) x_2(t-1) - \exp(-2/5) x_2(t-2) + u_2(t). \end{aligned} \quad (2.5.26)$$

Each signal is dependent only on its own past two values, and thus we expect PDC coefficients, in both directions, to be zero. We model the above time series by three different VAR models, with corresponding model orders $p = 2$, which is the optimal one, 50 and 200. The resulting PDC calculations and critical values are presented in Figure 2.1. Increasing the model order far beyond the optimal value introduces an increasing variability in PDC which makes it difficult to assess straightforwardly whether it is significantly nonzero. Yet, the critical value is always higher than PDC in almost all frequencies and for all model orders, suggesting the validity of the analysis of its derivation. Note that our results differ with the published ones only in that both our PDC and critical values are lower than those in [178] for model orders 50 and 200. This could be due to various factors such as the different VAR modelling technique implemented there.

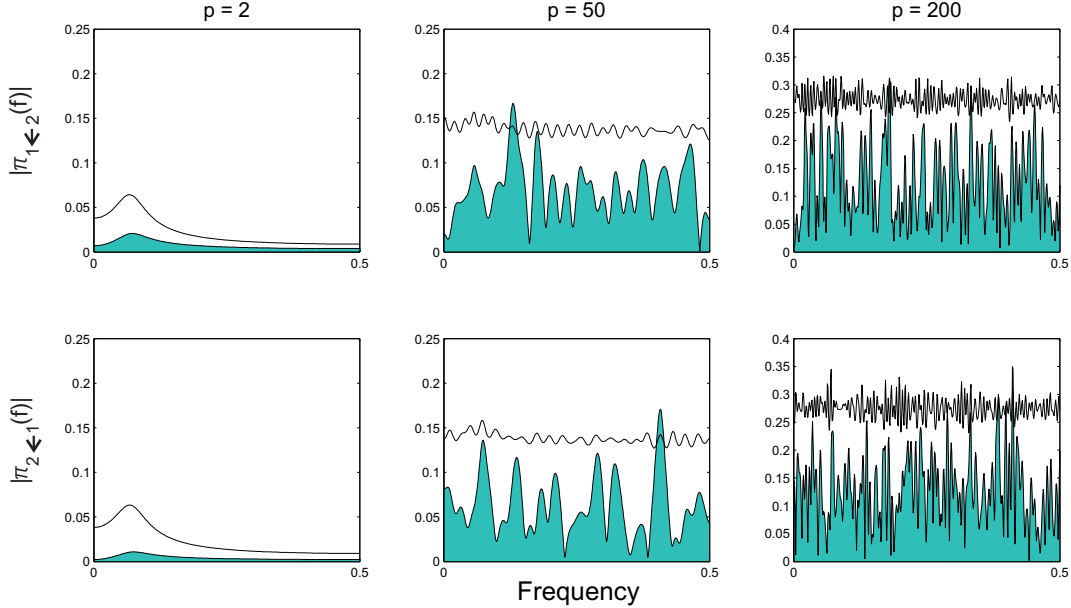


Figure 2.1: PDC (filled) and critical values (black lines) for a set of uncoupled oscillators given by (2.5.26), for increasing VAR model orders. The top row corresponds to $|\pi_{1\leftarrow 2}(f)|$ and the bottom one to $|\pi_{2\leftarrow 1}(f)|$.

2.5.8 Comparing Coherence, PDC, DTF and dDTF

To gain an insight on how these tools perform when analysing multivariate signals, they can be tested over simplistic cases where the connectivity of the signal generators is known in advance. We perform here one such test that was first used in [11] in order to compare the efficiency of PDC against DTF. We will use the same test to compare

these two measures, along with ordinary coherence and dDTF.

Consider the set of five coupled oscillators: x_1, \dots, x_5 , whose connectivity diagram is shown in Figure 2.2A. The signal propagates from the driving source $x_1(t)$ to $x_2(t)$ and $x_3(t)$ independently and also to a set of coupled oscillators $x_4(t)$ and $x_5(t)$. The equations generating the 5 time series are the following:

$$\begin{aligned} x_1(t) &= 0.95\sqrt{2}x_1(t-1) - 0.9025x_1(t-2) + u_1(t), \\ x_2(t) &= 0.5x_1(t-2) + u_2(t), \\ x_3(t) &= -0.4x_1(t-3) + u_3(t), \\ x_4(t) &= -0.5x_1(t-2) + 0.25\sqrt{2}x_4(t-1) + 0.25\sqrt{2}x_5(t-1) + u_4(t), \\ x_5(t) &= -0.25\sqrt{2}x_4(t-1) + 0.25\sqrt{2}x_5(t-1) + u_5(t). \end{aligned} \tag{2.5.27}$$

Note that the maximum lag between oscillator pairs is 3 timesteps. Following this set of equations, 10000 data points were created for each oscillator, with $u_i(t)$ being random numbers drawn from a standard normal distribution. The data were modelled with a VAR_p model, according to the ARFIT algorithm [152] (see next chapter for a description of our custom built numerical package FunCAT that was used for this analysis). The model order was detected correctly by both the FPE and BIC criteria as $p_{\text{opt}} = 3$ (Figure 2.2B). The resulting VAR model was stable since its eigenvalues satisfied the stability criterion (the maximum eigenvalue was 0.948). The VAR coefficients are shown in Figure 2.2C, where those corresponding to each pair of oscillators, in each direction, are connected by a line of different color. Their confidence intervals are also shown. Note that the only significantly non-zero elements are coefficients: $a_{2 \leftarrow 1,2}, a_{3 \leftarrow 1,3}, a_{4 \leftarrow 1,2}, a_{4 \leftarrow 5,1}$ and $a_{5 \leftarrow 4,1}$. Note also that indirect connections like $a_{5 \leftarrow 1,i}$ are insignificantly non-zero for any step i . Therefore, one can already detect the connectivity diagram and the lags between the signals simply by this figure, which corresponds to a Granger Causality test analysis.

The results for PDC, along with its critical value at 95% significance level, DTF, dDTF and ordinary coherence are shown in Figure 2.3. The diagonal diagrams contain the power spectrum of each oscillator. Coherence, apart from symmetrical (thus non-directional), is non-zero in all cases, detecting both direct and indirect interactions. The PDC and DTF plots are in full agreement with those presented in [11]. PDC is significantly nonzero only at panels $\{2, 1\}, \{3, 1\}, \{4, 1\}, \{4, 5\}, \{5, 4\}$, correctly revealing only the direct interactions between the oscillators, unlike DTF which detects indirect causalities as well, e.g. from x_1 to x_5 . Note that dDTF is also nonzero in the same

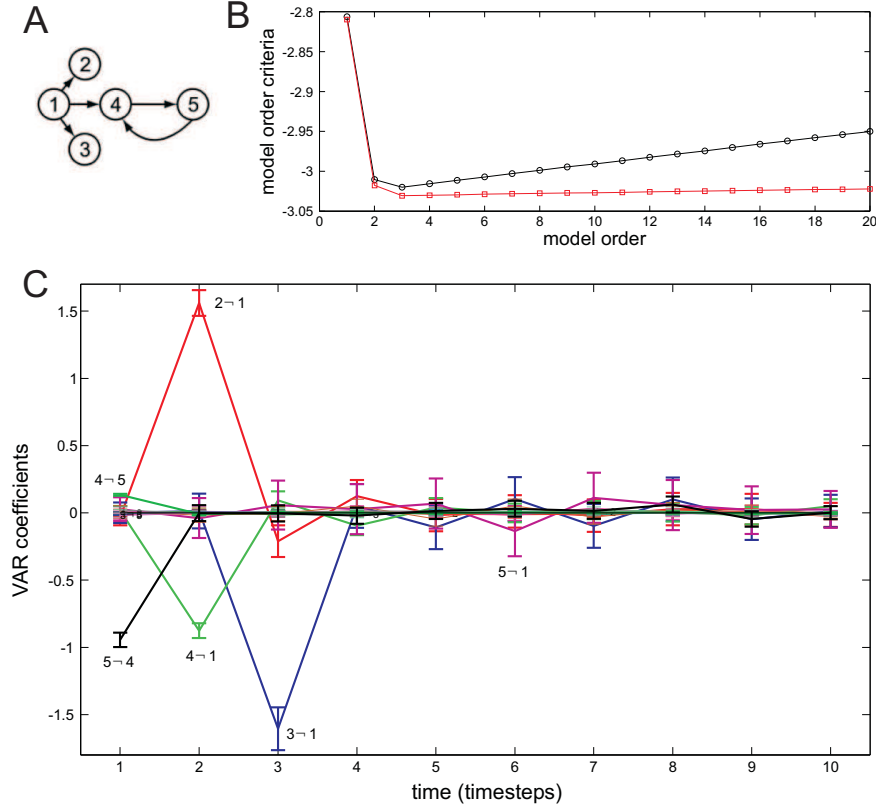


Figure 2.2: **A:** Diagram of the causality pattern of the oscillators (circled 1 to 5) given by (2.2). **B:** Model order criteria FPE (red) and BIC (black) for VAR modelling the time series generated by the oscillators. **C:** VAR coefficients of the optimal-order model and 95% confidence intervals for each directed pair of signals and each step. Note that only those corresponding to direct causal interactions are significantly non-zero.

diagram for a few frequencies, but the absence of a critical value for this measure makes it unclear whether this is significant. It acquires similar small values in the $\{4, 5\}$ panel as well, which is a direct connection, making it even harder to conclude whether it only detects direct causality or not (an issue also noted in [9]).

More tests similar to this, some with more complex connectivity, can be found in [11, 123, 178]. In all of them, PDC performs equally well or better than other measures in detecting direct causal interactions. Nevertheless, below we examine one case where PDC fails as well, and present a newer version that can tackle this problem.

2.5.9 Generalised Partial Directed Coherence

Although PDC appears to detect direct interactions only, it has been shown that large differences in the variances of the modelled time series can yield distortions in the

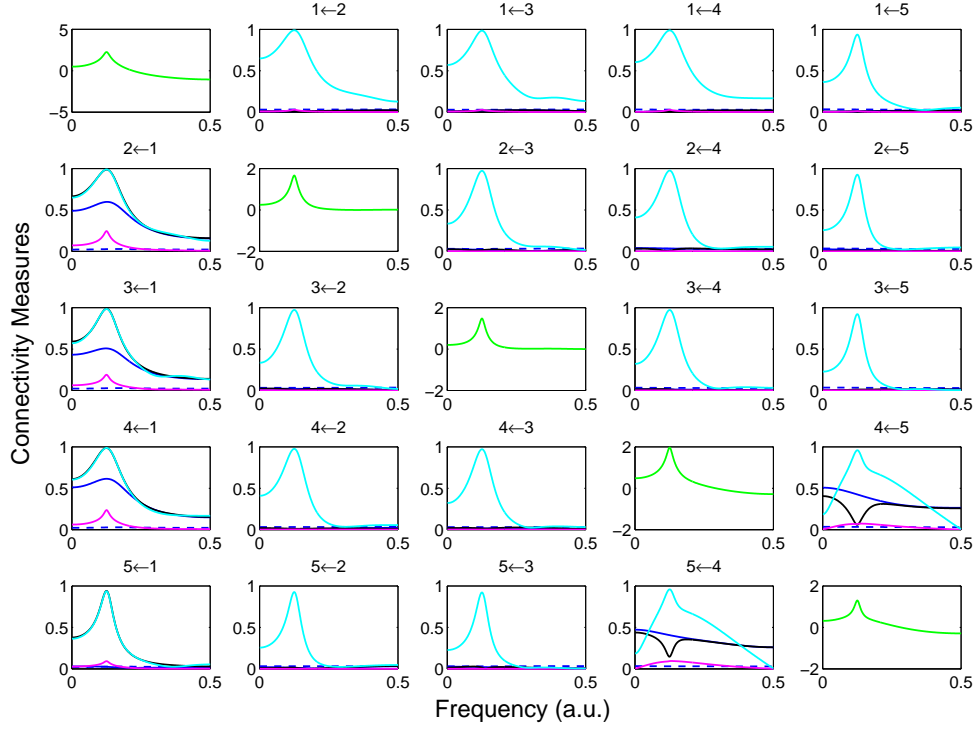


Figure 2.3: Plots of PDC (blue), DTF (black), dDTF (magenta) and ordinary coherence (cyan, often overlapping with DTF). The critical value of PDC is also included (dashed blue). The logarithm of each oscillator's power spectrum is plotted in the diagonal diagrams. Note that all measures, except PDC, indicate a causal interaction from signal 1 to 5. PDC detected the directed interactions only.

resulting PDC values. In a simple bivariate VAR model where there is causality only from x_2 to x_1 , amplifying x_2 by some constant α would lead to $\pi_{1 \leftarrow 2}(f)$ approaching 0 as α increases [12]. In another example, a set of three uncorrelated white noise processes, where two of them have much larger variance than the third, will produce a distorted connectivity profile, since PDC wrongly detects connections from the low-variance process to the other two [236]. To avoid such distortions from differences in time series scaling, a variation of the original PDC, the generalised PDC (gPDC) was introduced [12]. In gPDC, the coefficients $|\bar{A}_{ij}(f)|$ are normalised by the standard deviation, $\sigma_i = \sqrt{\Sigma_{ii}}$, of the u_i model residuals:

$$|\pi_{i \leftarrow j}^g(f)| = \frac{\frac{1}{\sigma_i} |\bar{A}_{ij}(f)|}{\sqrt{\sum_k \frac{1}{\sigma_k^2} |\bar{A}_{kj}(f)|^2}}. \quad (2.5.28)$$

It has been suggested that gPDC, when accompanied by a Z-score pre-normalization of

the time series, yields correct connectivities even for very short data samples of series with different scaling [12]. Indeed in the example of three oscillators we find that the lack of connections is correctly detected after replacing PDC with gPDC (Fig. 2.4).

In need of a critical value for gPDC, we followed the same analysis as in the derivation of (2.5.25) for an estimated VAR model. Under the null hypothesis of $\frac{1}{\hat{\sigma}_i^2}|\hat{A}_{ij}(f)|^2 = 0$, which is equivalent to $|\pi_{i \leftarrow j}^g(f)| = 0$, we can simply rewrite (2.5.22) as:

$$\frac{N}{\frac{\hat{C}_{ij}(f)}{\hat{\sigma}_i^2}} \frac{1}{\hat{\sigma}_i^2} |\hat{A}_{ij}(f)|^2, \quad (2.5.29)$$

whose asymptotic distribution is again a weighted average of two independent χ^2 -distributions with one degree of freedom. By following the same steps as in [178], we obtain a similar approximation to (2.5.25) for the critical value of gPDC:

$$\sqrt{\frac{\hat{C}_{ij}(f) \chi_{1,1-\alpha}^2}{N \hat{\sigma}_i^2 \sum_k \frac{1}{\hat{\sigma}_k^2} |\hat{A}_{kj}(f)|^2}}, \quad (2.5.30)$$

where $\hat{C}_{ij}(f)$, $\hat{\sigma}_i$ and $\hat{A}_{ij}(f)$ are estimated as before. This expression was tested in the aforementioned example, where all gPDC values are below the critical value (Fig. 2.4).

2.6 Discussion

We have reviewed here some basic elements of multivariate time series analysis and autoregressive modelling, as well as statistical tools that assess structural or functional interactions between time series.

Traditionally effective connectivity between pairs of signals has been estimated using cross-correlation analysis which gives an indication of a delay in coupling, but does not provide any inference to causality, i.e. which structure drives which. It could, for example, be the case that a common mechanism is simply driving both signals with different time delays (“common source” problem). The difference between these delays would in this case appear as the maximum-correlation lag between them. This is why high correlations between signals should be treated with some caution before making inferences about their functional or structural relations. Use of coherence provides a frequency-domain representation of the magnitude and phase relationship between two signals, but again provides no information on the directionality between them. However some directionality information can be obtained from the coherence phase spectrum

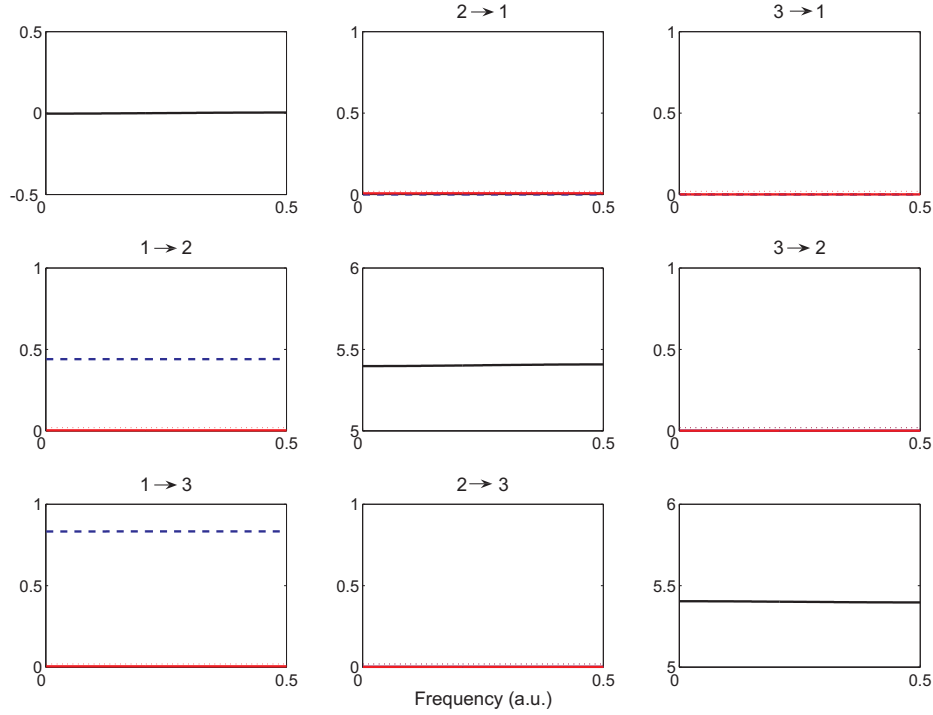


Figure 2.4: Comparing the effectiveness of PDC against gPDC in detecting non-causality. PDC (dashed blue) and gPDC (red) for three independent white noises with different variances. PDC erroneously indicates that process x_1 is driving the processes with higher variance, x_2 and x_3 . gPDC on the other hand, being practically zero and lying entirely below the critical value (2.5.30) (dotted red) restores the independence of the processes. The diagonal panels represent the logarithm of the power spectral density of each process.

which provides a measure of phase (time delay) between coherent signals [94] but it has been argued to be rather ineffective [125].

Apart from these measures, we have presented here a review of some more recent techniques for statistical analysis of time series that have been developed to provide tools for assessing the directional effective connectivity between signals. They are being increasingly applied in studies of directional interactions between neurons, neural ensembles and whole brain regions. Amongst them, the closely related Partial Directed Coherence (PDC) [11, 174] and Directed Transfer Function (DTF) [114, 115] have been implemented in numerous studies of either simplified tests [11, 85, 125, 178, 236] such as the comparative test presented here, in simple neuronal models of interconnected neurons [9, 114, 174] and in real neurophysiological data [72, 102, 198, 226, 227, 240]. Generalised PDC, was introduced to cope with time series of widely different variances [12] and is also being increasingly applied in neuroscience studies [71, 100, 176, 202]. These

techniques are mainly based on the notion of Granger causality [86], which was described here, and have the advantage that they can be applied to multivariate time series, not restricting the connectivity analysis to pairs of neural structures but extending it to whole sets of simultaneously recorded signals. This way, typical problems with pairwise analysis, such as serial connections or common sources [36, 125], can be avoided.

Nevertheless, since most of these measures rely on VAR modelling the sampled time series, caution is needed in the construction of the model. Issues such as the stationarity of the available data or their size can severely affect the VAR modelling procedure and consequently lead to incorrect connectivity profiles. For example, in the connectivity test presented in Section 2.5.8, reducing the size N of data points generated by the oscillators, leads to nonzero PDC between unconnected pairs. Although the rise of the critical value (being inversely proportional to N) renders these PDC profiles non-significant, reducing N to below ~ 500 points, leads to significantly nonzero PDC between unconnected oscillators (not shown). Although there are, to our knowledge, no standard tests for the adequacy of the data size, one should use as many (stationary) data as possible and check features of the resulting VAR model, such as its stability or the whiteness of the residuals, to gain a better understanding of the model's efficiency.

In the following chapter we put this analysis into practice by applying gPDC on a set of neurophysiological data. One issue that arises there, is how to transform binary spike trains into continuous stationary signals, which we examine as well.

Assessing Cortico-Hippocampal Connectivity During Anaesthesia

3.1 Introduction

Here we make use of the previous analysis on generalised PDC (gPDC) by applying it to neurophysiological data, recorded from anaesthetised rats, in order to assess directionality of information flow between prefrontal cortex and hippocampus, as well as intrahippocampally. Specifically, we applied gPDC on electrophysiologically recorded local field potentials (LFPs, recorded at 1 kHz) and unit spike trains from dorsal hippocampus (medial and lateral regions) and medial prefrontal cortex (mPFC) from isoflurane-anesthetised rats [50]. After a period of basal recordings, the excitatory amino-acid kainic acid (KA) was administered in order to induce increased excitatory activity. The experimental methods and recording procedures were described in detail in [202] (Figure 3.1). gPDC was selected from the tools presented in the previous chapter because, as demonstrated, it succeeds in detecting direct causal interactions without getting distorted by signal variance differences.

Each recorded spike train is considered as a time series in the form of a series of Dirac delta functions. Spectral analysis through the usual Fast Fourier Transform (FFT) algorithm cannot be performed in such binary signals since they cannot be regularly sampled [75]. Furthermore, VAR modelling cannot be successfully performed either since spike trains do not constitute stationary time series with relatively constant mean values and variances. Therefore, one needs to convert spike trains into continuous stationary signals. Before the final PDC results are presented, we delve into an analysis of two

different methods that can be used for the spike train preprocessing.

All subsequent analyses (as well as the tests described in the previous chapter) were performed with FunCAT (Functional Connectivity Analysis Toolbox), our custom built numerical package on autoregressive modelling and functional connectivity analysis, implemented in MATLAB, that we describe in Section 3.4. The biological implications of our findings on the effects of anaesthesia and kainic acid on brain connectivity, along with implementation issues of gPDC analysis on field potentials and spike trains, are discussed in Section 3.6 (part of which was co-written with Dr. Benjamin Coomber).

Parts of the work in this chapter were published in [202].

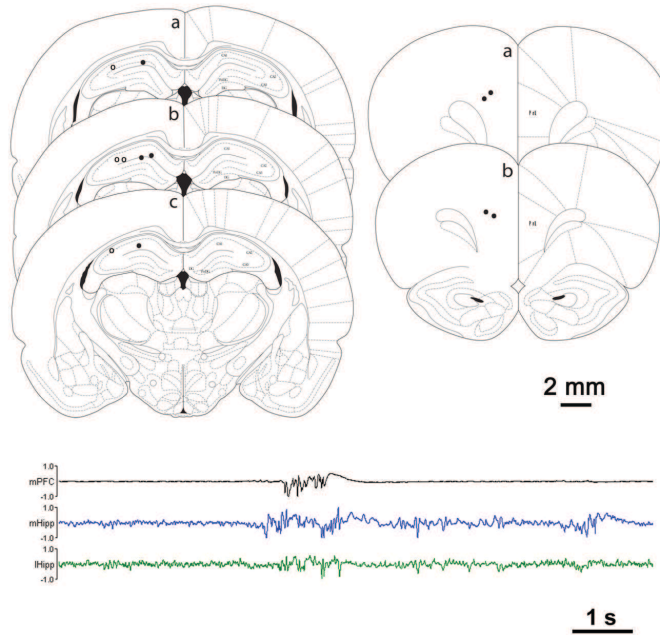


Figure 3.1: Schematic representation of histological verification of dorsal hippocampal (upper left) and mPFC (upper right) recording sites. The most medial (black circles) and lateral (white circles) electrode placements of the hippocampal array and an electrode from the bundle in mPFC (black dots) are shown in coronal brain images [160] for each experiment included in this study ($n = 4$ rats). The distance posterior to bregma in the sequence of images for hippocampal placements is: (a) -3.8 mm; (b) -3.6 mm; (c) -3.3 mm. The distance anterior to bregma in each image for mPFC placements is: (a) 3.2 mm; (b) 3.7 mm. The lower panel shows representative segments of LFP activity during basal conditions in one rat recorded simultaneously in mPFC, medial hippocampus (mHipp), and lateral hippocampus (lHipp) respectively over a 10 s period.

3.2 Methods of Spike Train Preprocessing

Two methods were examined for transforming spike trains into continuous stationary time series. Namely, spike trains were either substituted with their corresponding firing rates, or transformed into a continuous function through the kernel convolution method of the French-Holden algorithm [75]. Both methods have been used before in neurophysiological studies with PDC [72, 102, 174, 226, 227, 240].

3.2.1 Firing Rates

Firing rates are calculated for each neuron by splitting spike trains into time-bins and computing the rate in each bin. The choice of bin length and overlap between successive bins set a specific “binning scheme” which affects the shape of the resulting rates. Namely, the shorter the length of the bin and the larger the overlap between two neighbouring bins, the more detailed are the firing rates. The sampling rate of the resulting signal is also set through the binning scheme since each data point will correspond to the rate in the corresponding bin. Each bin can be uniquely represented by, for example, its starting point, so that the resulting sampling rate is equal to the inverse time between two consecutive such points. This time distance can be calculated straightforwardly as follows. If L is the time length of each bin and v is the percentage of overlap between two consecutive bins, then the distance T between two consecutive starting-bin-points is given by:

$$T = L - L \frac{v}{100} = L \frac{100 - v}{100}, \quad (3.2.1)$$

yielding the corresponding sampling rate f_S :

$$f_S = \frac{100}{L(100 - v)}. \quad (3.2.2)$$

According to the Nyquist Theorem, the maximum frequency f_{max} over which spectral analysis can be performed without aliasing is:

$$f_{max} = \frac{f_S}{2} = \frac{50}{L(100 - v)}. \quad (3.2.3)$$

In conclusion, apart from determining the amount of processed information, the binning scheme for the firing-rate processing of spike trains also determines the available range of frequencies for spectral analysis and the subsequent PDC calculation.

Below we test how the choice of bin length and overlap may affect PDC results by comparing them with results from LFP signals. Figure 3.2 displays PDC and its critical

value along with ordinary coherence, calculated over one LFP segment from the three recorded brain regions of one rat. The diagonal panels show the power spectra of the three regions. Both the BIC and FPE criteria, which were used for the VAR_p model order selection, decreased monotonically with increasing order, so p was set to 50 and the goodness of fit of the VAR models was established by visual comparison of the power spectra from the model with those of the original data (see Section 3.3). PDC indicates that there is, in general, a much stronger flow of information from mPFC to both medial and lateral hippocampus than in the opposite direction, more prominent over low frequencies from 0 to 20 Hz. Moreover, the lateral hippocampus appears to be leading the medial, which is anatomically consistent with excitatory (Schaffer collateral) projections from CA3 to CA1 (see Chapter 5). We examine whether PDC, calculated from firing rates, also yields similar results and importantly whether these results are severely affected by the implemented binning scheme.

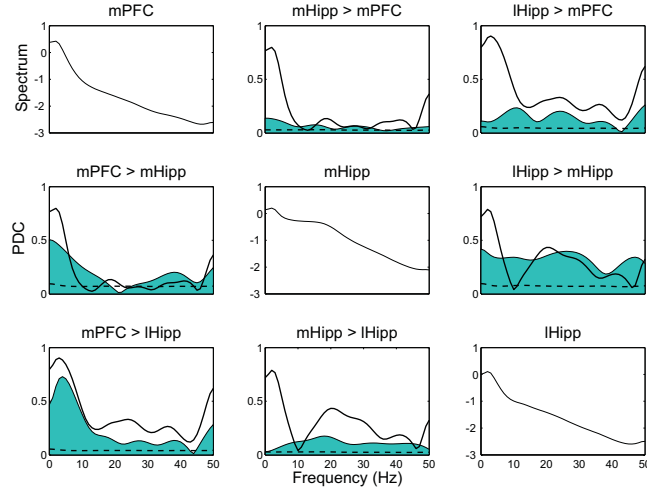


Figure 3.2: PDC (filled) between the three rat-brain regions, calculated over one segment of the LFP recordings from one rat. Critical values (dashed) and ordinary coherence (solid curves) are also displayed. The diagonal panels display the power spectra of the three regions.

3.2.2 Bin Length Effect

We computed the firing rates of all spike trains and grouped those belonging to each region by performing principal component analysis and keeping the first principal component (see Section 3.3). There was no overlap between bins and four different bin lengths were implemented: 10, 50, 100 and 200 msec. The resulting PDC, critical values and ordinary coherence from all cases are displayed in Figure 3.3. The model order was

fixed in all cases to avoid any effects that different orders may produce between models. Both BIC and FPE yielded particularly small optimal orders, so p was set to 17, producing accurate VAR models in all cases, which was again established by a spectral comparison between the model and the original data.

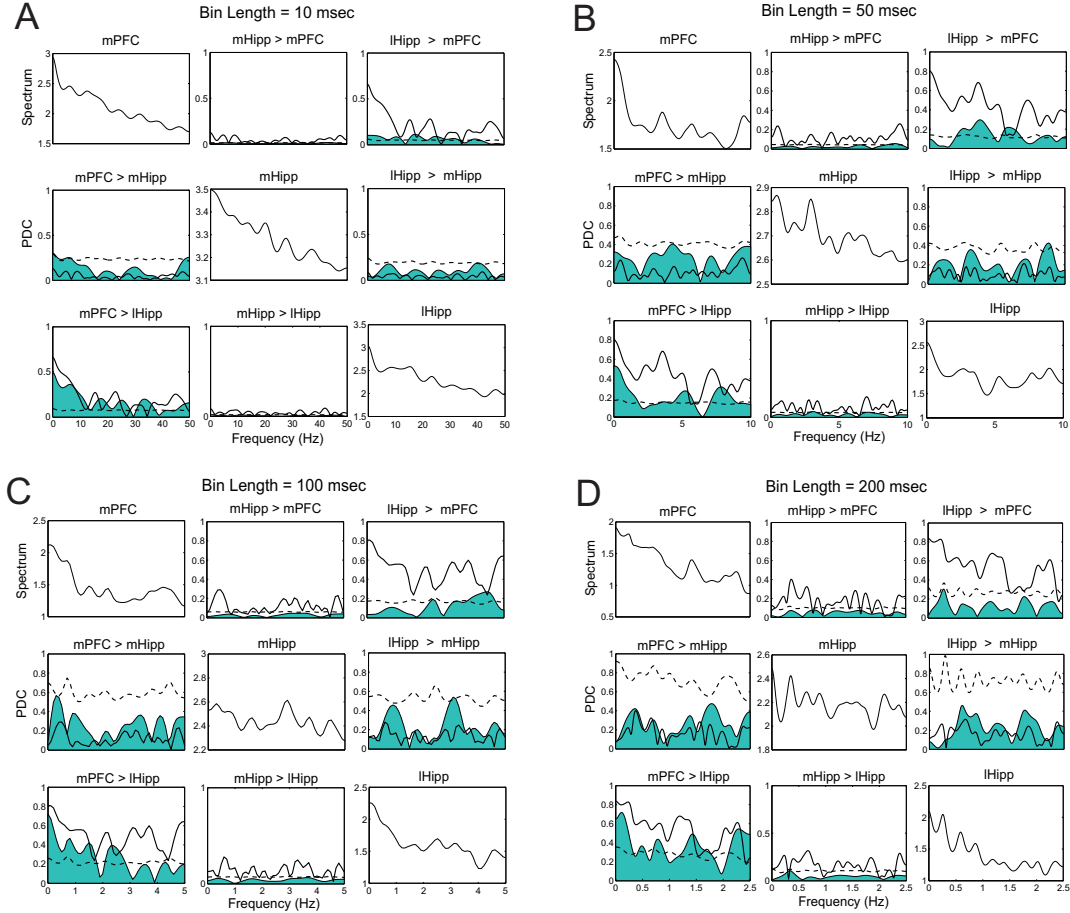


Figure 3.3: PDC (filled) between the three recorded brain regions calculated through firing rates of spike trains. Critical values (dashed) and ordinary coherence (solid curves) are also displayed. The bin lengths for the firing rate calculation are (A-D): 10, 50, 100, 200 msec. No overlap between bins was set.

The general connectivity profile from all 4 bin-length cases is qualitatively in agreement with that from LFP analysis. In particular, the case with the smallest bin size (Figure 3.3A) appears to be closest to the LFP results. Note though, that the critical value is much higher now and increases with the bin length, since the value is inversely proportional to the data length (thus proportional to the bin length, eq. (2.5.25)). So the critical value suggests that PDC is statistically insignificant in most cases. Therefore the connectivity is in accordance with the LFP results only if one chooses to disregard

the critical value for all binning schemes.

Another point to be noted is that all results were plotted up to the maximum (Nyquist) frequency corresponding to each bin length. So although one could expect the results from each bin length to be roughly a zoom into the lower frequencies PDC profile of the smaller length, this is hardly the case. We thus conclude that the amount of information lost by increasing the bin length strongly alters the frequency profile of PDC. A similar conclusion can be drawn for ordinary coherence as well.

So, as expected, the shorter the bin length the higher the amount of information contained in the data and the closer the resulting connectivity to that derived from a continuous signal. On the other hand, setting the length too low may result in non-stationary time series, even binary ones, in cases of sparsely spiking neurons.

3.2.3 Overlap Effect

We next tested the effects on PDC of the overlap between two consecutive bins. We processed the same spike trains as before with the same four choices of bin length, but with 50% overlap (Figure 3.4). The model order was set to 26 for all bin lengths, which was optimal, according to BIC, for the first case of 10 msec bin and also gave accurate models for all the other cases.

The general characteristics of these plots are similar to the previous ones and the comments in the previous section apply to this case as well. A more direct comparison can be made between cases of 100 and 200 msec bins with 50% overlap (Figure 3.4C,D), with those of 50 and 100 msec in Figure 3.3C,D respectively, since their sampling rates are equal. There, the overlap does not create significant differences in the PDC profiles over frequency, other than making the ones with overlap more detailed. This is somewhat intuitive since the 50% overlap essentially adds one more bin between any two adjacent bins in the corresponding no-overlap case. This extra bin will carry some extra information for the “transition” of PDC from one bin to the next, thus making the PDC profile more sensitive to frequency fluctuations.

Our results suggest that the bin length is the most crucial parameter of the binning scheme. A bin length of around 10 msec seems to be the best option in our case study as it gives results that closely resemble those of corresponding LFP data. Increasing the bin length substantially, will cause distorted PDC profiles over frequency and critical values, since large amounts of essential information will be lost in the spike grouping and

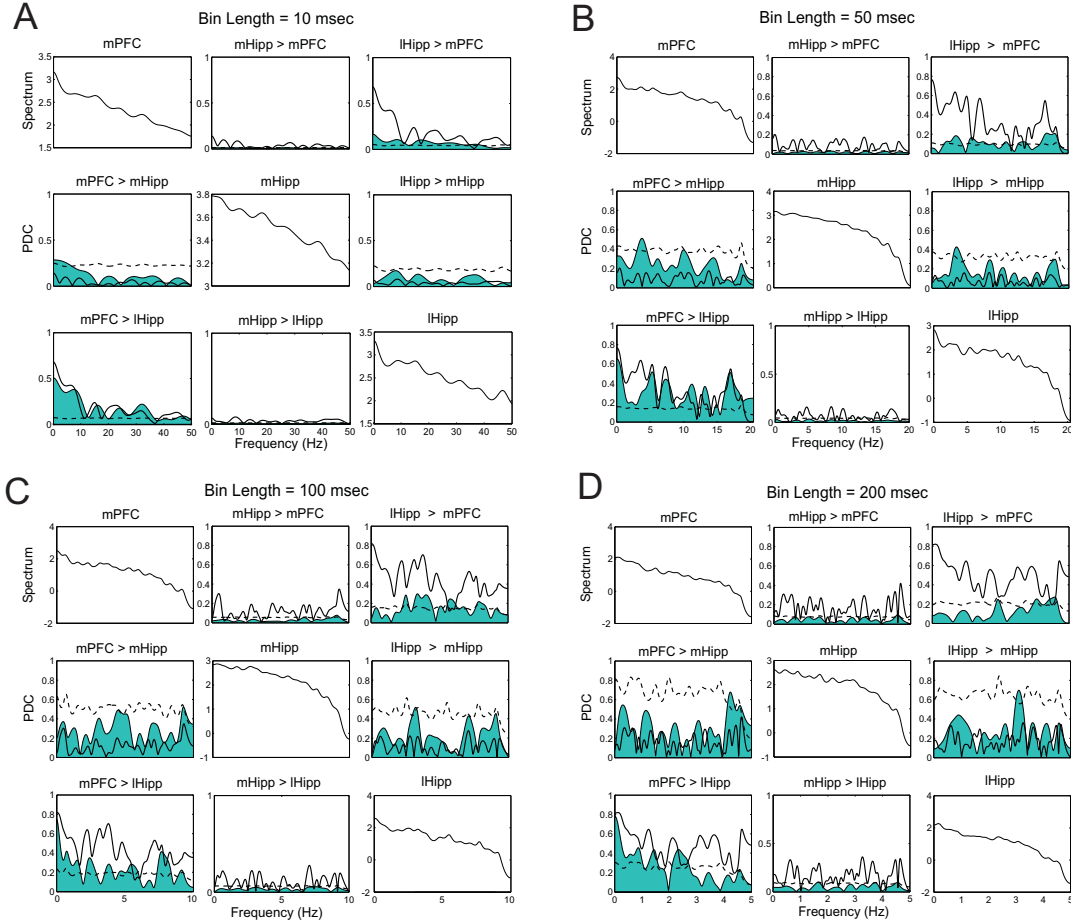


Figure 3.4: PDC (filled), critical values (dashed) and ordinary coherence (solid curves) between the three recorded brain regions, calculated through firing rates of spike trains as in Figure 3.3, but with a 50% overlap between adjacent bins.

the resulting data points will be too few. Decreasing it further, e.g. to 5 msec, does not substantially affect results (not shown), only making computations more demanding. The overlap between the bins appears to be of relatively minor importance. A large overlap will create smoother firing rates and thus more detailed (but also more noisy) PDC frequency profiles, but will be more computationally demanding. We thus suggest that a moderate ($\sim 50\%$) overlap is adequate.

3.2.4 Spike Times

French and Holden [75] described a technique which applies an ideal low-pass phaseless filter to the original spike train and transforms it to a discrete time series, sampled at equally spaced intervals. This filter is implemented by convolving the spike train with

a kernel function. Specifically, each spike is convolved with the function:

$$S(t) = \frac{\sin(2\pi f_N(t - t_i))}{2\pi f_N(t - t_i)}, \quad (3.2.4)$$

where t_i is the time of the i -th spike and f_N is the cut-off frequency of the filter. After the convolution, the continuous function is sampled at some sampling frequency f_S .

The steps of this “French-Holden Algorithm” (FHA) are the following:

1. Firstly, the effect of having only a finite window of recorded spikes (from a theoretically infinite in length spike train signal), is the convolution of the infinite signal’s power spectrum with the function:

$$S_P(f) = P \frac{\sin(\pi f P)}{\pi f P}, \quad (3.2.5)$$

where P is the duration of the recorded spike train.

2. Every spike of this finite signal is convolved with the kernel $S(t)$, yielding an infinite-length signal. This step corresponds to applying a low-pass filter with cut-off frequency f_N to the spike train’s spectrum. f_N will then be the Nyquist frequency of the spectrum resulting from this step.
3. From this infinite continuous signal, only the time-window of the original recorded spike train is kept. This results to a second convolution of the spectrum with $S_P(f)$ which, due to the one in the first step, will have no further effects.
4. Finally the continuous windowed function is sampled with some sampling rate f_S , yielding a new Nyquist frequency for the final spectrum, given by $f_S/2$.

French and Holden argued that this process yields a power spectrum that is unbiased and free of aliasing and distortion effects. Although the FHA is computationally implausible due to the infinite length of the kernel $S(t)$, the resulting signal can be obtained by just centering the kernel around each spike and calculating its value at each sampling point within the recorded time window. The results are then added over each sample point. To simplify the FHA further, the calculation of $S(t)$ at each point for each spike can be avoided by setting $f_S = 2f_N$. In this case, the FHA will give a regularly sampled function with values:

$$S(j\Delta t) = \sum_{i=1}^N \frac{(-1)^{j+1} \sin(2\pi f_N t_i)}{\pi(j - 2f_N t_i)}, \quad (3.2.6)$$

where the summation is over all N spikes in the spike train [159]. Hence, only one sinusoidal computation is required for each spike time and one summation of N components for each sample point.

The resulting signal is a continuous function whose magnitude represents an estimate of the instantaneous firing rate of the spike train [159]. The choice of the f_N parameter affects the shape of the final signal and the accuracy of the spike times. Figure 3.5 depicts an example spike train and the sampled function resulting from the FHA with $f_N = 10$ Hz and 100 Hz. It is clear from the figure that the higher the Nyquist frequency the more accurate the resulting spectrum.




Figure 3.5: The French-Holden algorithm applied on a series of spikes (top) with Nyquist frequency set to 10 Hz (middle) and 100 Hz (bottom).

3.3 Data Processing and Choice of VAR Model Order

We describe here the pre-processing of all types of data before the final gPDC analysis, along with the choice of model parameters for the VAR_p models.

LFP signals were analysed first. LFPs usually provide locally (almost) stationary series when examined in segments of only a few seconds [23, 64]. As proposed in [64], all LFP recording epochs were split in 10 sec segments corresponding to 10,000 data points (as in the comparison test of the previous chapter) with 80% overlap, and each segment was modelled and analysed separately. Smaller segments of 2 and 5 sec duration were also implemented but, although yielding qualitatively similar connectivity features, their short data length resulted in poor gPDC estimation and very high critical values (in agreement with our analysis on the effect of long bin lengths, i.e. short data lengths).

To further avoid possible distortions due to the different variances of the signals, all LFPs from each segment were Z-score normalised.

Following our conclusions from the binning scheme tests, the firing rates of spike data were calculated in 10 msec bins, overlapping by 50%. Though the maximum frequency to be analysed was 50 Hz, the Nyquist frequency for the FHA spike train processing, was set to 100 Hz to ensure accurate power spectral computations (see [159] for practical guidelines on the FHA implementation). Principal component analysis was performed to group either the firing rates or the FHA processed units of each brain region together. This is a well known technique for dimensionality reduction of high-dimensional data sets and it is based on the eigenvalue decomposition of the data covariance matrix [112]. The first principal component from each region was implemented for the VAR modelling and subsequent PDC analysis. In most cases presented below, the first principal component contained more than 75% of the processed units' variance, rendering it a good representative of the individual processed spike trains. The only exception was the third case study where the percentage dropped to approximately 45% for the lateral hippocampal neurons. Each recording epoch was split into segments of 10 seconds duration with 90% overlap. The resulting signals had only their mean values subtracted. A full Z-score normalization was avoided because the standard deviation in some segments was often too small to be used for normalization (due to the sparseness of spikes). Again VAR modelling and gPDC analysis were performed for each segment separately. Unfortunately, the very low firing rates of neurons under anaesthesia, especially those in the mPFC, led in most occasions to very few non-zero values in the firing rate time series, even after KA administration, rendering these series non-stationary. In fact, many segments had completely zero firing rates, making it impossible to perform VAR analysis on them. Moreover, the very sparse firing led to many highly non-stationary (but non-zero) segments in the FHA processed cases as well. All the non-stationary segments of both firing rates and FHA processed spikes yielded VAR models that did not satisfy the stability criterion (2.3.16) and their gPDC results were set to zero.

To perform a consistent analysis between all the data segments and be able to make appropriate comparisons between PDC results, the same model order should be set for all VAR models, as it represents the maximum delay of information transfer between the recorded brain regions. This is not a straightforward choice, as different data sets usually require different orders for the VAR model to accurately capture the data characteristics. The model order differences are potentially even more substantial between different types

of data, as in our case. Even though they correspond to the same experiment, each of the three types of neurophysiological data examined here carries a different type of information and the corresponding data have different form.

Both FPE and BIC criteria were applied. In the LFP cases they failed to yield an optimal order as they kept decreasing with increasing order. This is a typical situation for EEG and LFP data [64, 105]. It is customary to select a model order manually so that the power spectrum of the modelled data closely follows that of the original data [26, 85, 179]. For the spike train data, preprocessed either way, both selection criteria indicated different optimal orders that were strongly fluctuating between segments. For the LFPs a model order of 100 sample points was set, corresponding to 100 msec of history since the sampling rate was 1 KHz. This order yielded, on average, a good fit of the model spectrum on the data spectra (although it naturally varied between different segments and different recordings). For the spike data, a delay of 100 msec corresponded to a model order of 20 data points for both firing rates and FHA processed trains. This order appeared to be close to optimal, as by visual inspection it was found that significantly increasing or decreasing p diminished the fit between VAR and data spectra. The power spectra of the original data and those of the VAR models are presented for each case study in the following section as representations of the fit of the VAR models to the original data.

3.4 VAR Modelling and Structural Analysis with FunCAT

We developed a numerical code, named FunCAT (Functional Connectivity Analysis Toolbox) and written exclusively in MATLAB, to analyse sets of continuous neurophysiological signals, such as EEGs and LFPs, as well as spike train data, with respect to the potential functional connectivity of the neuronal ensembles that generated the signals. FunCAT performs VAR modelling of the input signals and applies most of the statistical tools, described in the previous chapter, to assess any directional interactions between them. In addition to our own custom written routines, many functions are taken from other numerical toolboxes that are available online. Particularly, FunCAT contains (often modified) algorithms from the ARFIT [152, 181] (<http://www.gps.caltech.edu/~tapio/arfit/>), Biosig (<http://biosig.sourceforge.net/documentation.html>) and Neurospec 20 (<http://www.neurospec.org/>) numerical toolboxes.

Our goal was to combine elements from such toolboxes, along with routines on spike processing, data handling and plotting, in one complete and particularly user-friendly package that can be used even by non-experts. Moreover, we believe that simultaneously examining various connectivity tools in a structural analysis study helps in giving a more complete picture of the causalities. This is why FunCAT combines PDC, gPDC, DTF, dDTF, GCI, ordinary coherence and time-domain correlations in an effort to assess such causalities.

The main function that is called to initiate the procedure is named FunCAT. When running it, the first main choice of the user is between:

1. performing an analysis on LPF data (or any similar continuous signals), or
2. an analysis on discrete spike-train data.

Alternatively, the user has the options of computing the average from a set of connectivity results, running a demo, or just plotting results from preprocessed data.

If the spike train option is chosen, then the user can choose between preprocessing the spike trains by:

- 2.a. computing their firing rates in a user-specified binning scheme, or
- 2.b. applying the FHA with user-specified Nyquist frequency.

In either case, if the processing has already been performed in an earlier run of the code, the user can chose to start the analysis on the first principal components of the processed trains.

To process spike trains through their firing rates, the following steps take place:

- 2.a.1. The user specifies which file to open and which spike train arrays, that belong to each brain area, will be grouped together.
- 2.a.2. The length and overlap of the time bins are specified and the firing rate of each train is calculated for each bin.
- 2.a.3. Principal component analysis is performed on the firing rates of all trains in a particular group and the first principal component (PC) is returned. The eigenvalues and weights of the three first PCs are displayed, so that the user can assess the variance representation by the first PC.

- 2.a.4. All firing rates and the first PC of each group are stored and plotted.
- 2.a.5. The sampling frequency resulting from the binning scheme is calculated through eq. (3.2.2) and displayed.

To process spike trains by the FHA, the first step is the same as step 2.a.1 followed by these steps:

- 2.b.1. The maximum frequency f_{max} , over which spectral estimation will take place, is specified. The Nyquist frequency is set to $f_{max}+10$ Hz, to avoid any spectral distortions [159].
- 2.b.2. The kernel convolution of the FHA takes place for each spike train, with a sampling rate of twice the Nyquist frequency.
- 2.b.3. PCA is performed to group the trains of each brain area and the first PCs are returned.
- 2.b.4. All processed trains and the first PCs are stored and plotted.

The user can now analyse either the LFPs from each brain area, or the first PCs of spike trains, processed either way, through the following procedure:

- 3. The names of the LFP- or PCs-containing arrays are specified.
- 4. The data are separated into user-specified segments.
- 5. The model order p is set. Alternatively, the optimal model order can be estimated for each segment according to the FPE and BIC criteria. In the latter case, the optimal orders are returned for each data segment and both criteria are plotted over p .
- 6. The sampling frequency of the signals is specified, along with the range of frequencies over which analysis will be performed.
- 7. The data of each segment are Z-score normalised.
- 8. A VAR_p model is created for each segment, following the ARFIT algorithm [152].
- 9. The stability of each model is checked and a warning is displayed if the model is unstable.

10. The 95% confidence intervals of the VAR coefficients are calculated (according to a Student's t-distribution) [152].
11. The “whiteness” of the VAR model residuals is examined by checking whether the correlation matrices of the residuals (up to lag p) are significantly non-zero. This step, along with the stability test, assesses the goodness of fit of the model to the data.
12. The VAR model is Fourier transformed to the frequency domain and spectral analysis takes place along with the calculation of the PDC, gPDC, DTF, dDTF and ordinary coherence over the given frequency range.
13. The critical values of PDC and gPDC are calculated.
14. The Granger Causality Index of each modelled time series is calculated.
15. Ordinary correlations between all pairs of signals are computed.
16. All results are stored in a user-specified file.
17. Results on connectivity are plotted in various ways. The code also plots the power spectra of the original signals along with those calculated through the model, to further assess the goodness of fit of the model to the data.

Finally, FunCAT has a demo routine containing tests with sets of coupled oscillators that help examine the performance of the code and that of the individual statistical measures of causality. The test presented in the previous chapter is among them and all the plots presented there were generated by FunCAT. Moreover, the following results on generalised PDC were also derived from our code.

3.5 Results of gPDC Analysis

Figures 3.6 and 3.8 display the results of VAR modelling and gPDC analysis on LFP data from all studied cases individually, referred to as Cases 1 to 4. A single plot of the gPDC results averaged over all case studies was avoided so that individual differences and similarities in connectivity profiles could be examined. For each one, the gPDC results are presented along with the spectral comparison between the original data and the VAR model. This comparison is done separately for spectra averaged over the first four recording epochs (344 segments in total) of basal recordings and for spectra

averaged over the last four epochs during which KA is expected to have taken full effect. The averaged power spectra are plotted in the top row in each case, for each of the three brain regions under study.

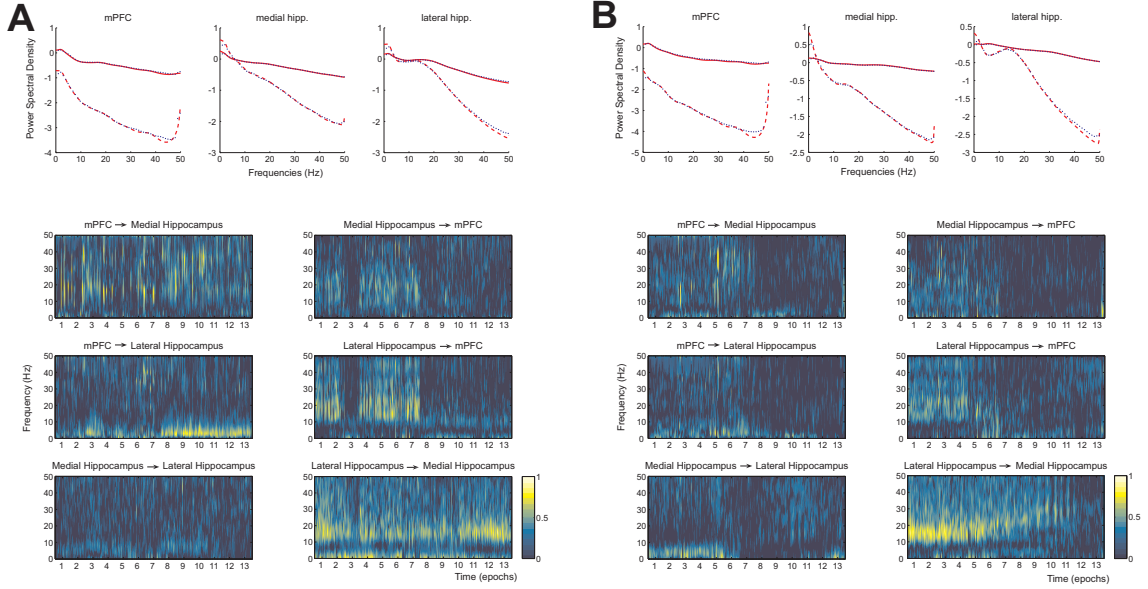


Figure 3.6: Results for LFPs from rat cases 1 (A) and 2 (B). Average power spectral densities of the VAR models, over the first four (basal) epochs (dashed lines) and over the last four epochs, when KA is expected to have taken effect (solid lines) are shown on top. LFP spectra are shown in dotted lines. Each panel represents one of the 3 brain regions (from left: mPFC, medial hippocampus, lateral hippocampus). gPDC results between all pairs of the three brain regions are shown below. Opposite directionalities are shown next to each other. False positive gPDC entries, according to the critical value (2.5.30) at 95% significance level, have been set to zero. The peaks at 50 Hz in some panels in this (and the following figures) represent 50 Hz mains noise artifacts.

In Case 1 (Figure 3.6A), gPDC indicates a bidirectional functional connection between the mPFC and the hippocampus. Information appears to be flowing continuously from the medial hippocampus to the mPFC, mostly in the beta frequency band (10-30 Hz), with the mPFC responding mainly during short bursts of activity where gPDC spans most of the frequency range. Also, the lateral hippocampus sends information to mPFC at frequencies >10 Hz, again peaking roughly around the beta band. mPFC seems to be influencing the lateral hippocampus mostly at the lower frequencies around the 0-10 Hz frequency band. gPDC also revealed a strong intra-hippocampal connection from the lateral to medial hippocampus, mostly concentrated on the beta (10-30 Hz) frequency band, with a gPDC maximum around ~ 15 Hz. Causality in the opposite direction is very low, peaking in the theta band (4-10 Hz). KA started taking effect

after the 7th epoch but did not disrupt the strong intra-hippocampal connection, while its effects were more evident in the cortico-hippocampal ones. Causal interactions from hippocampus to cortex was halted and reversed, with strong flow in the opposite direction, from mPFC to both medial hippocampus (covering most frequencies) and lateral hippocampus (localised to the 0-10 Hz band).

One distinctive feature of the basal activity epochs in Case 1 is a large gPDC decrease from hippocampus to mPFC during the third epoch with a simultaneous increase in the opposite direction, identical to the connectivity during the KA-induced enhanced-activity epochs. This reversal of directionality coincided with a period of continuous oscillatory transitions between active and quiet states in the cortex. Figure 3.7A displays gPDC between mPFC and hippocampus during epochs 2 to 4. Results from LFP analysis are shown on the left row, while the right one contains the corresponding results from FHA processed spike trains. Due to the very sparse firing, especially in the mPFC area, many segments gave unstable VAR models. gPDC from these segments was set to zero. For the same reason, VAR modelling of the firing rates was not possible since most segments yielded only zero values in the firing rate time series. Nevertheless, some segments gave significant non-zero gPDC. Figure 3.7B displays the LFP recordings from the same three epochs.

Clearly, the continuous hippocampus-to-mPFC directional interaction is interrupted and reversed only during short bursts of activity in mPFC (UP states). These bursts gave rise to high gPDC from mPFC towards mostly the medial hippocampal area. This reversed connectivity was sustained throughout the third epoch where mPFC exhibited a continuous oscillation between UP states and inactivity (DOWN states), sending information to hippocampus. Spike trains reveal the same scenario. Intuitively, significant gPDC results appeared only during the mPFC UP states when action potentials were actually produced. Again, it appears that those states coincided with causal interactions towards the hippocampus. Consequently, the third epoch gave rise to continuous high mPFC-to-hippocampus gPDC. It is noteworthy that, in some cases, results from spikes and LFPs reveal different hippocampal areas as targets of mPFC. For example, during the second UP state in epoch 2 (arrow in Figure 3.7B), spike train analysis indicated causal interactions only towards the lateral hippocampus, while gPDC from LFPs suggests mainly flow towards the medial hippocampus (arrows in Figure 3.7A).

In Case 2 (Figure 3.6B), gPDC yielded similar basal connectivity profile as Case 1. The cortico-hippocampal connection appears reciprocal, with the lateral hippocampus

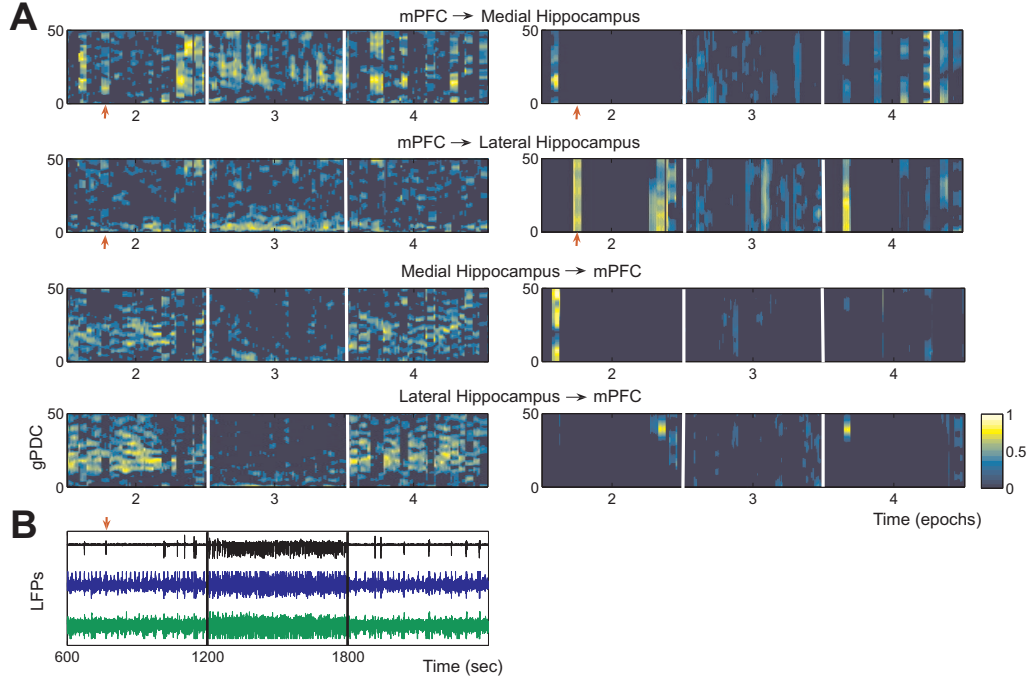


Figure 3.7: **A** gPDC results between the mPFC and the hippocampus of rat case 1 during the second to fourth recording epochs, plotted in similar fashion as in Figure 3.6. Left row displays gPDC from LFP recordings while the right one displays gPDC from FHA processed spike data. White lines indicate the limits between the epochs. **B:** LFP recordings from the mPFC (top), medial (middle) and lateral hippocampus (bottom) during the same epochs. The intervals of no recordings between the epochs have been removed. Black lines indicate the limits between epochs and time axis displays the starting time of each epoch (180 sec duration). The arrows indicate the second UP state in the mPFC signal during epoch 2 and the mPFC-to-hippocampus PDC results corresponding to that segment. These results are discussed in the text.

sending no information below ~ 10 Hz and the mPFC sending information back predominantly below this frequency boundary. Again, most of the intra-hippocampal directional interaction is found in the lateral-to-medial pathway, in the beta band, with flow below this range going only in the opposite direction, towards the lateral hippocampus. The disruptive effect of KA in the cortico-hippocampal connections is also evident here but this time with no apparent reversal of flow towards the hippocampus. This time KA also disrupted the intra-hippocampal connectivity as well, as there is a gradual drop of gPDC in the lateral-to-medial pathway and a dramatic reduction in the opposite direction. This reduction is followed by a small increase at higher frequencies suggesting a partial reversal of intra-hippocampal connectivity.

Case 3 (Figure 3.8A) presents a similar strong reciprocal mPFC-lateral hippocampal con-

nection as before. Here, the medial hippocampus-to-mPFC connection appears much less prominent, while the high gPDC values in the mPFC-to-hippocampus connections, present mostly in short bursts of activity, are more concentrated at low frequencies. Intra-hippocampal connections are again directed from the lateral to the medial although the strong peak on the beta band is now absent. KA took effect sooner in Case 3 (6th epoch compared to 8th and 7th in the two previous cases), again disrupting interactions from hippocampus to mPFC, but only for a transient period. After epoch 9 the basal connectivity is partly restored. Also, here the mPFC-to-medial hippocampus connection was greatly enhanced in the low frequency band (0-10 Hz), while the mPFC-to-lateral hippocampus was interrupted by KA. Finally, KA action also resulted in a partial reversal of the intra-hippocampal connectivity, with gPDC presenting a substantial drop of the medial-to-lateral hippocampus connection strength and a simultaneous rapid increase in the opposite direction, most prominently in the low frequency band.

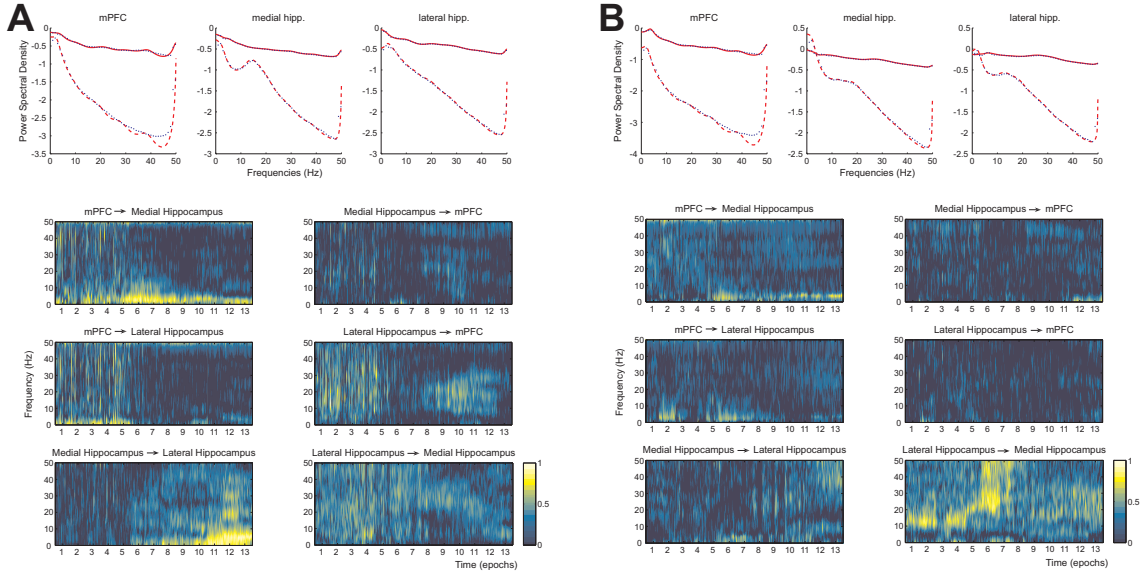


Figure 3.8: Same set of plots as in Figure 3.6 for the third (A) and fourth (B) rat case studies.

In Case 4 (Figure 3.8B), the hippocampus-to-mPFC connections are almost absent. From the mPFC, information flow is directed mostly towards the medial hippocampus, with only a transient increase towards the lateral, corresponding to a period of similar UP/DOWN-transitions in mPFC-activity as epoch 3 in Case 1. The most distinctive characteristic in the intra-hippocampal connectivity is again the very strong gPDC in the lateral-to-medial pathway on a frequency range around the beta band where gPDC is maximal while in the theta band it is almost absent. This strong connectivity is

again partially reversed under KA activity, while the mPFC-to-medial hippocampus connection is enhanced at low frequencies, as in Case 3.

Case 4 exhibited a firing rate after KA administration which was high enough to allow gPDC analysis to be performed on spike data. However, the very sparse firing during the initial basal epochs allowed VAR modelling of the firing rates to be performed only on the last seven epochs. It also resulted in many segments with unstable VAR models for the FHA processed spikes. In Figure 3.9, gPDC results from the FHA processed spikes are displayed in a similar fashion as for the LFP results. Power spectra comparisons are performed only for the KA epochs where the VAR models were mostly stable. Results from the firing rates are not shown as they were practically identical with the last seven epochs in Figure 3.9. The basal epochs exhibit no specific sustained connectivity. The most distinctive characteristic is the pronounced increase in medial-to-lateral hippocampal connectivity after the KA administration, sustained until the end of the recordings. This increase coincides temporally with the gPDC increase observed in the corresponding LFP results. Although it does not verify the reversal in hippocampal connectivity observed through LFPs, it supports the KA-induced increase of information flow from medial to lateral hippocampus.

3.6 Discussion

A significant challenge in neuroscience lies in determining how complex interactions between discrete populations of neurons elicit functional output. PDC analysis was applied to LFP and unit activity recorded at mPFC and medial (CA1) and lateral (CA3) hippocampal recording sites, to address interactions between these regions. Generalised PDC (gPDC) [12], was used in order to avoid distorted connectivity results due to different scaling of the data in cortex and hippocampus. Spike train data from each brain region were replaced by their firing rate or a continuous function generated through the FHA algorithm and grouped using principal component analysis. Before performing the gPDC analysis, we first examined two spike processing techniques and the effect of the binning scheme of firing rates in computing PDC.

Connectivity results were consistent between LFP recordings from different animals during basal recording conditions. Within the hippocampus, the predominant connectivity under isoflurane anaesthesia was clearly from lateral towards medial, concentrated in the beta (10-30 Hz) frequency band, peaking strongly around 15 Hz. Usually, little informa-

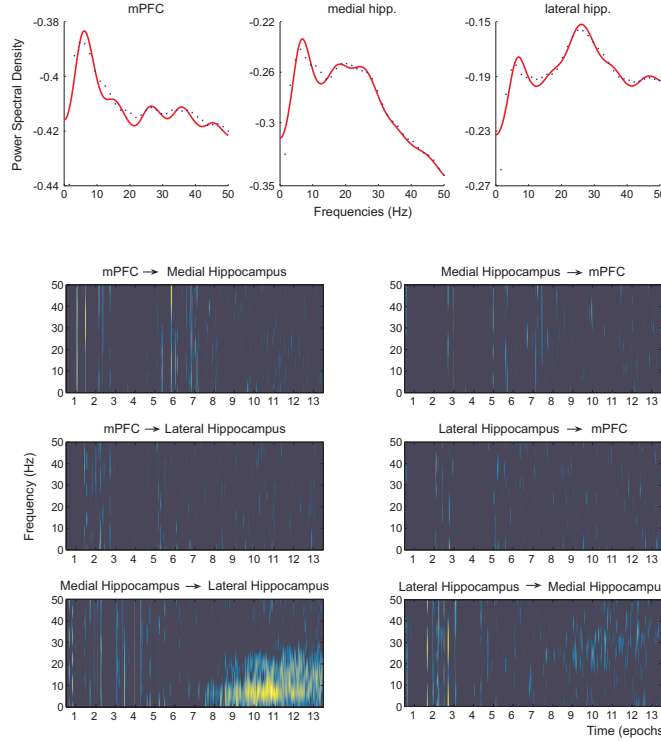


Figure 3.9: Results for the first principal components of FHA processed spike trains in rat case 4. The sets of plots are similar to the ones in Figures 3.6 and 3.8. Only the power spectra of the last four epochs are displayed as the basal epochs contained too many unstable VAR models. We note that gPDC analysis on the last seven epochs of firing-rates processed spike trains gave identical results as the ones presented here.

tion was sent in the opposite direction, mostly in the theta (4-10 Hz) frequency band. This lateral-to-medial directionality was consistent with the established neuroanatomy of the hippocampus, i.e. pyramidal cells located in CA3 projecting, via Schaffer collaterals, to cells in CA1 (see Chapter 5).

Analysis of LFP activity further revealed a reciprocal flow of information between the mPFC and hippocampus. In most cases, the hippocampus showed an almost continuous information flow towards the mPFC, with the lateral hippocampus-to-mPFC gPDC peaking predominantly between 10 and 30 Hz. In the opposite direction, the mPFC-to-hippocampal information flow was less strong during periods of mPFC inactivity, but interestingly peaked during short bursts of mPFC activity (Figure 3.10). These activity and inactivity periods are reminiscent of the so called UP and DOWN states observed during the cortical slow oscillation (< 1 Hz, see Chapter 4). This suggests that the hippocampus feeds the mPFC, while the latter is silent, in an almost-continuous fashion, with the mPFC responding in bursts towards both the medial and the lateral

hippocampus.

The hippocampus forms multiple connections with other structures in the brain, including the mPFC. Specifically, ventral CA1 (medial hippocampus) projects to mPFC, while no projections from CA2, CA3, or dentate gyrus to mPFC are known [38, 73, 83, 108, 110, 184] (see Chapter 7). This CA1-mPFC projection is primarily unidirectional, given that no direct return projections have been identified [28, 83, 103, 170, 184, 200]. The monosynaptic hippocampal-to-mPFC pathway innervates both pyramidal cells and interneurons in the rat [73, 107, 127, 204], and has been implicated in memory acquisition and consolidation, specifically for working spatial memory [82] (Figure 3.10).




Figure 3.10: Simplified schematic representation of the information flow, under basal conditions, between the structures recorded in the present study. Schaffer collaterals project from lateral (CA3) to the medial (CA1) hippocampal area (Hpc) and monosynaptic projections have been reported from medial hippocampus (CA1) to the mPFC. No monosynaptic pathways have been observed from mPFC back to the hippocampus, but it is possible that other areas can act as relay structures (see text). The arrows represent the directionality as assessed by gPDC.

Interestingly, the observed mPFC-to-hippocampus directionality appears inconsistent with an unidirectional monosynaptic projection from hippocampus to mPFC. This observation may be explained by the participation of a “relay” structure, possibly the entorhinal cortex (perforant pathway) or the nucleus reuniens of the thalamus, a structure that receives strong input from prelimbic and infralimbic cortices of the mPFC and projects to the hippocampus [222, 223].

Recent electrophysiological studies (which we review in Chapter 7) suggest that the neocortical slow oscillation [192, 195] which engages neurons in prefrontal and other cortical regions into synchronous transitions between depolarised (UP) and hyperpolarised (DOWN) states during deep sleep and anaesthesia, influences activity in hippocampal cells even though such activity lacks the neocortical bimodality of the slow oscillation

[91, 104]. Specifically, it has been suggested that cortical UP-states result in increased activity of dentate and most CA1 neurons, as well as a higher probability of ripple events in the hippocampus [104]. The delay between this cortical influence and its effect on hippocampus is in the order of tens of milliseconds [149, 188]. The high model order that was applied during the present study contains such delays, and the gPDC results presented in Figure 3.7 (epoch 3) seem to support this notion, implying that during slow oscillation-type cortical activity, information flows mainly from mPFC to hippocampus. Further, these results suggest that this information flow appears generally during the cortical UP state, even isolated ones.

It can be argued that the clear increase in mPFC-to-medial hippocampus gPDC observed during these UP states is a numerical artifact produced by poor VAR modelling, resulting from the non-stationarity that an UP state introduces in a modelled LFP segment. Nevertheless, the absence of similarly strong gPDC increase in the corresponding mPFC-to-lateral hippocampus panels supports the validity of our results.

In the present study, kainic acid administration resulted in a variety of effects on connectivity profiles, yet two major effects occurred in most case studies. Namely, KA caused a partial disruption of connectivity between the mPFC and hippocampus, in both directions and, following KA, there was a weakening in the lateral-to-medial connection, often accompanied with a strong increase of the opposite pathway, leading to a gradual and partial reversal of connectivity in the hippocampus.

The use of KA is well established as a model of status epilepticus (SE) and spontaneous seizures in the rat, and the hippocampus is particularly susceptible to pharmacological manipulation with KA [20, 29, 131, 151]. KA receptors are broadly distributed in the hippocampus, presynaptically regulating glutamate release at mossy fibre-CA3 synapses [180], and GABA release between interneurons [52, 150] and at interneuron-to-pyramidal cell synapses [42, 137, 169]. The mechanisms by which SE and spontaneous seizures are generated in animal models with KA are not fully understood. Our results confirm the particular susceptibility of the hippocampus to KA. The decline in lateral-to-medial hippocampal information flow, and the subsequent reversal of directionality, exhibited by some rats administered KA in these experiments, is intriguing, in light of both the described anatomy of the hippocampus, and the suggested mechanisms for KA-induced limbic seizure. The perforant path does not represent the sole pathway within the hippocampus (e.g. the entorhinal cortex projects directly to CA1 via the temporoammonic pathway); however reverse projections from CA1 to CA3 have not been reported. Fur-

thermore the susceptibility of CA3 mossy fibres to KA does not fit with medial-to-lateral directionality. One caveat is the influence of anaesthesia on KA-evoked activity which may underlie the hippocampal directionality observed in the present study. In addition, different sub-populations of neurones may have been recorded because of variations in the precise array-placement between experiments.

Our results further indicate that applying PDC on spike trains can be problematic due to their highly non-stationary nature. Preprocessing of the unit data is necessary in order to generate a stationary series. While the effects of such processing may not always be clear, they should be taken into account in subsequent analysis. Both of the preprocessing techniques, analysed and applied here, yielded similar results as binning parameters were chosen, according to our presented study, so that the resulting firing rate did not deviate much from the instantaneous firing rate given through the FHA. In general, processing parameters should be chosen carefully, and according to the specific case under study. The conclusions reached by our binning scheme study can only serve as merely guiding suggestions and we believe that there should not be a black box rule for the binning scheme. It will depend on the experiment and the experimenter's intentions. For example, grouping the spikes into large bins leads to a loss of information about the firing patterns of the neuron. Fast spiking neurons require small bins and large overlaps so that the detailed and rapidly changing characteristics of the firing rates can be captured. In contrast, small bin sizes may be inadequate for a sparsely firing neuron. In this particular study, anaesthesia resulted in low firing rates. Thus, spike train gPDC analysis did not reveal any significant connectivity, with the exception of Case 4, where spike train gPDC did parallel elements of the connectivity revealed by corresponding LFP analysis; namely the increase in medial-to-lateral hippocampus information flow under KA. The same conclusion can be drawn from the analysis of the three basal epochs of Case 1 presented in Figure 3.7. Despite the very low firing of mPFC, spike train analysis did support the mPFC-to-hippocampus connectivity during the mPFC UP-states. Thus gPDC analysis of spike trains can be a useful complementary tool along with analysis of LFPs.

PDC represents a novel technique for assessing interactions between brain regions and revealing their directionality. Its recent variation, gPDC, is a powerful tool which, combined with data normalization, can overcome distortions caused by differences in the electrophysiological signal scaling. Nevertheless, caution is necessary when applying such statistical tools in the analysis of neurophysiological data as inappropriate use

can lead to false conclusions. There are many issues, mainly related with the VAR modelling of the signals, such as the scaling and stationarity of time series data, the size of the sampled data and VAR model order that need to be taken into account when constructing the VAR models which is the most crucial part. We believe that our numerical package FunCAT can provide a useful tool in this analysis, by being easy to use, by including the spike train processing routines, and most importantly by combining various connectivity tools that help give a more complete picture of the underlying structure.

Part II

Modelling Cortical and Hippocampal Dynamics and Their Interactions

Cortical Sleep Rhythms and a Model of the Slow Oscillation

4.1 Introduction

Sleep is defined behaviorally by four criteria [116]:

1. reduced motor activity,
2. decreased response to stimulation,
3. stereotypical postures,
4. relatively easy reversibility (unlike coma or anaesthesia).

Unlike most body parts, the brain is busy during sleep as well. In fact, through EEG recordings, one can witness a rich ensemble of different brain rhythms throughout a human sleeping sequence. According to these rhythms, sleep can be divided into at least 5 stages, with different EEG characteristics and progressively higher waking thresholds [33]:

1. Stage 1 is the phase transition between wakefulness and sleep. It consists of a relatively low amplitude EEG with mixed frequencies, mainly slow alpha and theta activity.
2. Stage 2 is distinguished by the emergence of spindles and K-complexes (see below).
3. Stage 3 is a mixture of spindles and delta waves (20%-50%).

4. Stage 4 is dominated by delta activity with only traces of spindles.
5. REM-sleep is characterised by waking-type EEGs, rapid eye movements, loss of muscle tone and dreaming.

The first four stages together constitute the so-called “non-REM” sleep while stages 3 and 4 are often referred to as deep sleep or “slow wave sleep” (SWS).

Every night, at least four or five cycles of the above stages take place during sleep, with a period of 70-90 mins, progressively reaching less deep stages (Figure 4.1). Approximately half of sleeping time consists of stages 2 and 3, and only 5%-15% of stage 4. REM takes up 20%-25% of sleeping time and indicates the end of a non-REM/REM sleeping cycle [33].




Figure 4.1: Left: Sleep cycles for an adult. Right: EEG recordings from the four non-REM stages (spindles indicated by arrows). Taken from the scholarpedia article http://www.scholarpedia.org/article/Neurobiology_of_sleep_and_wakefulness.

Wakefulness and REM sleep are characterised by sustained beta (20-30 Hz) and gamma (30-80 Hz) rhythms whereas SWS is mainly defined by three rhythms: spindles (7-15 Hz), delta waves (1-4 Hz) and slow oscillation (0.5-1 Hz) [192]. The first two are examined below, while the slow oscillation is analysed in much more detail in section 4.3.

4.2 Cortical Sleep Rhythms

4.2.1 Spindles

The EEG spindles, appearing during sleep stage 2 in humans, are the landmark of sleep onset and have been extensively studied *in vivo* and *in vitro* and successfully modelled

computationally (see [60] for an extensive review). They are associated with loss of perceptual awareness and their intensity increases as non-REM sleep progresses into stage 3. They consist of waxing and waning field potentials, oscillating at 7-15 Hz, grouped in sequences that last for 1 to 3 secs and recur once every 3 to 10 secs [33, 58].

Spindles are synaptically generated in thalamic networks involving two cell types which oscillate in synchrony: reticular GABAergic thalamic neurons (RE) and excitatory thalamocortical cells (TC). RE cells act as pacemakers, displaying synchronous bursts at around spindle frequency as a result of incoming excitatory postsynaptic potentials (EPSPs). Their inhibitory output is received by TC cells that get hyperpolarised and start bursting intrinsically (often skipping cycles), further depolarising and synchronizing RE cells. Key elements in the TC bursts are the intrinsic hyperpolarisation-activated rebound current I_h and low-threshold Ca^{2+} current I_T [60, 191].

One characteristic of spindles in the thalamus of an intact brain is that they occur synchronously over large thalamic areas. Cortically recorded spindles are also coherent over relatively large areas [33]. This supports the idea that thalamic activity is synchronised by cortical long range connections. Indeed it has been shown that by removing a whole cortical hemisphere of a cat brain, thalamic spindles continue to emerge but are now completely independent and asynchronous [48] (Figure 4.2).




Figure 4.2: Multiple *in vivo* extracellular field potential recordings in the cat thalamus. In the intact thalamocortical system, spindle oscillations were synchronised over large distances (7 mm between Th1 and Th8). After removal of the cortex, the large-scale synchrony was largely abolished. Taken from [48].

4.2.2 Delta Waves

Delta waves are large amplitude EEG patterns (1-4 Hz) and have two components, a cortical and a thalamic one. The cortical component, shown to arise between cortical layers 2 to 3 and 5, is characterised by synchronous delta waves, throughout the entire

neocortex, which survive thalamectomy [191, 224].

Delta oscillations are also an intrinsic rhythm of hyperpolarised TC cells, resulting from the interplay between the two aforementioned intrinsic currents (I_T and I_h) [60]. In the isolated thalamus, these oscillations are asynchronous between TC cells. Delta waves appear during later stages of sleep than spindles and at more negative membrane potentials. These facts led to the postulation of a progressive hyperpolarisation of TC neurons as SWS moves to deeper stages [194].

4.3 Slow Oscillation

The slow oscillation (SO) brain rhythm, arising during SWS in humans and animals, was first introduced and analysed in a series of three seminal papers by Steriade et al. that appeared in the Journal of Neuroscience in 1993 [193, 195, 196]. Below we examine them along with the major characteristics of this oscillation, that they describe, and its grouping role for other sleep oscillations. We return to this grouping role, in relation with hippocampal activity, in Chapter 7.

4.3.1 General Characteristics of the Slow Oscillation

The SO was first described in [195], recorded intracellularly in the neocortex of anaesthetised cats. As displayed in Figure 4.3A (from naturally sleeping animals), it consists of step-like rhythmic transitions of neuronal membrane potentials between low depolarising envelopes, 0.8-1.5 sec long, with superimposed action potentials (UP states) alternating with hyperpolarisations where spikes are absent (DOWN states). Such sequences recur rhythmically at less than 1 Hz, with a prevailing oscillation between 0.3 and 0.4 Hz in cells of urethane-anesthetised cats, which varies between animals and between anaesthetics (e.g. under ketamine it increases to 0.6-1 Hz). In [195], the SO was described as a “massive population event” since it was observed in all recorded neocortical areas and all types of pyramidal cells and interneurons [49]. Silences and increased firing occurred simultaneously in units recorded at 7 mm distances throughout the cortex [59].

The depolarised UP states are due to synaptic excitatory inputs (mostly non-NMDA EPSPs) but the hyperpolarised DOWN states are not simply due to synaptic inhibition but due to disfacilitation (removal of synaptic inputs) [207]. The neuron can be in the

DOWN state for several seconds, much longer than expected by GABA inhibition or spike refractoriness. The almost flat membrane potential is another indication of the lack or extreme paucity of synaptic activity.

The form of the SO on the network level as represented through EEG signals was first described in [196]. Depth-negative (surface positive) half-waves of EEGs are associated with the cellular depolarised UP states, while depth-positive (surface negative) half-waves are associated with hyperpolarised DOWN states [49, 59, 192]. The other rhythms (spindles, delta) are superimposed on this slower oscillation of the EEG amplitude (Figure 4.3A).

Similarly to depth EEG, LFPs during SO consist of a positive envelope (delta wave), associated with the neuronal silence of the DOWN states, followed by a sharp negative peak, correlated with the increase of firing in UP states (Figure 4.3B). During the periods between the slow wave complexes, LFPs exhibit faster oscillations and the LFP-units relation resembled awake or REM activity.

A possible dependence of SO generation on the thalamus seems improbable since the slow rhythm survives extensive ipsilateral thalamic destruction [196]. Though the thalamus is not essentially implicated in generating cortical SO, thalamocortical pulse trains (at 10 Hz) are able to increase its frequency. Interactions between cortex and thalamus during SO were further explored in the third paper by Steriade et al. [193], in which they reported a SO in 65% of recorded thalamic reticular neurons and in 44% of thalamocortical neurons. The reticular SO was closely related to EEG wave complexes that had the same frequency. The SO in thalamocortical neurons coexisted on average with delta and spindle oscillations in the same neurons.

Survival of the cortical SO under thalamectomy suggests that it is generated within the cortex. This cortical nature was further shown by its absence in the thalamus of decorticated animals [208] and its presence in isolated (from subcortical inputs and long-range connections) cortical slabs [205], and even in cortical slices maintained *in vitro* [175]. Nevertheless the SO synchronization over the cortex was lost after disconnection of intracortical synaptic links indicating that such recurrent connections are necessary for cortical synchrony [7].

Finally the SO has been observed through EEG and MEG recordings in human SWS as well [139, 148]. Interestingly, it has been shown that each cycle of the SO is a traveling wave, originating mostly in frontal cortical regions and propagating over the scalp in an anterior-posterior direction at an estimated speed of 1.2-7 m/sec [139].

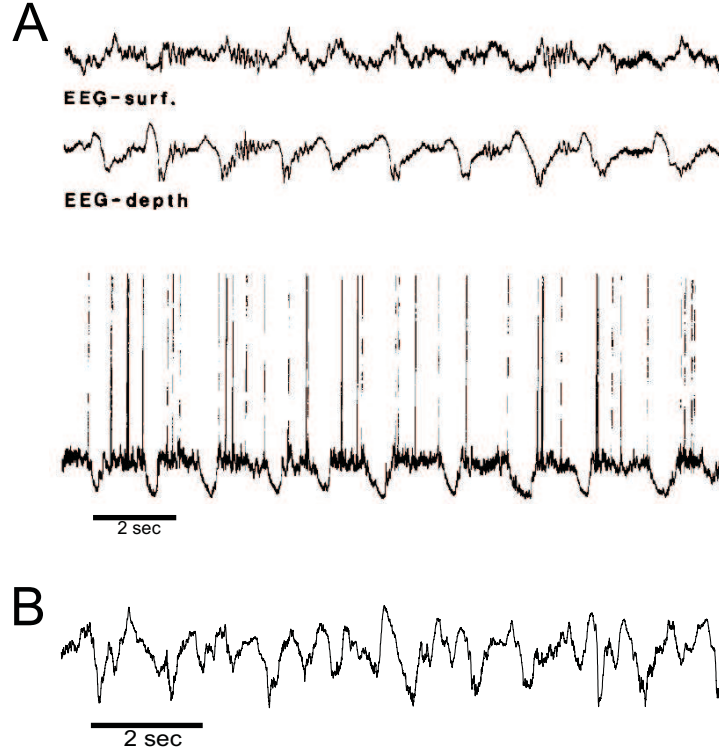


Figure 4.3: **A:** Slow oscillation in *in vivo* EEG signals and intracellular recordings from an anaesthetised cat. Raw data from a neuron from the left motor cortex, recorded simultaneously with surface and depth EEG from the same area. The cell membrane potential is oscillating at ~ 0.7 Hz. Note the clear UP and DOWN states and the presence of spindle sequences during UP states and depth-EEG positivities (delta waves) during the DOWN states. Taken from [49]. **B:** Slow oscillation in an *in vivo* LFP recording from the prefrontal cortex of a naturally sleeping rat [R. Mason, unpublished data].

4.3.2 Grouping of Spindles and Delta Waves by the Slow Oscillation

A possible relation between the SO and other SWS brain rhythms, specifically spindles and delta waves, was first studied in [196]. It was found that in many recorded cortical neurons the membrane potential SO was combined with delta oscillation with both intracellular rhythms phase locked to the slow and delta oscillations of the EEG signal. In some recorded neurons the SO was combined with spindles as well. The grouping role of the SO on these (and faster) rhythms has been extensively studied by Steriade et al. [192–194].

An SO cycle can be followed by a brief sequence of spindles in thalamocortical neurons and in human EEGs. Such EEGs display the so-called “K-complex” pattern, a sharp

depth-negative component followed by a slower depth-positive component and a few spindle waves [92]. Moreover, cortical hyperpolarised DOWN states are accompanied by delta waves recurring with the same frequency as the stereotypical intrinsic delta oscillations of the TC cells. The grouping role of the SO on spindles and delta waves is also suggested by the observation that such purely thalamic oscillations, without the cortex, become nonsynchronous.

The above observations on the three major SWS rhythms led to the postulation of a scheme for the generation and grouping of these rhythms under the SO [193]. The suggested steps, taking place during every SO cycle in corticothalamic circuits, are the following (Figure 4.4):

1. The activity generated by cortical cells during their DOWN-UP state transition is projected down to the thalamus, exciting thalamic neurons, both RE and TC cells.
2. RE cells display rhythmic depolarisations, effectively grouping their intrinsic spindle-frequency oscillations under the SO.
3. In parallel, TC cells are also activated by the DOWN-to-UP transition but fewer spikes are produced, due to the synchronous inhibition coming from the GABAergic RE cells. Thus the intrinsic delta-oscillation activity of the hyperpolarised (during deep sleep stages) TC cells gets interrupted by the excitation from the cortical UP state and the inhibition from the reticular spindles. Hence their internal delta oscillations dampen and get revived with SO frequency.
4. Finally TC delta oscillations are fed back to the cortex.

Faster rhythms are also grouped by the SO with episodes of beta/gamma frequency range oscillations (20-80 Hz) [191], and even ultra-fast ripple episodes (~ 200 Hz) [87], appearing during the UP states and disappearing in DOWN states (see Chapter 7).

4.4 Computational Models of the Slow Oscillation

The cortical SO during SWS has been modelled by many theoretical and computational approaches of different architectures and various levels of complexity [18, 47, 99, 101, 117, 141, 158, 225]. Here we review three of these computational models. The first two




Figure 4.4: The functional scheme proposed in [193] for the SO grouping of spindles and delta waves is based on this corticothalamic circuit. Direction of axons is indicated by arrows and excitatory or inhibitory synapses are indicated by + or -. Taken from [196].

are based on different potential mechanisms underlying the initiation and termination of the UP and DOWN states while the third includes a combination of both.

4.4.1 The Bazhenov et al. Model

Bazhenov et al. [18] implemented a computational model of a thalamocortical circuit with both cortical and thalamic networks being one-dimensional two-layer arrays (Figure 4.5). The thalamus is modelled by single-compartment Hodgkin-Huxley models of RE and TC cells, while the cortex consists of pyramidal cells (PY) and interneurons (IN) modelled by two-compartment (dendritic and axosomatic) Hodgkin-Huxley cells. PY cells contain AMPA, NMDA and GABA_A receptors, RE cells contain AMPA and GABA_A receptors and TC cells AMPA, GABA_A and GABA_B receptors. Importantly, short term depression of intracortical excitatory connections and Poisson-process miniature EPSPs and IPSPs (minis) of fixed amplitude are also included.

The results of their simulations indicate that:




Figure 4.5: Network geometry in the Bazhenov et al. model. Taken from [18].

- UP states are initiated by random summation of mini EPSPs in some PY cell. This cell spikes, spreading activity throughout the network via recurrent excitatory synapses only if the neighbouring cells are also relatively depolarised by minis.
- The UP state terminates through PY-PY synaptic depression.

A large network or a sufficiently high rate of mini EPSPs is required by this mechanism for consistent UP state transitions. The UP state spreads from the cortex down to the thalamic network. TC cells often exhibit waning spindle episodes which are terminated by non-synchronous PY-to-RE input. Input from TC to PY cells further extends the UP states and shortens the DOWN states.

4.4.2 The Compte et al. Model

Compte et al. [47] implemented a one-dimensional cortical model of 1024 PY and 256 IN cells. Single cell models consisted of two-compartment (dendritic and somatic) Hodgkin-Huxley cells for PY and single-compartment Hodgkin-Huxley cells for IN. AMPA, NMDA and GABA_A synapses were implemented and the connections of each cell were stochastically distributed among the network, following a Gaussian probability distribution centered around the cell (Figure 4.6). The standard deviation (SD) of this distribution and the number of connections between cell types were fixed, with PY cells having more extensive connections than IN cells.

The mechanism for SO generation in the model is the interplay between neuronal spontaneous firing, propagating via recurrent excitation, and a negative feedback due to slow activity-dependent K⁺ currents. Specifically:




Figure 4.6: **A:** Schematic representation of the spatial connectivity in the Compte et al. model [47]. **B:** Raster plot displaying the synchronous population oscillation (5 neighbouring cells per site, sites are spatially separated by $500 \mu\text{m}$). **C:** Intracellular somatic voltage (V_m) and intracellular sodium concentration $[Na^+]_i$ of 3 representative pyramidal neurons. **D:** Blow-up of an individual depolarised episode of the cells in C. In C and D, arrows point at -75 mV for voltage traces and at 10 mM for internal sodium concentration. Taken from [47].

- Occasional firing of a sufficient number of close spontaneously-active neurons triggers a cascade of recurrent excitation that brings the network in that area in the UP state. This activity propagates throughout the whole network in a wave-like fashion.
- During the UP state, K^+ currents and particularly the Na^+ -dependent K^+ current (I_{KNa}) accumulate in active PY cells reducing their excitability and finally bringing the network to the DOWN state. After the K^+ currents decay, the UP state can be reached again.

During the UP state, PY and IN cells fire in phase with IN cell spiking starting first, in response to the traveling activity-wave, and lasting longer.

The authors argue that the SO is due to an intrinsic bistability between a stable UP and a stable DOWN state of the cortical system, with I_{KNa} inducing the slow switching between them. This led to the prediction that without the I_{KNa} current, both UP and DOWN states are stable and the cortex can be made to switch between them by external stimulations, yet this has not been tested experimentally. Since the UP state propagates through the network, blocking recurrent excitation should abolish the bistability, while blocking GABA_A receptors should result in more intense discharge in the UP state and longer oscillation periods. Therefore they argue that anaesthetic-dependent changes in the network excitability could explain the variety in SO patterns (duration of UP or DOWN states and oscillation frequency), seen in the various recordings.

4.4.3 The Hill and Tononi Model

Hill and Tononi [99] created a much larger and more complex corticothalamic model with 65,400 neurons and 4,860,450 synapses. It consists of a primary and a secondary visual area, organised in 3 layers each, two dorsal thalamic nuclei and two reticular thalamic nuclei (Figure 4.7). All neurons are single-compartment Hodgkin-Huxley cells but action potentials are modelled by a hybrid between Hodgkin-Huxley and Integrate-and-Fire models. AMPA, NMDA, GABA_A and GABA_B synapses are included along with synaptic depression. Intrinsic ion-channel currents include the depolarisation-activated potassium current (I_{KNa}), while diffuse neuromodulatory projections to the thalamus, spontaneous stochastic activity in the form of optic track firing and mini PSPs are also implemented.

Transition to SO (considered to signify sleep) from an activated-awake state results from an induced increase in potassium leak conductance, g_K , corresponding to reduction of neuromodulatory input. The g_K -increase eventually causes synchronous propagating SO in both cortex and thalamus. Simulations revealed that the alteration between UP and DOWN states is caused by intrinsic (I_h , I_{KNa} , I_{NaP}) and synaptic currents. Specifically:

- For the transition to an UP state, the activation of a voltage-dependent persistent-sodium current I_{NaP} is necessary. This activation is accomplished by a variety of means, including minis, synaptic input by other cortical or thalamic areas, or intrinsic activation of hyperpolarisation-activated I_h .
- Currents I_{NaP} and I_{KNa} tend to covary in the UP state balancing each other. The UP state is maintained intrinsically by I_{NaP} , and synaptically through excitatory




Figure 4.7: Schematic diagram of the thalamocortical model of the Hill and Tononi model [99]. Primary thalamocortical circuit (left) including a 3-layered primary visual cortical area (Vp), reticular nucleus (Rp), and dorsal thalamus (Tp) and secondary visual area Vs (with its associated thalamic sectors Rs and Ts) (right). Visual inputs (left) excite the primary thalamus (Tp). Diffuse neuromodulatory systems project throughout the entire thalamocortical network. Taken from [99].

currents and their balance with $GABA_A$. Blocking any of these factors during an UP state disrupts it and terminates the SO.

- The DOWN state is initiated through both I_{KNa} (as in the Compte et al. model [47]) and synaptic depression (as in the Bazhenov et al. model [18]). I_{KNa} is essential for stopping the UP state but also for maintaining excitation-inhibition balance throughout its duration. Synaptic depression also aids in stopping an UP state and in the overall synchronization of the SO.

Finally, removal of inputs from thalamus does not stop the cortical oscillations, which supports the notion of a pure cortical nature of the SO.

4.5 Implementation of the Compte et al. Slow Oscillation Model in Brian

Due to its intermediate complexity, we chose to implement the Compte et al. model, and the SO it exhibits, within the framework of our eventual cortico-hippocampal model. Since we did not have access to the original numerical code of the model, written in C++, we needed to implement it from start, based on its description in the literature [46, 47]. In this section, we describe the implementation of this model in the Python-based spiking neural networks simulator “Brian” [84] and the resulting differences with the original model.

4.5.1 Single Cell Models

Pyramidal neurons are modelled by two-compartment Hodgkin-Huxley cells with a dendritic and an axosomatic compartment communicating through an internal current. The axosomatic compartment contains the typical sodium, potassium and leakage currents along with a fast A-type K^+ current (I_A), a slow non-inactivating K^+ current (I_{KS}) and a Na^+ -dependent K^+ current (I_{KNa}). The dendritic compartment contains a high-threshold Ca^{2+} current (I_{Ca}), a Ca^{2+} -dependent K^+ current (I_{KCa}), a persistent sodium current (I_{NaP}) and an inwardly rectifying K^+ current (I_{AR}). Parameter values are taken from various references, cited in [47]. Interneurons are modelled by the single-compartment Wang-Buzsàki model [228] (see Chapter 5) with all parameters having their default values.

According to Compte et al. [47], when a depolarising current of 0.25 nA is applied for 0.5 sec in the pyramidal cell model, it responds as a regular spiking neuron with firing frequency adaptation and average firing rate of ~ 22 Hz. To test our single cell implementation, we applied the same current on the dendritic compartment of our pyramidal cell model. The results are plotted in Figure 4.8, where the panels depict (from top) the axosomatic membrane potential, the intrinsic Na^+ concentration and the Ca^{2+} concentration. Indeed, the neuron exhibits regular spiking with adaptation and its mean firing rate is roughly 22 Hz. Intrinsic sodium and calcium concentrations rise with each spike, but note that the sodium concentration decays very slowly.

When the same current is applied to the modelled interneuron, it responds as a fast spiking neuron with firing rate of ~ 75 Hz. This holds for our implementation as well (Figure 4.9). The firing rate, computed through the power spectrum of the membrane




Figure 4.8: Firing pattern of the pyramidal cell model, described in [47], when injected with a depolarising current of 0.25 nA for 0.5 sec (100-600 msec). From top: axosomatic membrane potential (mV), intrinsic Na concentration and intrinsic Ca^{2+} concentration (both in mM). The cell responds with regular adaptive firing at approximately 22 Hz.

potential (not shown), is 78 Hz.

The only sources of noise in the network, apart from the randomness in the connectivity (see below), are some intrinsic neural properties. Namely the leakage resting potentials and maximum conductances in both pyramidal and interneuron cells and the conductance for the coupling current between dendritic and axosomatic compartments in the pyramidal cells. All these parameters are Gaussian-distributed throughout the network with fixed means and SD, as in [47].




Figure 4.9: Firing pattern of the interneuronal model, described in [47], when injected with a depolarising current of 0.25 nA for 0.5 sec (100-600 msec). The firing rate is 78 Hz.

The initial conditions in our simulations are also Gaussian distributed over the network, around the steady state of the corresponding single cell model, with deviations equal to 10% of the means. Note that the resting membrane potential of a pyramidal cell is slightly below -75 mV. This means that when an inhibitory spike reaches the axosomatic

compartment of a pyramidal cell in its steady state, it will give rise to an EPSP instead of an IPSP, exciting the cell (this becomes important at the beginning and end of an UP state, where most pyramidal cells are close to their steady states).

Finally, due to an error in the Compte et al. paper [47], the maximal conductance of the leakage current was actually Gaussian-distributed over the pyramidal cells with: $\bar{g}_L = 0.02 \pm 0.002$, instead of $\bar{g}_L = 0.0667 \pm 0.0067$ as reported in [47]. Additionally, each pyramidal cell dendritic compartment had the same leakage conductance as its axosomatic compartment, which was not mentioned [A. Compte, personal communication]. After the change in \bar{g}_L and the inclusion of the dendritic leakage current, pyramidal cells become slightly more excitable. Since we were informed of these errors after all simulations of this chapter were performed, and since they did not significantly affect any results at the network level, we omitted these changes in our version.

4.5.2 Modelling Synapses

In the original model all post-synaptic currents are given by:

$$I_{syn} = \bar{g}_{syn} s_{syn} (V - V_{syn}), \quad (4.5.1)$$

with \bar{g}_{syn} a fixed maximum conductance, $s_{syn}(t)$ the synaptic gating variable and V_{syn} the reversal potential. Fast AMPA- and slow NMDA-mediated excitatory synapses are implemented between pyramidal cells and from pyramidal to interneuronal cells. Recurrent and interneurons-to-pyramidal cells inhibition is implemented by fast GABA_A-mediated synapses. The synaptic variable for fast AMPA and GABA_A connections is given by:

$$\frac{ds_{syn}}{dt} = \alpha_{syn} f(V_{pre}) - \frac{s_{syn}}{\tau_{syn}}, \quad (4.5.2)$$

with $\tau_{AMPA} = 2$ msec, $\tau_{GABA_A} = 10$ msec, $\alpha_{AMPA} = 3.48$ and $\alpha_{GABA} = 0.5$ for AMPA and GABA_A PSPs respectively. V_{pre} is the membrane potential of the presynaptic cell. For the slow NMDA variable:

$$\frac{ds_{NMDA}}{dt} = \alpha_{NMDA} x (1 - s_{NMDA}) - \frac{s_{NMDA}}{\tau_{NMDA}}, \quad (4.5.3)$$

$$\frac{dx}{dt} = \alpha_x f(V_{pre}) - \frac{x}{\tau_x}, \quad (4.5.4)$$

with $\alpha_{NMDA} = 0.5$, $\alpha_x = 3.48$, $\tau_{NMDA} = 100$ msec and $\tau_x = 2$ msec. The reversal potentials are set at 0 mV for both excitatory synapses and at -70 mV for GABA_A. In all cases, the function $f(V_{pre})$ is given by:

$$f(V_{pre}) = \frac{1}{1 + e^{-(V_{pre} - 20)/2}}, \quad (4.5.5)$$

which has a sigmoidal shape (Figure 4.10).




Figure 4.10: Shape of $f(V_{pre})$ in eq. (4.5.5).

The above scheme implies that the PSP in a cell is not fixed but depends directly on the membrane potential V_{pre} of the presynaptic cell. Figure 4.11 displays the effects of two spike-like inputs with different peaks in the voltage of the presynaptic spike. We use two Gaussian functions for the presynaptic V_{pre} with a peak at 30 mV (Figure 4.11A) and at 10 mV (Figure 4.11B), to represent two different incoming spikes. $f(V_{pre})$ is calculated in each case and eqs. (4.5.2-4.5.4) are numerically solved to get the fast and slow postsynaptic variables. In the first case, where $f(V_{pre})$ reaches 1, the peak of the fast AMPA or GABA_A variable is higher than the slower NMDA peak. In the second case, $f(V_{pre})$ does not reach its plateau and s_{NMDA} acquires slightly higher values than the fast variable. Note though, that in the latter case the NMDA EPSP will again be smaller due to the very low value of \bar{g}_{NMDA} . It thus becomes clear that the relevant strength of the fast and slow PSPs vary as a function of the shape (peak, duration) of the presynaptic spike.

To make our numerical code less time and memory consuming and reduce the number of variables that need to be stored in each numerical step, we implemented a different scheme for the calculation of PSPs. Firstly, we replaced $f(V_{pre})$ in the computation of s_{syn} with a fixed spiking threshold. Since $f(V_{pre})$ becomes nonzero at 10 mV, we set a spike to be detected whenever the membrane potential of the presynaptic cell exceeds 10 mV. This way, all presynaptic spikes that would exert a PSP in the original model, are accounted for in our version too. Secondly, the $\alpha_{syn}f(V_{pre})$ term in ds_{syn}/dt is replaced by an instantaneous kick in s_{syn} . So, when a spike is detected, the synaptic variable in the postsynaptic cell is instantaneously increased by a fixed fraction of the corresponding parameter α . We set this fraction to be equal to 1/3 for excitatory connections and 1/2 for inhibitory ones. We avoided using the full value of α (corresponding to $f(V_{pre}) = 1$) since that would correspond only to spikes that peak at ≥ 30 mV, whereas we find by




Figure 4.11: Synaptic variables s_{syn} of AMPA and NMDA with two different spike-like inputs. From top, presynaptic voltage V_{pre} , $f(V_{pre})$ and s_{AMPA} (or s_{GABA_A} , solid line) with s_{NMDA} (dashed line). The peak of the presynaptic input V_{pre} is 30 mV (A) and 10 mV (B) respectively. Note that the postsynaptic response depends on the shape of the presynaptic spike, whereas in our version it is fixed.

visual inspection of individual membrane potentials that most spikes peak at 10-20 mV. Moreover, we replaced the variable x in eq. (4.5.4) with s_{AMPA} since both variables are described by the same equation.

Note that in the original model each synapse is modelled independently, i.e. there is an independent postsynaptic variable for each synapse in a cell [A. Compte, personal communication]. In our implementation, each cell has one “total” variable s_{syn} for each synapse type, instantaneously increasing whenever a spike arrives in any incoming synapse. Since eq. (4.5.2) is linear, s_{AMPA} (or s_{GABA}) from all incoming synapses can simply be added to give the total ds_{AMPA}/dt in the equation. Therefore, the resulting I_{syn} will be the same in both versions. Yet, the nonlinear eq. (4.5.3) for the s_{NMDA} variable is not additive, implying that s_{NMDA} is modelled slightly differently in our version. Consequently, in the original model, each NMDA synapse saturates independently at $s_{NMDA} = 1$ at high presynaptic firing rates, whereas in our version the total postsynaptic variable will saturate at the same value. In effect, there will be less NMDA current in

the postsynaptic cell of our implementation than in the original version of the model.

All the above differences in the calculation of the PSPs yield a different balance between excitation and inhibition over the network. As a result, the model produces too much spiking activity during the UP state compared to that reported in [47] (see also section 4.5.5). We therefore reduced the overall excitation by reducing the original model's maximal conductances of the excitatory synapses:

- $\tilde{g}_{\text{AMPA}}^{P-P} = g_{\text{AMPA}}^{P-P}/1.3$,
- $\tilde{g}_{\text{AMPA}}^{P-I} = g_{\text{AMPA}}^{P-I}/4.5$,
- $\tilde{g}_{\text{NMDA}}^{P-P} = g_{\text{NMDA}}^{P-P}/4$,
- $\tilde{g}_{\text{NMDA}}^{P-I} = g_{\text{NMDA}}^{P-I}/4.5$.

and we enhanced the inhibition by changing the GABA reversal potential from -70 mV to -75 mV. As we show in Section 4.5.5, with this combination of parameter values we manage to accurately reproduce most of the results reported in [47].

4.5.3 Modelling the Network Connectivity

Our cortical network comprises of a single layer of 1000 pyramidal and 250 interneuron cells (instead of 1024 and 256 cells respectively in [47]) distributed equidistantly, i.e. every four pyramidal cells there is an interneuron. The total length is considered to be 5 mm so the intercellular distance is $\sim 4 \mu\text{m}$. The number of connections of each cell are Gaussian distributed over the network, with each cell being connected to 20 ± 5 (mean \pm SD) other cells of the same type and 20 ± 5 cells of the opposite type. Multiple synapses are allowed, but no autapses. The probability of any cell being connected to another neuron lying at distance x is given by a Gaussian distribution, centered on the source cell, with a standard deviation equal to $250 \mu\text{m}$ (62 cells) for excitatory connections and $125 \mu\text{m}$ (~ 31 cells) for inhibitory ones.

This connectivity is implemented in our version of the model through the following steps:

1. The interneurons are equidistantly placed among the pyramidal cells in an array of length equal to the total size of the network. The indexes of the pyramidal cells on this array and those of the interneurons are stored separately.
2. For each cell i , its number of connections to pyramidal cells and to interneurons is a pair of random numbers drawn independently from a normal distribution: 20 ± 5 .

3. The distances of the cell from the other pyramidal/interneuron cells are calculated as distances of its index from the indexes of the cells.
4. Each distance x is transformed to a connection probability: $P(x) = e^{-x^2/2\sigma^2}/\sqrt{2\pi\sigma^2}$, with σ corresponding to the type of the i cell.
5. All probabilities are normalised so that they add up to 1.
6. The interval from 0 to 1 is separated into intervals with lengths equal to these probabilities.
7. A set of random numbers equal to the number of connections is uniformly drawn between 0 and 1. The interval on which each number falls, is linked with its corresponding cell j .
8. The entry $\{i, j\}$ on the connectivity matrix (for the specific type of synapses) is increased by +1 (to account for multiple connections).

The connectivity matrices between all pairs of neuron types, for a specific network realization, are shown in Figure 4.12.

4.5.4 Numerical Methods

As aforementioned, the model was implemented in the Python-based spiking neural networks simulator “Brian” (www.briansimulator.org) [84], using a second-order Runge-Kutta method with a time step of 0.05 ms. Spikes were recorded at every time step, while all other variables were recorded every 1 ms, yielding a maximum frequency of 500 Hz for spectral estimation.

4.5.5 Modelling Slow Oscillations

We combine the above single neuron models, synapses and connectivity profiles into a single model of a small single-layer cortical network, in order to try to verify the SO observed in the simulations of Compte et al. [47].

Results from one simulation are plotted in Figure 4.13, which can be compared with Figure 4.6 (taken from [47]). The raster plot is shown in the top panel (blue for pyramidal spikes and red for interneuronal ones). The SO in the model is obvious, with long DOWN states that last approximately 2-2.5 seconds and short UP states of roughly 1 second




Figure 4.12: Connectivity matrices derived from our implementation of the Compte et al. model [47]. Clockwise from top left: pyramidal-pyramidal, pyramidal-interneuron, interneuron-interneuron and interneuron-pyramidal connections.

duration, yielding an oscillation frequency 0.4 Hz as in [47]. Also in accordance with [47], the UP state is triggered in a nonspecific area of the network, and propagates to neighbouring sites in the form of a wave, eventually recruiting the whole network into intense firing. Typically, the most active part of the network acts as the initiation site, but it is not unique and varies between simulations. In some cases the UP state is initiated at more than one site in the network. In that case, multiple waves are generated which eventually collapse into a single wave (not shown). The wave propagation speed in our simulation is roughly 9 mm/sec, whereas in [47] it is <3 mm/sec (note though that in the experimental recordings, reported in [47], it is 10 mm/sec). Also note that, again in agreement with [47], interneurons appear to start firing before the pyramidal cells, at a population level, and that their firing lasts longer (due to the slow-decaying NMDA postsynaptic currents in interneurons from pyramidal spikes).

Membrane potentials of three individual neurons along with their corresponding intrinsic Na^+ concentrations are plotted in Figure 4.13B-D. There is a clear bimodal distribution of the membrane potentials between the two states, in parallel with waxing and waning of the intracellular Na^+ concentration. As in Figure 2 of [47], the three selected neurons

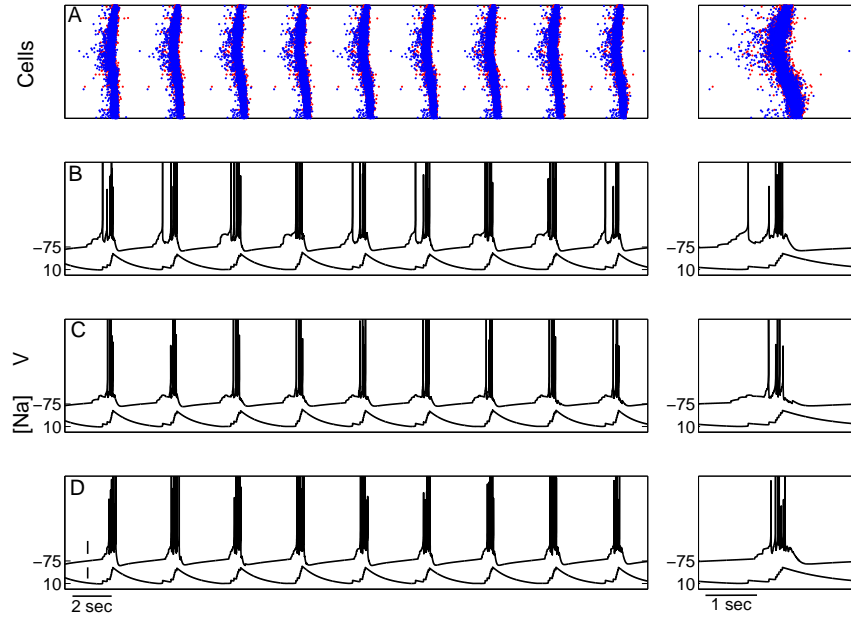


Figure 4.13: Slow oscillations in the cortical network model. **A:** Raster plot of spikes from pyramidal cells (blue) and interneurons (red, lying in the background). **B-D:** Examples of individual axosomatic membrane potentials from three pyramidal cells. Each cell represents different levels of intrinsic excitability (see text) with the top being among the most excitable cells and the bottom among the least. Below each potential, the corresponding intrinsic Na^+ concentration is plotted. The calibration bars in **D** correspond to 10 mV and 2 mM respectively and correspond to all panels. Next to each panel is a zoom into the last of the displayed UP states. This figure can be compared with Figure 4.6 or with Figures 2-3 in [47]

represent three different intrinsic excitability levels, with the one in panel B representing the most excitable cells and in D the least excitable. The most excitable cells fire action potentials before the actual onset of the UP state, acting as seeds of the wave propagation, and show little after-hyperpolarisation. Less excitable cells show a clear bimodal distribution of membrane potential with no firing during the DOWN state and distinct, slowly-decreasing hyperpolarisation after the end of the UP state. The two states are separated by ~ 10 mV and ~ 2 mM in intrinsic Na^+ concentrations in the least excitable cells, as in [47].

Figure 4.14A displays the average firing rates of pyramidal cells and interneurons for the last 5 UP states of the same simulation, calculated over 10 msec non-overlapping bins. It appears that the average maximal firing rate is ~ 15 Hz for pyramidal cells and ~ 30 Hz

for interneurons during the UP states. This is similar to the firing rates in [47], since the corresponding values reported there are ~ 10 Hz and ~ 20 Hz respectively. The power spectra of the membrane potential of one pyramidal cell and one interneuron display peaks at similar frequencies (Figure 4.14C-D. Note that a second peak in the pyramidal spectrum reflects the interneuronal firing). The same spectral peaks (but for spiking activity) are reported via simulations of the model described in [46]. When calculating the firing rate over one UP state using the same scheme as in Figure 3A in [47], the two firing rate peaks shift to ~ 20 and ~ 80 Hz respectively (Figure 4.14B). Again from these rates it becomes clear that interneuronal firing precedes and exceeds pyramidal firing.

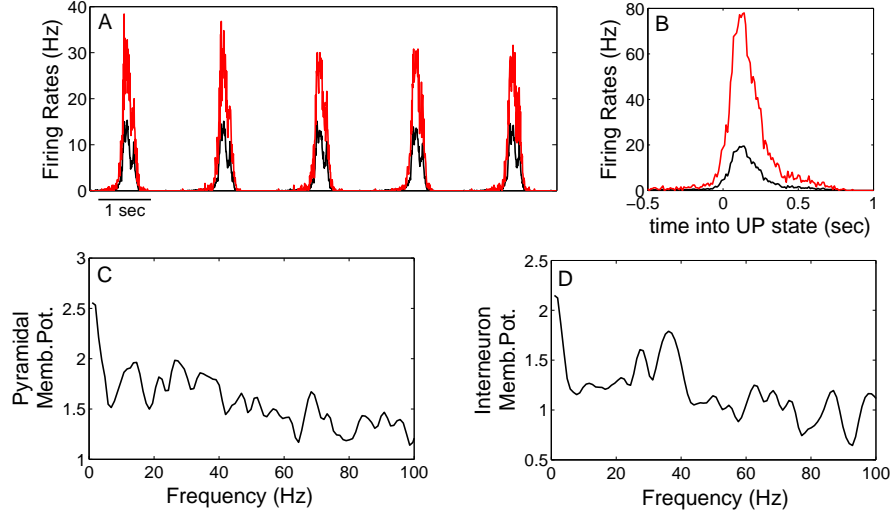


Figure 4.14: **A:** Average firing rates of pyramidal cells (black) and interneurons (red) in the cortical SO model, calculated over 10 msec non-overlapping bins. **B:** Firing rates over the last UP state in **A**, calculated in the same fashion as in Figure 3A in [47]. Namely, within the time window, we subtract from the spike times for each excitatory (inhibitory) neuron the time of the first spike of the closest lying inhibitory (excitatory) neuron, and then we construct the time histogram of those intervals in the black (red) curve. **C-D:** Power spectra of the membrane potential of one pyramidal cell (**C**) and one interneuron (**D**) over a 30 sec simulation.

We also checked various ionic conductance contributions in the course of the SO, in an effort to partially reproduce Figure 5 in [47]. We follow the authors in reducing g_{KNa} to 0.5 mS/cm^2 to prolong the duration of UP states, (a reduction to 0.33 mS/cm^2 as in [47] resulted in practically no DOWN states in our version). The membrane potential of a pyramidal cell and all intrinsic conductances are presented in Figure 4.15. Most plots are similar to the published ones. Nevertheless, the inhibitory conductance is much lower than the excitatory one in our model, although they are practically the same in

[47]. This is a result of the different implementation of synaptic inputs in our version.

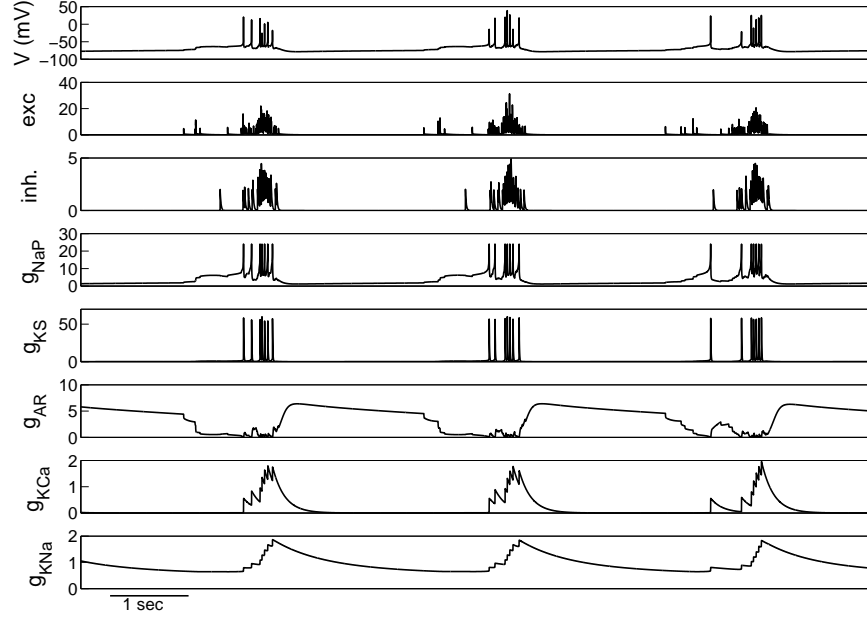


Figure 4.15: Various ionic conductances of a modelled pyramidal cell during slow oscillations. From Top: Membrane potential of one pyramidal cell, total excitatory conductance, inhibitory conductance, g_{NaP} , g_{KS} , g_{AR} , g_{KCa} , g_{KNa} . All conductances are shown in nS. For clarity, g_{KNa} was reduced to 0.5 mS/cm² to prolong the duration of UP states. This figure can be compared with Figure 5 in [47]

Finally, in order to reproduce the effects of blocking particular synaptic types, as in Figure 9 in [47], we blocked each of the three synaptic conductances separately and produced raster plots of the whole array for each case (Figure 4.16). Our results are qualitatively similar to the published ones. AMPA blockade results in random firing of most cells without any modularity. NMDA blockade has a similar result, but occasional clusters of spiking activity emerge transiently (resembling an UP state). GABA blockade, on the contrary, leads to UP states that propagate faster throughout the network, and longer DOWN states (~ 3.5 sec).

4.6 Discussion

We have briefly reviewed here some of the most prominent EEG cortical rhythms observed during slow wave sleep. The slow oscillation (SO) was described in more detail




Figure 4.16: Raster plots of the whole model when either AMPA, NMDA or GABA connectivity is blocked. This figure can be compared with Figure 9 in [47]

and some of the computational models developed were presented.

The cellular and/or synaptic mechanisms that generate this intracellular UP and DOWN bimodality in cortical tissue are not well understood. The models discussed above illustrate that different networks can generate such states, alternating at low frequencies, via mechanisms that include various different key elements. Experimental recordings could shed more light by excluding some of these alternatives or hinting at new crucial components. On the other hand, different experimental set ups and procedures can lead to different conclusions.

Although most data and computational models agree on the sustainability of the UP state through recurrent synaptic interactions, one of the most interesting open questions on the SO is how the UP state is initiated and terminated. The Bazhenov et al. model [18] suggests that UP states are initiated through the summation of random mini EPSPs arriving in PY cells during the DOWN states. This approach was based on experiments on a cortical slab preparation, synaptically isolated from the rest of the cortical network but kept with a natural blood supply [205]. While the cortex exhibited SO, neurons in the slab exhibited small depolarising potentials, represented by the minis in the model, occasionally interrupted by bursts of large depolarising events consisting of EPSPs and IPSPs [206]. Much larger slabs exhibited normal SO, suggesting that random summations of such synaptic miniature EPSPs could be spontaneously initiating the UP states in sufficiently large cortical networks.

In a different approach, the Compte et al. model [47] postulates the UP state initiation through intrinsic spike discharges of cortical cells. Their analysis supports *in vitro* experimental recordings on isolated cortical slices [175], where SO appeared after the

traditional slice bathing medium was replaced by a higher concentration of K^+ , but no miniature EPSPs were applied to the slice. Instead, the UP state was intrinsically generated in the recurrent excitatory network of layer 5 of the slice, propagated to the other layers through excitatory pathways and was regulated by inhibition and K^+ currents. Moreover, during blockade of NMDA and AMPA receptors in the slices, spontaneous activity in layer 5 persisted. According to the model these activities can trigger UP states, although in the recordings they fail since they cannot spread without excitatory synapses.

The above models also disagree on the terminating mechanism of the UP state, with the former placing it on synaptic depression in the cortical network and the latter on hyperpolarisation buildup through the I_{KNa} current.

The Hill and Tononi model [99] includes all the above components in a very complex corticothalamic network, bridging the two models. It suggests that the UP state can be initiated either intrinsically by spontaneous firing of cortical cells or by external minis. What is important is that either of these two mechanisms can induce firing in a small area of the cortical network and activate the persistent sodium current I_{NaP} , which will prolong the depolarisation and firing of the cells involved (the importance of I_{NaP} is also emphasised by the other two models). This will enable the activity to propagate through the rest of the network. The termination of the UP state can again be intrinsic, by I_{KNa} activation, or synaptic, by depression, but the hyperpolarisation needed to sustain a DOWN state seems to be dependent of the I_{KNa} current.

Furthermore, the Compte et al. model suggests that both UP and DOWN states are stable states of the network and the SO is the result of periodical switching between the two, regulated by I_{KNa} . Such a bistability was shown for a theoretical model of a purely excitatory network with synaptic depression [101]. Within a specific parameter regime, the phase space of the system contains two attractors (UP and DOWN states) and transitions between the two is achieved by synaptic noise. Interestingly, it has been shown that *in vivo* external stimulations can shift the system between the two states. For example, electrical stimulation in the DOWN state can induce an UP state shift and the same stimulus can terminate the UP state. Hence, the same cortical input can produce diametrically opposite changes in the network state, depending on the short-term history of the network [33]. Computational and theoretical models that address these and other issues on SO include [117, 141, 158, 225].

The three models described in this Chapter provide different approaches to modelling

a cortical network (in some cases including a thalamic network as well) that exhibits SO with various mechanisms. The cortical network of the first model is a very simple circuit with fixed connectivity and relies in synaptic depression, while the last model is highly complex including numerous elements that are not essential in producing the SO. Here we reproduced the Compte et al. model [47], implemented for the first time in the Python neural simulator “Brian”. We chose this model because it is of moderate complexity, with simple enough single cell models, relatively few variables and contains stochastic elements within a simple but realistic connectivity scheme. Moreover it has been described and analysed in detail [46, 47] and its network architecture will also be used in our CA3 and CA1 hippocampal networks (see next chapter).

We have demonstrated that after making some modifications in the synaptic interactions, due to the different implementation of PSPs in Brian compared to the original model, our version accurately reproduces the simulation results presented in [47]. One advantage of our implementation, owing to the various components of Brian, is that many new features could be easily implemented in the model. For example, it is rather straightforward to include synaptic depression, spike-time dependent plasticity, or axonal delays in the network, in an effort to study the effect of such elements on SO.

This modified version of the Compte et al. model will be coupled to our CA3-CA1 network in order to model cortico-hippocampal interactions during sleep and correlations between cortical SO and hippocampal ripples (see Chapter 7).

Network Models of the Hippocampal CA3 and CA1 Areas

5.1 Introduction

The previous chapter dealt with EEG rhythms of the cortex during deep stages of sleep, focusing on the cortical SO. We now shift our attention to hippocampal EEG rhythms and patterns, emerging during sleep.

The rodent hippocampus EEG has been traditionally divided into two categories [237], based on their spectral characteristics and occurrence in specific sleep states:

- The activated state, consisting of prominent theta oscillations (4-10 Hz). Theta is the main hippocampal rhythm, appearing during wakefulness and REM sleep and is coexpressed with a faster gamma (20-80 Hz) rhythm, superimposed on the theta cycles (see [33] for an extensive analysis of the theta rhythm and its role in memory and learning).
- The deactivated state, consisting of a broadband signal known as large-amplitude irregular activity (LIA), appearing during non-exploratory behavior, deep sleep and under most general anaesthetics [237]. During this state, three basic patterns alternate in the hippocampus: gamma oscillations, silent periods and the hippocampus-specific pattern called sharp wave-ripple complex.

In contrast with the relatively uniform gamma oscillations modulated by theta-oscillations

during the activated state, gamma activity during LIA is much more irregular, varying both in frequency and amplitude [33]. This chapter is focused on hippocampal dynamics during the deactivated state, while in the next chapter we focus particularly on sharp wave-ripple complexes.

Here, we briefly describe the basic morphology of the hippocampus and we review three single cell models for hippocampal neurons. We then describe in detail the computational models of the CA3 and CA1 areas that we have developed. We show how these models reproduce various of the characteristic anatomical and/or functional properties of the corresponding areas, and we discuss the underlying mechanisms. These models will be coupled in the next chapter in order to yield sharp wave-ripple complexes.

Parts of this chapter were published in the *Hippocampus* journal [203].

5.2 Hippocampal Morphology

The hippocampus is a neural structure in the medial temporal lobe that belongs to the limbic system along with the amygdala, entorhinal cortex and hypothalamus. It is a one-layer cortex comprising of the tooth-shaped dentate gyrus, the C-shaped hippocampus proper or cornu ammonis (CA), which is further divided into the CA1 and CA3 regions, and the subiculum (Figure 5.1). A discrete CA2 region has also been described [8] but we will not be dealing with it here. Like most cortical formations, it comprises pyramidal cells and GABAergic interneurons, with specific layer and wiring organizations, while the dentate gyrus has a radically different cell type called the granule cell [6, 8, 56].

If unfolded, the hippocampus is actually a large sheet with a single layer of granule cells and pyramidal neurons. The dentate gyrus acts as a major entry point from the entorhinal cortex (perforant path). The axons of its granule cells (mossy fibers) reach the pyramidal cells of the CA3 region which in turn excite those of the CA1 region. The main conduit axons from the CA3 to the CA1 cells are called Schaffer collaterals. The excitatory synaptic chain:

1. entorhinal cortex to dentate gyrus through the perforant path,
2. dentate gyrus to CA3 through mossy fibers,
3. CA3 to CA1 through Schaffer collaterals,

is the so-called trisynaptic circuit [8]. The cells in the CA3 area form a strongly recursive

network of short and long distance connections, whereas CA1 pyramidal cells are sparsely connected with each other. Direct connections also exist from the entorhinal cortex (layers 2 and 3) to the distal dendritic layers of the CA3 and CA1 areas (temporoammonic pathway). We return to this circuitry in Chapter 7. Finally, CA1 cells project to the subiculum area and out to the entorhinal cortex (Figure 5.1).




Figure 5.1: Left: Position of the hippocampal formation in the rat brain. The hippocampus is an elongated, C-shaped structure with the longitudinal (septotemporal) axis running from the septal nuclei rostrally (S) to the temporal cortex (T) ventrocaudally. The transverse axis (TRANS) is oriented perpendicular to the septotemporal axis. Taken from [6]. Right: Basic circuit of the hippocampus in one transverse slice, consisting of the dentate gyrus (DG), the CA3 and CA1 regions and the subiculum (Sub). The entorhinal cortex (EC) is its basic input and output. Taken from the wikipedia article: <http://en.wikipedia.org/wiki/Hippocampus>.

About 90% of the cells in the CA areas of the hippocampus are pyramidal cells. Hippocampal pyramidal cells, when depolarised, can fire isolated spikes or bursts of action potentials both *in vivo* and *in vitro*. Bursts usually last about 30-50 msec and are followed by long after-hyperpolarisations (AHP).

It has been observed that the firing patterns of CA3 pyramidal cells depend on the intensity of the depolarising stimulation. When they are slightly depolarised by somatic current injection, they fire isolated bursts whose frequency increases with the level of depolarisation. Increasing the current further, causes the bursts to be replaced by spike doublets, then tonic firing of continuous action potentials and eventually by membrane oscillations of ~ 20 Hz without any firing. On the contrary, hyperpolarising stimulation causes suppression of the bursts [213, 214].

These bursts are intrinsic properties of the cells and not a collective network phenomenon since they can be seen in completely isolated hippocampal cells. They depend on calcium currents which are mostly found in the cells apical and basilar dendrites but can be

found in the soma as well. Summation of excitatory postsynaptic potentials leads to the generation of one or more calcium spikes which contribute to the bursts. Their termination is caused by a slow AHP mediated by one or more calcium dependent potassium currents. Blocking these calcium currents replaces the bursts with single spikes [214].

CA1 cells have different firing properties. Somatic injection of depolarising current usually evokes tonic firing of fast spikes with frequency adaptation, instead of bursts. Burst firing can be seen in the dendrites after a strong current injection or synaptic activation with inhibition blocked. These dendritic bursts appear to be a mixture of calcium and sodium spikes. Somatic bursts can be seen usually when inhibition is blocked and the cells are driven maximally by synchronised CA3 firing [8, 134].

In contrast with the rather uniform population of pyramidal cells, there is a large variety of interneuron cells in the hippocampus, with different shapes, axonal distributions, locations, firing patterns and intrinsic currents. Examples of hippocampal interneurons include basket cells, chandelier (or axo-axonic) cells, dendritic inhibitory cells and many more. Extensive reviews of the different classes of hippocampal interneurons can be found in [8, 56, 76]. The basket cells, which are included in our models, are fast (or regular) spiking parvalbumin-positive perisomatic-inhibitory interneurons, lying on the pyramidal cell layer.

5.3 Models of Hippocampal Cells

Before describing the two hippocampal network models of CA3 and CA1 areas that we developed, we first need to examine the firing properties of single cell models. Below we review two computational models that have been developed for pyramidal CA3 and CA1 neurons and one model for fast spiking interneurons. Since two of these, namely the Pinsky-Rinzel and the Wang-Buzsáki models, will be used as single cell models in our hippocampal networks, we examine their individual behavior and show that they accurately reproduce many of the aforementioned firing properties.

5.3.1 The Traub Model

A complex mathematical model of a guinea pig CA3 pyramidal neuron was introduced by Traub et al. [221] and was used in a series of studies of hippocampal properties [214].

The “Traub model”, a modification of an earlier version [210], is a Hodgkin-Huxley-kinetics cable model consisting of 19 compartments: 8 basilar dendritic, 10 apical dendritic and 1 axosomatic. Throughout these compartments, the conductances of 6 ionic currents with compartment-specific densities are distributed: a leakage current, a sodium current, a high-voltage activated calcium current and 4 potassium currents: a delayed rectifier ($I_{K(DR)}$), an A-type transient ($I_{K(A)}$), a slow Ca^{2+} -dependent after-hyperpolarisation ($I_{K(AHP)}$) and a fast Ca^{2+} -dependent current ($I_{K(\text{Ca}^{2+})}$). The model kinetics for all the above currents, except for $I_{K(\text{Ca}^{2+})}$ were obtained from voltage-clamp data from isolated hippocampal pyramidal neurons. Calcium channels are mostly located in dendritic compartments while Na^+ channels were mostly in perisomatic compartments. The intracellular Ca^{2+} concentration is given by:

$$d[\text{Ca}^{2+}]/dt = -\alpha I_{\text{Ca}^{2+}} - \beta[\text{Ca}^{2+}] \quad (5.3.1)$$

with α and β fixed parameters.

The model succeeds in reproducing three different types of firing pattern of CA3 neurons, when stimulated by depolarising constant currents or brief current pulses (Figure 5.2A-B):

- When the model neuron is depolarised by a low current, it fires in a low-frequency bursting mode. The bursting frequency increases with the current magnitude.
- For large enough depolarising current, the firing mode switches to high-frequency tonic firing.
- In between these modes, a mixture of bursts and tonic spikes exists.

Therefore, the model manages to reproduce the two characteristic CA3 firing patterns mentioned in Section 5.2. A third firing mode, tonic dendritic Ca^{2+} spiking, was also observed when the cell was stimulated by a large depolarising current applied in the middle or distal dendrites. Namely, the large depolarisation caused tonic Na^+ spikes in the soma that initiated respective Ca^{2+} currents in the dendrites.

By increasing the maximal conductance of the $I_{K(DR)}$ current, decreasing those of the dendritic $I_{\text{Ca}^{2+}}$ and $I_{K(C)}$ currents and adjusting the conductances distributions, the CA3 neuron model could be transformed into a model of a CA1 pyramidal cell. After these alterations, somatic depolarisation leads to tonic firing with adaptation, whereas dendritic depolarisation leads to bursting.

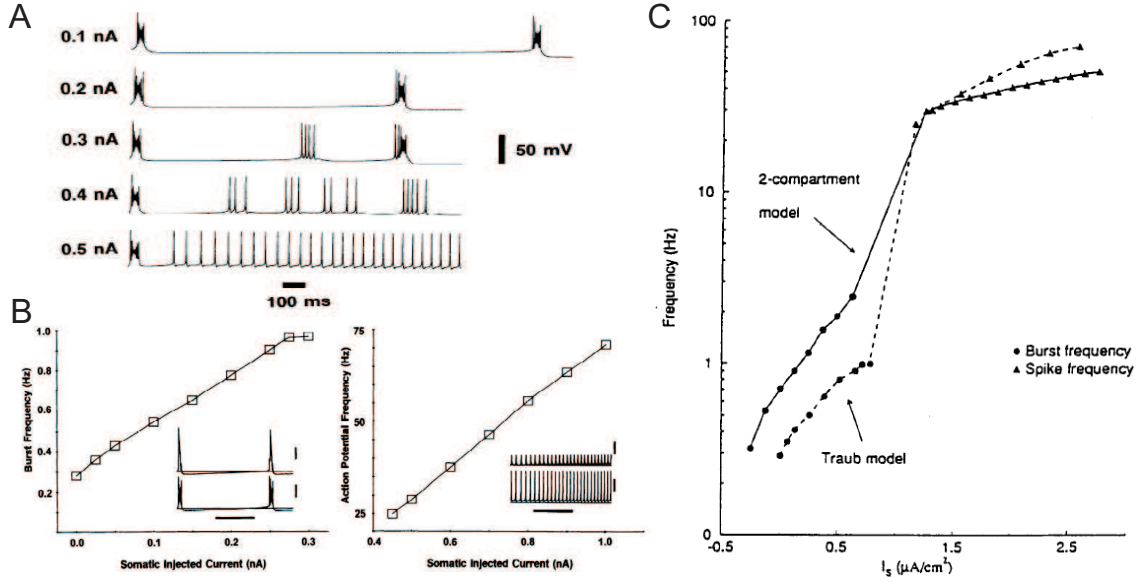


Figure 5.2: **A:** The three firing patterns of a CA3 pyramidal cell as reproduced by the Traub model. **B:** Frequency-current curve for the bursting (left) and tonic spiking (right) firing modes of the Traub model. The insets show representative dendritic (top) and somatic (bottom) voltage traces of the 2 modes. Taken from [221]. **C:** Comparison of frequency-current curves for the Traub and the Pinsky-Rinzel models. Taken from [165].

An extended version of the CA3 Traub model was later developed [212], containing 64 somatic and dendritic compartments with branching dendrites, an axon initial segment and 4 axonal compartments. The kinetics and the distributions of the ionic currents remained similar as before and so did the single cell firing patterns.

5.3.2 The Pinsky-Rinzel Model

The same year that the extended version of the Traub model was published, a much simplified version of this CA3 pyramidal cell model was developed by Pinsky and Rinzel [165]. Their goal was to reduce the original Traub model while retaining its main characteristics and properties. Their approach was to keep the same currents and gating kinetics but reduce the number of compartments.

The Pinsky-Rinzel model is a two compartment Hodgkin-Huxley cell with one dendritic compartment where all the calcium and calcium-dependent currents are located, and one axosomatic compartment with the sodium and the delayed rectifier potassium currents. The connection between the two compartments is regulated by a simple coupling

conductance whose value was shown to be critical in achieving the same behavior as the Traub model. Again, the intracellular calcium concentration is given by an ODE of the same form as in eq. (5.3.1). Altogether, this model consists of only 8 variables, compared to 120 variables for the original Traub model.

The basic firing modes of the hippocampal cell were successfully reproduced by this model as well:

- The model exhibits periodic bursting with a low somatic current or with a stronger constant dendritic input. The burst duration is primarily determined by the time course of dendritic Ca^{2+} build up while the interburst interval depends on the time courses of the hyperpolarising currents $I_{K(\text{Ca}^{2+})}$ and $I_{K(\text{AHP})}$. The burst is initiated by a sodium spike in the soma which triggers a slower dendritic calcium spike and the burst is formed by a “ping-pong type” current flow between the somatic and dendritic compartments [51].
- Greater depolarisation leads to periodic somatic spiking without active dendritic spikes. For specific coupling conductance and compartment sizes the resulting firing frequency as a function of the applied current (f-I curve) resembles that of Traub’s (Figure 5.2C). In between the bursting and tonic firing modes, aperiodic behavior and spike doublets are observed, as in Traub’s model.
- Even higher current injection eventually leads to annihilation of spikes and a steady depolarised state is reached.
- Dendritic input results in bursting with frequency increasing along with the current amplitude.

5.3.3 A Modification of the Pinsky-Rinzel Model

In order to simulate the firing properties of CA1 pyramidal cells, we introduce a slight modification in the Pinsky-Rinzel model’s parameter values. We aim to replace the typical CA3 intrinsic bursting, under a somatic-injected depolarising current, with tonic firing exhibiting frequency accommodation [8, 134]. Since the somatic bursts are superimposed on a longer dendritic calcium spike, we keep all parameter values the same as in the original model with the only exception being the Ca^{2+} maximum conductance $g_{\text{Ca}^{2+}}$ which is reduced to 7 mS/cm² (from 10 mS/cm²).

We simulated the original Pinsky-Rinzel model in “Brian” and, after accurately reproducing its firing properties (not shown), we implemented the $g_{Ca^{2+}}$ parameter alteration. Depolarising the somatic compartment with low currents now results in tonic firing with a frequency that quickly attenuates to a constant value. We calculated the firing rate of the cell through the first, second and last interspike interval after a 2 sec run of the model. As shown in Figure 5.3, the firing rate increases linearly with increasing current and the accommodation becomes faster. Greater depolarisations result in an initial burst that is followed by accommodating tonic firing. The slope of the f-I curve, calculated through the last interspike interval is approximately 30.35 Hz/nA which is in good agreement with the 34.1 Hz/nA slope estimated through single cell recordings [126].

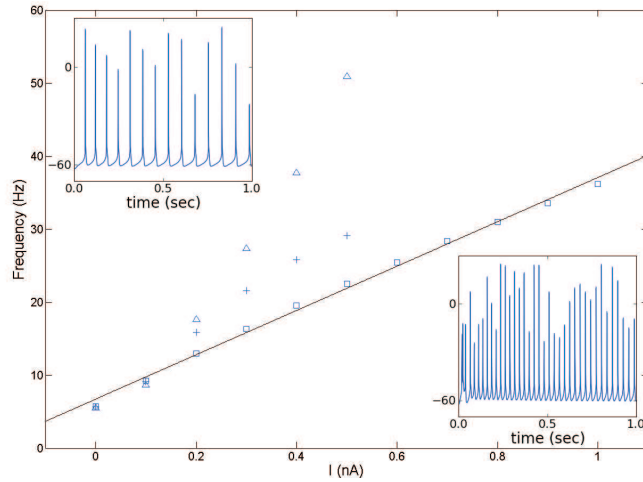


Figure 5.3: f-I curve for the Pinsky-Rinzel model of the CA1 pyramidal cell, after the parameter alteration described in Section 5.3.3. Frequencies are calculated by the inverse of the first (triangles), second (crosses) and the last interspike interval (squares) after a 2 sec simulation. The slope of the least squares fit on the last interspike intervals is 30.35 Hz/nA. The insets show the membrane potential of the cell when depolarised in the dendrite by constant current of 0.2 nA (upper) and 0.8 nA (lower).

5.3.4 The Wang-Buzsáki Model

The Wang-Buzsáki model is a simple single-compartment Hodgkin-Huxley cell with only the basic spike-generating currents (leakage, sodium and delayed rectifier potassium currents) [228]. It is well suited for modelling fast spiking interneurons, such as hippocampal basket cells, since it can fire repetitive spikes at high frequencies and each spike is followed by a very brief refractory period. Its firing rate increases very steeply as

a function of the applied current and can reach high frequencies of several hundred Hz even with depolarising pulses of a few $\mu\text{A}/\text{cm}^2$ (see Figure 4.9 in the previous chapter).

5.4 Description of a CA3 and a CA1 Network Model

Below we describe in detail the two computational network models we developed to simulate characteristic properties of the CA3 and CA1 areas. The anatomical and/or functional justification for our choices of parameter values is also explained.

5.4.1 General Model Architecture and Single Cell Models

The two models represent the CA3 and CA1 rat hippocampal areas respectively, in the longitudinal (septotemporal) direction. They are one-dimensional arrays containing 1000 pyramidal cells and 100 interneurons, keeping the ratio of 10:1 pyramidal cells-to-interneurons that has been estimated for the hippocampus [8, 221]. The intercellular distance is set to 10 μm in both cases, yielding a total length 11 mm, close to the 10 mm length of a typical longitudinal CA3 slice [221]. The interneurons are equidistantly distributed among the pyramidal cells, with one interneuron every ten cells. The total surface area of each pyramidal cell is set to 50,000 μm^2 , close to the 50,400 μm^2 reported for CA3 cells [37] but a slight overestimate for CA1 cells with 36,000 μm^2 surfaces [37]. The total surface area of each interneuron is set to 20,000 μm^2 , which is of similar order of magnitude as the 12,000 μm^2 for dentate gyrus basket cells reported by [14]. The two arrays are considered parallel and the distance between them was set to 100 μm .

Pyramidal cells in CA3 are modelled by the two-compartmental Pinsky-Rinzel model [165], and in CA1 by the parameter-modified version of the model, presented in section 5.3.3. Although many different types of interneurons exist in the hippocampus [56, 76], all interneurons here are considered to be perisomatic basket cells, i.e. targeting only the axosomatic compartment of pyramidal cells. In contrast, all excitatory connections target the pyramidal dendritic compartment. Interneurons are modelled by the Wang-Buzsáki model [228].

Following Compte et al. [47], heterogeneity in the system is included by varying the reversal potential and the maximum conductance of the leakage current throughout the network. These are distributed over the cells according to a Gaussian distribution centered on the value for the corresponding single cell model with standard deviation

(SD) equal to 0.5% of that value. Additionally, the maximum conductance of the coupling current between the two compartments of the Pinsky-Rinzel model is varied in pyramidal cells following the same scheme. Since firing properties of both single cell models depend heavily on these parameters, the small SD introduces variability in the excitability of the cells.

The initial conditions in all simulations were also Gaussian-distributed over both pyramidal and interneuronal populations. The means of the distributions were equal to the steady states of the corresponding single cell models. Their SD was 10% in all cases.

5.4.2 Synaptic Interactions

Only fast AMPA and GABA_A-mediated synaptic interactions are considered in both models, with synaptic currents given by:

$$I_{syn} = \bar{g}_{syn} s_{syn} (V - V_{syn}), \quad (5.4.1)$$

where \bar{g}_{syn} is the maximum conductance, V is the membrane potential of the postsynaptic cell, V_{syn} is the reversal potential of the synapse and s_{syn} is an exponentially decreasing gating variable:

$$\frac{ds_{syn}}{dt} = -\frac{s_{syn}}{\tau_{syn}}, \quad (5.4.2)$$

with τ_{syn} the decay constant. When an action potential arrives at a synapse, s_{syn} in the postsynaptic cell is increased by a fixed value α_{syn} . For simplicity, we have set $\bar{g}_{syn} = 1$ nS for all synapses, so that the postsynaptic potential (PSP) amplitudes are controlled by α_{syn} . The values of α_{syn} for all types of synapses are given in Table 5.1, along with the PSPs they yield and the corresponding recorded PSPs that were used as reference. Note that the inhibitory PSPs (IPSPs) between CA1 interneurons somewhat deviate from the cited value (0.25 mV) because in these particular connections we replicated the 5 nS peak conductance reported by [15].

The AMPA decay time constant is set to 2 msec [79]. For GABA_A synapses between interneurons we set an equally short time constant of 2 msec, following [15, 16]. For GABA_A synapses targeting pyramidal cells the decay time constant was set to 7 msec which is between the ~ 4 msec reported by [15] and the 10 msec time constant usually implemented in similar simulations [217, 219, 228].

The conduction velocity of the CA3 pyramidal cell axon has been shown to be approximately 0.5 mm/msec [144]. We use this velocity in both our CA3 and CA1 models. The

		α_{syn}	PSP (mV)	Reported PSP (mV)	Reference
CA3	PY-PY	15	1	0.6-1.3	[145]
	PY-IN	3	1.6	1.9	[142]
	IN-PY	50	1.2	1.2	[143]
CA1	PY-IN	2.5	1.3	1.4	[4]
	IN-PY	15	0.5	0.45	[30]
	IN-IN	5	0.5	0.25	[44]

Table 5.1: Parameter values for the synaptic strength of all types of synapses in the CA3 and CA1 models along with the resulting PSPs and the ones reported through physiological recordings. All PSPs in the model were measured from a background membrane potential of -65.3 mV for CA3 pyramidal cells, -62.6 mV for CA1 pyramidal cells and -62 mV for all interneurons.

corresponding velocity for CA1 interneurons was set to 0.1 mm/msec [173]. We treat propagation over CA3 interneuronal axons as instantaneous.

5.4.3 General Connectivity

The CA3 model is characterised by many recurrent excitatory connections while connections between CA3 interneurons are omitted. The opposite holds for CA1 where interneurons are strongly connected with each other while recurrent excitatory connections are absent.

The connectivity scheme in both models is similar and follows the one applied in the Compte et al. SO cortical model [47] described in the previous chapter (Section 4.5.3). The number of connections from cells of type i to cells of type j , where i and j can stand for pyramidal cells or interneurons, is drawn from a Gaussian distribution with mean k_{ij} and standard deviation $k_{ij}/20$. The assigned connections are distributed to the j -cells according to a Gaussian probability distribution, centered on the source cell, with a fixed SD equal to σ_i :

$$P_i(x) = \frac{e^{-x^2/2\sigma_i^2}}{\sqrt{2\pi\sigma_i^2}}, \quad (5.4.3)$$

where x is the distance of each j -cell from the source cell. In both CA3 and CA1 models, all excitatory connections have a SD of $\sigma_{PY} = 1 \text{ mm} = 100 \text{ cells}$ while interneurons have $\sigma_{IN} = 100 \text{ } \mu\text{m} = 10 \text{ cells}$. Multiple connections are allowed but not autapses.

Our connectivity scheme means that the further away a cell is from the source, the smaller the probability that it will get connected. The only exception to this scheme is connections from CA3 interneurons to CA3 pyramidal cells, for which all connections are uniformly distributed around each interneuron to cells within distance $3\sigma_{IN}$.

Since about 99.7% of values drawn from a Gaussian distribution are within three standard deviations away from the mean, 99.7% of the distributed connections will lie within a distance of $3\sigma_i$ from the source cell. Therefore almost all connections are expected to be distributed to cells within a “connectivity cluster” of total size $6\sigma_i$. For example, the pyramidal-pyramidal connections in CA3 can reach out to distances of $3\sigma_{PY} = 3$ mm, covering 600 cells. These clusters are in broad agreement with data from neurophysiological recordings. For recurrent excitatory connections in CA3, it has been shown that *in vivo* CA3 axons can extend at least one-half the length of the hippocampus in the longitudinal direction [213] but that the probability of synaptic contact between two pyramidal cells drops as a function of their distance, reaching out to 3-5 mm [144, 214].

In practice, the i -to- j connectivity cluster will be smaller than $6\sigma_i$, since cells of the opposite type of j will also lie within it and since the closer the source cell is to the array boundaries the fewer the available cells. Moreover, multiple connections in the case of small clusters may result in the average number of contacted cells being reduced from k_{ij} to a value \tilde{k}_{ij} . The connectivity probability between any pair of cell types is estimated as the average number of contacted cells \tilde{k}_{ij} over the average cluster size measured in numbers of cells.

The implemented average number of connections k_{ij} along with the average number of contacted cells \tilde{k}_{ij} are summarised in Table 5.2 for every pair of cell types. The average cluster sizes (estimated through 20 repetitions of the connectivity algorithms) and the corresponding connectivity probabilities are also listed. A schematic diagram of the model is illustrated in Figure 5.4 while an example of the connectivity matrix of pyramidal-pyramidal connections in CA3 is shown in Figure 5.5.

5.4.4 CA3 Connectivity

The average number of connections a CA3 pyramidal cell makes was set to 190. Of these, $97.4\% = 185$ connections target other pyramidal cells and the remaining $2.6\% \approx 5$ connections target CA3 interneurons, in accordance with the observation that approximately 2.1% of the boutons of a studied CA3 pyramidal cell targeted parvalbumin-




Figure 5.4: Schematic diagram of the full CA3-CA1 model. The bottom array represents CA3 and the top CA1. Blue spheres are pyramidal cells and red ones are interneurons (1:4 cell ratio is shown for clarity, instead of the real 1:10). A qualitative picture of the connectivity is illustrated through four example cells. The two Gaussian distributions represent the connectivity probability as a function of distance from a pyramidal cell (bottom) and an interneuron (top). Connections from CA3 to CA1 represent Schaffer collaterals and are discussed in the following chapter.

positive interneurons [187]. It has also been estimated that most of the pyramidal cells each CA3 cell contacts are located in the CA1 area, the rest forming recurrent excitatory connections within CA3 [130, 213]. Accordingly, the 185 connections targeting other pyramidal cells in the model are further divided into $70\% = 130$ that will later form Schaffer collaterals to CA1 (see next chapter) while the remaining $30\% = 55$ connections target other pyramidal cells within CA3.

Pair recordings of pyramidal cells in transverse CA3 slices indicate that the probability two such cells are connected is roughly 2% [133]. However, it has been argued that the actual probability inside the intact CA3, in the longitudinal direction, can be much higher, even up to 15% [22]. The connectivity probability in our model is on average 10.62%, well within the estimated limit for the intact CA3.

The average cluster size of CA3 pyramidal-to-interneuron connections is approximately 47 cells in the model. Since this cluster is large enough for the 5 distributed connections, very few polysynaptic connections arise ($\tilde{k}_{PY-IN} = 4.66$), in accordance with the observation that excitatory connections to interneurons are mostly monosynaptic [213]. Moreover, the estimated probability of finding a pyramidal-to-interneuron connection within a transverse CA3 slice is about 10% [144, 214]. Similarly, the corresponding probability in our model is 9.88%.

		Average Cluster Size	k_{ij}	\tilde{k}_{ij}	Connectivity probability
CA3	PY-PY	471.77	55	50.1	10.62%
	PY-IN	47.15	5	4.66	9.88%
	IN-PY	55.2	68	39.16	71.3%
CA1	PY-IN	47.15	20	14.67	31.11%
	IN-PY	55.2	400	46.35	83.97%
	IN-IN	4.94	100	4.49	90.89%

Table 5.2: Connectivity parameter values for every connected pair of cell types. The average cluster size is the number of target cells within a distance of $3\sigma_i$ from an average neuron, k_{ij} is the average number of connections each cell makes with the target cells while \tilde{k}_{ij} is the average number of different cells it contacts.




Figure 5.5: Connectivity matrix of the pyramidal-to-pyramidal connections in CA3. The color scale represents the number of multiple connections.

The interneuronal axons deviate widely in a transverse CA3 slice but their outputs extend only to around $300 \mu\text{m}$ [22, 213]. In a series of CA3 computational models [214, 215], the number of interneuronal outputs are tenfold the pyramidal ones but excitatory connections are 5 times more extensive. In our model, each interneuron has 68 connections on average, distributed uniformly around the cell over a maximum distance of $3\sigma_{IN} = 300 \mu\text{m}$. Hence excitatory connections are 10 times more extensive than inhibitory ones. As expected, many polysynaptic connections arise ($\tilde{k}_{IN-PY} = 39.16$) in accordance with the observation that perisomatic interneurons form multiple connections on each pyramidal cell [213]. The resulting connection probability is approximately 71.3%, close to the estimated 60% in slices [144, 214].

5.4.5 CA1 Connectivity

Since, to our knowledge, there have not been many recordings in CA1 that clearly reveal the anatomical features of the circuit, connectivity in our CA1 network is more loosely based on anatomical data and more focused on reproducing features of the intrinsic CA1 gamma oscillations through a strongly interconnected inhibitory network.

The extent of pyramidal connections is the same as in CA3, but now they are set to target only interneurons, since various studies have shown that there is a distinct preference for CA1 pyramidal cells to target interneurons while recurrent excitation is very sparse [8]. The connection probability from pyramidal cells to stratum pyramidale interneurons was calculated to be roughly 30% [122]. Although different interneuronal subclasses have been shown to have different probabilities of receiving excitatory inputs, we chose to follow the general stratum pyramidale estimate, by setting $k_{PY-IN} = 20$, which yields a connection probability of 31.11%. Since the actual average number of interneurons targeted by each pyramidal cell is $\tilde{k}_{PY-IN} = 14.67$, few polysynaptic connections arise, again in agreement with pair recordings [88].

Unlike the relatively few excitatory connections, interneurons form many polysynaptic connections to neighbouring cells. Basket cell axons in CA1 were found to generally span less than 1 mm [186]. In a CA1 modelling study this span was set to 500 μm [218]. In our model the corresponding span is of the same order and can reach 300 μm . Furthermore, the CA1 basket-to-pyramidal cell connection probability was found to be roughly 30%, dropping abruptly with distance from the soma, from 54% for immediate neighbours to 5% for distant ones [8, 122]. In the model the average connection probability is larger, at 83.97%. Nevertheless, this probability yields that the average number of contacts each interneuron makes on a pyramidal cell is 8.54 with the corresponding experimentally estimated average being 6 contacts [31, 89].

Although little information is available on the characteristics of recurrent inhibitory connections in CA1, it is established through *in vitro* pair recordings in slices that basket cells are highly interconnected [16] with each basket cell contacting at least 60 others [186]. We have set an average of 100 connections each interneuron makes to others, but the fact that they are equidistantly distributed on the array, 110 μm apart, yields a very small connectivity cluster (~ 5 cells), and consequently a very high connection probability of 90.89% within the cluster.

5.4.6 Numerical Methods

As with the cortical model in the previous chapter, both hippocampal models were implemented in Brian [84], using a second-order Runge-Kutta method with 0.05 msec time step. Spikes were recorded at every time step and all other variables were recorded every 1 ms.

5.5 Simulation Results

Here we test the two individual models separately by reproducing the features observed in neurophysiological recordings or simulations of more sophisticated models. This was done in an effort to validate the two models by showing that they are realistic representations of the corresponding hippocampal areas and can reproduce some of their basic intrinsic characteristics and behaviors. These tests will also give an insight into the mechanism of sharp wave-ripple complexes generation by the full CA3-CA1 model, described in the next chapter.

5.5.1 Validation of the CA3 Connectivity

Before looking into the actual behavior of the CA3 model, we tested its connectivity against some characteristics seen in a more sophisticated two dimensional CA3 model by Traub et al., which accurately reproduces a large number of CA3 features [214].

The pyramidal-pyramidal connectivity was tested by counting the number of pyramidal cells reached through excitatory connections, monosynaptically, disynaptically and trisynaptically from a pyramidal cell in the middle of the array. The particular cell contacts 4.4% of all pyramidal cells directly, 59.3% through two steps and 99.7% through three steps. Therefore, virtually the whole network can be reached within three excitatory connections, which holds in the aforementioned model as well (see Figure 5.4 in [214]).

In order to test the full CA3 connectivity, we counted the pyramidal cells that are inhibited by the previous central pyramidal cell through one intermediate interneuron (one step), those that are inhibited by the interneurons contacted by its pyramidal connections (two steps) and so on. We found that 12.3% of all pyramidal cells are inhibited through one step, 67.9% through two and 99.7% through three steps. Therefore, virtually the whole network can be reached within three excitatory connections from a central pyramidal cell or can be indirectly inhibited by its activity. This is again in agreement

with the two dimensional model by Traub et al. (see Figure 5.6 in [214]).

5.5.2 Spread of Activity Over the CA3 Array

Another test of the implemented connectivity, and particularly the balance between excitation and inhibition, was performed by examining the network's response to a stimulation on one end of the array. We applied a depolarising stimulus of 8 nA for 2 msec in the leftmost 10 cells of the array and recorded the membrane potentials of pyramidal cells lying at distances up to 5 mm away (Figure 5.6A). Responses at distances below 4 mm consist of an excitatory component followed by inhibition, while no responses are observed at greater distances.

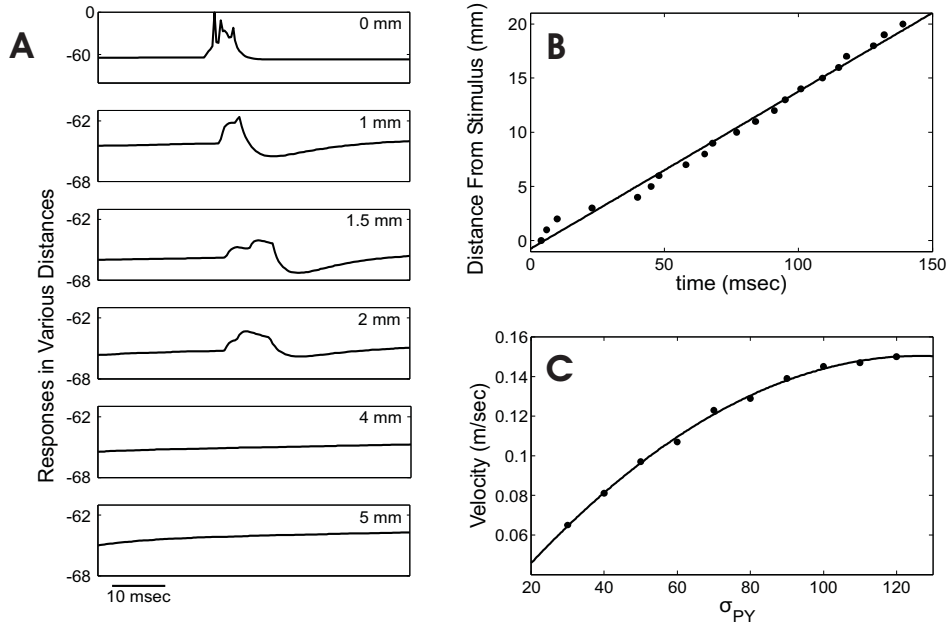


Figure 5.6: **A:** Responses of pyramidal cells at different distances from a stimulation at one end of the CA3 array. The membrane potential of one of the stimulated cells is shown in the top panel. The other panels correspond to cells at different distances from the stimulation site. **B:** Distance from the stimulation at one end of the disinhibited CA3 array, against the time the propagating activity reaches that distance. The line represents the least squares fit and yields a propagation velocity of approximately 0.145 m/sec. **C:** The propagation velocity increases as a function of the extent of excitatory connections.

We then examined the spread of bursting activity over the disinhibited network by disconnecting all interneurons and stimulating the leftmost 10 cells of the array with a 10 nA depolarising current for 2 msec. The model exhibited a wavelike propagation

of burst responses. The time the propagating activity reached a particular distance from the stimulus, was taken to be the moment the average AMPA-current over a block of ten cells, centered at that distance, became nonzero. The propagation velocity, calculated through a least squares fit to distance-time data, was approximately 0.145 m/sec (Figure 5.6B). Note that the size of the array was doubled to 2000 pyramidal cells in this simulation to get a better estimate of the propagation velocity.

Finally, we tested the influence on the propagation velocity of varying the extent of excitatory connections. We gradually increased σ_{PY} from a range of 300 μm (30 cells) to 1.2 mm (120 cells) and each time repeated the above experiment (Figure 5.6C). The propagation velocity increased along with the spread of excitatory connections, eventually reaching a plateau.

All the above features have also been observed in recordings in rat CA3 longitudinal slices and simulations of a more sophisticated CA3 computational model [144], and validate the connectivity scheme implemented here.

5.5.3 Rhythmic Population Bursts in the CA3 Model

One characteristic of CA3 activity is that population discharges can occur and quickly spread throughout the network [214]. Since such discharges are considered to be the cause of the large depolarisations that form sharp waves in CA1 (see next chapter), it is essential that they can be reproduced by our model. To examine whether the full CA3 model exhibits such population activity, we restored inhibition in the network and excited all pyramidal cells on their dendritic compartment by a current whose magnitude was Gaussian distributed over the population: $I = 0.3 \pm 0.03$ nA (mean \pm SD). To increase the input heterogeneity, it was redrawn for every cell every 1 msec of the simulation. No current was applied to the interneurons.

The whole CA3 network exhibited population oscillations that involved both pyramidal cells and interneurons (Figure 5.7A). Bursts were initiated in non-specific locations on the array and quickly spread throughout the whole network via recurrent excitatory connections. They were terminated by strong interneuronal inhibition and by the afterhyperpolarisation of pyramidal cells. The population oscillation frequency was estimated, from the autocorrelation of the average membrane potential of the array (not shown), to be approximately 7.5 Hz, which lies within the theta frequency range (4-10 Hz). Membrane potentials of individual cells (Figure 5.7B-C) indicate that the pyrami-

dal cells fire mostly in bursts but very rarely, without a clear rhythmicity. In contrast the interneurons fire in almost all cycles at high frequencies and skip cycles only if their neighbouring pyramidal cells did not fire. The shape of synaptic events, seen through the membrane potentials, varies between cycles and the underlying population oscillation is not always apparent through inspection of individual pyramidal cell potentials. One needs to examine the interneuronal activity or the total raster plot to see the overall rhythmical activity.




Figure 5.7: Quasi-synchronous theta-frequency oscillations in the full CA3 model. **A:** Raster plot (pyramidal cell spikes are shown in black and interneuronal in gray) and average membrane potential of the whole model. **B-C:** Examples of membrane potentials from one pyramidal cell (**B**) and one interneuron (**C**). Some spikes have been truncated in the figure and others appear as spikelets due to the 1 kHz sampling frequency of membrane potentials. **D:** Relation between the frequency of the population bursts in CA3 and the depolarising current, applied on the dendritic compartments of pyramidal cells. Note that even for zero current the oscillations are within the theta-frequency band.

Similar population events and membrane potential characteristics were reported via field and intracellular recordings of pyramidal cells in CA3 slices *in vitro* [68, 216], and simulations of a two-dimensional CA3 model [216]. The sparse participation of individual pyramidal cells to a population burst (only $\sim 27\%$ of cells fire in a burst on average) is a result of the strong inhibition they receive. By removing interneurons, the

model exhibits epileptiform bursts at much lower frequencies involving all pyramidal cells (not shown), similar to what was shown in disinhibited CA3 slices *in vitro* [68], and to population oscillations seen under the inhibition-suppressing cholinergic agent carbachol [215].

Finally, to examine whether the oscillation frequency of the population bursts is bounded within the theta range or depends strongly on the applied current we plotted the f-I curve for various mean amplitudes of depolarising current (Figure 5.7D). Each time, the standard deviation of the distributed currents was set to 10% of the mean. The frequency never dropped below the theta range and monotonically increased with increasing depolarisation. Although for currents higher than $I = 0.8$ nA the frequency is already slightly above the theta range, the characteristics mentioned above are distorted, as all pyramidal cells burst in every cycle and exhibit a clear oscillation. Therefore the quasi-synchronous oscillatory bursts appear to be bounded within the theta frequency range.

5.5.4 Interneuronal Gamma Oscillations in the CA1 Model

To the best of our knowledge, unlike the CA3 case, there have not been many recordings in CA1 that clearly reveal the anatomical features of the circuit connectivity. We thus targeted our CA1 model validation on reproducing functional properties of this hippocampal area, focusing specifically on gamma oscillations that have been extensively studied in CA1 [213].

We first focused on the purely interneuronal network in CA1 by removing all excitatory interactions. Figure 5.8A illustrates the model behavior when all interneurons are depolarised by a current $I = 0.3 \pm 0.003$ nA (mean \pm SD) Gaussian distributed over the population and redrawn every 1 msec (the current's SD is lower than in the CA3 model because interneuronal firing depends very strongly on the current magnitude; higher SD reduces the population synchrony). The raster plot indicates that there is rhythmical activity with local synchrony but the network lacks an overall synchronous oscillation. The average membrane potential shows a noisy rhythmicity and its autocorrelation reveals a strong peak at 14 msec corresponding to a gamma oscillation of 71.43 Hz. When increasing the mean current tenfold (3 nA) with SD again at 1% of the mean (0.03 nA) the oscillation frequency increases beyond the gamma range, reaching approximately 167 Hz (Figure 5.8B) which is within the “ripple range” (150-200 Hz). The overall synchrony appears to be much higher as the average membrane potential

exhibits clear oscillations.

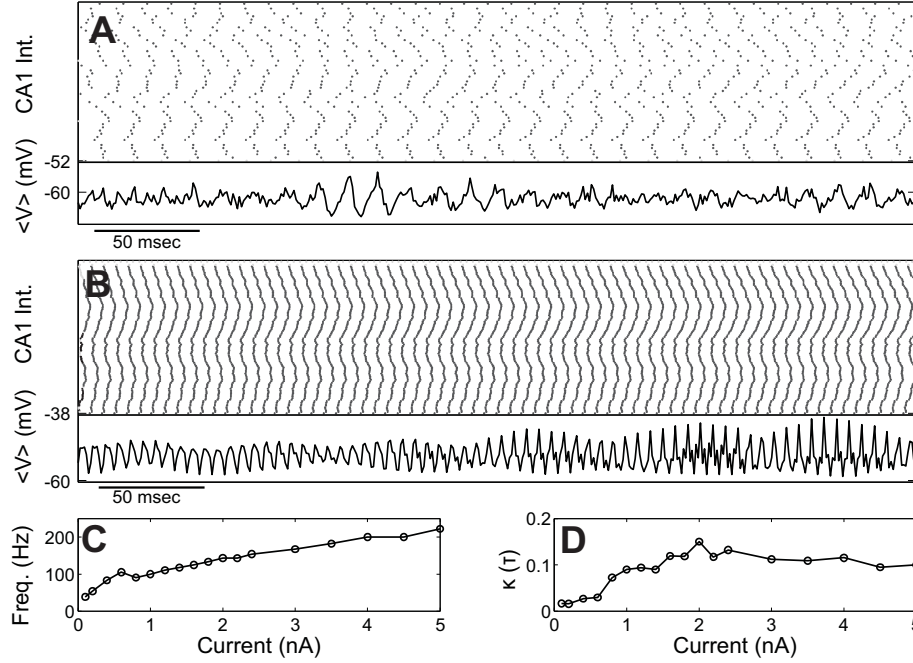


Figure 5.8: Gamma oscillation in the CA1 interneuronal network, where pyramidal cells have been removed. **A:** The interneuronal raster plot indicates weak local synchrony and the average membrane potential reveals a population oscillation of approximately 71.43 Hz, within the gamma frequency range. All interneurons are depolarised by a current that is Gaussian-distributed over the network, with 0.3 ± 0.003 nA (mean \pm SD) and is redrawn every 1 msec. **B:** When the gaussian distributed current is increased to 3 ± 0.03 nA, a faster and more synchronous population oscillation emerges, with frequency 167 Hz, within the ripple range. **C-D:** Oscillation frequency (**C**) and average synchrony measure $\kappa(\tau)$ (**D**) as a function of the mean current. Note that the population synchrony is higher at frequencies above the gamma range.

To quantify the network's overall synchrony we used the average coherence measure $\kappa(\tau)$, introduced in [228], which ranges from 0 to 1, with 1 corresponding to total synchrony. After fixing the time window τ to $0.1/f$, where f is the population frequency, $\kappa(\tau)$ was calculated to be approximately 0.022 for the first case but increased to about 0.112 for the second one, verifying that the population oscillations were much more synchronous at a frequency well beyond the gamma range.

The dependence of the oscillation frequency and the population synchrony on the magnitude of the applied current was examined by depolarising the network with currents of increasing mean amplitudes. The current's SD was fixed at 1% of the mean. Figure 5.8C reveals an almost linear dependence of the frequency on the mean current. Moreover,

the frequency does not drop below ~ 40 Hz, indicating that the network can only exhibit oscillations within the gamma range or above it. The coherence measure $\kappa(\tau)$ initially increases, reaching a peak for a mean current of 2 nA, beyond which it slowly drops, remaining at relatively high values even for corresponding frequencies above the ripple range (Figure 5.8D). As will be shown in the next chapter (Figure 6.13), this coherence at high frequencies depends strongly on the decay time of inhibition whereas the strength of inhibitory connections mainly affects the oscillation frequency. Note that, for all currents, the population frequencies are far below the range of frequencies (~ 100 -400 Hz) that interneurons with the implemented heterogeneous intrinsic properties would fire at, if they were uncoupled and the same current was applied to them [228].

5.5.5 Gamma Oscillations in the Full CA1 Model

We then restored excitatory connections and interneuron-to-pyramidal inhibition, and depolarised pyramidal cells by the same scheme that was used in the interneuronal network, with a mean current of 6 nA. No current was injected in the interneurons.

The raster plot and the average membrane potential, shown in Figure 5.9A, illustrate a stable oscillatory activity, with pyramidal cells firing in synchrony and interneurons firing mostly in spike doublets, with an inter-spike interval of 3-4 msec. The first spike of the doublets follows the pyramidal spike by 2 msec on average. The autocorrelation of the average membrane potential has a peak at 21 msec, indicating that the oscillation frequency is roughly 48 Hz, which is again within the gamma range, although well below the range of frequencies exhibited by the interneuronal network for tenfold weaker currents.

Similar gamma oscillation characteristics have been seen *in vitro* in tetanically stimulated CA1 slices [220, 231] and in computational CA1 models [217, 220, 231].

5.6 Discussion

We presented here separate network models of the CA3 and CA1 areas in the rat hippocampus, along the longitudinal direction. The two CA models reproduce a series of neurophysiological observations related to anatomical features or functional properties of the two areas.

Both models are relatively simple, one-dimensional, with similar architecture, parameter




Figure 5.9: Gamma oscillations in the full CA1 model. **A:** Raster plot of the whole CA1 model (pyramidal cells are shown in black and interneurons in gray) and average membrane potential of the whole array. The population oscillation is approximately 48 Hz with interneurons firing mostly in spike doublets with the first spike following the quasi-synchronous pyramidal spikes by 2 msec on average. **B-C:** Examples of membrane potentials from one pyramidal cell (**B**) and one interneuron (**C**).

values and connectivity schemes, and single cell models that accurately capture the firing properties of the corresponding neurons, without any detailed cell anatomy. The main emphasis was on the implemented connectivity being physiologically realistic with connectivity extents and probabilities close to reported ones and synaptic strengths following reported PSPs.

5.6.1 The CA3 Model

Recordings have shown that when applying a stimulus to one end of a rat CA3 longitudinal slice, responses consisting of an excitatory component followed by inhibition can be recorded in the stratum pyramidale at distances within 5 mm from the stimulation [144]. It has also been shown that when the stimulus is presented at one end of the longitudinal CA3 slice in which inhibition is blocked with picrotoxin, bursts propagate, without any decrease, in a wavelike fashion throughout the slice with an average propagation velocity of 0.14 ± 0.04 m/sec [144]. Simulations with a spatially exponentially decaying recurrent connectivity revealed that the propagation velocity increases when

increasing the spatial extent of these recurrent connections [144]. All the above three features were successfully reproduced by our CA3 model (Figure 5.6), validating the implemented connectivity of the recurrent excitatory connections, the balance between excitation and inhibition and the physiologically-driven choice of $\sigma_{PY} = 1$ mm which was important in getting the correct propagation velocity in the disinhibited array.

Similarly to intracellular recordings in CA3 slices *in vitro* [68, 216], and computational modelling [216], the full CA3 model produces population bursts initiating at non-specific sites and quickly propagating through recurrent excitation with pyramidal cells firing rarely during a population event and interneurons participating in every cycle with multiple spikes (Figure 5.7). Strong inhibition is necessary to keep the number of participating pyramidal cells low, thus increasing the recurrence frequency of the bursts by decreasing the overall after-hyperpolarisation in the network. The theta frequency range (4-10 Hz) occurrence of these population bursts is not to be confused with the classical theta oscillations in the hippocampus during wakefulness or REM sleep, described and studied extensively elsewhere (for a review see [33]). Moreover, all cells receive a variety of excitatory and inhibitory synaptic inputs in each population burst, unlike in CA3 slices *in vitro* where their synaptic input is almost exclusively inhibitory [68]. This inconsistency could be due to severing of long range excitatory connections in the slice preparation that diminishes the effect of excitatory interactions, which could also explain the small spatial extent of the bursts in slices [68], unlike in our model where each population burst travels throughout the whole network.

We can thus conclude that our simplistic one-dimensional CA3 model is able to reproduce, with sufficient accuracy, various characteristics and behaviors seen in experimental recordings, and in simulations of a much more complex two-dimensional CA3 model [214]. Note that various features of CA3 circuits such as transduction probabilities, PSP rise times, synaptic depression or facilitation [8] have not been included in the model. For example it has been estimated that the single-spike-transduction probability in a pyramidal-pyramidal connection is 5% but the burst-transduction probability can rise up to 50% [214]. In our model, neither a single EPSP nor a single burst can cause a post-synaptic pyramidal cell to fire. Nevertheless such features appear not to be necessary for the simulation of the functional and anatomical features presented here.

5.6.2 The CA1 Model

Interneuronal gamma oscillations have been detected in the CA1 region *in vitro* under tetanic stimulation, in the presence of glutamate receptor blockers [233]. This and other studies indicated that transient gamma oscillations can rely solely on GABA-mediated inhibition, with fast somatic inhibition, namely basket cells, being the most likely candidate for generating and sustaining such oscillations [16]. Moreover, various computational models of inhibitory networks reproduced gamma oscillations, usually assuming slow and weak inhibition and homogeneous or global network connectivity [219, 228, 230, 231]. The population synchronization in such models was shown to be sensitive to heterogeneities in connectivity and in the driving current, and changes in the synaptic strength and kinetics [228, 230]. Even slight inhomogeneities or a limited spatial spread of connections can shift the network from global synchrony to local [213]. This very limited spatial synchrony is also encountered in our CA1 inhibitory network during gamma oscillations. Tonic depolarisation, albeit with a noisy component, short-range heterogeneous connectivity and noisy intrinsic cell parameters do not allow an overall synchronization of firing during the oscillations (Figure 5.8A).

This sensitivity to heterogeneity could be ameliorated to some extent by incorporating fast, strong, shunting inhibition along with synaptic delays [15, 16]. Though lacking shunting inhibition, synchrony is enhanced in our network by incorporating multiple connections between neighbouring interneuronal pairs, with strong and fast decaying IPSPs along with axonal delays. In fact, axonal conductance delays were crucial, since, by removing them, any synchrony of oscillations was completely abolished (not shown). The importance of incorporating long synaptic delays, longer than the inhibitory decay time, in establishing fast synchronous interneuronal oscillations has been shown through mathematical and computational analysis in a similar set up [27]. In our model, the conductance delay between two connected interneurons can be longer than the decay time, reaching up to 3 msec (0.1 mm/msec conductance velocity with connections reaching out to $\sim 300 \mu\text{m}$, Table 5.2), thus aiding the oscillation synchrony.

Although the frequency of the population oscillations does not drop below the gamma range, with stronger depolarisation it reaches the ripple frequency range. The model reproduced the linear relationship between the applied current amplitude and the oscillation frequency (Figure 5.8C-D) seen in other inhibitory network simulations [219, 228]. Interestingly, we found that global synchrony is considerably stronger at those higher frequencies, with cells firing in a clear rhythmical pattern (Figure 5.8B). A similar re-

lationship between synchrony at ripple frequencies and the external drive was reported in simulations of interneuronal networks with global random connectivity and fixed delays [27]. As will be shown in the following chapter, synchrony at such high frequencies depends strongly on the small decay time of inhibition. The condition for this time constant to be smaller than ~ 5 msec for synchronous oscillations at ~ 200 Hz to occur, was shown through mathematical analysis and computational modelling of heterogeneous interneuronal networks [40, 230]. It appears that even with our scheme of spatially restricted connectivity and distance-dependent delays, a similar condition and synchrony-drive dependence still hold. Note that our unusual choice of 2 msec for the decay time is based on recent electrophysiological studies which report similar GABA_A time constants in CA1 [15, 16].

When depolarising the pyramidal cells in the full CA1 model they fired synchronously followed by interneuronal spike doublets (Figure 5.9). The network oscillated, still within the gamma range, but at considerably lower frequencies than the interneuronal ones, even though the applied current was much stronger. The frequency decrease is a result of the extra delays introduced by long range excitation in the pyramidal-interneuron-pyramidal loop of the full network [27]. Similar characteristics of gamma oscillations have also been seen *in vitro* in electrophysiological recordings from tetanically stimulated CA1 slices [220, 231] and have been reproduced in computational models with a simplistic chain-like architecture [220, 231] or more realistic two-dimensional connectivity [217]. In these recordings, when the stimulation was performed simultaneously on two sites 4 mm apart, very tightly synchronised transient gamma oscillations appeared in both sites after a short latent period. This long range synchrony that extended several millimeters consisted of pyramidal cells firing single spikes and interneurons firing mostly spike doublets with the first spike being in phase with the pyramidal spikes [220, 231]. Doublets, along with long range synchrony, would disappear when the network was stimulated only locally implying that doublet firing is a result of the increased excitation the interneurons received due to the long-range axonal connections of further away pyramidal cells [220, 231]. These results led to the suggestion that doublet firing is the necessary driving mechanism for long range synchrony of gamma oscillations in CA1 (for a qualitative analysis see [213] which was also demonstrated with analytical methods [70]. It has been shown that for gamma oscillations in a network with long range non-recurrent excitation and local inhibition, doublets would arise when the first interneuronal spike was generated by local excitatory input whereas the second one by the delayed input from more distant cells. The resulting inhibition determines the time

of the next spike of nearby pyramidal cells and consequently the next oscillatory cycle. Under this regime, the network undergoes synchronous oscillations, robust to mild heterogeneities [70]. In our model, the lag between the pyramidal spikes and the two spikes of the doublets (2 and 5-6 msec respectively) implies that the first spike is due to excitatory input from local cells at distances within σ_{PY} whereas the second one is from distant cells up to $3\sigma_{PY}$ away, verifying the above analysis.

The ratio of excitation to inhibition and the synaptic delays are important in this mechanism [27, 70]. Specifically, removing conductance delays, either excitatory or inhibitory, completely abolished any oscillations. Moreover, if pyramidal-to-interneuron excitation is not strong enough to evoke doublets, or if feedback inhibition is not strong enough to time the next firing of pyramidal cells, then long range synchrony can not be sustained. In support of this, administration of AMPA-receptor antagonists [231] or GABA-blocking morphine [232] to CA1 slices disrupts long-range synchrony of the oscillations. Reducing the unitary AMPA-conductance in our model by 75% or the GABA-conductance to pyramidal cells by 75% and to interneurons by 85%, as in simulations in [232], abolished long-range oscillations as well, although some gamma rhythmicity persisted in the former case and small clusters of locally synchronous oscillations still existed out of phase (also reported by [213]). The fact that our full CA1 model reproduces these features indicates a correct balance between inhibition and excitation in the network.

Having validated both CA3 and CA1 models by establishing that they accurately reproduce a series of anatomical and functional properties of the corresponding areas, in the next chapter we couple them in a full CA3-CA1 network in order to model sharp wave-ripple complexes. Through our analysis of these complexes, various modelling assumptions made here, such as the inclusion of fast synapses and fast somatic inhibition exclusively, will be justified. The mechanism behind the interneuronal gamma- and ripple-frequency activity is also further explored within the framework of the whole CA3-CA1 network.

A Computational Model of Sharp Wave-Ripple Complexes

6.1 Introduction

Sharp Wave-Ripple Complexes (SWRs) are a self-organised endogenous hippocampal event, occurring when the brain has minimal or no interaction with the environment. They are the first and only organised population activity in the developing hippocampus, present even when the hippocampus is completely isolated [33].

Below we review the two components of this complex separately and we briefly outline some of their basic characteristics. We describe one existing modelling approach to ripples and we then describe how our CA3 and CA1 models are coupled together in a realistic CA3-CA1 network which exhibits SWRs and suggests a novel mechanism for their generation. We present the characteristics of these complexes and compare them with those reported in literature. We finally discuss the mechanisms behind them and their possible implications in memory consolidation.

Parts of this chapter were published in the *Hippocampus* journal [203].

6.2 The Sharp Wave-Ripple Complex

6.2.1 Sharp Waves

One major hippocampal EEG pattern, observed virtually simultaneously in different regions of both hippocampi during deep sleep and anaesthesia, is the so called “sharp

wave” (SW). It is also present during awake rest, consummatory behavior and immobility in the rat brain and disappears during sensory stimulations [32]. It consists of a large negative EEG peak followed by small positive peak, resembling an “EEG spike” (Figure 6.1, bottom trace) and is located mostly in the pyramidal cell layer of the CA regions and specifically in the apical dendritic layer of the CA1 region where its amplitude reaches up to 3.5 mV [32, 241] (Figure 6.2, left). Its duration varies between 30 and 120 msec and its frequency of occurrence ranges between 0.01-3 Hz [32, 136]. Under anaesthesia SW periods alternate spontaneously with theta oscillation periods in a 20-40 min cycle [241].




Figure 6.1: Simultaneous recordings from the CA1 rat pyramidal layer (electrode 1) and stratum radiatum (electrode 2) exhibiting an SWR. Electrode 2 displays an example of a dendritic SW. Uppermost trace of electrode 1 is wide-band recording (1 Hz to 10 kHz). Second and third traces are digitally filtered traces showing pyramidal layer unit activity (500 Hz to 10 kHz) and ripple oscillation (100-400 Hz) respectively. Calibrations: 0.5 mV (top trace), 0.25 mV (middle traces) and 1.0 mV (bottom trace). Taken from [34].

SWs can be detected in EEG and LFP hippocampal recordings by low-pass filtering the raw data at 50 Hz and setting a threshold for event detection at a few (4-5) times the standard deviation of event-free baseline data [136]. This threshold (and the subsequent one for ripples) is not a standard one and is usually set by visual inspection of the available data.

Various studies have shown, through inspection of raw traces, correlation analysis and current source density analysis, that SWs emerging in the CA1 region are the result of strong depolarisation by the CA3 Schaffer collaterals caused by the synchronous bursting

of CA3 pyramidal cells [32, 54, 136, 241]. Specifically, population bursts were found to emerge first in the CA3a-b subregions and then spread to the CA3c [54]. Both *in vivo* [41] and *in vitro* [136], SWs propagate from the CA3 region towards the CA1 and subiculum regions and even backwards to the dentate gyrus, with a propagation velocity within CA1 of around 2-7 cm/sec [32, 136]. In one study, at least 10% of the recorded CA3 neurons should discharge synchronously within a 100 msec window to exert a detectable influence on CA1 cells [54]. Disconnecting the Schaffer collaterals or isolating the CA1 results in total or great suppression of SW generation in CA1 [136]. Therefore it appears that SWs are generated in the CA3 recurrent excitatory network which can trigger synchronised activation of the CA3 pyramidal cells (see previous chapter). The synchronous bursting creates a barrage of EPSPs that travel through the Schaffer collaterals to the CA1 region where recurrent connections between pyramidal cells are sparse (possibly explaining why SWs are so rare in isolated CA1 slices). Naturally these CA3 discharges activate not only CA1 pyramidal cells but interneurons as well, possibly giving rise to the second component of the SWR complex.

6.2.2 Ripples

Hippocampal SWs typically appear in EEGs and LFPs in conjunction with another characteristic oscillatory pattern of the hippocampus, the so called fast-field oscillations or “ripples”. Ripples are observed during the negative phase of the SW and together they form the SWR (Figure 6.1), present both in *in vivo* recordings and *in vitro* [34, 136]. Ripples are spindle-shaped oscillatory patterns, of ~ 30 -120 msec duration consisting of 5 to 15 sinusoidal waves with ~ 150 -200 Hz intraburst frequency [34, 136, 241]. They are generated synchronously mostly throughout the CA1 area as a result of the depolarisation from CA3 synchronous bursting during SWs [34]. The main current source of ripples in CA1 was shown to be located in the pyramidal cell layer, where their amplitude is maximal and decreases sharply towards the stratum radiatum and stratum oriens [34, 136, 241] (Figure 6.2, left).

Ripples can be detected by band-pass filtering the raw EEG or LFP data (150-250 Hz) and setting a threshold of detection at a few (3-4) times the standard deviation of the event-free baseline noise [136, 149]. Further criteria can also be applied, e.g. that ripples must last roughly 25-75 msec [149], that they must contain at least six successive peaks/troughs [147], or that they have to consist of at least three consecutive negative peaks below threshold, with interpeak interval within 2-11 msec (90-500 Hz) [136].

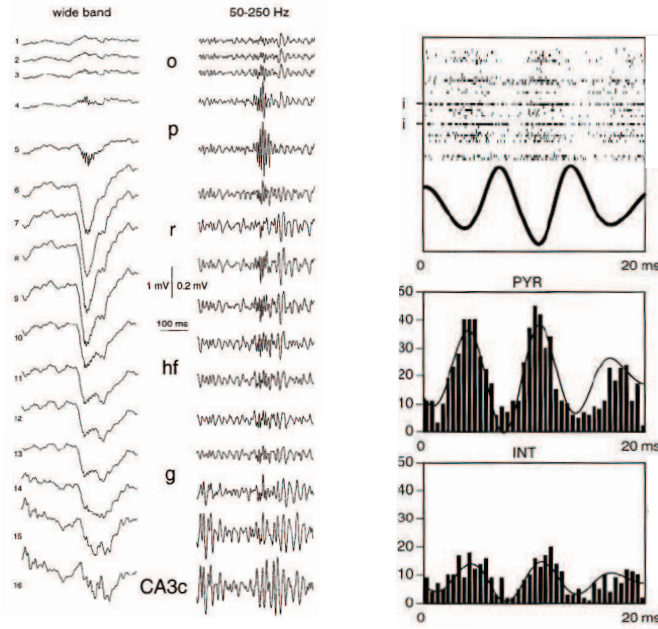


Figure 6.2: Left: Depth profile of a single SWR event. Wide band (1 Hz to 5 kHz, left column) and bandpass filtered (50-250 Hz, right column) simultaneous recordings of a SWR event from 16 sites in the CA1-dentate gyrus axis (o: stratum oriens, p: pyramidal cell layer, r: stratum radiatum, hf: hippocampal fissure, g: granule cell layer, CA3c: hilar portion of the CA3 pyramidal layer). Note high-frequency oscillation (ripple) associated with the SW peak, with an amplitude maximum at the CA1 pyramidal cell layer. Right: Relationship between ripple oscillation and neuronal spike activity. Averaged field event and superimposed spike events from seven recording sites in the CA1 pyramidal layer (300 μ m horizontal tip separation, 40 pyramidal cells and two putative interneurons). Isolated pyramidal cells and interneurons (top panel) and cross-correlogram between field oscillation and summed pyramidal cell activity (PYR) and interneuronal firing (INT). Taken from [241].

Firing rates in the pyramidal cell layer increase significantly during the ripples [53]. Yet, only a small percentage of pyramidal cells actually fire during the ripples, with their spikes phase-locked to the ripple troughs [34, 241] (Figure 6.2 right), and with each cell firing only a few times (one or two usually) per ripple, so not at the population frequency [53, 136]. Interestingly, the participation of different cells in ripple spiking is not homogeneous, with extracellular recordings indicating 10% of all recorded pyramidal cells participating in an average ripple with some cells participating in up to 40% of successive episodes and others firing very rarely [241]. It thus appears that only a minority of cells fire during the majority of ripples. Moreover, intracellular recordings of individual pyramidal cells during ripples revealed a fast, 1-5 mV membrane voltage

oscillation that coexisted with the ripple [241].

Unlike pyramidal cells, the different classes of interneurons in CA1 differentiate their participation in ripples with some classes increasing their firing rate during ripples (mostly basket cells) while others show no change or even reduce firing [120, 121]. Perisomatic-inhibitory basket cell discharge rates increase severalfold, reaching high frequencies but rarely ripple frequencies, with multiple spikes per event [53]. This lack of a one-to-one correspondence between unit spikes and ripples in pyramidal and interneuronal cells indicates that ripples do not reflect the activity of single cells, but instead are clearly a field rhythm. Interneuronal firing is also phase-locked to the ripples, slightly following the negative ripple peaks by 1-2 msec [241] (Figure 6.2, right). This intense activity causes a pronounced inhibition in most pyramidal cells [68, 136].

An important feature of SWRs is that they are widespread, extending throughout the CA3-CA1-subiculum-entorhinal cortex axis, involving the synchronised participation of tens of thousands of cells [41]. This synchronous activity that quickly builds up during the ripples is not consistently initiated from the same site but emerges simultaneously from several locations [241]. The enhancement of population synchrony of the CA1 pyramidal cells (defined as percentage of cells firing) during ripples, is more than 2 times greater than that of CA3 pyramidal cells, while interneuronal population synchrony is similar between the two regions [54]. The strong transient output, produced during such massive population discharges, is very likely to affect neocortical targets, rendering SWRs a possible means for memory transfer from the hippocampus to the neocortex for long-term storage, during deep stages of sleep. Correlations between ripples and cortical slow oscillations during deep sleep [17, 104, 185, 188] and the fast-scale replay, during SWRs, of temporal patterns of correlated place cells in forward [129, 189, 235] or reversed order [62, 74] support the role of SWRs in memory consolidation. This role is discussed in the final section of this chapter and further studied in the next chapter.

6.2.3 Possible Mechanisms for Ripple Generation

Although SWRs have been observed in rodents, monkeys and humans [33] and have been widely studied *in vivo* [32, 34, 41, 53, 120, 241] and *in vitro* [19, 24, 135, 136, 153, 238] in rats and mice, the exact cellular and synaptic mechanisms for ripple generation in CA1 and the source of the high population synchrony during ripples still remain unclear. Sparse networks of CA1 pyramidal cells coupled together with axo-axonal gap junctions have been proposed to underly high frequency hippocampal oscillations

[65, 180]. This proposed mechanism is supported by the suppression of ripples under gap-junction blockers [136, 241], but contradicted by lack of SWR generation from antidromic stimulation of CA1 pyramidal axons and by preservation of SWRs after removal of the distal part of these axons where the gap junctions are expected to lie [24].

Another suggested mechanism involves the cross-talk between pyramidal cells and interneurons. It has been proposed that the excitatory barrage by discharging CA3 and CA1 pyramidal neurons results in a relatively steady depolarisation of the target interneurons causing fast rhythmic oscillation of their membrane potential and sustained firing. When a critical number of spatially distributed interneurons fires at ~ 200 Hz, they will coordinate the firing of their pyramidal cell targets which in turn will phase reset the next firing of their interneuronal contacts. Thus spatially distributed synchrony could be established through pyramidal-interneuron interactions [34, 53, 241]. It is noteworthy that the involvement of fast perisomatic inhibition was recently shown to be critical for the *in vitro* generation of SWRs in CA3 slices [68].

6.3 A Model for Ripples Based on Axo-Axonal Gap Junctions

An existing computational model of ripple oscillations is based on random networks of pyramidal cells coupled by axo-axonal gap junctions [217]. All pyramidal cells were modelled by the extended Traub model [212] (see Section 5.3.1) that contains multiple axonal segments. Random networks of such modelled neurons were implemented, containing only excitatory AMPA connections, to simulate a CA1 network. Gap junctions were placed between the penultimate axonal compartments of some pairs of cells. These connections were sparse and localised, so that a cell could have on average less than two gap junctions with other cells, and only with cells that lie within $200 \mu\text{m}$.

Simulations demonstrated that such networks can generate transient 100-200 Hz synchronised oscillations. These oscillations appeared irrespective of whether the network included recurrent excitatory synapses or consisted purely of gap junctions. In fact it was shown that these network bursts required only two elements in the structure of the network:

1. a sparse but strong axo-axonal gap junction connectivity,
2. a background of spontaneous current pulses to the axons which are capable of

inducing action potentials provided the soma is not too hyperpolarised.

Thus, the mechanism for the high frequency field oscillations appeared to be determined not by membrane and synaptic parameters, but by topological structure and by the time for an axonal spike to pass to another axon through a gap junction.

A further study on this subject involved an extended version of the above network, including interneurons which were divided into four classes (perisomatic basket cells, axo-axonic cells inhibiting initial segments of pyramidal axons, and two types of dendrite contacting interneurons) [211]. Pyramidal cell synaptic connectivity was global in the network but interneurons could contact only cells within 500 μm . Noise was again implemented in the form of spontaneous Poisson spikes applied in the axons of cells. SW-like inputs were also implemented by applying sinusoidal excitatory conductance pulses on the dendrites of all cells, emulating the Schaffer collateral input. A block diagram of the network model is given in Figure 6.3. Simulations yielded a ~ 140 Hz coherent oscillation during the SW-input, with interneurons increasing their firing to almost ripple frequency. The SW depolarisation helps the spontaneous induced spikes to propagate in the network through the axo-axonal gap junctions, generating the high frequency network oscillation. Since most spikes did not backpropagate from the axons, the pyramidal cell somata appear to fire at lower rates than interneurons, but their axons fire at more than ripple frequencies. Again the oscillation frequency depended strongly on the topological characteristics and the parameters of the gap junctions and on the spontaneously induced spikes.

Both these models are based on axo-axonal gap junctions in order to achieve large scale firing synchrony during ripple oscillations. There do exist some biological indications that gap junctions might indeed be the mechanism that drives fast synchrony, e.g. under the gap-junction blocking anaesthetic halothane, ripples are suppressed [241]. Nevertheless, there has not been a direct and definitive morphological evidence that axo-axonal gap junctions in hippocampal neurons exist. If their density in the hippocampus is as sparse as the theory predicts (less than two junctions per axon suffice to get ripples), then it is natural that such evidence is extremely hard to obtain. Nevertheless, some recent *in vitro* experiments in hippocampal slices do support the scenario that some pairs of pyramidal neurons can be electrically coupled not far from the axon's initial segment [180].

Below we describe the results of our combined CA3-CA1 network model which we show to exhibit SWRs sharing many characteristics with those reported via recordings. Our




Figure 6.3: Block diagram of the gap junction-based network model for ripples in [211]. The model contains 3072 pyramidal cells and 384 interneurons (basket cells, axo-axonic cells, and dendrite-contacting cells). The figure highlights the conceptual distinction between the axons of the pyramidal cells and the somatodendritic compartment. Interactions between components of the model are shown. Afferent inputs in the model are purely excitatory. Only AMPA and GABA_A receptors are simulated. Taken from [211].

model proposes a ripple-generating mechanism based on high interneuronal activity and strong recurrent inhibition in CA1. Gap junctions in this model are not required for ripples but could potentially aid in their long-range synchrony.

6.4 Connecting the CA3 and CA1 Models and Detecting Ripples

6.4.1 Schaffer Collateral Connectivity

Schaffer collaterals are the major input to CA1 pyramidal cells with each CA3 pyramidal cell contacting 30,000-60,000 cells and its axonal projections extending to more than two thirds the length of CA1 [130]. Schaffer collaterals are expected to excite CA1 interneurons as well, although to our knowledge the characteristics of such connections have not been studied yet.

We connect the two CA models, described in the previous chapter, in a feedforward fashion, with each CA3 pyramidal cell contacting 130 CA1 neurons on average (see also Section 5.4.4). Tables 6.1 display the synaptic and connectivity parameters of the

		α_{syn}	PSP (mV)	Reported PSP (mV)	Reference
Schaffer	PY-PY	1.5	0.13	0.13	[177]
	PY-IN	0.8	0.4		

		Average Cluster Size	k_{ij}	\tilde{k}_{ij}	Connectivity probability
Schaffer	PY-PY	548.02	130	99.81	18.21%
	PY-IN	54.84		10.01	18.25%

Table 6.1: Parameter values for the synaptic strength and connectivity of the modelled Schaffer collaterals between the CA3 and CA1 models, similarly to Tables 5.1 and 5.2 of the previous chapter respectively.

implemented Schaffer collaterals as in Tables 5.1 and 5.2 in the previous chapter. Since, to our knowledge, there have been no reported values for Schaffer collateral excitatory PSPs (EPSPs) on CA1 interneurons, we set the EPSP to 0.4 mV. This gives a ratio of Schaffer collaterals' EPSPs to EPSPs from CA1 pyramidal cells equal to $0.4/1.4 \approx 28\%$. This is comparable to the corresponding ratio for CA1 pyramidal cells which is roughly $0.13/0.7 \approx 19\%$ (given an average EPSP amplitude between CA1 pyramidal cells of roughly 0.7 mV [61]). Similarly to the connectivity scheme in the previous chapter, the Schaffer collateral connections of each CA3 cell are distributed according to a Gaussian probability distribution centered on the corresponding pyramidal cell in CA1. However, the connections are distributed in CA1 irrespective of whether the contacted cell is pyramidal or interneuron. The SD was set to $\sigma_{CA3-CA1} = 1.2$ mm, so that connectivity clusters can reach 7.2 mm in size, roughly two thirds the total extent of the array.

Consequently, each CA1 cell receives weak input from approximately 100 CA3 pyramidal cells, which is a great underestimate of the total amount of excitation each cell is expected to receive from CA3 [8]. To counterbalance the small number of pyramidal cells in our CA3 array and have the necessary strong input in CA1 we assign multiple synapses to each CA3-to-CA1 connection. Although, to our knowledge, the number of contacts between a CA3 and CA1 pyramidal cell has not been thoroughly studied (see [8] for an analysis of the difficulty of such a task), it has been estimated that it lies between two and ten contacts [190]. For each CA1 interneuron we fix the number of synapses it receives from each contact to 13. The corresponding number for the pyramidal cells varies. For each cell, it is drawn from a Gaussian distribution: 13 ± 13 (mean

\pm SD). To avoid negative numbers, the absolute value of each drawn number is taken. This scheme results in most cells receiving relatively few synapses (~ 15 on average) from each assigned CA3-to-CA1 connection. Nevertheless, a minority of pyramidal cells receives many more synapses than average, yielding a subset of pyramidal cells that will be driven by much stronger input from CA3 than the rest. These cells are assumed to represent real neurons in the rat CA1 that are targeted by a larger portion of CA3 or receive stronger EPSPs than average. An example of the histogram of this heterogeneous “excitability” is shown in Figure 6.4.




Figure 6.4: Histogram of the synapses-per-Schaffer collateral connection of CA1 pyramidal cells. The dotted line represents the mean of the distribution (~ 15 synapses).

6.4.2 Simulations of Extracellular Recordings

Most studies on SWRs are made with depth-EEG and LFP recordings, detecting extracellular activity from an area around the tip of the electrode. To simulate such local activity recordings in CA1, we avoided storing data from the whole array and focused on a specific site of length $560 \mu\text{m}$ in the middle of the CA1 array, containing 50 pyramidal cells and 6 interneurons. We will refer to this as the “default recording site”.

We model extracellular synaptic activity using two similar measures:

1. the total post-synaptic currents $I_{syn}(t)$ and
2. the total conductances $\bar{g}_{syn}s_{syn}(t)$,

summed over all synaptic connections over the whole recording site. The first measure was used to reproduce raw extracellular recordings visually and for the study of post-synaptic currents developing in the average cell during a SWR. However, its definition contains the intracellular membrane potential of cells (eq. 5.4.1), so it will be affected by any phase delays between population synaptic activity and membrane potential responses. Since the second measure takes only non-negative values it does not reproduce

the raw recordings of an EEG or LFP probe, but is appropriate for fine timescale analysis since it will not be affected by any of the aforementioned delays. We will refer to these two measures as “total synaptic currents” and “total synaptic conductances” throughout the text.

In order to simulate CA1 dendritic layer and pyramidal cell layer recordings separately, the aforementioned measures are split into two components:

1. Dendritic layer recordings are assumed to correspond to the excitatory synaptic currents/conductances from Schaffer collaterals that act on the dendritic compartment of pyramidal cells and on the single-compartment interneurons (since the dendrites of basket cells lie in the CA1 dendritic layer [76]).
2. Somatic (pyramidal cell) layer recordings are assumed to correspond to synaptic currents/conductances from pyramidal-to-interneuron, interneuron-to-pyramidal and interneuron-to-interneuron connections.

6.4.3 Ripple Detection

Ripples are detected from the 150-200 Hz bandpass filtered version of the total synaptic conductances of the somatic layer (Figure 6.5). The algorithm for their detection is the following: The root mean square (RMS) of the measure is calculated in non-overlapping bins of 10 msec duration. The SD of the RMS is derived over the whole signal. Ripples are detected when the RMS exceeds a $2 \times \text{SD}$ threshold. Their boundaries are set where the RMS drops below $1 \times \text{SD}$ around the ripple. Ripples with less than 20 msec total duration are discarded and neighbouring ripples less than 10 ms apart are taken as one event. Similar algorithms were used in [53, 54, 120].

In a typical 30 sec simulation, 170 ripple episodes are detected, with average duration of 61 msec (similar to the 50 ms duration in [241]) with a range between 20-90 msec (30-80 msec in [136]).

For the averaging over ripples, all detected events are aligned using their minimum point as reference (point 0 msec).




Figure 6.5: Ripple detection. The root mean square (RMS, thick solid line) of the filtered total synaptic conductance (grey) is calculated. Ripples are detected when $\text{RMS} > 2 \times \text{SD}$ (dashed line). Their boundaries (solid lines) are set where RMS drops to $1 \times \text{SD}$ (line) around their peak. If their total duration is less than 40 msec they are discarded. Detected ripples that lie less than 10 msec apart are taken as one.

6.5 Simulation Results

6.5.1 Sharp Wave-Ripple Complexes in the Full CA3-CA1 Model

All pyramidal cells in CA3 are depolarised by implementing the same variable Gaussian-distributed current of 0.3 ± 0.03 nA as in Figure 5.7. No external input is given to interneurons or CA1 cells.

Figure 6.6A illustrates the activity of the full model as depicted by the raster plots of both areas and the average membrane potential and total synaptic currents of all CA1 cells within the recording site. As described in the previous chapter, the CA3 area exhibits theta-periodic population bursts (4-10 Hz) that initiate in non-specific locations and spread over the whole CA3 network. Each population burst produces a corresponding burst of activity in CA1 by exciting both pyramidal cells and interneurons through the Schaffer collaterals. Note that interneurons participate very strongly in the CA1 burst, causing the average membrane potential of the recording site to decrease significantly on every burst. The synaptic currents reveal a fast oscillatory activity during each burst. The individual membrane potentials of three random CA1 pyramidal cells and three interneurons (Figure 6.6C-D) imply that although both pyramidal cells and interneurons receive strong excitation from CA3, only the latter fire almost on every burst and usually with a large number of spikes, while pyramidal cells fire rarely with mostly just one spike per cycle. Similar spiking characteristics have been observed in CA1 ripple recordings [34, 53, 241]. Finally both cell types show rich synaptic activity that fluctuates strongly with a very fast alteration between excitatory and inhibitory synaptic inputs.

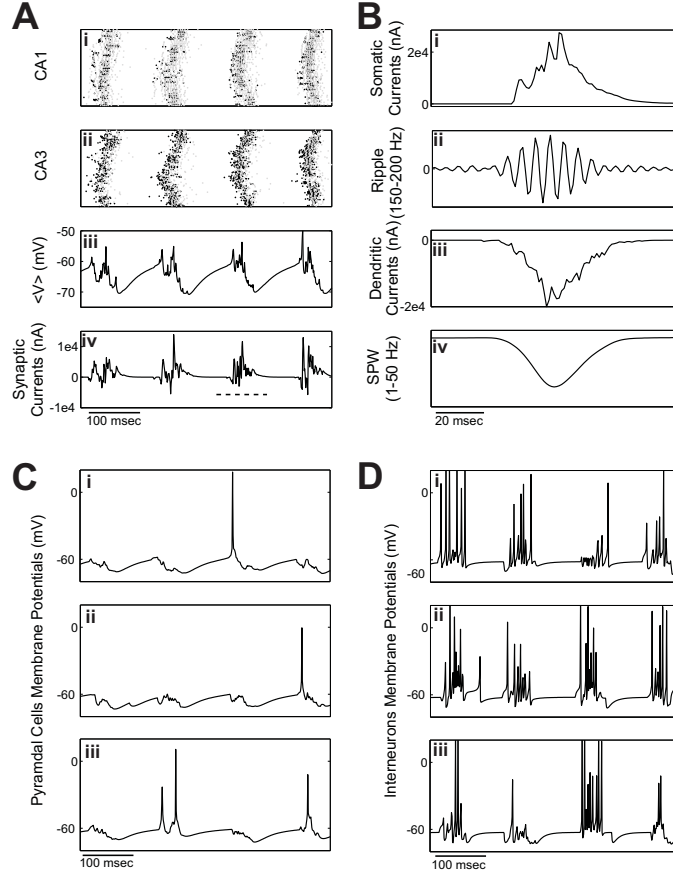


Figure 6.6: **A:** Activity of the full CA3-CA1 model. (i) Raster plot of the whole CA1 array, (ii) raster plot of the whole CA3 array (pyramidal cell spikes are shown in black and interneuronal in gray), (iii) average membrane potential and (iv) total synaptic currents of the 560 μm -long recording site in CA1. **B:** (i) CA1 synaptic currents corresponding to the pyramidal cell layer during the underlined burst and (ii) their 150-200 Hz bandpass filtered version. (iii) Schaffer collateral excitatory currents corresponding to the dendritic layer and (iv) their 1-50 Hz lowpass filtered version. **C-D:** Membrane potentials from (C.i-iii) three CA1 pyramidal cells and (D.i-iii) three CA1 interneurons.

To examine the activity in CA1 during an individual burst we zoom into one of them (underlined in Figure 6.6A.iv). In Figure 6.6B, the total synaptic currents of the recording site were split into those corresponding to the pyramidal cell layer (Figure 6.6B.i) and those corresponding to the dendritic layer (Figure 6.6B.iii). Bandpass filtering the somatic currents in the 150-200 Hz frequency range (Figure 6.6B.ii) yields a clear ripple-like oscillatory pattern. Similarly, lowpass filtering the dendritic currents in the 1-50 Hz range (Figure 6.6B.iv) yields an SW curve that peaks in parallel with the ripple. This

activity, exhibited by the model, closely resembles corresponding raw and filtered CA1 extracellular recordings during SWRs [34, 53, 120, 136, 241].

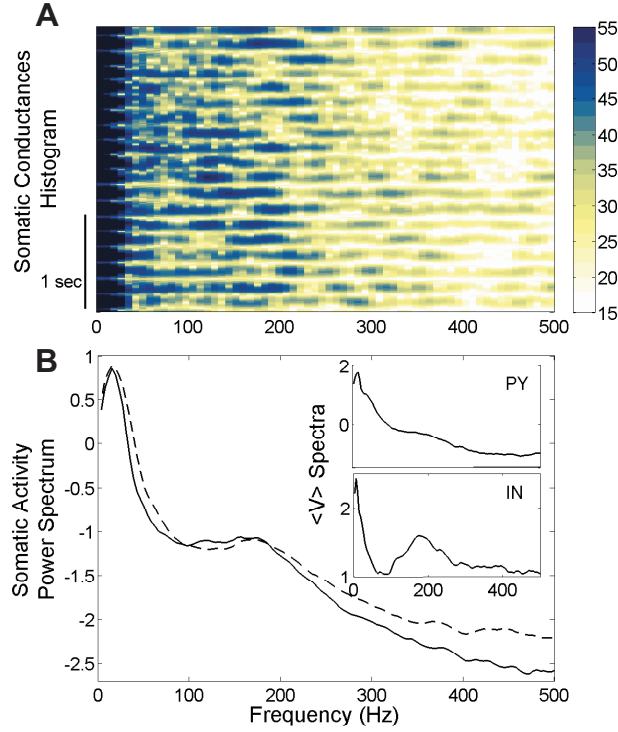


Figure 6.7: **A:** Spectrogram of the total synaptic conductance over the recording site in the full CA3-CA1 model. Bands of high power at ripple frequencies (~ 150 - 200 Hz) are visible. **B:** Power Spectrum of the total synaptic conductance (solid line) and the total synaptic currents (dashed line) in the site over a 30 sec-long simulation. Both measures were normalised by subtracting their means and dividing by their standard deviations. The strong theta-peaks are due to the theta-periodic recurrence of the input from the CA3 bursts. Insets: Power spectra of the average membrane potentials of pyramidal cells (top) and interneurons (bottom) in the recording site over the same simulation.

The power spectral profile of the synaptic activity in the CA1 somatic layer was examined in order to ensure that the observed fast oscillatory modulation lies within the ~ 150 - 200 Hz ripple-frequency range. The spectrogram of the total synaptic conductance (Figure 6.7A) reveals narrow bands of high power bounded in a range of frequencies roughly around the ripple range but varying between cycles. Figure 6.7B contains the power spectrum of both the total synaptic conductance and total synaptic currents in the recording site over a 30 sec-long simulation. A clear peak appears in the spectral power of both measures around the ripple frequency range. The insets show the power spectra

of the average membrane potential of pyramidal cells (top) and interneurons (bottom) in the recording site. The interneurons show a very clear intracellular ripple-frequency oscillation whereas the pyramidal cells exhibit a much less pronounced peak, suggesting the interneurons as a likely candidate for generating ripple oscillations.




Figure 6.8: Field oscillation correlation with unit firing and average membrane potential. **A:** All detected ripples (gray lines) are aligned according to their minimum point and the average ripple is drawn (thick black line). Spikes from pyramidal cells (gray) and interneurons (black) are aligned around the corresponding ripples. **B:** Spike histogram of pyramidal cells aligned with the average ripple. Pyramidal cells are split into two groups: the highly excitable ones that receive more than 19.5 synapses for each Schaffer collateral connection ($1.5 \times$ the average, 13; gray bars) and those that receive less than 19.5 (black bars). **C:** Interneuronal spike histogram. **D:** Average membrane potential of pyramidal cells (solid line) and interneurons (dashed line) aligned with the average ripple.

We next examined correlations between spike activity and the average ripple. First we isolated all detected ripple events and calculated the average ripple as described in Section 6.4.3. The spikes of all pyramidal cells and interneurons in the recorded site during each detected ripple were also aligned with its minimal point as reference (0 msec) and were superimposed on the average ripple (Figure 6.8A). We calculated the spike histogram for all cells, aligned with the average ripple (Figure 6.8B-C). At this

point we separated the pyramidal cells into two groups, those that receive many synapses for each Schaffer collateral connection and those that receive a close-to-average number or below. We arbitrarily set the threshold to 19.5 synapses per input connection (1.5 times the mean number of synapses per connection). The first group contains 32% of all pyramidal cells, those that are very excitable due to the stronger input they receive. The histograms of both groups are shown in Figure 6.8B (grey bars for the first group and black for the second). The interneuronal histogram (Figure 6.8C) shows a very clear rhythmic spiking, with maximum activity phase locked to the negative half-wave of the average ripple, its peak slightly following the ripple's negative peak (~ 1 msec). In the pyramidal cell histogram the number of spikes is much lower, and the activity differs between the two cell groups. Specifically, cells that belong to the group receiving weaker excitation appear to have almost no ripple modularity in their firing. Most of their firing activity is concentrated on the beginning of the event, before significant synchronous inhibitory spiking begins. In contrast, the cells receiving stronger input exhibit clear rhythmic spiking similar to the interneurons. Their activity is also phase locked to the average ripple and its peak coincides with the ripple's negative peak, indicating that on average they fire slightly before the interneurons. As mentioned earlier, similar pyramidal and interneuronal phase lockings between spike activity and field ripples have been established through recordings [34, 53, 241].

The membrane potentials of all recorded cells during the detected ripples were also isolated, highpass filtered (40-500 Hz) to remove any slow components, and aligned with the minimum point of their corresponding ripple as reference. The average potentials of pyramidal cells and interneurons are plotted with the average field ripple in Figure 6.8D. It appears that the average pyramidal cell membrane potential (solid line) oscillates in phase with the interneuronal one (dashed line) during the ripples and they are both correlated with the interneuronal spike histogram. Consequently, they exhibit a small phase precession relative to the average field oscillation, with their peaks coinciding with the rising phase of the ripple, similar to intracellular potentials observed in anaesthetised rats [241].

The histograms in Figure 6.8 support the idea that the ripple oscillation is actually a result of the interneuronal activity since only very few pyramidal cells, those receiving stronger input than average, manage to fire during a ripple episode. Figure 6.9A shows the percentage of the recorded pyramidal cells (solid line) and interneurons (dashed line) that fire on every ripple. Nearly all interneurons fire on every cycle, while the percentage

of pyramidal cells that fire during an episode is 21.9% on average. Figure 6.9B shows that the percentage of ripples during which a pyramidal cell fires increases with the cell's excitability, quantified by the number of synapses it receives per connection. An almost linear relationship arises close to our implemented threshold between the two cell groups (solid line).




Figure 6.9: **A:** Percentage of pyramidal cells (solid line) and interneurons (dashed line) that fire on each detected ripple. **B:** Percentage of ripples during which each pyramidal cell fires, versus the cell's excitability (as number of synapses it receives per CA3 incoming connection). The vertical line represents the threshold (19.5 synapses per CA3 incoming connection) between the two pyramidal cell groups.

The activity of an average cell during a ripple was studied by examining the synaptic input it receives and its resulting firing rate. Figure 6.10A depicts the average postsynaptic current from all types of synapses received by a cell during a ripple, averaged over all detected ripples and all relevant cells. All currents start to rise early during the ripple and peak along with the peak of the excitatory input from CA3 received by the average pyramidal cell (solid line) or interneuron (dashed line). The inhibitory currents in both pyramidal cells and interneurons exhibit a ripple oscillation phase locked to the average ripple with the same phase delay as seen in Figure 6.8D. In contrast, excitatory inputs are much more smooth and they appear to be in close balance with inhibitory inputs in both the average pyramidal cell and interneuron. Note that the average excitatory current an interneuron receives during a ripple, from both CA3 inputs and CA1 pyramidal cells, is ~ 3.7 nA. With a steady depolarising current of this amplitude, an isolated non-inhibited cell would fire with 400 Hz frequency [228]. The corresponding frequency for the purely inhibitory network would be 180 Hz (Figure 5.8C), which is

roughly in the middle of the ripple range.

The average firing rates in the recorded site are depicted in Figure 6.10B for both pyramidal cells (solid line) and interneurons (dashed line). They were calculated using 5 msec non-overlapping bins. Interneuronal firing increases slightly earlier and is sustained longer than the pyramidal, while both rates peak at roughly 5 msec before the ripple power peak (0 msec), in agreement with firing rates in [53, 54]. Most interestingly, the interneuronal rate barely reaches the ripple frequency range, indicating that the ripple oscillation is an interneuronal population oscillation while the average interneuron fires at lower frequencies.

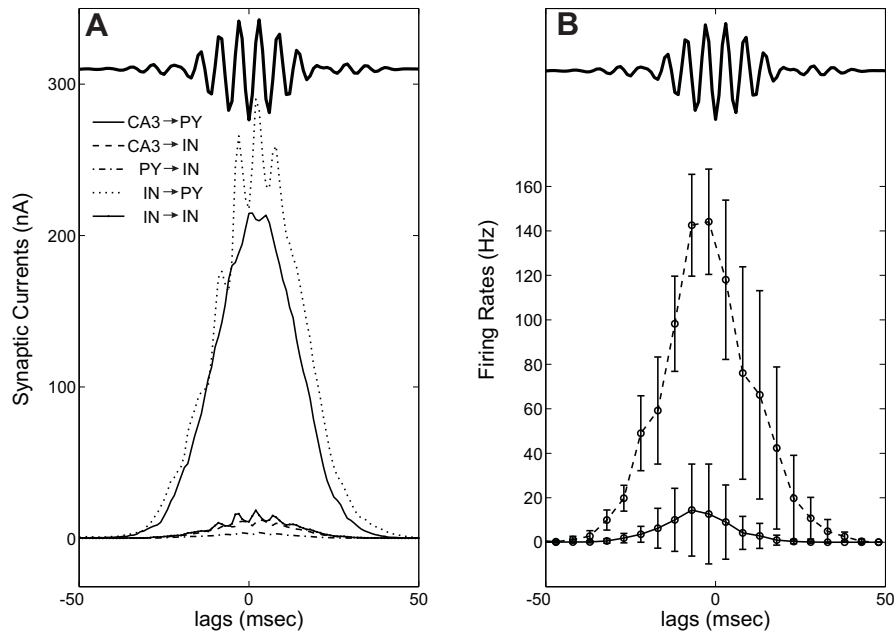


Figure 6.10: **A:** Postsynaptic currents received by the average pyramidal cell or interneuron during an average ripple (thick line). The ripple power correlates with the rise of excitatory input a CA1 cell receives from CA3. Note the close balance between the excitation and inhibition a CA1 pyramidal cell receives during the ripple (solid and dotted lines respectively). The same balance holds for the average interneuron as well. **B:** Average firing rates of pyramidal cells (solid line) and interneurons (dashed line) in relation to the average ripple (thick line), calculated over 5 msec bins. Error bars correspond to the SD. The interneuronal rate is below the ripple frequency and both rates peak roughly 5 msec before the ripple's power peak (0 msec).

Another major characteristic of SWRs is that they appear to be highly synchronous throughout the hippocampus [41, 241]. To check whether the ripples in our model are close to synchronous, we simultaneously recorded the total synaptic conductance from

three more sites of the same length as the default ($560\ \mu\text{m}$). The respective centers of two of them lay $300\ \mu\text{m}$ away from that of the default recording site, and the center of the third site is $500\ \mu\text{m}$ away. Figure 6.11A displays one example of the same ripple episode recorded from all four sites. The filtered total synaptic conductances appear to oscillate nearly synchronously throughout all four sites, with very small jitter, as was also seen for similar space scales in [241]. This synchrony is disrupted for larger distances of the order of several mm (not shown). We computed the average ripple on each site during the same 30 sec simulation, isolating the bandpass filtered data fragments from all sites that coincide with each detected ripple in the default site and aligning them with its minimum value. The averages from each site plotted along with the average ripple on the default site (Figure 6.11B) appear to be coherent, with zero time lag between them. Therefore, averaging over a large number of events abolishes the jitter seen in individual ripples, again in agreement with [241].




Figure 6.11: Long range synchrony of modelled ripple events in CA1. **A:** One ripple recorded simultaneously from four different sites: the default recording site (solid line), two sites lying $300\ \mu\text{m}$ away from it (dashed and dashed dotted lines), and one lying $500\ \mu\text{m}$ away (dotted line). **B:** Similar for the average ripple calculated from each site.

To check that CA1 ripples are not a direct consequence of a similar oscillation in CA3, we examined the correlation between the CA1 average ripple and the CA3 pyramidal cell firing. If the CA1 ripples are just a CA3 rhythm that travels through the Schaffer collaterals to CA1, a correlation between the CA3 pyramidal spike histogram and the average CA1 ripple, similar to that in Figure 6.8B should exist, with some phase difference due to the conductance delay of the Schaffer collaterals. Figure 6.12A shows that

this is not the case. In fact, there appears to be no ripple-like modularity in the firing of pyramidal cells in CA3, to correlate with CA1 activity. Note that we have taken all CA3 pyramidal cells into account here since the Schaffer collaterals are so widespread that a CA1 cell receives input from the majority of CA3 cells. Nevertheless similar results are acquired even if we take only the cells in the CA3 part of the array that corresponds to the CA1 recording site (not shown).




Figure 6.12: **A:** CA1 average ripple correlation with CA3 pyramidal cell firing histogram. Spikes from all CA3 pyramidal cells were included in the histogram. **B:** Mean firing rate of all CA1 pyramidal cells during each CA3 burst, plotted over the percentage of CA3 pyramidal cells that fired during the burst. An almost linear relation arises between the magnitude of a CA3 burst and the frequency of the CA1 response. The solid line is a least squares fit with a slope of 0.145.

The varying frequency ranges of the power peaks in the spectrogram of Figure 6.7A imply that the ~ 150 -200 Hz range is not intrinsic in the CA1 model but depends on the input from CA3. This was verified by examining the relationship between the magnitude of CA3 bursts and the magnitude of CA1 responses. CA3 bursts were detected by a threshold equal to the SD of the RMS of the total synaptic currents recorded in the CA3 equivalent of the CA1 default site. The threshold for setting the burst limits was half the SD. The burst magnitude was quantified by the percentage of CA3 pyramidal cells that fired during the burst while the magnitude of CA1 responses was measured as the average firing rate of all pyramidal cells in CA1 during the burst. Plotting the CA1 responses over the extent of the corresponding CA3 bursts (Figure 6.12B) reveals an overall monotonic relation, resembling qualitatively a corresponding relation seen in

[54]. Similar results were found for the average firing rate of CA1 interneurons (not shown).

6.5.2 The Role of Inhibition and Excitation in Ripples

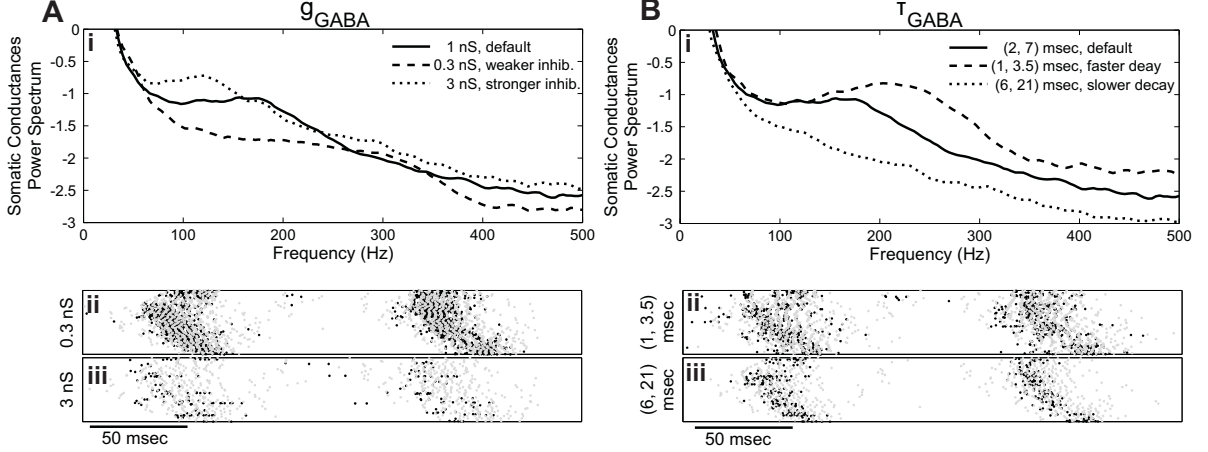


Figure 6.13: Comparison of CA1 somatic conductances power spectra between the default model and cases where the following synaptic parameters have been modified: **A:** the maximum unitary GABA conductance \bar{g}_{GABA} , **B:** the decay time of inhibitory synaptic variables τ_{GABA} . Below each spectral profile, examples of CA1 raster plots from the two non-default cases are plotted and can be compared with Figure 6.6A.i. The parameter values in each case are given in the figure. Changes in the strength of inhibitory synapses alter mainly the oscillation frequency, while prolonging the decay time of inhibition abolishes the spectral peak implying a disruption of synchronous oscillations in any particular frequency. Qualitatively similar results hold for the ripple range oscillations of the purely interneuronal network of CA1.

The role of inhibition and excitation in CA1 on the generation and synchronization of ripples was studied by modifying the parameters of inhibitory and excitatory interactions. Parameters in CA3 were unaltered to retain the same input to CA1 as before.

We altered the strength of all inhibitory connections in CA1 by first reducing the maximum synaptic conductance \bar{g}_{GABA} in eq. (5.4.1) by 70%, setting $\bar{g}_{\text{GABA}} = 0.3$ nS, and then increasing it to $\bar{g}_{\text{GABA}} = 3$ nS. The same was done in another pair of simulations for excitatory connections, setting \bar{g}_{AMPA} to the previous two values respectively. We also examined the effect of the synaptic decay times by decreasing and increasing τ_{GABA} and τ_{AMPA} separately. Since τ_{GABA} is different for pyramidal cells and interneurons (7 msec and 2 msec respectively) both values were altered accordingly. Each case was simulated for 30 sec. Figure 6.13 presents the power spectrum of the somatic conductances for

each case, along with the default value case for comparison, and excerpts of CA1 raster plots for the non-default cases.

The effect of altering these parameters is partly reflected in the way the high-frequency peak (150-200 Hz) of the default spectral profile changes. For weaker inhibition the peak widens and shifts to higher frequencies (~ 300 Hz) indicating that there is still some activity at ripple frequencies but on average the oscillations are much faster (Figure 6.13A). The corresponding raster plot indicates that interneurons now fire at higher rates along with an increased pyramidal activity. On the contrary, stronger inhibition results in shifting a narrower peak to lower frequencies and the raster plot clearly shows weak interneuronal and pyramidal activity.

Changing the inhibitory decay time has a different effect (Figure 6.13B). For faster inhibition decay, spectral peaks are sustained but lie at higher frequencies (~ 250 Hz) as interneurons can now fire at much higher rates. Slower inhibition though does not shift the peak but instead just attenuates it, indicating the loss of coherent oscillations in any particular frequency range.

Changing the excitation strength affects the input all neurons receive from CA3 and the extra excitation of interneurons from nearby pyramidal cells. It has the opposite effect to that seen for \bar{g}_{GABA} changes, with weaker (stronger) excitation decreasing (increasing) the oscillation. Similar results hold for the decay time with faster decay leading to slower oscillations and slower decay to faster ones (not shown).

In summary, changing the level of inhibition or excitation in the network, or even the decay time of excitation, alters mainly the oscillation frequency. On the other hand, prolonging the decay time of inhibition disrupts the population synchrony. Note that the same results hold for the ripple frequency oscillations of the purely interneuronal network of CA1, described in the previous chapter (Figure 5.8B), another indication of the relatively small contribution of pyramidal firing in ripple generation.

6.6 Discussion

The two network models described in the previous chapter were coupled together in a Schaffer collateral-like manner with feedforward excitation from CA3 to CA1. We have shown that quasi-synchronous population bursts in the CA3 model evoke responses in the CA1 model that exhibit numerous characteristic features observed in real SWRs. Below, we discuss these characteristics, in relation with the architecture and the features

of the model, along with the underlying mechanisms controlling them and their potential functional implications.

6.6.1 Model Architecture

Although a large variety of inhibitory cell types exists in the hippocampus [56, 76], we only implemented fast somatic inhibition in the models. In addition to the obvious model simplification that this allows, this choice was based on the observation that, during SWRs, different interneuronal types exhibit diverse behaviours, with only basket and bistratified cells increasing their discharges, while other interneurons stop spiking [120, 121]. Our model suggests that the increased activity of basket cells suffices to model many basic ripple characteristics. The mechanisms that generate such an unequal participation of interneuronal classes in ripples are beyond the scope of this work.

Only fast synaptic interactions were implemented (AMPA and GABA_A) since pharmacological disruption of NMDA or GABA_B receptors in CA1 did not affect SWRs [136], implying that they are not involved in ripple generation. On the other hand application of the AMPA-receptor antagonist CNQX completely abolished ripples [19, 136, 238]. Our results support this observation since fast synaptic time scales appear to be key for ripple generation (Figure 6.13B).

6.6.2 SWRs in the Full CA3-CA1 Model

CA1 responded to the strong input from CA3 with intense synaptic activity on fast time scales involving both excitatory and inhibitory populations (Figure 6.6A-B), with characteristics very similar to SWRs from extracellular recordings in CA1 [34, 53, 120, 136, 241]. Note that the SWs in the model do not show a small inhibitory component after the negative peak [32, 34]. Perhaps this afterhyperpolarisation component is caused by dendritic inhibition, which is omitted in the model. It could also be due to somatic inhibition captured by dendritic-layer recordings. Spectral analysis of synaptic activity in the CA1 somatic layer revealed a clear peak around 150-200 Hz (Figure 6.7), similar to spectral profiles in [24, 34, 136], verifying that on average the CA1 somatic fast oscillations were indeed within the ripple frequency range.

Interestingly, the high frequency of the ripple oscillations was not reflected in the spiking activity of individual neurons. Membrane potentials of pyramidal cells showed that they fired extremely rarely during a ripple, and mostly with a single spike, while interneurons

were much more active, firing in every episode with a large number of spikes (Figure 6.6C-D). These features are characteristic of CA1 recordings during ripple activity [34, 53, 241] and hint at the interneurons as the driving force in ripple oscillations. Indeed, average firing rates during ripple episodes in the model showed that interneurons increase their spiking slightly earlier and their firing lasts longer than pyramidal cells (Figure 6.10B), seen also in [53, 54]. This conclusion is also supported by the clear modularity and phase locking of the interneuronal firing histogram to the negative half wave of the average ripple (Figure 6.8C). Nevertheless, firing rates were on average below the 150-200 Hz band, implying that ripples are an overall population oscillation, not an effect of the average firing of individual interneurons. In short, the external drive from CA3 produces synchronous ripple-frequency oscillations of the CA1 interneuronal network in the model, similar to those in Figure 5.8B, only now cells can skip cycles due to heterogeneities in the input and recurrent inhibition.

Pyramidal cells, embedded in the strong background interneuronal oscillation, spike much more rarely, mostly right before the interneurons, resulting in a phase locking of their spiking to the negative peak of the ripple on the onset of the inhibitory barrage. As aforementioned, these phase correlations are also established via intra- and extracellular recordings [34, 53, 120, 241]. In the model, only the cells receiving stronger input from CA3 can synchronize their firing with the overall field oscillation. The rest get too strongly inhibited to show rhythmical activity and fire randomly mainly in the beginning of the ripple episode, before the peak of interneuronal firing (Figures 6.8B and 6.9B). Therefore, the excitatory spikes come mostly from a minority of pyramidal cells, which fire on many more ripples than the rest. Specifically, only about 22% of pyramidal cells fired on average during a ripple (Figure 6.9A). This feature agrees with extracellular recordings where approximately only 10% of all recorded pyramidal cells participated in an average ripple with some cells participating in up to 40% of successive episodes and others firing very rarely [241].

At the peak of the CA3 input, interneuronal activity reaches its maximum intensity, reducing further both pyramidal and inhibitory spikes. This is reflected in the average firing rates which peak approximately 5 msec before the ripple peak (Figure 6.10B), again in agreement with a similar time difference (5-10 msec) seen in average rates from extracellular recordings [J. Csicsvari, personal communication].

Average membrane potentials of both cell types appear to be in phase with the interneuronal spike histogram, exhibiting a similar small phase shift compared to the average

extracellular ripple, with their peaks coinciding with the rising portions of the ripple waves (Figure 6.8D). This phase difference is also seen in intracellular recordings from anaesthetised rats [241], and is probably due to the slow interneuronal conduction velocities. Inhibitory postsynaptic conductances (which constitute the vast majority of postsynaptic activity within CA1, Figure 6.10A) appear to lag behind spikes and membrane potentials by ~ 2 msec. The time lag corresponds to a distance of $200\text{ }\mu\text{m}$ or $2\sigma_{IN}$ of inhibitory synaptic inputs from interneuronal spikes, which would imply that the delay is mostly due to the time it takes for inhibition to reach cells lying further than distance σ_{IN} from an average interneuron. The delay indicates the importance of using the synaptic conductance as a measure of the extracellular synaptic activity, particularly in fine timescale studies. Processing synaptic currents instead, which contain the membrane potential in their definition (eq. 5.4.1), will result in an average ripple that is slightly shifted backward in relation to the spikes. This will artificially produce a similar phase difference between the field oscillation and the unit firing (with pyramidal cells firing during the rising phase of the ripple and interneurons firing at its peaks), which contradicts the aforementioned phase lockings seen in recordings.

The fact that the fast oscillations were generated within CA1, and were not contained within the CA3 input, was shown by the lack of any correlation between CA3 firing and CA1 oscillations (Figure 6.12A). According to our simulations, the actual frequency range of the ripples is not an intrinsic property of the CA1 circuit architecture but a direct result of the magnitude of the excitatory input from CA3. Due to the variability in the magnitude of CA3 bursts, the frequency range also varies, with a clear relationship between the number of CA3 pyramidal cells firing during a burst episode and the mean firing rate of CA1 pyramidal cells during the evoked response (Figure 6.12B). A similar relationship between these two measures was reported through extracellular recordings [54] and implies that the more extensive and synchronous a CA3 burst is, the higher the frequency of the field oscillation. This is also supported by the shift in the spectral ripple-peak when excitation parameters are altered, with stronger or more slowly decaying EPSPs leading to higher frequency oscillations (not shown). According to this dependence of the ripple frequency on the magnitude of the Schaffer collateral input, the model predicts that an *in vitro* partial cut of Schaffer collateral axons in a longitudinal CA3-CA1 slice would decrease the observed ripple oscillation frequency.

Finally, note that SWRs have also been reported in CA3 slices *in vitro* [68, 136]. In our CA3 model, although the population bursts may account for SWs, ripples are absent.

This is probably due to the lack of recurrent inhibition in the model. The inclusion of CA1-like connections between CA3 interneurons with an anatomically realistic connectivity scheme, in a future version of the model, could synchronize inhibitory firing and potentially generate ripple oscillations in CA3 as well. Moreover, we expect that long range excitation is reduced in CA3 slices making the connectivity more similar to CA1, which may help explain the generation of ripple-like oscillations in the slices.

6.6.3 Pharmacological Effects on Ripples

An important factor in the mechanism of ripple generation appears to be the relation between excitation and inhibition in the CA1 network. The postsynaptic currents generated in the average cell during ripples reveal a tight balance between smooth excitatory input from CA3 and fast-oscillating inhibition (Figure 6.10A). The average interneuron receives enough excitation to fire at frequencies twice as fast as ripples, but cross-inhibition retains the interneuronal network within ripple frequency range. A disruption of this sensitive balance has a corresponding effect on the ensuing oscillation frequency, shifting the spectral peak to different frequency ranges (Figure 6.13).

This balance disruption can be induced pharmacologically by various anaesthetics. Ketamine for example has no significant effect on AMPA excitation (only depresses NMDA receptors which are not relevant to SWRs) but enhances GABA_A inhibition [124, 171]. According to Figure 6.13A, this inhibition enhancement would shift the oscillations to lower frequencies. A similar inhibition enhancement in CA3 would further diminish the excitatory input to CA1, reducing the frequency further. This shift is in agreement with the observation that population oscillations dropped to 90-150 Hz, in *in vivo* extracellular recordings on ketamine-anaesthetised rats [241].

The opposite shift, to higher frequencies, occurs when the inhibition-to-excitation ratio is decreased, although ripple-range frequencies still have significant power, even when IPSP amplitudes are reduced by 70% (Figure 6.13A). This is supported by recordings on CA1 minislices *in vitro* under the GABA_A receptor antagonist gabazine, where ~200 Hz ripple oscillations persisted to some degree [153]. The reduction of IPSP amplitudes under 0.3 μ M of gabazine was roughly 70%, similar to our corresponding simulation with $\bar{g}_{\text{GABA}} = 0.3$ nS. It is noteworthy that, in CA3-CA1 slices administered with bicuculline and gabazine, abolishing GABA_A inhibition, ripples were replaced by large epileptiform bursts that still carried a ~200 Hz oscillation component [136].

Moreover, fast synaptic time scales are also crucial, with faster decay times leading to faster oscillations, while with slower IPSP decay times, the oscillations do not become slower but rather get attenuated overall (Figure 6.13B). The importance of fast decay inhibition could be linked with reported effects on ripples of the gap junction-blocker anaesthetic halothane. Halothane has been shown to disrupt CA1 ripple activity in *in vivo* recordings of anaesthetised rats [241]. Under deep halothane-induced anaesthesia, irregular SW activity was maintained but 200 Hz oscillations were almost completely abolished. Other gap junction blockers (carbenoxolone and octanol) had similar effects in CA1 minislices *in vitro*, reducing the frequency of spontaneous SW occurrence and even more so that of ripple activity [136]. This has led to the suggestion that gap junctions are a key factor in the SWR mechanism, since, as described in section 6.3, they can yield fast and synchronous oscillations. Nevertheless, gap junction blockers are quite non-specific, affecting chemical synapses as well. Specifically, halothane has a combined effect on both chemical excitation and inhibition [154], depressing glutamate-mediated synaptic EPSP responses in both pyramidal cells and interneurons [119, 162, 163], in combination with an approximately 2.5-fold prolongation of GABA_A IPSP decay time [78, 154]. A similar prolongation of τ_{GABA} in our model also resulted in severe attenuation of the ripple spectral peak indicating the loss of any particular high frequency oscillation (Figure 6.13B). The effect would be even stronger with the combined reduced excitation in CA3, shifting oscillations to even lower frequencies. Thus the model suggests that halothane effects on chemical synapses alone could explain the loss of ripple-like oscillations.

6.6.4 SWRs and Memory Consolidation

Another very important feature of SWRs is that they exhibit long range synchrony over the whole longitudinal extent of both CA areas [41, 241]. In fact SWRs extend to the rat subicular complex and even the enthorinal cortex, involving the synchronised participation of tens of thousands of cells [41]. Individual ripples, recorded simultaneously with multi-site probes 300 μm apart along the longitudinal CA1 axis, show a very small jitter between them (~ 1 msec) but on average their time lag drops to zero [241]. Synchrony over such space scales of hundreds of μm is successfully reproduced by our model (Figure 6.11). Apart from the wide distribution of the Schaffer collaterals, which induce a homogeneity in the input from CA3, this synchrony appears to be induced mainly by the recurrent inhibition in CA1 which locks the inhibitory spikes in global oscillations

and is enhanced by the interneuronal phase resetting caused by sparse firing of nearby pyramidal cells, as was also suggested in [241]. Nevertheless, synchrony is disrupted over larger space scales of several mm (not shown).

The strong transient output that is produced during such massive population discharges is very likely to affect cortical targets, making SWRs a possible vehicle for transferring neuronal information from the hippocampus to the neocortex for long-term storage, during deep stages of sleep. Various studies of the correlations between hippocampal ripples and slow oscillations in the cortex during deep sleep provide further support for the role of ripples in memory [17, 104, 185, 188]. This subject is studied in detail in the following chapter.

One basic assumption of the model is that interneurons receive direct excitation by Schaffer collaterals. The model can produce similar SWRs without such an input to interneurons, with all excitation coming from their neighbouring pyramidal cells. Nevertheless, this set up would abolish the particular group of pyramidal cells that discharge during ripples and would result in all pyramidal cells firing in every cycle of the ripple before evoking the subsequent interneuronal firing. This contradicts the observation of very sparse pyramidal firing [53, 54, 241] and interneuronal discharge rising before the pyramidal [53, 54]. Most importantly, it would contradict the fact that neuronal participation in ripples is not random, with only a particular subset of pyramidal cells discharging during most ripple episodes [241].

The heterogeneity in the response of CA1 pyramidal cells to the CA3 input was implemented through a variability in the number of synapses Schaffer collaterals make on each cell. The small number of cells that fire during ripples receive more synapses per incoming connection than average, and can be thought of as representing neurons in the rat CA1 that are targeted by a larger portion of CA3 or receive stronger EPSPs than average. Although such a highly excitable subset is necessary in the model for reproducing the selective participation reported in [241], its existence in the rat CA1 area has not yet been established. Nevertheless, we believe it supports the memory consolidation scenario since the cell subset discharging in experimental ripples appears to be formed by pre-existing experiences. Specifically, correlated discharges of rat hippocampal pyramidal cells during sleep have been shown to reflect the correlated activity of the same cells during earlier spatial exploration [235]. In fact the actual temporal sequences of those cells during wakeful training are preserved during sleep [189] but replayed at faster timescales during the short window of SWRs, in the same order [129]

or reversed [62, 74], (for an extensive discussion of the possible role of SWRs in memory consolidation see [33]).

In support of these observations, the subset of participating neurons in the model, instead of being implemented *a priori*, could emerge intrinsically in CA1 through a combination of repetitive firing of highly correlated CA3 cells and synaptic plasticity of the Schaffer collateral outputs. By implementing ensembles of strongly connected neurons in CA3, representing correlated place cells, the spontaneous CA3 population bursts could trigger their sequential firing and in combination with synaptic plasticity could potentiate synapses with their targeted CA1 cells. These CA1 cells would eventually form the strongly-driven subset, firing during ripples. A similar enhancement of single CA1 pyramidal cells' participation in ripples was found after their long-term potentiating stimulation [118]. The replay of temporal patterns of the correlated CA3 ensemble during a burst would perhaps form a similar temporal pattern in CA1 and the whole process could take place either in a forward or reversed direction, shedding light to the bidirectional replay mechanism of place cell patterns observed during SWRs [62]. Changing the correlated CA3 ensembles would eventually change the CA1 subset. Such a scheme would provide computational support to the observed SWR generation in CA1 after long-term potentiation inducing stimulation of CA3 [19] and even the place selective firing of CA1 cells in SWRs during exploration [155]. Coupling of such an extended model with a cortical network model in a feedback loop could help the study of temporal pattern transfer to the neocortex for long-term storage and the correlations between SWRs and cortical slow oscillations [104, 185, 188]. On the other hand, the model suggests that, in a situation where such a subset in CA1 is abolished, pyramidal spiking during ripples would be almost terminated, without severely affecting any other ripple characteristics (frequency, synchrony etc), but potentially affecting the memory consolidation process to the cortex through ripples.

6.6.5 Gap Junctions

As described earlier in this chapter, models based on sparse axo-axonal gap junction connectivity give rise to fast 200 Hz population oscillations even in the absence of chemical synapses [211, 217]. Nevertheless, the lack of SWR generation from antidromic stimulation of CA1 pyramidal axons and their preservation after removal of the distal part of these axons [24] suggests that axo-axonal gap junctions may not be necessary for SWRs.

Our model offers a different approach to the ripple generation mechanism that relies on inhibitory activity and solely on chemical synaptic interactions, avoiding the need for gap junctions. Since the model reproduces a large number of basic SWR characteristics and offers an interpretation of such characteristics, we propose that gap junctions are not a necessary driving component for the SWRs but rather an additional mechanism that helps the overall ripple generation. Such an “aiding” role of axo-axonal gap junctions is also hinted by the fact that SWRs were not eliminated in mice lacking gap junction protein Cx36, only their frequency of occurrence decreased [135]. Most importantly, gap junctions between pyramidal cell axons could serve as the mechanism for long range fast synchrony of the ripples over spatial scales of many mm. The future inclusion of an underlying gap junction network between pyramidal cells in our CA1 model would help the study of such a long-range synchronizing mechanism.

6.6.6 A Proposed Mechanism for SWRs

In summary, our simulation results, combined with the discussed neurophysiological findings, lead us to formulate a mechanism for the generation of SWRs that is based purely on chemical synaptic interactions between excitatory and inhibitory populations. The proposed mechanism can be summarised by the following steps:

- CA3 quasi-synchronous population bursts produce a strong depolarising input to both pyramidal cells and interneurons in CA1. The large number of outgoing connections from each CA3 pyramidal cell and their wide spatial extent helps render this input more homogeneous throughout CA1.
- Strong excitatory AMPA-synaptic currents generated in the dendrites of both cell types during the CA3 input produce an SW in the dendritic layer of CA1.
- Fast spiking interneurons, such as the basket cells modelled here, get directly depolarised by the CA3 input, enough to intrinsically start spiking at a wide range of high frequencies ($\sim 100\text{-}400$ Hz).
- The strong recurrent inhibitory connectivity in CA1, combined with the fast timescales of IPSPs, confines interneurons to a frequency range of about 150-200 Hz, which varies with the extent of the CA3 burst, and helps synchronize their membrane oscillations within a range of hundreds of μm . Recurrent inhibition also causes interneurons to skip cycles within the ripple resulting in a firing

rate which is lower than the average membrane potential oscillation frequency of interneurons.

- Pyramidal cells have a much more passive role during the ripple episode. The strong inhibitory barrage they receive during the synchronous interneuronal firing cancels on average the depolarisation from CA3. As a result they rarely manage to fire. Their sparse spikes appear mostly just before the interneuronal barrage and increase even further the depolarisation of neighbouring interneurons, helping enhance the population synchrony.
- A minority of pyramidal cells, which receives more CA3 input than average, manages to overcome the vast inhibition and produces the majority of pyramidal spikes in most events.
- The termination of the CA3 burst results in the consequent termination of CA1 activity. Interneurons continue to fire for a short amount of time due to delayed input from distant CA3 pyramidal cells.

In the following chapter we couple our SWR-exhibiting CA3-CA1 model and the cortical model described in Chapter 4, implementing an anatomically realistic connectivity, in an effort to examine the temporal and functional interactions between SWRs and UP/DOWN states and uncover some key parameters on which these interactions depend.

A Model of Cortico-Hippocampal Interactions

7.1 Introduction

The standard model for memory consolidation assumes that new memories are stored temporarily in the hippocampal area and in closely connected medial temporal lobe regions, whereas remote memories are stored in the neocortex [33]. The transferring of memories from hippocampus to cortex for long-term storage is generally thought to take place during slow wave sleep (SWS). The previously discussed observations of firing pattern temporal replay in the hippocampus during subsequent sleep sessions, along with recent reports of similar replays in the prefrontal cortex [111, 140, 164], support the notion of memory consolidation taking place during SWS (for a discussion on the details of such a procedure see [33]).

One related issue is how the hippocampus communicates with the neocortex during SWS and how information flows between the two structures. An influential suggestion has been that information travels from cortex to hippocampus during wakefulness and in the opposite direction during SWS [35]. As discussed in the previous chapter, a series of recent observations indicate that hippocampal ripples could be the vehicle for information transfer to the cortex. This is supported by various studies on the correlations between hippocampal ripples and slow oscillations in the cortex during SWS. These studies are reviewed here.

In Chapter 3, we examined cortico-hippocampal connectivity under anaesthesia, which shares many characteristics with SWS. Our results hinted that during short cortical UP-

like states information is sent to the hippocampus from mPFC and that the process is reversed during longer cortical DOWN states. Here we tackle this issue from a modelling perspective. We combine the cortical slow-oscillation (SO) model, described in Chapter 4, with the ripple-producing CA3-CA1 network, described in Chapters 5 and 6, in an effort to validate the ripple-SO correlation studies, reviewed below, and uncover the important parameters and mechanisms that yield these correlations.

7.2 Slow Oscillation-Ripple Correlations in Electrophysiological Recordings

A lot of research on memory consolidation during SWS has recently been focused on the temporal relationships between hippocampal SWRs and rhythmic cortical events during SO, such as spindles and delta waves (see Chapter 4). Since these events are also present under anaesthesia, the analysis has been extended in many cases to anaesthetised animals as well.

Correlations between cortical spindles and hippocampal ripples during SO, within temporal windows of ~ 1 -2 sec, were first reported by Siapas and Wilson [185] who analysed LFPs and multiunit activity from the mPFC and CA1 of sleeping rats. Correlation asymmetries further implied a tendency for ripples to precede spindles and general hippocampal firing to precede prefrontal firing.

Sirota et al. [188] also observed strong neocortical and hippocampal activity coupling during SO, in LFPs and unit activity from the somatosensory cortex and CA1 of rats and mice. Ripples were found to be temporally associated with a power increase of spindles and delta waves in the cortex, again in a time scale of ~ 1 -2 sec. This co-occurrence was modulated by SO, with more events happening during the UP state and less in the DOWN state. A similar spiking correlation between CA1 and cortex was found, with cortical activity preceding hippocampal by ~ 50 msec. Based on the phase locking of CA1 spikes to ripple troughs [53, 241] and cortical spikes to troughs of spindles and delta waves, they argued that the maximum probability of ripple troughs occurred ~ 50 msec after the troughs of spindles and delta waves. In other words, the hippocampal output becomes temporally “sandwiched” between the cyclic discharges of spindle-activated neocortical neurons (Figure 7.1).

Cross-correlation analysis by Mölle et al. [147, 149], in a similar set of recordings, supported the temporal coupling of hippocampal ripples and cortical spindles to SO. They

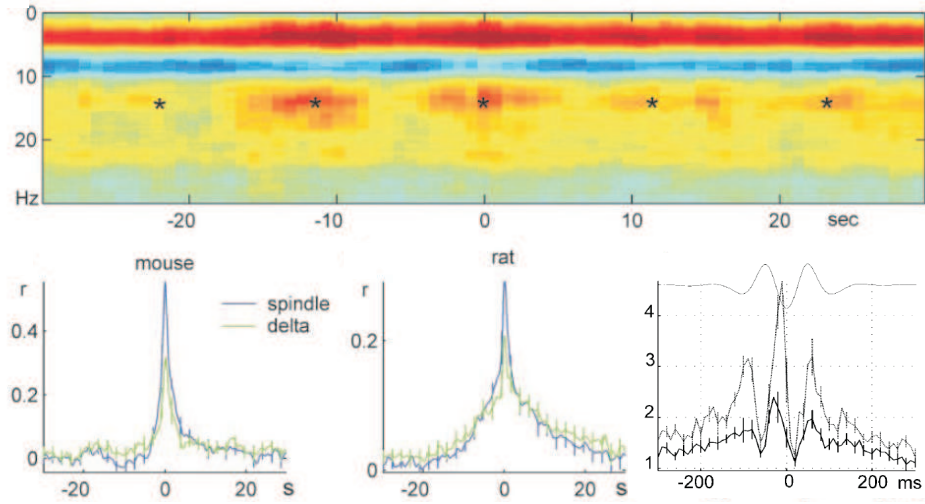


Figure 7.1: Slow time scale correlation between neocortical and hippocampal patterns during slow wave sleep. Top: Hippocampal ripple-triggered neocortical spectrogram in a mouse. Power spectrograms, centered on ripples (time 0 s), were averaged and normalised by the mean power over the entire recording session. Note increased correlation of power in the delta (1-4 Hz) and sleep spindle (10-20 Hz) bands with hippocampal ripples. Asterisks indicate slow comodulation of neocortical and hippocampal activity. Below from left: Group averages of cross-correlograms of ripple power (time 0 s) to spindle and delta power (blue and green lines, respectively) in mice and rats. Cross-correlogram between troughs of sleep spindles and CA1 hippocampal unit activity (solid line, mean \pm SD) and between troughs of sleep spindles and ripple troughs (dotted). Modified from [188].

also detected weak ripple-spindle correlations with spindle activity increasing 200 msec before ripples and lasting around 2 sec more. They concluded that UP/DOWN states, defined by spiking activity, precede ripples by ~ 30 msec, in accordance with Sirota et al. [188].

Battaglia et al. [17] examined peri-event time histograms around SWRs in CA1 and also found them temporally correlated with spindles and delta waves in the cortex. They further reported an increased SWR probability of occurrence during the DOWN-to-UP transition. Surprisingly, they observed more SWRs during the DOWN states than UP. Nevertheless, they defined the UP/DOWN states, not based on LFP/EEG recordings (as usually done) but based on spike activity (which actually did not exhibit a clear UP/DOWN bimodality).

Perhaps the most broad and detailed study on this subject was made by Isomura et al. [104] who performed extra- and intra-cellular recordings on the neocortex, entorhinal

cortex and hippocampus of both anaesthetised and naturally sleeping rats. The cortical SO was shown to extend to all neocortical areas, engaging cells even at the entorhinal cortex which followed the oscillation with a delay of ~ 150 -200 msec. Hippocampal cells did not exhibit a bimodal membrane potential distribution but their spiking activity was modulated by the cortical SO (Figure 7.2A). Specifically, neurons in the dentate region were shown to participate in the SO, firing mostly in the UP state, but with less clear UP-DOWN state bimodality, while most CA3 neurons were preferentially active during the DOWN state firing at gamma frequencies (Figure 7.2B). In CA1 cells, three distinct types of behavior were described. Most cells fired with higher probability in the UP state, yielding a correlation of the overall CA1 activity with the SO (that was described as a hippocampal SO by Wolansky et al. [237]). Few CA1 cells fired preferentially in the DOWN state and some in both. Finally, most ripples occurred in the UP state ~ 200 msec after the neocortical DOWN-UP transition, with less ripples at the onset of the DOWN state (Figure 7.2C-G).

Hahn et al. [91] also observed SO modulation in the dentate gyrus. In contrast with the observations of Isomura et al. [104], some CA3 pyramidal cell membrane potentials were correlated with the UP state, while others were anticorrelated. Most CA1 pyramidal cells were anticorrelated with the SO, being hyperpolarised in the UP state. However, the fact that these recordings come from anaesthetised animals rather than naturally asleep may have a substantial effect. Moreover, whole-cell recordings from CA1 R-LM interneurons (type of hippocampal interneurons lying at the stratum lacunosum moleculare [76]), revealed a weak SO modulation of their membrane potentials, correlated with cortical LFP DOWN-UP transitions at ~ 150 msec delay [90].

All the aforementioned reports clearly show the grouping role of the cortical SO on hippocampal activity, but lack any evidence for feedback of information flow from the hippocampus to cortex, for memory consolidation [146, 209]. Wierzynski et al. [234] calculated cross-covariances between firing activity of cell pairs from the mPFC and CA1 and found that some mPFC cells tended to fire around 36 msec after hippocampal cells. These strong covariances of firing activity appeared during ripples and were abolished during REM sleep.

Although some of the above results are contradictory, various fundamental conclusions can be reached. One is the grouping role of the SO on cortical and hippocampal unit activity and on the probability of spindle and ripple generation [90, 104, 147, 149, 188]. These two rhythmic events appear to be strongly phase-modulated by the UP/DOWN

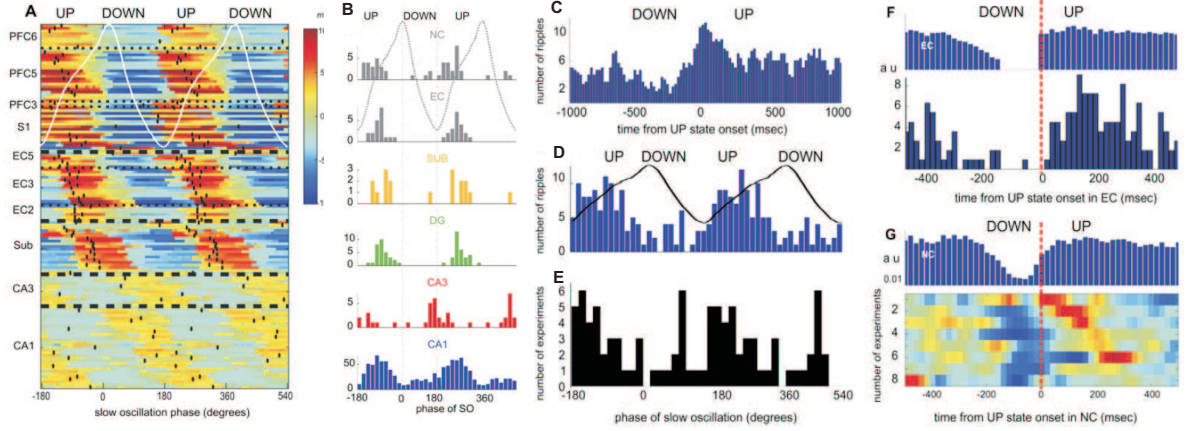


Figure 7.2: Relationship between cortical SO, hippocampal firing and SWRs. **A:** Group representation of the SO-related modulation of distribution of the membrane potential. Each line corresponds to a single intracellularly recorded neuron. Color code: magnitude of the most likely values of the membrane potential, beyond what is expected from independence, for each phase. White trace: averaged field trace from CA1 pyramidal layer. PFC, prefrontal cortex; S1, somatosensory cortex; EC, entorhinal cortex; Sub, subiculum. **B:** Distribution of preferred discharge phases of single neurons in various regions. Only neurons that were significantly modulated by slow oscillation are shown (643 cells from 130 animals). **C:** Temporal distribution of ripple occurrence relative to DOWN-UP transition (0 msec) of the membrane potential in entorhinal neurons. **D:** Distribution of ripple occurrences as a function of phase of neocortical SO in LFP in a representative experiment. **E:** Distribution of preferred phases of ripples across experiments with significant modulation of ripples occurrence by the SO. **F:** Relationship between DOWN-UP transition of population firing in the EC (top) and the occurrence of hippocampal ripples (bottom). **G:** Relationship between DOWN-UP transition of population firing in the neocortex (top) and the occurrence of hippocampal ripples (bottom). Each line is a color-coded cross-correlogram. Modified from [104].

states of the SO, with their probability of occurrence increasing during the cortical UP state and decreasing during the DOWN state. Thus the grouping effect is not limited only in the corticothalamic areas but extends to the hippocampus, although hippocampal cells do not exhibit a clear bimodal distribution in their firing rates. It appears that ripple events (and consequently hippocampal firing) generally follow the SO with a delay of the order of tens of milliseconds [90, 104, 149, 188]. Additionally, this modulation results in strong temporal correlations between the spindle and ripple events within time scales of 1-2 sec [17, 149, 185, 188]. These results support a general model of memory consolidation through a closed cortico-hippocampal loop that exists during SWS. Namely, the cortical SO exerts a feed-forward control over hippocampal activity,

and during its UP states (perhaps within spindles) it causes neuronal reactivation of newly encoded information in hippocampal circuitry. This “memory reactivation” is fed back to the cortical networks through ripples, gradually integrating the new memories to the cortex for long-term storage [111, 140].

7.3 Data Analysis on Slow Oscillation-Ripples Correlations

We analyse here a set of electrophysiological recordings from a naturally sleeping rat in an effort to reproduce some of the aforementioned observations.

LFPs (sampling rate 1 kHz) and multi-unit activity (40 kHz) were recorded from a freely-moving male Lister hooded rat over a total period of 60 min. The animal was previously implanted with two eight-channel micro-wire Teflon-coated stainless steel electrode arrays (NB Labs, USA), one array in the mPFC area and one in the dorsal hippocampus (recording details in [50]). Recording channels 1-8 correspond to mPFC LFPs and channels 9-16 to hippocampal LFPs, with channel 9 corresponding to the most lateral location (CA3) and channel 16 the most medial location (CA1) [Drs Rob Mason and Yvonne Mbaki, unpublished data]. All data segments that corresponded to awake behavior of the rat were removed by visual inspection of the LFPs and the concurrent video recording of the experiment. Figure 7.3A displays a segment of the data, recorded during a sleep session of the rat, with two of the mPFC and four of the hippocampal LFPs.

To isolate the data segments corresponding to SO cycles (pairs of DOWN-UP states) in the mPFC LFP (channel 2) we followed an algorithm resembling ones used elsewhere [139, 148]:

- The signal was lowpass filtered to 5 Hz with a 1st order Butterworth filter.
- Segments that lay above 0.05 mV were taken as DOWN states and those that lay below -0.15 mV as UP states.
- Any consecutive DOWN states less than 200 msec (5 Hz) apart were taken as one. Same for UP states.
- If a detected DOWN state was followed by more than one or no detected UP state before the next DOWN state, it was discarded.

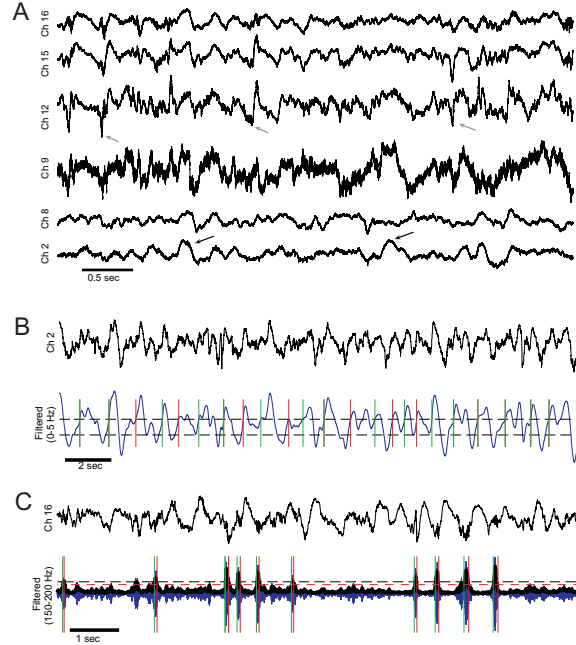


Figure 7.3: **A:** Rat LFP recordings from two mPFC locations (Channels 2 and 8) and four hippocampal locations (Channels 9 to 16) during one sleep session. Black arrows indicate distinct delta waves (DOWN states) in the cortical signals. Grey arrows indicate sharp waves in the hippocampus. **B-C:** Data segments from the mPFC (**B**) and the medial hippocampal signal (**C**). Below each signal, the filtered versions are depicted. The cortical signal was lowpass filtered (5 Hz) and the hippocampal one was bandpassed (150-200 Hz) and its RMS was calculated (black line). The green lines indicate the beginning of a SO cycle or a ripple accordingly. The red lines indicate their respective end times (in many SO cycles, these overlap with the beginning of the next cycle). The dotted lines indicate the detection criteria for the delta waves (DOWN states), the UP states and the ripple peaks accordingly. The red dashed line represents the detection criterion for the beginning and end of ripples.

This process returned only data segments that contained one detected DOWN state followed by one detected UP state, thus one SO cycle with a robust DOWN-UP transition. In total, 1162 SO cycles of average duration 767.42 msec (~ 1.3 Hz) and maximum duration 2430 msec ($SD = 272.6$ msec) were detected.

To detect any ripple episodes in the LFP from the medial hippocampus (channel 16), the same algorithm as in our CA3-CA1 model (but with different event detection thresholds) was followed:

- The signal was bandpass filtered at 150-200 Hz with a 10th order Butterworth filter.

- Its root mean square (RMS) was calculated over 10 msec non-overlapping bins. The SD of the RMS was calculated over the whole signal.
- Ripples were detected when the RMS exceeded a $4 \times \text{SD}$ threshold. Their boundaries were set where the RMS drops below $3 \times \text{SD}$ around the ripple.
- Ripples with less than 20 msec total duration were discarded and neighbouring ripples less than 10 ms apart were taken as one event.

The algorithm returned 1478 detected ripples of average duration 61.45 msec and maximum duration 151 msec ($\text{SD} = 20.23$ msec). Figures 7.3B-C display characteristic segments of the original LFPs from mPFC and medial hippocampus (CA1) and the filtered signals with detected SO cycles and ripples respectively.

We aligned all SO cycles over the minimum point of the UP state and calculated the average DOWN-UP transition (Figure 7.4A). We then calculated the firing histogram from all the spikes recorded close to the LFP signal (channels 2 and 3). Figure 7.4B demonstrates that most spikes occur during the average UP state, thus verifying our SO-detecting algorithm. We also calculated the histogram of all the spikes close to the medial hippocampal LFP (channels 15 and 16) and correlated it with the average DOWN-UP transition. Figure 7.4C reveals a similar tendency of (putative pyramidal) cells in CA1 to fire during UP states, though with less modulation than the cortical cells, since spikes appear in the average DOWN state as well. This preliminary result is in agreement with the observation of higher probability of CA1 firing during cortical UP states [104]. However, a larger sample of unit activity is needed to reach more robust conclusions.

We also superimposed all the detected ripples on the average DOWN-UP transition. Figures 7.4D-E reveal a tendency of ripples to occur during the UP state and particularly during its onset (negative peak). This result is again in support of the various studies on the coupling of hippocampal rippling by the cortical SO UP states [104, 149, 188].

Finally, we superimposed all the detected ripples over their minimum point and calculated the firing histogram of the spikes close to the hippocampal LFP (channels 15 and 16) in relation with the average ripple (Figure 7.4F). It appears that the spikes are phase locked to the average ripple, similar to previous observations [53, 241] and in agreement with our computational model (Figure 6.8), verifying our ripple-detecting algorithm. Nevertheless, the maximal firing precedes the ripple troughs by ~ 1 -2 msec instead of being superimposed (for pyramidal cells) or even slightly delayed (for interneurons) as

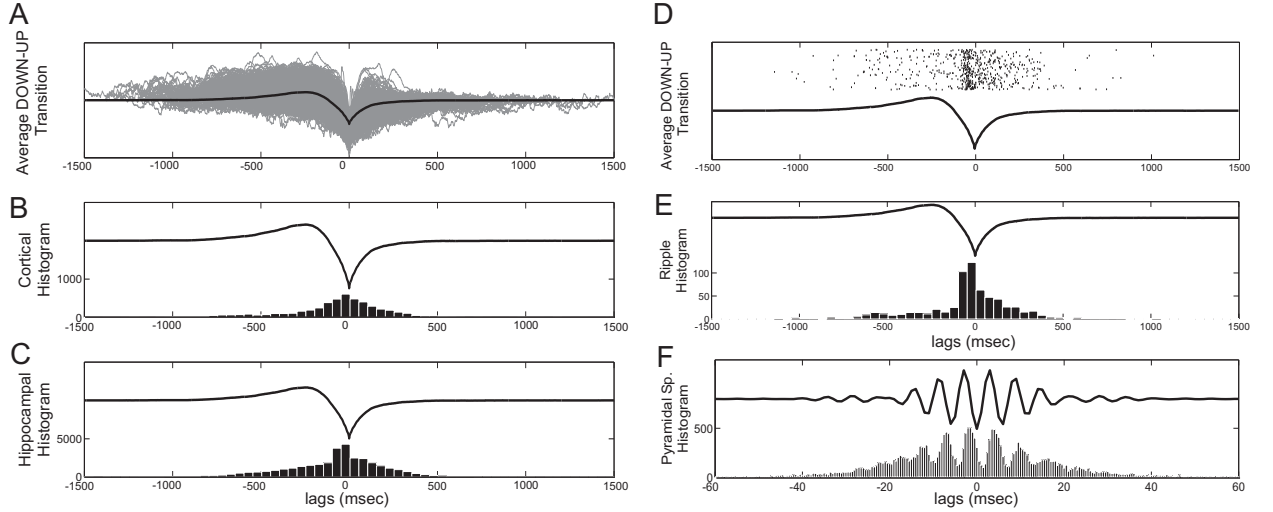


Figure 7.4: **A:** Average DOWN-UP transition (black) derived from superimposing all SO cycles over their minimum point (grey). **B:** Correlation of the average DOWN-UP transition with cortical spikes (Channels 2 and 3). **C:** Correlation with hippocampal spikes (Channels 15 and 16). Both cortex and hippocampus increase their firing at the UP state onset (minimum point). **D:** Average DOWN-UP transition and raster of detected ripples corresponding to each cycle, aligned over the the UP state onset. **E:** Ripple histogram correlated with the average DOWN-UP transition. Note that most ripples occur during the UP state. **F:** Correlation of the average ripple with the firing histogram of the spikes.

was expected, according to our discussion in the previous chapter. This systematic time shift of the recorded spikes could be due to an artifact in the recording procedure, e.g. in the algorithmic detection of a spike time stamp or in the intrinsic filtering of the LFP signal during recording, or even due to the very small number of spikes available (coming from only 2 cells).

In conclusion, this preliminary data analysis (albeit from a very small sample size from one rat) supports the aforementioned grouping role of cortical SO on hippocampal ripples and the higher probability of hippocampal firing during cortical UP states compared to DOWN states.

7.4 A Model of Cortico-Hippocampal Connectivity

7.4.1 General Connectivity and PSPs

Figure 7.5 shows the basic wiring diagram we use to couple the cortical and hippocampal models from Chapters 4 and 5 respectively. The prefrontal cortex (PFC) sends output to both CA3 and CA1 areas. Since the SO has been shown to extend even to the entorhinal cortex [104], we assume that the cortical output represents the output produced by the SO of the entorhinal cortex that reaches the hippocampus. CA1 receives direct input from the entorhinal layer 3 (EC3) via the temporoammonic pathway (TA), so the PFC-to-CA1 connection represents multistep connections originating from PFC and reaching CA1 through EC3. The afferent connections in CA3 are the mossy fibres (MF) originating from the dentate gyrus (DG) which has also been shown to partially participate in the SO [91, 104]. We therefore assume that the cortical output that is fed into CA3 represents the polysynaptic pathway through the entorhinal cortex and the dentate gyrus. Direct monosynaptic projections from the entorhinal cortex to CA3 are ignored. Finally, pyramidal cells in CA1 send monosynaptic projections to mPFC [38, 73, 83, 108, 110, 184], completing a feedback loop between PFC and hippocampus in the model.




Figure 7.5: Schematic diagram of the implemented cortico-hippocampal connectivity. The cortex provides indirect input to both CA3 and CA1 during the UP state, through the dentate gyrus (DG) mossy fibre connections (MF) and the entorhinal cortex layer 3 (EC3) temporoammonic pathway (TA). CA1 projects back to the cortex with monosynaptic connections.

Although our goal is for both CA areas to receive an input that is somewhat realistic, the cortex-to-hippocampus connections in the model represent polysynaptic pathways

		α_{syn}	EPSP (mV)	EPSC (nA)	Reported EPSP (mV)	Reported EPSC (nA)	Reference
PFC-CA3	PY	70	5	4.5	2-12	0.2	[113, 239]
	IN	0.5	0.25	0.03		0.03	[128]
PFC-CA1	PY	2	0.16	0.13		0.16	[156]
	IN	2	1	0.13	1		[69]
CA1-PFC	PY	8	9.5	3.5	10		[57, 157]
	IN	13	6	0.4	8		[204]

Table 7.1: Parameter values for the synaptic strength of all types of cortico-hippocampal synapses in the model along with the resulting EPSPs and EPSCs and the ones reported through physiological recordings (where available). All PSPs in the model were measured from a background membrane potential of -74.5 mV for cortical pyramidal cells, -65.3 mV for CA3 pyramidal cells, -62.6 mV for CA1 pyramidal cells and -62 mV for all interneurons.

involving many types of cells and therefore they are oversimplified. Both connections are modelled according to the connectivity scheme applied in the cortical and hippocampal models examined in chapters 4 and 5, i.e. Gaussian probability distribution of connections around each cell, with number of connections also gaussian distributed throughout the network. The same holds for the feedback via CA1-to-mPFC projections.

Mossy fibres are mostly oriented transverse to the longitudinal axis of CA3, with very limited overlap in the septotemporal direction [6, 8]. Temporoammonic connections are more widespread in CA1 and are topographically organised, with each EC3 axon covering about one quarter of the septotemporal CA3 axis [8]. For simplification, we assign Gaussian probability distributions for both types of connection with standard deviation (SD) of 500 cells (5 mm), which implies that both outputs will be spread over a cluster that covers the entire extent of the CA areas. This is so that the wave propagation of UP states in the cortical model is not reproduced in the hippocampus.

Both mossy fibres and the temporoammonic pathway excite pyramidal cells and interneurons in their respective target areas [1, 69, 77, 204], so we implemented connections from cortical cells to both cell types in CA3 and CA1. To reduce further the number of free parameters, we fix the magnitude of the unitary EPSPs to pyramidal cells and interneurons for all cortico-hippocampal connections, so that finally only the number of connections varies. The EPSP and/or EPSC levels, summarised in Table 7.1 and explained below, are taken from the literature (partly summarised in [56]).

The implemented mossy fibre EPSP to pyramidal cells is 5 mV, within the reported range of 2-12 mV [239], although the EPSC is much higher than reported (4.5 nA versus 0.2 nA. However, EPSCs up to 1 nA have been recorded [56]). The implemented synapse triggers spikes in the CA3 pyramidal cell after a train of 4 spikes at 40 Hz, in agreement with recordings [97]. The implemented EPSC in CA3 interneurons (30 pA) is similar to the recorded quantal EPSC [128] (a corresponding average EPSP has not been reported).

The TA synapses lie in the stratum lacunosum moleculare of CA1 exciting the distal dendrites of pyramidal cells [45]. As a result, they are relatively weak and get attenuated along the dendrites, producing very small depolarisations in the somatic compartment that have little effect on the firing properties of the cell [45, 69, 106]. We thus assign a small somatic EPSP magnitude to these connections, 0.16 mV which is comparable to the Schaffer collateral EPSP the cells receive from CA3 [177]. A similar EPSP has also been used in a more complex simulation of the temporoammonic effect on pyramidal cells [106]. The corresponding EPSC peaks at 0.13 nA which is close to the reported 0.16 nA [156]. Moreover, the temporoammonic input exerts a powerful feedforward inhibition in CA1, with EC3 cells possibly targeting inhibitory neurons as well, lying in the strata lacunosum moleculare and radiatum [45, 69]. We model this inhibition by assigning direct connections to CA1 interneurons in the model, although these represent pyramidal-layer somatic-inhibitory neurons. The implemented EPSP is 1 mV, as suggested by recordings [69], and the peak EPSC is the same as in pyramidal cells.

Unidirectional monosynaptic projections from hippocampal cells in CA1 and the subiculum to cells in the PFC have been identified in rats, cats and primates [38, 73, 83, 108, 110, 184, 199]. These projections target mostly the mPFC and orbitofrontal cortex and are topographically organised along the longitudinal and transverse axis of the hippocampus. Pharmacological studies have revealed that the glutamatergic synapses involved are primarily AMPA-receptor mediated [109], though NMDA receptors are also present, affecting LTP induction [107]. We implement these in the model by connecting all CA1 pyramidal cells with the mPFC cortical network, whose cells contain both AMPA and NMDA synaptic currents. The hippocampal afferents have been shown to target both cortical pyramidal cells [38] and interneurons [77, 204], thus we assign connections to both cell types in mPFC. Afferents stimulation to CA1 revealed synaptic responses in almost all ($\sim 90\%$) recorded mPFC pyramidal cells [57] and most PFC interneurons ($\sim 80\%$, [204]). These responses consisted mostly of early EPSPs of roughly

5-10 mV amplitude in pyramidal cells [57, 157] and one or two action potentials in interneurons [204], followed by inhibition. They were delayed by 14-18 msec from the stimulus [57, 73, 109] which is of the same order as the delay (36 msec) in the cell-pair covariance peaks measured by Wierzynski et al. [234]. To simplify this connectivity, we assign two connections from each CA1 pyramidal cell, one targeting pyramidal cells and one interneurons. These are spread throughout the whole cortical network by setting $SD = 500$ cells in the Gaussian distribution. The EPSPs are fixed (Table 7.1) and in the modified cortical network (where \bar{g}_{KNa} has been reduced to 0.5 mS/cm^2) can eventually produce an action potential in both cell types.

Finally, note that the maximum sodium-dependent potassium conductance \bar{g}_{KNa} , in the dendritic compartment of cortical pyramidal cells, is reduced from 1.33 to 0.5 mS/cm^2 , as in Figure 4.15 of Chapter 4. This is done to increase the excitability of pyramidal cells, prolonging(shortening) the duration of UP(DOWN) states, so that CA3 population bursts (and consequently ripples) can also take place during UP states. The resulting approximate equality in the duration of UP and DOWN states allows an unbiased temporal analysis of ripple occurrence in correlation with SO states.

All results presented below are taken from long-run simulations producing 40 sec of data. These contained 32 SO cycles of average duration ~ 1210 msec (0.83 Hz) when the CA1-to-PFC connections are omitted. In the final model version where they are included (Section 7.4.5) the SO frequency increases.

7.4.2 PFC-to-CA3: Dentate Gyrus Mossy Fiber Input

We first examine the effect of the cortical oscillatory input on CA3 which, when isolated, exhibits quasi-synchronised population bursts at theta frequencies.

As discussed earlier, mossy fibres have a very limited spatial extent along the longitudinal CA3 axis, with each granule cell having on average 10-18 large mossy terminals that target only 10-15 different nearby CA3 pyramidal cells [5]. To simplify the almost one-to-one mossy fibre projections we assign only one connection from each pyramidal cell in the cortex (which here represents the dentate gyrus as well) to CA3 pyramidal cells. Nevertheless, this output can reach almost any cell in CA3 due to the large SD of the Gaussian distribution of connections. Since the same number of pyramidal cells (1000) is implemented in both cortical and CA3 models, each CA3 cell receives on average one input.

Mossy fibres have been shown to excite almost 10 times more CA3 interneurons, but with much weaker connections involving small *en passant* boutons and very thin (filopodial) extensions of the mossy terminals [1]. This can result in an overall feedforward inhibitory input to CA3 from the dentate gyrus discharges, although the details of the corresponding circuits are still not well understood. Here, we present results from simulations with different average number of connections from each cortical (dentate) cell to CA3 interneurons (k_{DG-IN}) and examine the effect of the amount of feedforward inhibition on CA3 bursts.

Figure 7.6 displays the raster plot of the cortico-CA3 model for three different cases of k_{DG-IN} . Figure 7.6A shows the cortical SO, with the UP states providing the hippocampal input, and the bottom three panels display the CA3 raster plot for (from top) no connections to interneurons, one connection per cortical cell or 10 connections. Note that there are 1000 cortical pyramidal cells for 100 CA3 interneurons so in the latter two cases, each interneuron receives on average 10 and 100 connections respectively. When interneurons receive no cortical input, pyramidal cells get depolarised producing very clear population bursts in the UP states that involve most cells. In the DOWN state, CA3 is free to self-organize its bursting activity but the population events are less synchronised and coherent, involving fewer cells. As soon as the interneurons receive even a mild excitation the UP state bursts become less coherent since early interneuronal firing de-synchronizes the pyramidal cells. When feedforward inhibition on CA3 pyramidal cells becomes too strong for them to organize their firing in population bursts during UP states, they fire sparsely in disorganised patterns. On the contrary, the less robust the UP state bursts, the more synchronised population bursts appear in the DOWN state. When inhibitory firing dominates the UP state, the DOWN state population bursts are more clear and synchronised. It appears that the increased inhibition pyramidal cells receive, and the resulting lack of pyramidal spikes, synchronizes their post-inhibitory spiking, resulting in population bursts. The tendency of interneuronal activity to promote such post-inhibitory “rebound” bursts was also demonstrated in hippocampal electrophysiological experiments [43, 68, 96]. The difference between the case where interneurons are weakly excited or strongly driven by the external UP state input becomes more prominent when driving both cell populations with more connections (not shown).

A clear counterbalance arises, within the model, between population bursts in the UP and in the DOWN states. The stronger the cortical input to interneurons, the less

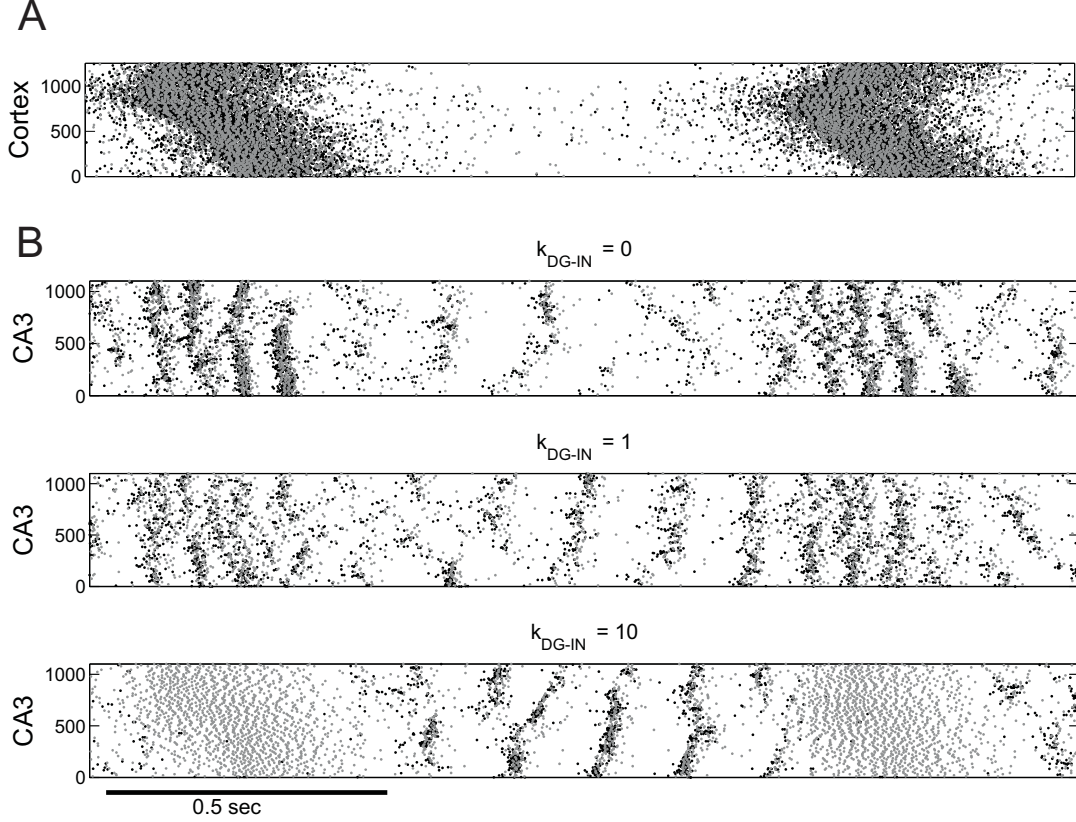


Figure 7.6: Raster plots of the PFC-CA3 model. Pyramidal cell spikes are shown in black and interneuronal ones in grey. **A:** SO in the cortical model acting as input to the CA3 model. **B:** Each cortical cell connects to one CA3 pyramidal cell, while the number of connections to interneurons varies between simulations. Three responses of CA3 are displayed corresponding to 0, 1 and 10 connections to interneurons per cortical cell. In both cases, the larger the feedforward inhibition in CA3, the less(more) population bursts during the UP(DOWN) state.

pyramidal firing in the UP state but the more synchronised the bursts in the DOWN state. In a full model including CA1, ripple responses to the CA3 bursts, will also be indirectly dependent on the SO input in CA3. If that is mostly inhibitory, then most ripples are eventually expected to arise during DOWN states, which contradicts most of the studies reviewed above.

Since, by setting $k_{DG-PY} = k_{DG-IN} = 1$, interneurons receive on average 10 times more connections than pyramidal cells, which is an anatomically realistic ratio [1], we keep this regime for the cortico-CA3 connectivity. After a long-run simulation we detect all UP states (whenever the total cortical synaptic activity exceeds $0.5 \times SD$) and superimpose all SO cycles (DOWN-UP pairs) aligned over the peak synaptic activity of the UP state

(0 msec). CA3 pyramidal and interneuronal spikes during the same data segments are also superimposed accordingly and the firing histograms are calculated for both cell types (Figure 7.7). As expected, spiking activity increases during the UP states, since the population bursts are more concentrated there and they engage more cells.




Figure 7.7: **A:** Total cortical synaptic activity (grey) in the PFC-CA3 model. All UP states are detected when the activity exceeds $0.5 \times \text{SD}$ and the window between the end of each UP state until the end of the next is taken as a DOWN-UP pair (black lines). **B:** Average SO cycle. **C-D:** Histograms of CA3 pyramidal (**C**) and interneuronal (**D**) cell spikes correlated with the average SO cycle.

We also examined the correlation of individual CA3 pyramidal firing rates with the SO. The firing rate of each pyramidal cell was calculated in 30 msec non-overlapping bins, and was averaged over the SO cycle, detected through the total firing rate of the cortical network (with the same detection criterion as before). Figure 7.8 displays the Z-score normalised average firing rate of each cell over the SO cycle. Cells were stacked according to the time of the peak of their normalised rate in descending order. It appears that roughly half the cells fire preferentially in the DOWN state bursts while the rest fire mostly in the UP state. The former get too inhibited by the UP state's increased feedforward inhibition to participate in population events. The excitation that some of the latter cells receive from their neighbours during the DOWN state bursts is not enough to make them burst, but the additional cortical input in the UP states drives

them to fire. Finally, some cells that are active in the DOWN state and still receive more excitation than the additional inhibition in the UP, participate in bursts during both states. This distinction of pyramidal cell firing agrees with intracellular recordings in the CA3 of anaesthetised mice, reporting a mixed modulation of pyramidal cell membrane potentials, with some correlated with the UP state and some anticorrelated [91].

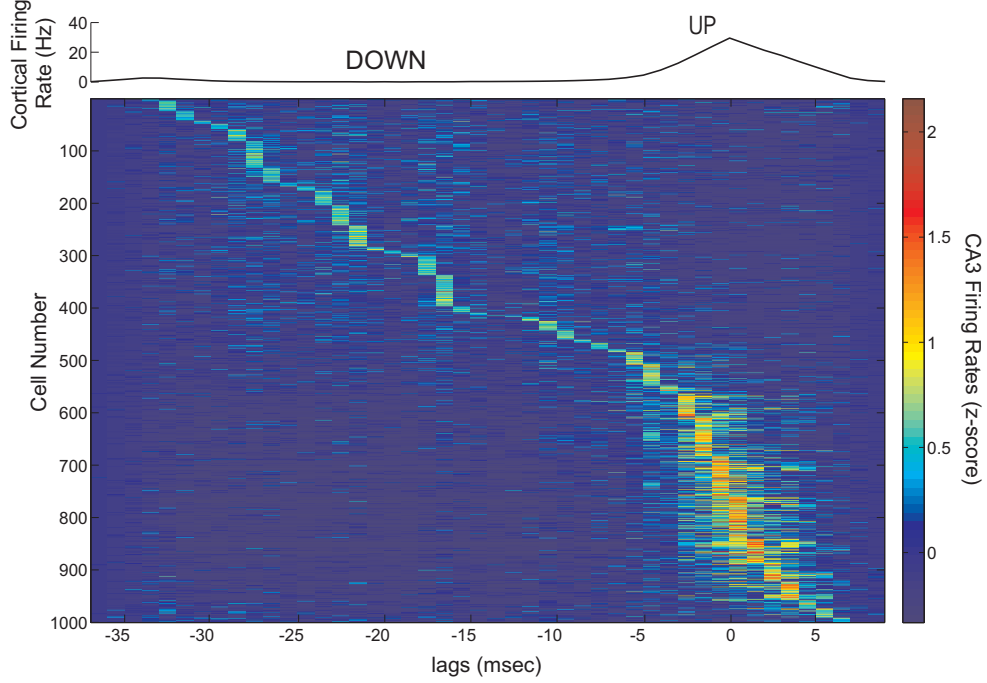


Figure 7.8: Top: Average SO cycle calculated from the average firing rate of the whole cortical network. **Bottom:** Colour coded Z-score normalised average firing rate of each CA3 pyramidal cell (rows), arranged by their peak in descending order and aligned with the DOWN-UP transition. Note that roughly half of the cells fire preferentially in the DOWN state and the other half in the UP state.

7.4.3 CA1 Response to the PFC-to-CA3 Model

We first look at the characteristics of the CA1 ripples arising from the population bursts in CA3 during the SO. We thus omit the TA input to CA1, so that the Schaffer collaterals are its only input. Since fewer CA3 cells participate in the population bursts, in comparison with the pure CA3-CA1 model where all CA3 pyramidal cells were driven by an external current, we increase the drive from CA3 to CA1 by setting the mean and SD of the Gaussian distribution of synapses per Schaffer connection to 15 (from 13).

We examined the spiking activity in CA1 in relation with the average SO cycle, estimated through the total synaptic activity of the cortical network. As in the previous

chapter, pyramidal cells are again split into those that receive less than 22.5 synapses for each Schaffer collateral connection ($1.5 \times$ the average of the Gaussian, 15) and those that receive more than 22.5 synapses. We refer to the former set as the “moderately-driven subset” and the latter as the “strongly-driven subset”. Figure 7.9 depicts the spike histograms for interneurons and pyramidal cells, both in the strongly-driven subset (grey) and the moderately-driven one (black). Both the former subset and the interneurons increase their spiking activity during the UP state, where the CA3 bursting is more intense. In contrast, the majority of pyramidal cells that receive an average (or below) CA3 drive do not participate in the UP state ripples due to intense interneuronal activity. They fire mostly in the DOWN states, when the CA3 input is weaker and they can overcome the feedforward inhibition. This is a contradiction with the findings of [104], where most recorded CA1 pyramidal cells tended to fire in the UP state, but agrees with anaesthesia-induced results in [91] where CA1 cells were anti-correlated with the cortical SO, being hyperpolarised during UP states.




Figure 7.9: Spiking activity in CA1, when the cortical input targets only CA3, correlated with the average SO cycle (thick line). **Middle:** Spike histogram of the strongly-driven CA1 pyramidal cells (grey) and the moderately-driven subset (black). **Bottom:** Spike histogram of interneurons

In the pure CA3-CA1 network, the power spectral density of the CA1 interneuronal average membrane potential peaks in the range of 150-200 Hz (see previous chapter). The frequency of that peak was shown to depend partly on the total excitatory input that CA1 receives through the Schaffer collaterals. Now that there is no input to CA3 during DOWN states and the UP states produce both excitation and inhibition, the drive to CA1 is weaker and thus the peak is broadened to a frequency range of 100-200 Hz, (Figure 7.10A).

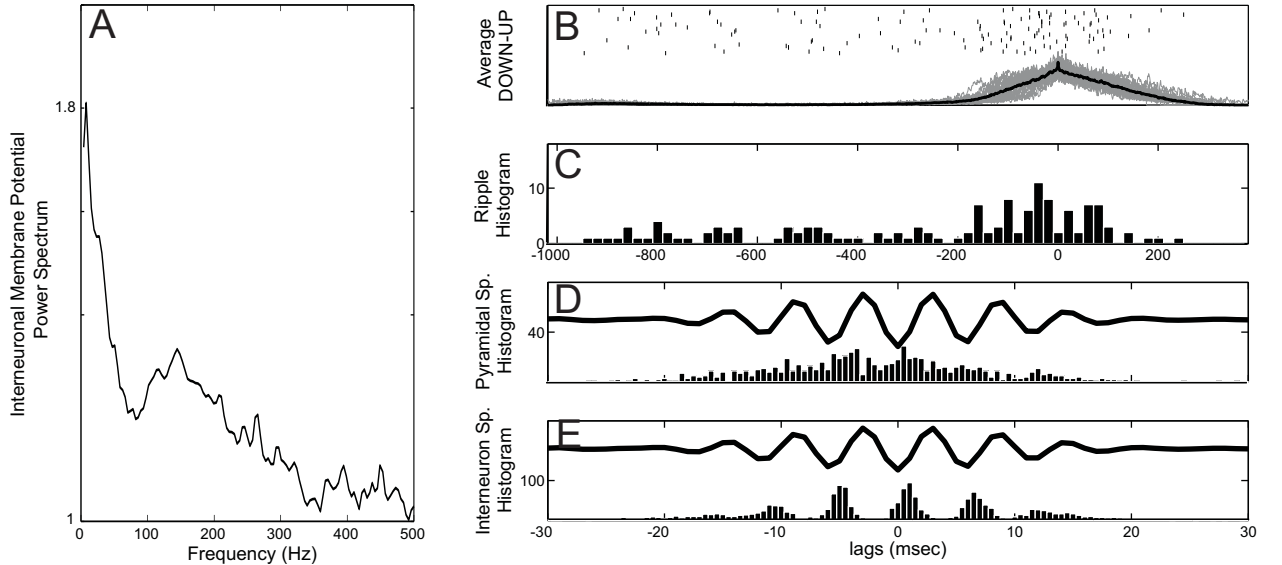


Figure 7.10: **A:** Power spectral density of the average membrane potential of CA1 interneurons when the TA input is omitted. The ripple peak corresponds to frequencies 100-200, slightly below those of the pure CA3-CA1 network. **B-C:** All detected ripples superimposed on the average SO cycle and ripple histogram. **D-E:** Pyramidal and interneuronal spike histograms correlated with the average ripple.

Figures 7.10B-C display the detected ripples and the ripple histogram relative to the average SO cycle. The ripple detection algorithm is similar to the previous chapter (with higher detection thresholds), i.e. ripples are detected when the RMS of the synaptic activity in the default recording site exceeds $2.5 \times \text{SD}$ (the ripple limits are detected at $1.8 \times \text{SD}$). As expected, more ripples appear during the UP state. The firing histograms of pyramidal cells and interneurons during the average ripple are also displayed. Although the interneuronal firing is phase locked to the ripple troughs, rhythmical spiking in pyramidal cells is less prominent than in the isolated CA3-CA1 model. This is again due to the low excitation the cells receive from CA3, particularly in DOWN states. Note that the cells were not split into two groups as the strongly-driven ones fired almost exclusively. The majority of cells produced practically no spikes during the ripple activity.

Therefore, additional excitation to CA1 is needed to induce robust ripple activity, i.e. interneurons to fire at higher frequencies and more pyramidal cells to fire during UP states with spikes phase locked to ripples. Assigning more synapses per Schaffer collateral connection to CA1 slightly increases the interneuronal oscillation frequency, but again most pyramidal cells were inhibited during UP states (not shown). An alternative approach

would be to assign a direct cortical input to CA1 through the TA pathway. The importance of the TA input for the coupling of CA1 with the cortical SO has been implied by *in vivo* studies [63, 237] and indirectly by the role of TA on CA1 spiking and plasticity [167] and on memory consolidation [168].

7.4.4 PFC-to-CA1: Entorhinal Cortex Temporoammonic Input

Although a topographic organization of the EC3 connections to the distal dendrites in the stratum lacunosum moleculare in CA1 has been described [8], the exact details of this circuit are not well understood. For example, an estimate of how many CA1 pyramidal cells or interneurons are contacted by a single EC3 axon is unknown. It has been shown though, that the TA input has little effect on the firing of postsynaptic pyramidal cells, but exerts a powerful inhibitory effect in CA1 activity [45, 69]. The interneurons implemented in our model are perisomatic, lying in the stratum pyramidale, for which limited connections are expected to arise via the TA pathway. But since the fixed pyramidal cell EPSP, in the model, is much weaker than the interneuronal one (Table 7.1), inhibition can easily dominate and prevent pyramidal spikes. Therefore a large number of connections from cortical cells to CA1 pyramidal cells is needed in order for them to overcome the strong feedforward inhibition.

We implemented various values for the numbers of pyramidal and interneuronal cells that each cortical cell targets (k_{EC-PY} and k_{EC-IN} respectively), in order to test their effect on the resulting CA1 responses. Firstly, we removed the CA3 input to isolate the effect of the TA input on CA1. Figure 7.11 contains spike raster plots with different sets of k_{EC-PY} and k_{EC-IN} (shown on top of each panel) from simulations where the CA3 network was omitted. Note that an interneuron receives on average $10 \times k_{EC-IN}$ connections. In all cases, cortical DOWN states are accompanied by sparse spontaneous firing of CA1 pyramidal cells, depolarising neighbouring interneurons. When $k_{EC-PY} = 50$ and $k_{EC-IN} = 1$, the pyramidal activity is practically the same during UP and DOWN states. Doubling k_{EC-PY} leads to more pyramidal spikes in the UP states and consequently more interneuronal spikes as well. Increasing k_{EC-IN} promotes interneuronal spiking and the suppression of pyramidal spikes. A further increase of k_{EC-PY} and k_{EC-IN} , while keeping their ratio constant, appears to favour interneuronal activity over pyramidal. Thus, in the last case, only few pyramidal cells fire during UP states, with multiple spikes. The majority are suppressed by feedforward inhibition.

After restoring the Schaffer collateral connections, we examined whether the TA input

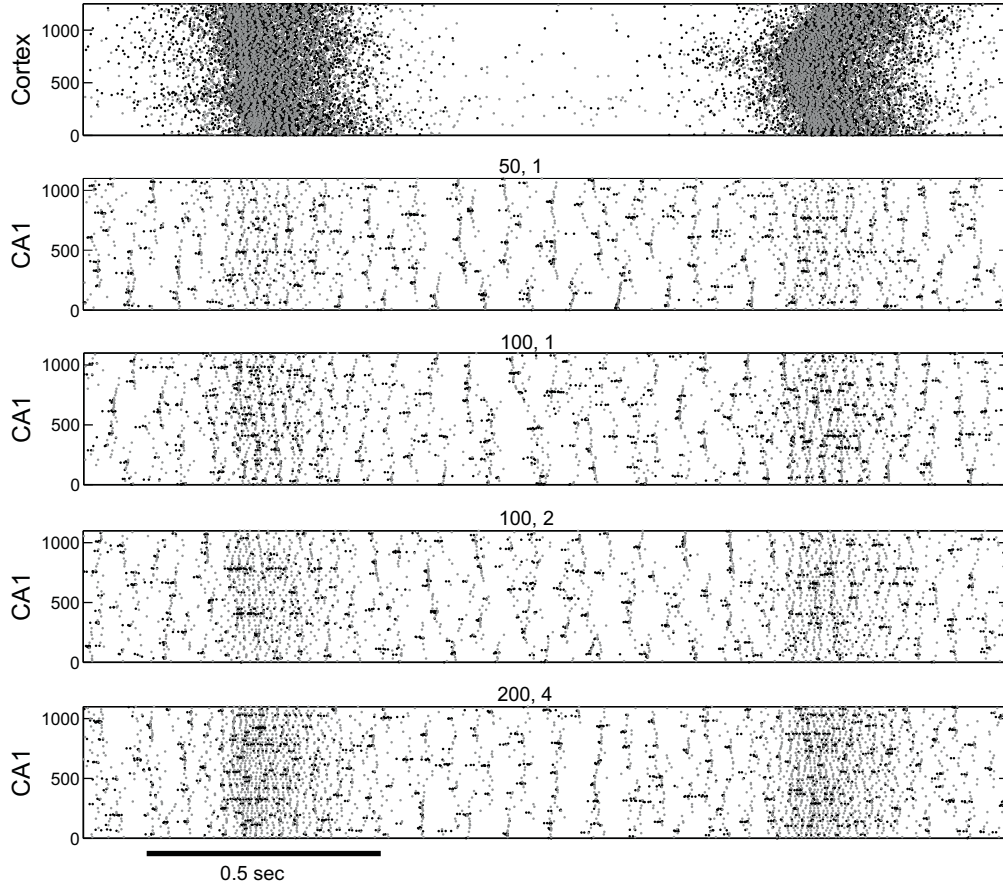


Figure 7.11: Spike raster plot of the cortical slow oscillation (top) and the CA1 response through only the TA input. Pyramidal cell spikes are shown in black and interneuronal ones in grey. The number of CA1 pyramidal and interneuronal cells each cortical cell targets, on average, is given on top of each panel.

can influence CA1 in producing ripple-frequency responses; namely whether TA can shift the spectral peak of the interneuronal average membrane potential back to 150 – 200 Hz. The corresponding spectral densities from simulations with various combinations of k_{EC-PY} and k_{EC-IN} (not shown) are similar to Figure 7.10A, indicating that the TA input alone cannot drive CA1 interneurons to oscillate at higher frequencies during CA3 population bursts.

Since the TA input affects CA1 spiking, we next examined how pyramidal cell firing rates depend on it. In Figure 7.12, the Z-score normalised firing rates of pyramidal cells are plotted, same as in Figure 7.8, for various combinations of k_{EC-PY} and k_{EC-IN} . Similar to CA3, most cells fire preferentially either in the UP or the DOWN states. As expected, increasing only the pyramidal cell drive results in more UP-state active cells. Increasing the interneuronal drive, even modestly, has the opposite effect. Moreover, an increase

in the TA drive to pyramidal cells does not significantly improve the phase locking of their spiking with ripple troughs (not shown). So the pyramidal spike histogram aligned with the average ripple remains similar to Figure 7.10.

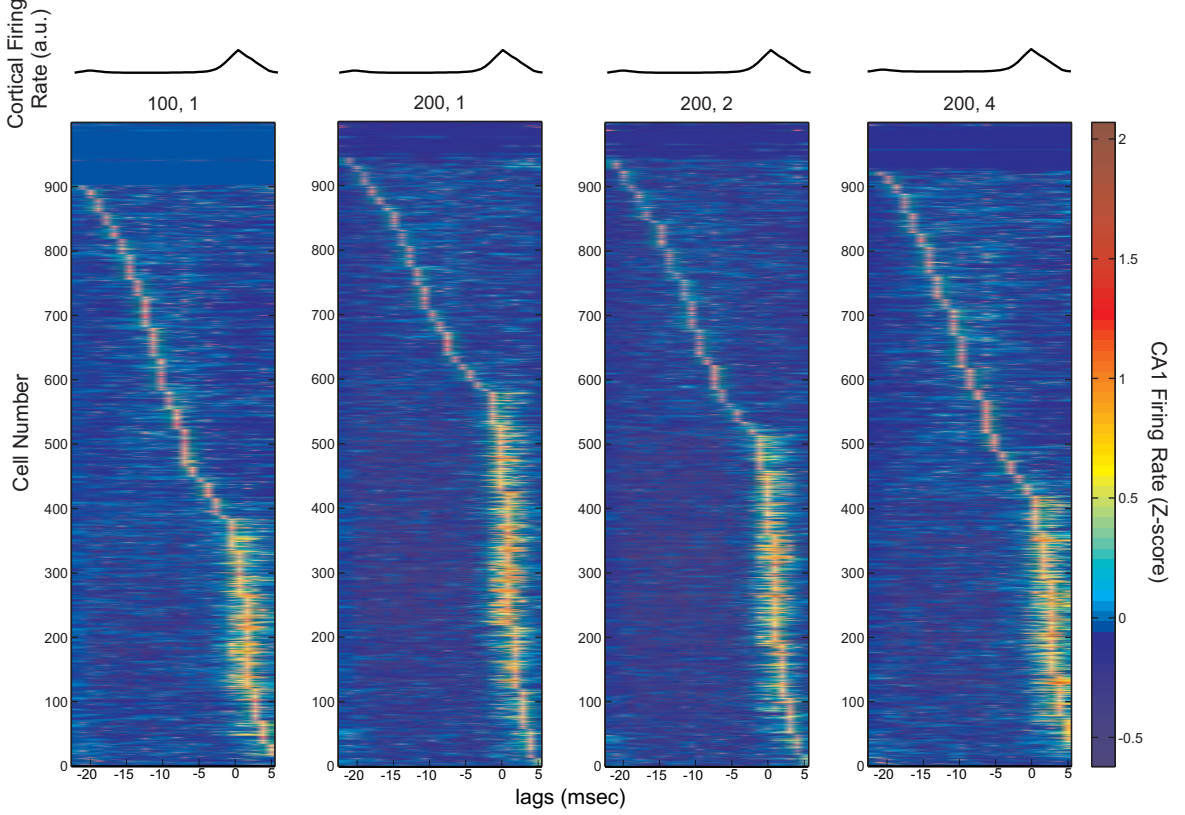


Figure 7.12: Colour coded Z-score normalised firing rates of CA1 pyramidal cells, for various TA inputs. The rates are averaged over all SO cycles (shown on top) and arranged by their peak in descending order. The values of k_{EC-PY} and k_{EC-IN} are shown on top of each panel. Note how the increase in the TA-induced excitation-to-inhibition ratio favours UP state pyramidal spiking and vice versa.

It thus becomes clear that, although the TA input does not particularly affect the frequency of the interneuronal responses to CA3 bursts, it does determine, to a large extent, the fraction of pyramidal cells that will overcome the UP state-elevated inhibition. The model suggests that a strong TA-depolarisation of CA1 pyramidal cells, compared to interneurons, aids in reproducing the CA1 spiking behavior described in [104]. We thus fix the TA input to $k_{EC-PY} = 200$ and $k_{EC-IN} = 2$. Moreover, since the Schaffer collateral input affects the ripple-frequency responses in CA1, we further increase the average number and SD of synapses per Schaffer connection to 18.

Figure 7.13A contains a segment of the raster plot from all three networks with these new

parameter values. We calculated the firing histogram of all pyramidal cells belonging either to the strongly-driven subset, or the moderately-driven one, in correlation with the average SO cycle (Figure 7.13B). This time, the additional excitation from the TA drives many more cells above spiking thresholds, hence spiking activity increases during UP states in both groups. As expected, interneurons also increase their firing during UP states.

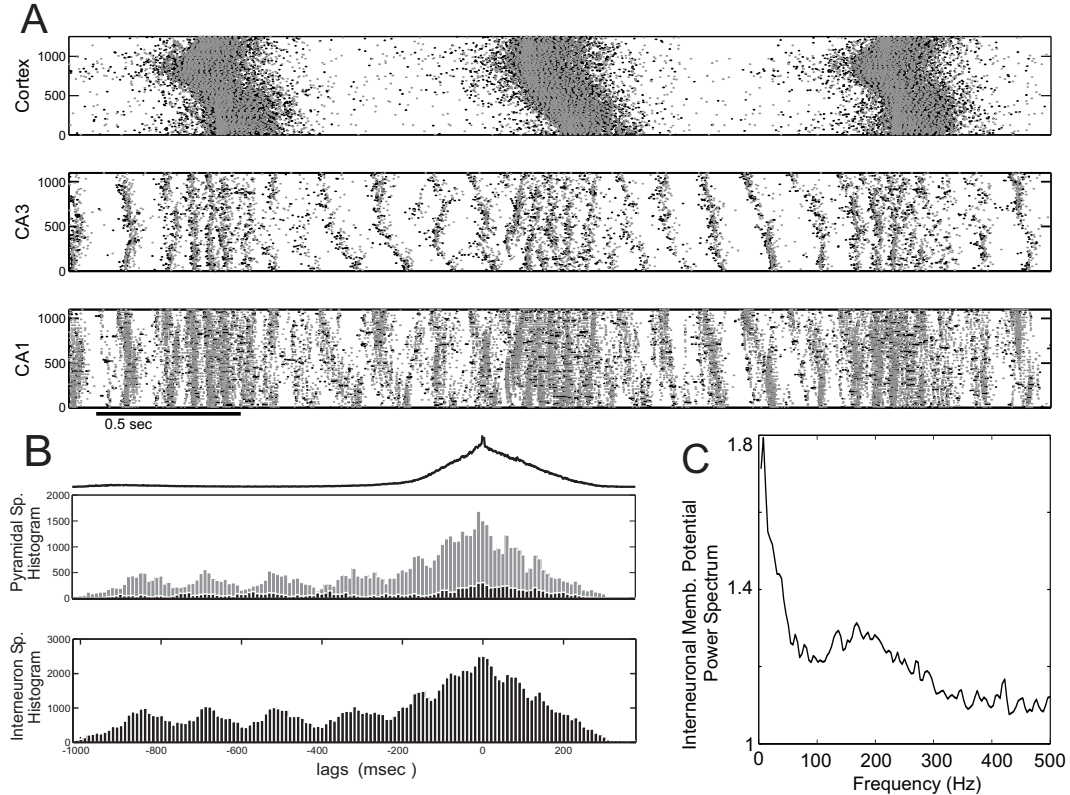


Figure 7.13: **A:** Spike raster plot of the cortical, CA3 and CA1 networks, when each cortical cell excites 1 CA3 pyramidal cell and 1 CA3 interneuron through the dentate gyrus input and 200 pyramidal cells and 2 interneurons through the TA pathway. **B:** Spike histogram of the strongly-driven CA1 pyramidal cells (top grey), the moderately-driven ones (top black) and the interneurons (bottom), correlated with the average SO cycle. **C:** Power spectral density of the average membrane potential of CA1 interneurons after increasing the CA3-to-CA1 input.

The power spectral density of the average membrane potential of CA1 interneurons is shown in Figure 7.13C. The additional drive the interneurons receive, through the increased Schaffer input, has shifted the peak of their oscillation back to 150-200 Hz, bringing the CA1 oscillatory responses again within the ripple frequency range.

All ripples detected in the recording site of CA1 were superimposed on the average SO

cycle. Approximately 12% of CA3 pyramidal cells fire on the average population burst. The same percentage of CA1 cells fires in the corresponding ripples. This is in close agreement with the recorded percentage (10%) in naturally sleeping rats [241]. Therefore, it is again only a minority of cells firing in each ripple episode. Figures 7.14A,B display a segment of the cortical synaptic activity and the corresponding CA1 activity bandpass filtered at 150-200 Hz. Detected SO cycles and ripples are also displayed. The ripple histogram over the average SO cycle (Figures 7.14C,D) indicates again that more ripples appear during the UP states. Moreover, the firing histogram of CA1 pyramidal cells in correlation with the average ripple (Figures 7.14E), reveals that the phase locking of spikes from the strongly-driven subset to the ripple troughs is now restored to a large extent. The fact that histogram peaks do not completely overlap with the ripple troughs is due to the ripples in DOWN states where the CA3 drive is weaker. In fact, after subtracting DOWN state ripples by increasing the ripple detection threshold, the prominent phase locking is restored (not shown).

Finally, to examine the SO effect on the spiking of individual CA1 pyramidal cells, we repeated the Z-score normalised firing rate calculation for each cell, and arranged them according to either the time of the firing rate peak (Figure 7.15A), or the number of synapses per Schaffer input each cell receives (Figure 7.15B). The number of synapses per Schaffer-connection that each cell receives is plotted next to each panel. Figure 7.15A indicates that, with the implemented combination of TA- and CA3-drive, roughly half of the CA1 pyramidal cells fire preferentially during UP states. The corresponding Schaffer-input is, on average, stronger for these cells. In fact, Figure 7.15B suggests that the stronger the CA3 drive, the higher the probability that the cell will fire mostly in the UP states. Therefore, both plots indicate that generally the strongly-driven cells fire preferentially during UP states, where they receive higher depolarising input. Note however, that some cells with strong CA3-drive get too inhibited in UP states and tend to fire more in DOWN states. In contrast, many of the cells that receive average or weak CA3 input, and are silent during DOWN states, get strongly depolarised from the TA input, and thus fire in UP states. In conclusion, the correlation of a cell's firing rate with the SO appears to depend on a combination of both the CA3 and the TA drive.

7.4.5 CA1-to-PFC: Monosynaptic Projections

We finally implemented the monosynaptic feedback projections from CA1 to pyramidal cells and interneurons of the cortical network (Figure 7.5) and examined its effect on

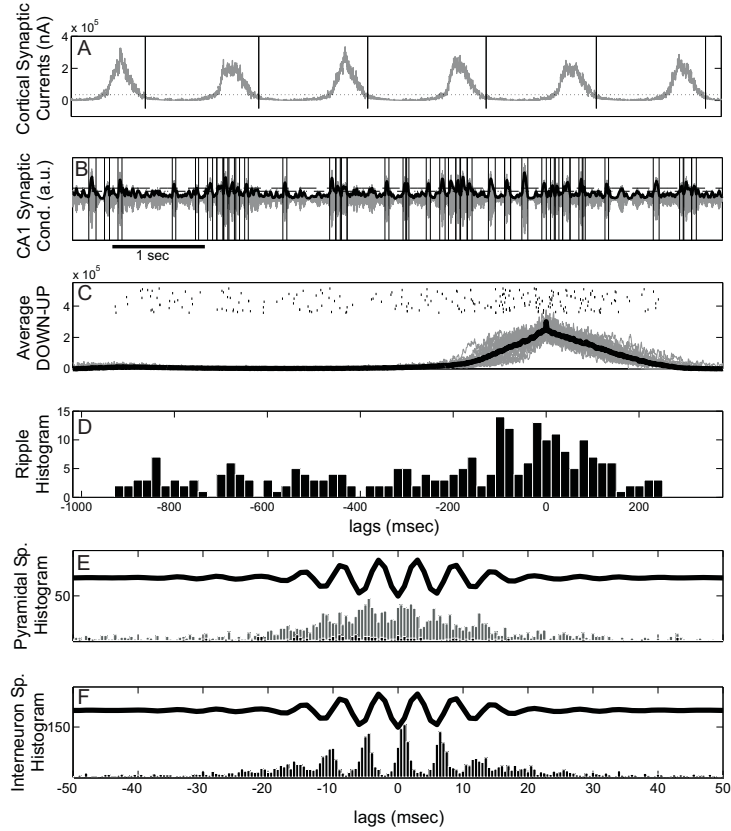


Figure 7.14: Ripple activity in CA1 when the cortical input targets both CA3 and CA1. **A:** Total synaptic currents in the cortical network. Vertical lines show the detected SO cycles. **B:** Filtered synaptic conductance in the CA1 recording site (150-200 Hz, grey) and RMS (black). Vertical lines show the time windows of detected ripples. **C-D:** All detected ripples superimposed on the average SO cycle and ripple histogram. **E-F:** Pyramidal and interneuronal spike histograms correlated with the average ripple. The subset of strongly driven pyramidal cells is shown in grey.

the SO.

Figure 7.16 displays the same segment of the cortical raster plot, without and with the inclusion of the feedback connections. The cortical excitation through CA1 activity appears to have two principal results. It increases the SO frequency and the spiking activity during DOWN states. Blocking the CA1 input on interneurons reveals that the SO frequency increases due to the additional depolarisation of the pyramidal cells during DOWN states, promoting the onset of the next UP state (which now propagates faster). Stimulating only interneurons results in longer DOWN states and reduced wave speed propagation of UP states, due to the hyperpolarisation of pyramidal cells.

The increased UP-state hippocampal activity also results in increased spiking in the

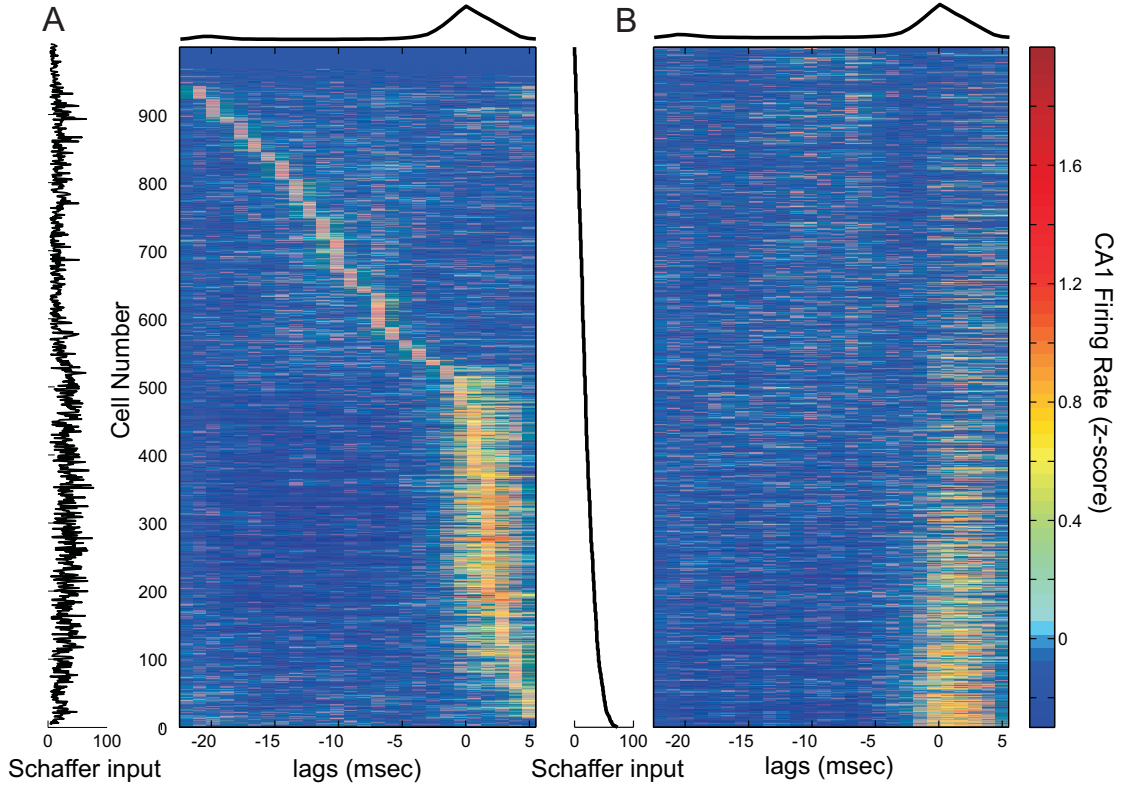


Figure 7.15: **A:** Colour coded Z-score firing rates of CA1 pyramidal cells during the average SO cycle. Cells are stacked according to the peak firing rate, in descending order. The number of synapses per Schaffer collateral-connection that each cell receives, is plotted on the left. **B:** Same, for cells stacked according to this number of synapses. The average SO cycle of the cortical firing rate is shown on top of each panel.

cortex. The average firing rates of both pyramidal cells and interneurons have doubled compared to the original cortical SO (not shown). These additional cortical spikes are again fed back to CA3 and CA1, further enhancing hippocampal activity. Figure 7.17A displays the raster plot from all three networks, similar to Figure 7.13A (note the smaller time scale) after the inclusion of the feedback mechanism. The elevated activity of the cortex during DOWN states, induces less coherent population bursts in CA3. As a consequence, ripple episodes appear to be more strongly correlated with the UP states than before (compare Figure 7.17B with Figures 7.14C,D). The stronger input that CA1 now receives through both CA3 and cortex induces increased CA1 firing as well. Although the firing rates of CA1 pyramidal cells have a similar form to Figure 7.15, the proportion of cells that fire preferentially during UP states is now roughly 60% (not shown). Finally, as can be seen in Figure 7.17C, the subset of strongly driven pyramidal cells exhibits a more prominent phase locking of spiking with ripple troughs

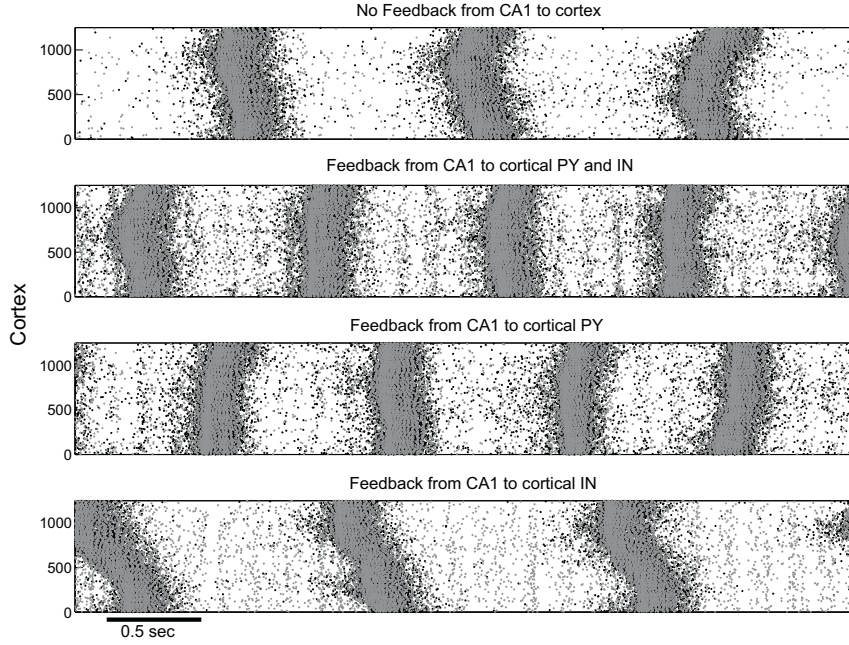


Figure 7.16: Raster plot of the cortical network, depending on the CA1 feedback projections. From top: With no CA1-to-mPFC feedback; with a feedback to both mPFC pyramidal cells and interneurons; with no connections to mPFC interneurons; with no connections to mPFC pyramidal cells. The same segment from all simulations is displayed.

than when feedback projections are omitted (Figure 7.14E). However, note that the almost continuous CA1 firing leads to particularly noisy (150-200 Hz filtered) synaptic activity which does not go low enough, between UP-state ripple episodes, to separate them, given the previously implemented detection threshold. This results in the need to increase the threshold, causing some DOWN-state (lower amplitude) ripples to be missed, but we do not expect the effect to be critical for our conclusions.

In conclusion, closing the cortico-hippocampal loop, through the CA1-to-cortex feedback projections, results in increased activity over the whole network, which affects the individual features of all three components. The frequency of the SO changes along with spiking during both UP and DOWN states. This additional firing results in stronger input arriving to the hippocampus, promoting pyramidal cell CA1 firing during ripples. On the other hand, the noisy and weak input arriving during DOWN states desynchronises CA3 population bursts, favouring ripples during UP states.

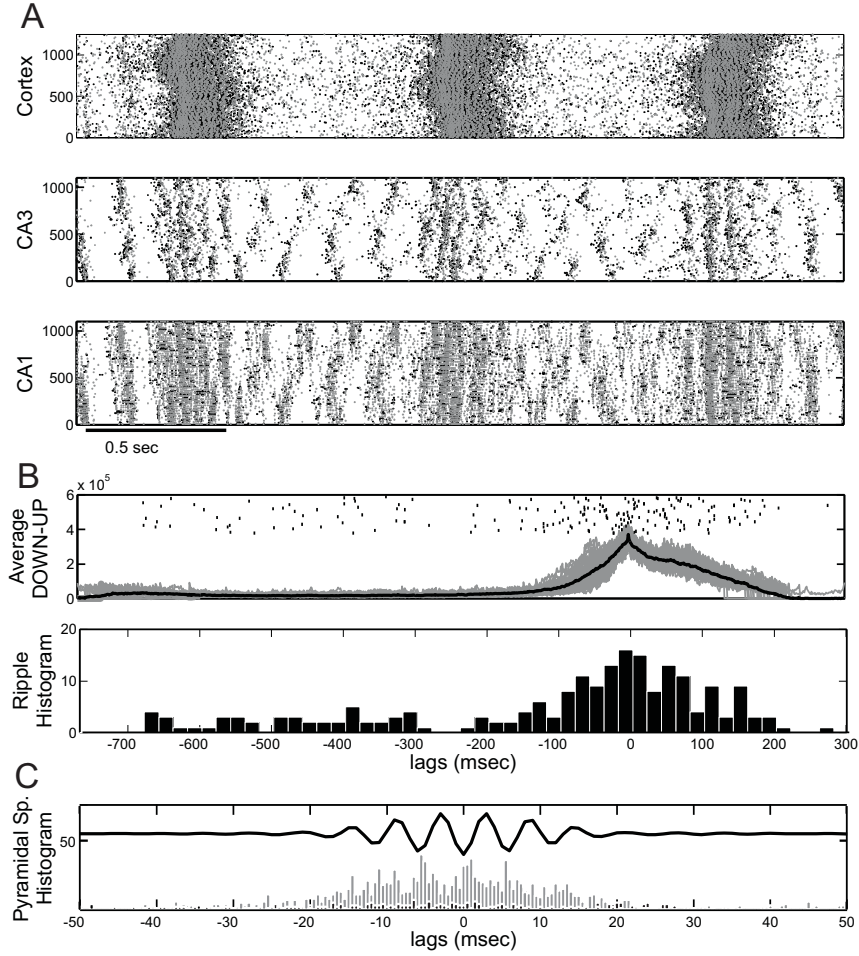


Figure 7.17: **A:** Raster plot of the cortical, CA3 and CA1 networks, after implementation of feedback projections from CA1 to cortical pyramidal cells and interneurons. **B:** Ripple histogram correlated with the average SO cycle. **C:** Pyramidal cell firing histogram correlated with the average ripple. Note that the phase locking of spikes from the strongly-driven subset is more prominent.

7.5 Discussion

Much research has recently been devoted to the role of the cortical slow oscillation (SO) impact on other brain areas including the thalamus and hippocampus. Correlations between rhythmic events in these structures during SO have been studied, along with their correlation with the actual SO [17, 90, 91, 104, 147, 149, 185, 188, 237]. All these studies support the fundamental role of the SO in coupling rhythmical activity of thalamic and hippocampal circuits, with most of them agreeing that the intense cortical spiking in UP states exerts a strong feedforward drive on these areas, promoting the occurrence of their intrinsic oscillatory events. In particular, thalamic spindles

and hippocampal ripples were shown to be grouped by UP states and moreover to be correlated with each other during those [149, 185, 188]. The spiking activity of the various hippocampal areas during the SO has also been recently studied, often with contradictory results [91, 104, 234].

We have reviewed here most of these studies, focusing on the correlation between ripples and cortical SO. Analysis of a small data set from a naturally sleeping rat, showed that CA1 spiking activity increased during UP states, verifying the entrainment of CA1 to the cortical SO [104, 237]. Ripple episodes also increased substantially during UP states, in agreement with the reviewed grouping role of the SO on SWRs [104, 149, 185, 188].

To study the temporal and functional relations between UP states and hippocampal activity, we coupled our previously analysed CA3-CA1 network with the SO-exhibiting cortical network of Compte et al. [46, 47], modified and described in Chapter 4, representing here mPFC. The implemented connectivity scheme (Figure 7.5) is a simplification of the direct and indirect projections between mPFC and hippocampus, with the cortical output functioning as the dentate gyrus and the entorhinal cortex input to the hippocampus [8, 56], and monosynaptic projections from CA1 to cortex [38, 73, 83, 108, 110, 184] constituting a feedback mechanism which is thought to be a pathway for memory consolidation [140]. Our simulations suggest that the correlation of CA3 firing and population bursts with the cortical SO, depends on the ratio of excitation-to-inhibition induced by the mossy fibre input. The spiking of CA1 cells and the occurrence of ripples in relation with UP/DOWN states appears to be controlled by the excitation-to-inhibition ratio induced via a combination of the Schaffer collateral and the temporoammonic input. CA1 projections to the cortex affect the SO frequency in our model, resulting to enhancement of UP state-ripple correlations.

Below we discuss our main findings for each implemented cortico-hippocampal connection, and compare them with relevant literature. Possible extensions to the model are also discussed.

7.5.1 Mossy Fibre Input

In Chapter 5 we showed that the CA3 network exhibits quasi-synchronous population bursts when isolated or under a constant depolarisation. Our analysis here revealed that the SO input to CA3 affects the occurrence of the bursts. When the pyramidal cells receive strong enough depolarisation to overcome inhibition, most population bursts ap-

pear during the UP states, while those in DOWN states are sparser and less synchronised over the network. On the contrary, when interneurons receive enough input to prevent pyramidal spiking during UP states, then well-synchronised bursts appear in DOWN states. The tendency of interneuronal activity to synchronize post-inhibitory pyramidal spiking has also been demonstrated in CA3 and CA1 *in vitro* slices [43, 68], and *in vivo* recordings [96]. Since these population bursts will produce ripple responses in CA1, the ratio of excitation to inhibition, produced within CA3 by the SO input, will determine whether the probability of ripple occurrence will be correlated or anticorrelated with UP states. When fixing the SO input in CA3 so that most bursts occur in UP states, the model reproduces the widely reported correlation of ripple activity with cortical UP states [17, 104, 149, 188]. Therefore the model predicts that a selective pharmacological blockade of the mossy fibre input on CA3 pyramidal cells will reduce the UP state-ripple correlation or even reverse it to DOWN states.

The average firing rates of individual pyramidal cells revealed that the SO input separates the CA3 population into two roughly equal size groups, one firing preferentially in UP states and the other in DOWN states. The former are driven by the UP state input, while the latter are silenced by the enhanced inhibition from their neighbouring interneurons. The relative sizes of these two groups is again determined by the excitation-inhibition ratio that CA3 pyramidal cells receive through the SO. Stronger(weaker) input to interneurons by UP states promotes bursts in DOWN(UP) states, increasing the group that fires mostly during those. This distinction in CA3 cell subgroups was also found through *in vivo* intracellular recordings in anaesthetised mice [91], where individual pyramidal cells exhibited a mixed modulation, some being correlated with the UP state and some anticorrelated. However, the Isomura et al. study on anaesthetised and naturally sleeping rats, reported increased depolarisation and preferential firing of most cells during DOWN states [104]. It is conceivable that species-, anaesthetic- or preparation-specific differences between the two studies, could result in differences on the input that CA3 receives through the SO, thus affecting the size of the two subgroups. Alternatively, it could be that most cells in the Isomura et al. study happen to be amongst those that get overall attenuated by the increased inhibition in UP states.

It is noteworthy that, although Isomura et al. report increased CA3 activity during DOWN states, most ripples still occur during UP states [104]. In fact they find that DOWN-activity is dominated by gamma oscillations which could produce more overall CA3 firing than population bursts (each burst involves only $\sim 30\%$ of cells in the

model), resulting in a preferential DOWN state firing. Since the current version of our CA3 network does not support gamma frequency oscillations (possibly due to the omission of recurrent inhibition), this remains to be examined. The lack of more bursts during DOWN states in the experiments, can be due to the short DOWN-state duration. Potentially, the synchronisation of CA3 cells to a self-paced population burst is already on the way during DOWN-state gamma activity, but often the UP state arrives first, accelerating this synchronisation, so that the UP-state onset triggers production of SWRs [A. Sirota, personal communication].

7.5.2 Temporoammonic Input

Since the average CA3 population burst involves less cells ($\sim 12\%$) than in the previous chapters where a tonic depolarisation was applied, CA1 receives a weaker input which is not enough to drive ripple-frequency responses. In the absence of any other input to CA1, the overall interneuronal oscillation frequency peak is below the 150-200 Hz range and the phase locking of pyramidal firing to ripple troughs is also attenuated (Figure 7.10). Although more ripple-like episodes occur during UP states, the majority of pyramidal cells fire preferentially during DOWN states, where the CA3-produced inhibition is weaker (Figure 7.9). The preferential firing of CA1 cells during DOWN states was also found in [91], where UP states coincided with an overall hyperpolarisation of CA1 pyramidal cells, but contradicts [104].

We implemented the TA pathway to CA1 to examine how a direct cortical input can affect the SO entrainment on CA1. The tendency of most CA1 cells to fire during UP states in naturally sleeping rats [104] and the correlation of TA-driven R-LM interneurons lying on stratum lacunosum moleculare with the SO [90] implies that the TA could have an active role in promoting the SO entrainment to the hippocampus [91, 104, 237]. The ability of the TA-CA1 synapse to exhibit both long-term depression and potentiation [167], and the dependence of Schaffer-evoked CA1 spikes on the timing difference between the TA and the Schaffer inputs [69, 106, 167], support the importance of the TA input during SWS. Finally, the deficiency that TA-lesioned rats exhibited in memory retrieval also indicates an active role of the TA input in memory consolidation.

We implemented a stronger TA input to CA1 interneurons so that the model reproduces the inhibitory effect of TA on CA1 [69] and the SO-modulation of R-LM interneurons [90]. The inclusion of the TA input did not increase the oscillation frequency of interneurons during ripple episodes or improve the phase locking of pyramidal firing to

ripple troughs. Additional CA3 drive was necessary for the CA1 responses to lie in the ripple frequency range and exhibit the expected spiking properties.

However, the TA input affected the number of pyramidal cells firing during UP states (Figure 7.15). The fraction of cells that manage to overcome the UP-state enhanced inhibition in CA1 depended on the ratio between excitation and inhibition produced by TA, with stronger input on pyramidal cells leading to more cells firing preferentially during UP states. Moreover, the alteration of the SO states results in four CA1 pyramidal cell subgroups, according to when they preferentially fire and the strength of their Schaffer collateral input:

1. Pyramidal cells strongly driven by CA3 that fire preferentially during UP states
2. Pyramidal cells strongly driven by CA3 that fire preferentially during DOWN states
3. Pyramidal cells receiving average/weak CA3 drive that fire preferentially during UP states
4. Pyramidal cells receiving average/weak CA3 drive that fire preferentially during DOWN states

Finally there are some cells that show very little variation in their firing during the SO. The above segregation is brought about by the combination of direct TA- and indirect dentate gyrus-CA3-input from the cortex. In the absence of TA, subgroups 2 and 3 largely disappear, as indicated by Figure 7.9. Although TA increased the number of UP state spiking cells, note that the actual ripples still involve only the few strongly CA3-driven cells (Figure 7.14E), indicating that the cells of subgroup 3 contribute mostly to the overall CA1 spiking during UP states.

A similar distinction between CA1 cells firing preferentially either at UP or DOWN states, or being uncorrelated with SO, was reported in Isomura et al. [104]. There, the majority of cells fired mostly at UP states. In our model, the two relevant subgroups can increase in size by stronger drives from Schaffer collaterals and the TA. Therefore, the model suggests that lesioning the TA pathway could reverse the Isomura et al. observation and result in most cells being hyperpolarised during UP states and firing during DOWN states, as observed in anaesthetised mice by [91]. Potentially, the anaesthetic could have altered the default TA input, in favor of interneuronal activity, or even have

suppressed it. In this case, the model predicts that a specific minority of cells would still fire in ripples, mostly during UP states, in contrast with the rest of the population.

7.5.3 Feedback to the Cortex and Memory Consolidation

Connecting the CA1 output with the cortical network had a substantial effect on the SO. The ratio of excitation-inhibition produced by these connections determined whether the oscillation frequency would increase or decrease, with higher pyramidal cell excitation leading to faster oscillations. This is due to the hippocampal activity received by the cortex during DOWN states. Since the effect of severing these projections on the cortical SO activity has not been studied yet, the model predicts that it could prolong DOWN states, particularly in prefrontal cortical regions where SO cycles are usually initiated [139].

One issue is the increased excitability of the cortical pyramidal cells in our model, resulting from the decrease in \bar{g}_{KNa} , introduced to prolong UP states (see Section 7.4.1). Even though the implemented CA1-to-mPFC EPSPs are inspired from literature (Table 7.1), this increased excitability leads to a somewhat exaggerated CA1-to-mPFC feedback, where all projections evoke spikes in their cortical targets (lower EPSPs were also applied but yielded similar results, not shown). We thus expect the effect of the projections to be more subtle in reality. Furthermore, we do not expect that the effects of hippocampal output on cortical cells will be directly fed back to CA3 or CA1, as in the model, since it would need to be sustained throughout polysynaptic pathways, in the same way as the SO. Currently, the only observed effect of these projections during SWS is in the level of single cell-pairs between CA1 pyramidal cells and targeted mPFC cells, whose firing appears correlated, particularly within SWRs [234]. These spike correlations, along with the replay of temporal patterns during SWRs in CA1 and the neocortex [111, 164] are currently the only evidence of information flow from the hippocampus to the cortex during SWS. Their modulation by ripples implies the potential role of this feedback circuit in transferring memories to the cortex through ripple episodes, for long-term storage. Our model supports the postulation that this consolidation process takes place mostly during UP states [188].

The application of a Granger-causality based analysis to a set of modelled or recorded data from mPFC, CA3 and CA1 would seem to be a promising idea to assess changes in cortico-hippocampal information flow between UP/DOWN states and ripples. This would be, to our knowledge, a novel approach. Nevertheless, the short duration of indi-

vidual SO states (<1 sec) and ripples (~ 100 msec), prevents the effective autoregressive modelling of the corresponding data segments (and the comparison of directionality between SO states). Furthermore, as aforementioned, the hippocampal effects on cortical spiking are directly fed back to both CA3 and CA1, in the model. This closed loop also prevents a potential gPDC analysis that could reveal any reversal of information flow between UP and DOWN states in support of our earlier findings (Chapter 3).

7.5.4 Future Steps

Clearly, the implemented model architecture ignores many components that might be crucial in the cortico-hippocampal loop, leaving out various other inputs that affect the hippocampus and many brain areas that filter the cortical signal before it reaches the hippocampus. Expanding the existing model with the inclusion of neuronal networks from other brain areas, such as the thalamus, would shed further light on the general cortico-hippocampal cross-talk during SWS. Thalamic spindle oscillations, for example, have been shown to closely correlate with ripples [149, 185], with hippocampal discharges often being entrained within spindles, implying a potential role of spindles in fine-tuning the timing of ripple occurrence while keeping cortical cells activated, and making them targets of ripple output [188]. Inclusion of plasticity in TA-CA1, CA3-CA1 and particularly CA1-mPFC synapses would also help a more detailed study of the process of temporal pattern transfer from CA3 to CA1, and eventually to mPFC, during memory consolidation.

However, even the relatively simplistic model developed and described here, offers a first approach to the problem of SO-ripple correlations and can help shed light on the role of certain important parameters and key components. To the best of our knowledge, this is the first modelling study on the ripple-SO correlation.

Summary - Future Directions

8.1 Summary of Results

A significant challenge in neuroscience lies in determining how complex interactions between discrete populations of neurones elicit functional output. One of the major questions in modern research, that lies within this framework, is where memories are stored and how are they retrieved. The standard model for memory consolidation assumes that new memories are stored temporarily in the hippocampal area and in closely connected medial temporal lobe regions, whereas remote memories are stored in the neocortex [33]. The transfer of memories from hippocampus to cortex for long-term storage is generally thought to take place during slow wave sleep (SWS). One related issue is how the hippocampus communicates with the neocortex and how information flows between the two structures. A series of observations indicate that hippocampal ripples could be the vehicle for information transfer to the cortex [111, 129, 234, 235]. This postulation is supported by various studies on the correlations between hippocampal ripples and slow oscillations in the cortex during SWS.

Having these studies as the main motivation, we have tackled the issue of cortico-hippocampal connectivity from two different perspectives. The first part of the thesis focused on data analysis from electrophysiological recordings, whereas in the second part we developed a computational modelling framework. In both cases we narrowed our analysis to SWS and anaesthesia, where a wealth of research on the subject of cortico-hippocampal interactions has been undertaken. Towards the development of a computational model on interactions between cortical and hippocampal outputs, we also delved into hippocampal dynamics and intrinsic rhythms, focusing on ripple generation.

Our goal was to reproduce part of the related studies on hippocampal dynamics and ripples and to expand and interpret their findings, either by revealing some of the key parameters, components and mechanisms, or by developing techniques for analysis of relevant neurophysiological data.

In the first part of this work we described mathematical tools and statistical techniques for assessing functional connectivity and directional interactions between neural ensembles and brain areas through electrophysiologically recorded signals. An introduction was given into some fundamental notions of multivariate time series analysis and autoregressive modelling of random processes, along with a review on some of the main statistical measures used to detect structural or functional connectivities between multivariate time series. We focused on the notion of Granger Causality and the related measures of Partial Directed Coherence (PDC) and generalised PDC (gPDC). For the latter, we implemented a novel critical value to detect the statistical significance of non-zero gPDC, based on a corresponding critical value for PDC [178]. To apply these measures, we developed a custom built numerical package in MATLAB, named FunCAT, which performs autoregressive modelling and functional connectivity analysis. We used it on LFP and unit spike data, recorded from anaesthetised rats, in order to assess directionality of information flow between medial prefrontal cortex (mPFC) and hippocampus, as well as intrahippocampally [202]. We also presented methods for spike train preprocessing and their conversion to continuous signals. The effects of such preprocessing and the relevant problems arising from it, were discussed.

Our gPDC-based connectivity analysis revealed a clear lateral-to-medial hippocampus directionality, consistent with the established neuroanatomy of Schaffer collaterals, projecting from CA3 pyramidal cells to CA1. A reciprocal flow of information between the mPFC and hippocampus was also detected, with the hippocampus showing an almost continuous information flow towards the mPFC, while the mPFC-to-hippocampal information flow peaked during short bursts of mPFC activity, reminiscent of the UP states observed during the cortical slow oscillation. When seizure response-evoking kainic acid was administered, a partial disruption of connectivity between the mPFC and hippocampus was observed, accompanied by a partial reversal of connectivity in the hippocampus. Our results led to the postulation of a general feedback scheme for the information flow between the three studied brain areas, which we aimed at reproducing through the computational modelling approach. Moreover, the reciprocal information flow between mPFC and hippocampus and its prominent alteration by cortical UP states, motivated

us to study the related problem of interaction between UP and DOWN states and hippocampal activity.

For this reason, in the modelling part of this work, we aimed at reproducing studies on the temporal relationships between intrinsic SWS rhythms of the neocortex and the hippocampus. Observed correlations between cortical slow oscillations and hippocampal ripples are one such example, and are thought to be important for the cortico-hippocampal-based memory consolidation scenario. In order to tackle this from a modelling perspective, we decided to couple a cortical network, exhibiting slow oscillations, and a hippocampal model exhibiting sharp wave-ripple complexes (SWRs).

For the cortical network, we reproduced the Compte et al. model [46, 47], implemented for the first time in the Python neural simulator “Brian” [84]. We chose this model because of its moderate complexity, simple single cell models, relatively few variables and inclusion of stochastic elements within a simple but realistic connectivity scheme. After implementing a modified version in Brian, we managed to reproduce a series of its reported features and characteristics.

In the absence of an established mechanism for ripple generation, we developed a new computational model of the rat hippocampus, consisting of two networks that represent the CA3 and CA1 hippocampal areas with network architecture similar to the Compte et al. model. We described in detail the set up of the model and showed that both networks reproduce various of the characteristic anatomical and/or functional properties of the corresponding hippocampal areas. The two models were then coupled in a realistic CA3-CA1 network, in order to yield SWRs, which are thought to emerge from the interaction between CA3 and CA1. We presented the characteristics of these complexes, compared them with those reported in literature and discussed the mechanisms behind them and their possible implications in memory consolidation. Our model proposes a novel mechanism for the emergence of SWRs, based on population bursts in CA3 and strong fast-decaying recurrent inhibition in CA1.

Finally we combined the cortical slow-oscillation exhibiting model, with the ripple-producing CA3-CA1 network. The cortical oscillatory output entrained the hippocampal model via connections to CA3, representing the mossy fibre input, and to CA1, via the temporoammonic input, while monosynaptic projections from CA1 to the cortex were also included. We showed how the spiking activities of CA3 and CA1 depend on the inhibition-excitation balance, induced by the two hippocampal inputs, and how these can affect the observed correlations between the cortical slow oscillation and ripples.

8.2 Some Future Directions

Perhaps one of the most interesting characteristics of SWRs is that neuronal participation in ripples is not random, with only a particular subset of pyramidal cells discharging during most ripple episodes [241]. This inhomogeneity further supports the memory consolidation scenario since the cell subset appears to be formed by preexisting experiences. Specifically, temporal sequences between correlated cells during wakeful training are preserved during sleep [189, 235], replayed at faster timescales during the short window of SWRs [129]. Such findings render the generation of the particular subset of cells, that does fire in ripples, crucial in understanding how the ripples could potentially send information to the neocortex. Most importantly, the question is how this subset changes in novel environments. One could study these mechanisms through a combination of electrophysiological recording analysis and modelling work. The inclusion of synaptic plasticity in the hippocampal model we developed, in combination with the implementation of specific firing patterns as input, could yield very interesting insights into the process. Moreover, in the full cortico-hippocampal model, one could also study the effect of such firing patterns as they reach the cortex through ripples. Another interesting question is whether this effect depends on the cortex being in the DOWN or UP state and whether the information transfer is more efficient in the latter case, as has been suggested [188]. A combination of modelling work with *in vivo* or *in vitro* electrophysiological recordings, which would support or reject the model's predictions, could address these complicated issues and give a better understanding of the processes involved.

Moreover, by examining in parallel the activity of multiple areas that are involved in memory consolidation, one can get a more complete image of the process. Apart from demanding parallel recordings over the involved areas, computational modelling of multiple areas, even in its most simplistic form, can also prove to be a very powerful additional tool in understanding how different formations contribute to such complex procedures as memory formation and consolidation. To this aim, a possible expansion of the existing cortico-hippocampal model could include neuronal networks from brain areas such as the thalamus or the amygdala that are involved in and affect the general cross-talk between cortex and hippocampus in deep sleep and anaesthesia. Thalamic spindle oscillations, for example, have been shown to closely correlate with ripples [185]. Hippocampal ripple discharges are often entrained within spindles, implying a potential role of spindles in fine-tuned timing of hippocampal ripples, while keeping cortical cells activated to be the targets of ripple output [188]. The temporal coupling of these rhythms under the group-

ing role of the slow oscillation is still open for investigation through a computational approach in combination with experimental work. Moreover, a Granger-causality based multivariate statistical analysis of the continuous or spike data, recorded from all these areas, could prove an important complementary tool in assessing neuronal functional interactions.

References

- [1] L. Acsady, A. Kamondi, A. Sik, T. Freund, and G. Buzsaki. GABAergic cells are the major postsynaptic targets of mossy fibers in the rat hippocampus. *Journal of Neuroscience*, 18(9):3386–3403, 1998.
- [2] H. Akaike. Fitting autoregressive models for prediction. *Annals of the Institute of Statistical Mathematics*, 21(1):243–247, 1969.
- [3] H. Akaike. A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6):716–723, 1974.
- [4] A.B. Ali, J. Deuchars, H. Pawelzik, and A.M. Thomson. CA1 pyramidal to basket and bistratified cell EPSPs: dual intracellular recordings in rat hippocampal slices. *The Journal of Physiology*, 507(1):201–217, 1998.
- [5] D.G. Amaral, N. Ishizuka, and B. Claiborne. Neurons, numbers and the hippocampal network. *Progress in brain research*, 83:1–11, 1990.
- [6] D.G. Amaral and M.P. Witter. The three-dimensional organization of the hippocampal formation: a review of anatomical data. *Neuroscience*, 31(3):571–91, 1989.
- [7] F. Amzica and M. Steriade. Disconnection of intracortical synaptic linkages disrupts synchronization of a slow oscillation. *Journal of Neuroscience*, 15(6):4658–4677, 1995.
- [8] P. Andersen. *The hippocampus book*. Oxford University Press, USA, 2007.
- [9] L. Astolfi, F. Cincotti, D. Mattia, M.G. Marciani, L.A. Baccala, F. de Vico Fallani, S. Salinari, M. Ursino, M. Zavaglia, L. Ding, et al. Comparison of different cortical connectivity estimators for high-resolution EEG recordings. *Human Brain Mapping*, 28(2):143, 2007.

- [10] L.A. Baccala and K. Sameshima. Directed Coherence: a tool for exploring functional interactions among brain structures. *Methods for neural ensemble recordings*, pages 179–192, 1998.
- [11] L.A. Baccalá and K. Sameshima. Partial directed coherence: a new concept in neural structure determination. *Biological Cybernetics*, 84(6):463–474, 2001.
- [12] L.A. Baccala, D.Y. Takahashi, and K. Sameshima. Generalized partial directed coherence. *15th International Conference on Digital Signal Processing*, pages 163–166, 2007.
- [13] M.S. Bartlett. Smoothing Periodograms from Time-Series with Continuous Spectra. *Nature*, 161(4096):686–687, 1948.
- [14] M. Bartos, I. Vida, M. Frotscher, J.R.P. Geiger, and P. Jonas. Rapid signaling at inhibitory synapses in a dentate gyrus interneuron network. *Journal of Neuroscience*, 21(8):2687–2698, 2001.
- [15] M. Bartos, I. Vida, M. Frotscher, A. Meyer, H. Monyer, J.R.P. Geiger, and P. Jonas. Fast synaptic inhibition promotes synchronized gamma oscillations in hippocampal interneuron networks. *Proceedings of the National Academy of Sciences of the United States of America*, 99(20):13222–13227, 2002.
- [16] M. Bartos, I. Vida, and P. Jonas. Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nature Reviews Neuroscience*, 8(1):45–56, 2007.
- [17] F.P. Battaglia, G.R. Sutherland, and B.L. McNaughton. Hippocampal sharp wave bursts coincide with neocortical “up-state” transitions. *Learning and Memory*, 11(6):697–704, 2004.
- [18] M. Bazhenov, I. Timofeev, M. Steriade, and T.J. Sejnowski. Model of thalamo-cortical slow-wave sleep oscillations and transitions to activated states. *Journal of Neuroscience*, 22(19):8691–8704, 2002.
- [19] C.J. Behrens, L.P. van den Boom, L. de Hoz, A. Friedman, and U. Heinemann. Induction of sharp wave–ripple complexes in vitro and reorganization of hippocampal networks. *Nature neuroscience*, 8(11):1560–1567, 2005.

- [20] Y. Ben-Ari. Limbic seizure and brain damage produced by kainic acid: mechanisms and relevance to human temporal lobe epilepsy. *Neuroscience*, 14(2):375–403, 1985.
- [21] J.S. Bendat and A.G. Piersol. *Random Data: Measurement and Analysis Procedures*. Wiley, 1986.
- [22] C. Bernard and H.V. Wheal. Model of local connectivity patterns in CA3 and CA1 areas of the hippocampus. *Hippocampus*, 4(5):497–529, 1994.
- [23] C. Bernasconi and P. KoÈnig. On the directionality of cortical interactions studied by structural analysis of electrophysiological recordings. *Biological Cybernetics*, 81(3):199–210, 1999.
- [24] M. Both, F. Böhner, O.B. und Halbach, and A. Draguhn. Propagation of specific network patterns through the mouse hippocampus. *Hippocampus*, 18(9):899–908, 2008.
- [25] P.J. Brockwell, R.A. Davis, and I. NetLibrary. *Introduction to time series and forecasting*. Springer New York, 2002.
- [26] A. Brovelli, M. Ding, A. Ledberg, Y. Chen, R. Nakamura, and S.L. Bressler. Beta oscillations in a large-scale sensorimotor cortical network: Directional influences revealed by Granger causality. *Proceedings of the National Academy of Sciences*, 101(26):9849–9854, 2004.
- [27] N. Brunel and X.J. Wang. What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *Journal of Neurophysiology*, 90(1):415–430, 2003.
- [28] S.L. Buchanan, R.H. Thompson, B.L. Maxwell, and D.A. Powell. Efferent connections of the medial prefrontal cortex in the rabbit. *Experimental Brain Research*, 79(2):469–483, 1994.
- [29] P.S. Buckmaster and F.E. Dudek. In vivo intracellular analysis of granule cell axon reorganization in epileptic rats. *Journal of Neurophysiology*, 81(2):712–721, 1999.
- [30] E.H. Buhl, S.R. Cobb, K. Halasy, and P. Somogyi. Properties of unitary IPSPs evoked by anatomically identified basket cells in the rat hippocampus. *European Journal of Neuroscience*, 7(9):1989–2004, 1995.

- [31] E.H. Buhl, K. Halasy, and P. Somogyi. Diverse sources of hippocampal unitary inhibitory postsynaptic potentials and the number of synaptic release sites. *Nature*, 368(6474):823–828, 1994.
- [32] G. Buzsaki. Hippocampal sharp waves: their origin and significance. *Brain Research*, 398(2):242–252, 1986.
- [33] G. Buzsaki. *Rhythms of the Brain*. Oxford University Press, USA, 2006.
- [34] G. Buzsaki, Z. Horvath, R. Urioste, J. Hetke, and K. Wise. High-frequency network oscillation in the hippocampus. *Science*, 256(5059):1025–1027, 1992.
- [35] G.O. Buzsaki. Memory consolidation during sleep: a neurophysiological perspective. *Journal of Sleep Research*, 7(S1):17–23, 1998.
- [36] A.J. Cadotte, T.B. DeMarse, P. He, and M. Ding. Causal Measures of Structure and Plasticity in Simulated and Living Neural Networks. *PLoS ONE*, 3(10):e3355, 2008.
- [37] R.C. Cannon, H.V. Wheal, and D.A. Turner. Dendrites of classes of hippocampal neurons differ in structural complexity and branching patterns. *The Journal of Comparative Neurology*, 413(4):619–633, 1999.
- [38] D.B. Carr and S.R. Sesack. Hippocampal afferents to the rat prefrontal cortex: synaptic targets and relation to dopamine terminals. *The Journal of Comparative Neurology*, 369(1):1–15, 1996.
- [39] C. Chatfield. *The Analysis of Time Series: An Introduction*. Chapman & Hall/CRC, 2004.
- [40] C.C. Chow, J.A. White, J. Ritt, and N. Kopell. Frequency control in synchronized networks of inhibitory neurons. *Journal of Computational Neuroscience*, 5(4):407–420, 1998.
- [41] J.J. Chrobak and G. Buzsaki. High-frequency oscillations in the output networks of the hippocampal-entorhinal axis of the freely behaving rat. *Journal of Neuroscience*, 16(9):3056–3066, 1996.
- [42] V.R.J. Clarke, B.A. Ballyk, K.H. Hoo, A. Mandelzys, A. Pellizzari, C.P. Bath, J. Thomas, E.F. Sharpe, C.H. Davies, P.L. Ornstein, et al. A hippocampal GluR5 kainate receptor regulating inhibitory synaptic transmission. *Nature*, 389(6651):599–603, 1997.

- [43] S.R. Cobb, E.H. Buhl, K. Halasy, O. Paulsen, and P. Somogyi. Synchronization of Neuronal-Activity in Hippocampus by Individual GABAergic Interneurons. *Nature*, 378(6552):75–78, 1995.
- [44] S.R. Cobb, K. Halasy, I. Vida, G. Nyiri, G. Tamas, E.H. Buhl, and P. Somogyi. Synaptic effects of identified interneurons innervating both interneurons and pyramidal cells in the rat hippocampus. *Neuroscience*, 79(3):629–648, 1997.
- [45] C.M. Colbert and W.B. Levy. Electrophysiological and pharmacological characterization of perforant path synapses in CA1: mediation by glutamate receptors. *Journal of Neurophysiology*, 68(1):1–8, 1992.
- [46] A. Compte, R. Reig, V.F. Descalzo, M.A. Harvey, G.D. Puccini, and M.V. Sanchez-Vives. Spontaneous high-frequency (10-80 Hz) oscillations during up states in the cerebral cortex in vitro. *Journal of Neuroscience*, 28(51):13828–13844, 2008.
- [47] A. Compte, M.V. Sanchez-Vives, D.A. McCormick, and X.J. Wang. Cellular and network mechanisms of slow oscillatory activity (< 1 Hz) and wave propagations in a cortical network model. *Journal of Neurophysiology*, 89(5):2707–2725, 2003.
- [48] D. Contreras, A. Destexhe, T.J. Sejnowski, and M. Steriade. Control of spatiotemporal coherence of a thalamic oscillation by corticothalamic feedback. *Science*, 274(5288):771–774, 1996.
- [49] D. Contreras and M. Steriade. Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. *Journal of Neuroscience*, 15(1):604–622, 1995.
- [50] B. Coomber, M.F. O’Donoghue, and R. Mason. Inhibition of endocannabinoid metabolism attenuates enhanced hippocampal neuronal activity induced by kainic acid. *Synapse*, 62(10):746–755, 2008.
- [51] S. Coombes and P.C. Bressloff. *Bursting: The genesis of rhythm in the nervous system*. World Scientific Pub Co Inc, 2005.
- [52] R. Cossart, R. Tyzio, C. Dinocourt, M. Esclapez, JC Hirsch, Y. Ben-Ari, and C. Bernard. Presynaptic kainate receptors that enhance the release of GABA on CA1 hippocampal interneurons. *Neuron*, 29(2):497–508, 2001.

- [53] J. Csicsvari, H. Hirase, A. Czurko, A. Mamiya, and G. Buzsaki. Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *Journal of Neuroscience*, 19(1):274–287, 1999.
- [54] J. Csicsvari, H. Hirase, A. Mamiya, and G. Buzsáki. Ensemble patterns of hippocampal CA3-CA1 neurons during sharp wave-associated population events. *Neuron*, 28(2):585–594, 2000.
- [55] J. Cui, L. Xu, S.L. Bressler, M. Ding, and H. Liang. BSMART: a Matlab/C toolbox for analysis of multichannel neural time series. *Neural Networks*, 21(8):1094–1104, 2008.
- [56] V. Cutsuridis, B. Graham, S. Cobb, and I. Vida. *Hippocampal Microcircuits: A Computational Modeler’s Resource Book*. Springer Verlag, 2010.
- [57] E. Degenetais, A.M. Thierry, J. Glowinski, and Y. Gioanni. Synaptic influence of hippocampus on pyramidal cells of the rat prefrontal cortex: an in vivo intracellular recording study. *Cerebral Cortex*, 13(7):782–792, 2003.
- [58] A. Destexhe. Sleep Oscillations. *Encyclopedia of Neuroscience*, 8:1037–1044, 2009.
- [59] A. Destexhe, D. Contreras, and M. Steriade. Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *Journal of Neuroscience*, 19(11):4595–4608, 1999.
- [60] A. Destexhe and TJ Sejnowski. Interactions between membrane conductances underlying thalamocortical slow-wave oscillations. *Physiological reviews*, 83(4):1401–1453, 2003.
- [61] J. Deuchars and A.M. Thomson. CA1 pyramid-pyramid connections in rat hippocampus in vitro: dual intracellular recordings with biocytin filling. *Neuroscience*, 74(4):1009–1018, 1996.
- [62] K. Diba and G. Buzsáki. Forward and reverse hippocampal place-cell sequences during ripples. *Nature neuroscience*, 10(10):1241–1242, 2007.
- [63] C.T. Dickson, T.D. Wolansky, and J.W. Kerber. Neocortical modulation of the hippocampal slow oscillation via the entorhinal cortex. *Soc Neurosci Abstr* 31:275.273, 2005.

- [64] M. Ding, S.L. Bressler, W. Yang, and H. Liang. Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling: data preprocessing, model validation, and variability assessment. *Biological Cybernetics*, 83(1):35–45, 2000.
- [65] A. Draguhn, R.D. Traub, D. Schmitz, and J.G.R. Jefferys. Electrical coupling underlies high-frequency oscillations in the hippocampus in vitro. *Nature*, 394(6689):189–192, 1998.
- [66] H. Eichenbaum. A cortical–hippocampal system for declarative memory. *Nature Reviews Neuroscience*, 1(1):41–50, 2000.
- [67] M. Eichler. On the Evaluation of Information Flow in Multivariate Systems by the Directed Transfer Function. *Biological Cybernetics*, 94(6):469–482, 2006.
- [68] T.J. Ellender, W. Nissen, L.L. Colgin, E.O. Mann, and O. Paulsen. Priming of hippocampal population bursts by individual perisomatic-targeting interneurons. *Journal of Neuroscience*, 30(17):5979–5991, 2010.
- [69] R.M. Empson and U. Heinemann. The perforant path projection to hippocampal area CA1 in the rat hippocampal-entorhinal cortex combined slice. *The Journal of Physiology*, 484(Pt 3):707–720, 1995.
- [70] G.B. Ermentrout and N. Kopell. Fine structure of neural spiking and synchronization in the presence of conduction delays. *Proceedings of the National Academy of Sciences of the United States of America*, 95(3):1259–1264, 1998.
- [71] L. Faes and G. Nollo. Extended causal modeling to assess Partial Directed Coherence in multiple time series with significant instantaneous interactions. *Biological Cybernetics*, 103:387–400, 2010.
- [72] E.E. Faselow, K. Sameshima, L.A. Baccala, and M.A.L. Nicolelis. Thalamic bursting in rats during different awake behavioral states. *Proceedings of the National Academy of Sciences*, 98(26):15330–15335, 2001.
- [73] F. Ferino, A.M. Thierry, and J. Glowinski. Anatomical and electrophysiological evidence for a direct projection from Ammon’s horn to the medial prefrontal cortex in the rat. *Experimental Brain Research*, 65(2):421–426, 1987.
- [74] D.J. Foster and M.A. Wilson. Reverse replay of behavioural sequences in hippocampal place cells during the awake state. *Nature*, 440(7084):680–683, 2006.

- [75] A.S. French and A.V. Holden. Alias-free sampling of neuronal spike trains. *Biological Cybernetics*, 8(5):165–171, 1971.
- [76] T.F. Freund and G. Buzsaki. Interneurons of the hippocampus. *Hippocampus*, 6(4):347–470, 1996.
- [77] P. Gabbott, A. Headlam, and S. Busby. Morphological evidence that CA1 hippocampal afferents monosynaptically innervate PV-containing neurons and NADPH-diaphorase reactive cells in the medial prefrontal cortex (Areas 25/32) of the rat. *Brain Research*, 946(2):314–322, 2002.
- [78] P.W. Gage and B. Robertson. Prolongation of inhibitory postsynaptic currents by pentobarbitone, halothane and ketamine in CA1 pyramidal cells in rat hippocampus. *British Journal of Pharmacology*, 85(3):675–681, 1985.
- [79] J.R.P. Geiger, T. Melcher, D.S. Koh, B. Sakmann, P.H. Seeburg, P. Jonas, and H. Monyer. Relative abundance of subunit mRNAs determines gating and Ca²⁺ permeability of AMPA receptors in principal neurons and interneurons in rat CNS. *Neuron*, 15(1):193–204, 1995.
- [80] J. Geweke. Measurement of Linear Dependence and Feedback Between Multiple Time Series. *Journal of the American Statistical Association*, 77(378):304–13, 1982.
- [81] J. Geweke, R. Meese, and W. Dent. Comparing Alternative Tests of Causality in Temporal Systems: Analytic Results and Experimental Evidence. *Journal of Econometrics*, 21(2):161–94, 1983.
- [82] P.S. Goldman-Rakic. Cellular basis of working memory. *Neuron*, 14(3):477–485, 1995.
- [83] P.S. Goldman-Rakic, L.D. Selemon, and M.L. Schwartz. Dual pathways connecting the dorsolateral prefrontal cortex with the hippocampal formation and parahippocampal cortex in the rhesus monkey. *Neuroscience*, 12(3):719–743, 1984.
- [84] D. Goodman and R. Brette. Brian: a simulator for spiking neural networks in Python. *Frontiers in Neuroinformatics*, 2(5):1–10, 2008.
- [85] B. Gourévitch, R.L. Bouquin-Jeannès, and G. Faucon. Linear and nonlinear causality between signals: methods, examples and neurophysiological applications. *Biological Cybernetics*, 95(4):349–369, 2006.

- [86] C.W.J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [87] F. Grenier, I. Timofeev, and M. Steriade. Focal synchronization of ripples (80–200 Hz) in neocortex and their neuronal correlates. *Journal of Neurophysiology*, 86(4):1884–1898, 2001.
- [88] A.I. Gulyás, R. Miles, A. Sik, K. Tóth, N. Tamamaki, and T.F. Freund. Hippocampal pyramidal cells excite inhibitory neurons through a single release site. *Nature*, 366(6456):683–687, 1993.
- [89] A.I. Gulyás, R. Miles, N. Hájos, and T.F. Freund. Precision and variability in postsynaptic target selection of inhibitory cells in the hippocampal CA3 region. *European Journal of Neuroscience*, 5(12):1729–1751, 1993.
- [90] T.T.G. Hahn, B. Sakmann, and M.R. Mehta. Phase-locking of hippocampal interneurons’ membrane potential to neocortical up-down states. *Nature neuroscience*, 9(11):1359–1361, 2006.
- [91] T.T.G. Hahn, B. Sakmann, and M.R. Mehta. Differential responses of hippocampal subfields to cortical up-down states. *Proceedings of the National Academy of Sciences*, 104(12):5169–5174, 2007.
- [92] P. Halász, M. Terzano, L. Parrino, and R. Bódizs. The nature of arousal in sleep. *Journal of Sleep Research*, 13(1):1–23, 2004.
- [93] D.M. Halliday and J.R. Rosenberg. Time and frequency domain analysis of spike train and time series data. *Modern Techniques in Neuroscience Research*, pages 503–543, 1999.
- [94] D.M. Halliday, J.R. Rosenberg, A.M. Amjad, P. Breeze, B.A. Conway, and S.F. Farmer. A framework for the analysis of mixed time series/point process data Theory and application to the study of physiological tremor, single motor unit discharges and electromyograms. *Progress in Biophysics and Molecular Biology*, 64(2-3):237–278, 1995.
- [95] E.J. Hannan and B.G. Quinn. The determination of the order of an autoregression. *Journal of the Royal Statistical Society*, 41(2):190–195, 1979.

- [96] K.D. Harris, H. Hirase, X. Leinekugel, D.A. Henze, and G. Buzsáki. Temporal interaction between single spikes and complex spike bursts in hippocampal pyramidal cells. *Neuron*, 32(1):141–149, 2001.
- [97] D.A. Henze, L. Wittner, and G. Buzsáki. Single granule cells reliably discharge targets in the hippocampal CA3 network in vivo. *nature neuroscience*, 5(8):790–795, 2002.
- [98] W. Hesse, E. Möller, M. Arnold, and B. Schack. The use of time-variant EEG Granger causality for inspecting directed interdependencies of neural assemblies. *Journal of Neuroscience Methods*, 124(1):27–44, 2003.
- [99] S. Hill and G. Tononi. Modeling sleep and wakefulness in the thalamocortical system. *Journal of Neurophysiology*, 93(3):1671–1698, 2005.
- [100] G.M. Hoerzer, S. Liebe, A. Schloegl, N.K. Logothetis, and G. Rainer. Directed Coupling in Local Field Potentials of Macaque V4 During Visual Short-Term Memory Revealed by Multivariate Autoregressive Models. *Frontiers in Computational Neuroscience*, 4:4, 2010.
- [101] D. Holcman and M. Tsodyks. The emergence of up and down states in cortical networks. *PLoS Comput. Biol*, 2(3):174–181, 2006.
- [102] J. Huang, J.Y. Chang, D.J. Woodward, L.A. Baccalá, J.S. Han, J.Y. Wang, and F. Luo. Dynamic neuronal responses in cortical and thalamic areas during different phases of formalin test in rats. *Experimental Neurology*, 200(1):124–134, 2006.
- [103] K.M. Hurley, H. Herbert, M.M. Moga, and C.B. Saper. Efferent projections of the infralimbic cortex of the rat. *The Journal of Comparative Neurology*, 308(2):249–276, 1991.
- [104] Y. Isomura, A. Sirota, S. Özen, S. Montgomery, K. Mizuseki, D.A. Henze, and G. Buzsáki. Integration and Segregation of Activity in Entorhinal-Hippocampal Subregions by Neocortical Slow Oscillations. *Neuron*, 52(5):871–882, 2006.
- [105] B.H. Jansen. Time series analysis by means of linear modelling. *Techniques in the Behavioral and Neural Sciences*, 5:157–180, 1991.
- [106] T. Jarsky, A. Roxin, W.L. Kath, and N. Spruston. Conditional dendritic spike propagation following distal synaptic activation of hippocampal CA1 pyramidal neurons. *Nature neuroscience*, 8(12):1667–1676, 2005.

- [107] T.M. Jay, F. Burette, and S. Laroche. NMDA Receptor-dependent Long-term Potentiation in the Hippocampal Afferent Fibre System to the Prefrontal Cortex in the Rat. *European Journal of Neuroscience*, 7(2):247–250, 1995.
- [108] T.M. Jay, J. Glowinski, and A.M. Thierry. Selectivity of the hippocampal projection to the prelimbic area of the prefrontal cortex in the rat* 1. *Brain Research*, 505(2):337–340, 1989.
- [109] T.M. Jay, A.M. Thierry, L. Wiklund, and J. Glowinski. Excitatory amino acid pathway from the hippocampus to the prefrontal cortex. Contribution of AMPA receptors in hippocampo-prefrontal cortex transmission. *European journal of neuroscience*, 4(12):1285–1295, 1992.
- [110] T.M. Jay and M.P. Witter. Distribution of hippocampal CA1 and subicular efferents in the prefrontal cortex of the rat studied by means of anterograde transport of Phaseolus vulgaris-leucoagglutinin. *The Journal of Comparative Neurology*, 313(4):574–586, 1991.
- [111] D. Ji and M.A. Wilson. Firing rate dynamics in the hippocampus induced by trajectory learning. *Journal of Neuroscience*, 28(18):4679–4689, 2008.
- [112] I.T. Jolliffe. *Principal component analysis*. Springer New York, 2002.
- [113] P. Jonas, G. Major, and B. Sakmann. Quantal components of unitary EPSCs at the mossy fibre synapse on CA3 pyramidal cells of rat hippocampus. *The Journal of Physiology*, 472(1):615–663, 1993.
- [114] M. Kamiński, M. Ding, W.A. Truccolo, and S.L. Bressler. Evaluating causal relations in neural systems: Granger causality, directed transfer function and statistical assessment of significance. *Biological Cybernetics*, 85(2):145–157, 2001.
- [115] M.J. Kaminski and K.J. Blinowska. A new method of the description of the information flow in the brain structures. *Biological Cybernetics*, 65(3):203–210, 1991.
- [116] E.R. Kandel, J.H. Schwartz, and T.M. Jessell. *Principles of neural science*. Appleton & Lange, 2000.
- [117] S. Kang, K. Kitano, and T. Fukai. Self-organized two-state membrane potential transitions in a network of realistically modeled cortical neurons. *Neural Networks*, 17(3):307–312, 2004.

- [118] C. King, D.A. Henze, X. Leinekugel, and G. Buzsáki. Hebbian modification of a hippocampal population pattern in the rat. *The Journal of Physiology*, 521(1):159–167, 1999.
- [119] E.D. Kirson, Y. Yaari, and M. Perouansky. Presynaptic and postsynaptic actions of halothane at glutamatergic synapses in the mouse hippocampus. *British Journal of Pharmacology*, 124(8):1607–1614, 1998.
- [120] T. Klausberger, P.J. Magill, L.F. Márton, J.D.B. Roberts, P.M. Cobden, G. Buzsáki, and P. Somogyi. Brain-state-and cell-type-specific firing of hippocampal interneurons in vivo. *Nature*, 421(6925):844–848, 2003.
- [121] T. Klausberger and P. Somogyi. Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. *Science*, 321(5885):53–57, 2008.
- [122] W.D. Knowles and P.A. Schwartzkroin. Local circuit synaptic interactions in hippocampal brain slices. *Journal of Neuroscience*, 1(3):318–322, 1981.
- [123] A. Korzeniewska, M. Mańczak, M. Kamiński, K.J. Blinowska, and S. Kasicki. Determination of information flow direction among brain structures by a modified directed transfer function (dDTF) method. *Journal of Neuroscience Methods*, 125(1-2):195–207, 2003.
- [124] M.D. Krasowski and N.L. Harrison. General anaesthetic actions on ligand-gated ion channels. *Cellular and Molecular Life Sciences*, 55(10):1278–1303, 1999.
- [125] R. Kus, M. Kaminski, and KJ Blinowska. Determination of EEG activity propagation: pair-wise versus multichannel estimate. *Biomedical Engineering, IEEE Transactions on*, 51(9):1501–1510, 2004.
- [126] T. Lanthorn, J. Storm, and P. Andersen. Current-to-frequency transduction in CA1 hippocampal pyramidal cells: slow prepotentials dominate the primary range firing. *Experimental brain research*, 53(2):431–443, 1984.
- [127] S. Laroche, T.M. Jay, and A.M. Thierry. Long-term potentiation in the prefrontal cortex following stimulation of the hippocampal CA1/subicular region. *Neuroscience Letters*, 114(2):184–190, 1990.
- [128] J.J. Lawrence, Z.M. Grinspan, and C.J. McBain. Quantal transmission at mossy fibre targets in the CA3 region of the rat hippocampus. *The Journal of Physiology*, 554(1):175–193, 2004.

- [129] A.K. Lee and M.A. Wilson. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron*, 36(6):1183–1194, 2002.
- [130] X.G. Li, P. Somogyi, A. Ylinen, and G. Buzsaki. The hippocampal CA3 network: an in vivo intracellular labeling study. *Journal of Comparative Neurology*, 339(2):181–208, 1994.
- [131] E.W. Lothman, R.C. Collins, and J.A. Ferrendelli. Kainic acid-induced limbic seizures: electrophysiologic studies. *Neurology*, 31(7):806–812, 1981.
- [132] H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, 2005.
- [133] B.A. MacVicar and F.E. Dudek. Local synaptic circuits in rat hippocampus: interactions between pyramidal cells. *Brain Research*, 184(1):220–223, 1980.
- [134] D.V. Madison and R.A. Nicoll. Control of the repetitive discharge of rat CA1 pyramidal neurones in vitro. *The Journal of Physiology*, 354(1):319–331, 1984.
- [135] N. Maier, M. Guldenagel, G. Sohl, H. Siegmund, K. Willecke, and A. Draguhn. Reduction of high-frequency network oscillations (ripples) and pathological network discharges in hippocampal slices from connexin 36-deficient mice. *The Journal of Physiology*, 541(2):521–528, 2002.
- [136] N. Maier, V. Nimmrich, and A. Draguhn. Cellular and network mechanisms underlying spontaneous sharp wave-ripple complexes in mouse hippocampal slices. *The Journal of Physiology*, 550(3):873–887, 2003.
- [137] F. Maingret, S.E. Lauri, T. Taira, and J.T.R. Isaac. Profound regulation of neonatal CA1 rat hippocampal GABAergic transmission by functionally distinct kainate receptor populations. *The Journal of Physiology*, 567(1):131–142, 2005.
- [138] S.L. Marple. *Digital spectral analysis*. Prentice-Hall Englewood Cliffs, NJ, 1987.
- [139] M. Massimini, R. Huber, F. Ferrarelli, S. Hill, and G. Tononi. The sleep slow oscillation as a traveling wave. *Journal of Neuroscience*, 24(31):6862–6870, 2004.
- [140] M.R. Mehta. Cortico-hippocampal interaction during up-down states and memory consolidation. *Journal of Neuroscience*, 10(1):13–15, 2007.
- [141] O. Melamed, O. Barak, G. Silberberg, H. Markram, and M. Tsodyks. Slow oscillations in neural networks with facilitating synapses. *Journal of Computational Neuroscience*, 25(2):308–316, 2008.

- [142] R. Miles. Synaptic excitation of inhibitory cells by single CA3 hippocampal pyramidal cells of the guinea-pig in vitro. *The Journal of Physiology*, 428(1):61–77, 1990.
- [143] R. Miles, K. Tóth, A.I. Gulyás, N. Hájos, and T.F. Freund. Differences between somatic and dendritic inhibition in the hippocampus. *Neuron*, 16(4):815–824, 1996.
- [144] R. Miles, R.D. Traub, and R.K. Wong. Spread of synchronous firing in longitudinal slices from the CA3 region of the hippocampus. *Journal of Neurophysiology*, 60(4):1481–1496, 1988.
- [145] R. Miles and R.K. Wong. Excitatory synaptic interactions between CA3 neurones in the guinea-pig hippocampus. *The Journal of Physiology*, 373(1):397–418, 1986.
- [146] M. Mölle and J. Born. Hippocampus Whispering in Deep Sleep to Prefrontal Cortex For Good Memories? *Neuron*, 61(4):496–498, 2009.
- [147] M. Molle, O. Eschenko, S. Gais, S.J. Sara, and J. Born. The influence of learning on sleep slow oscillations and associated spindles and ripples in humans and rats. *European Journal of Neuroscience*, 29(5):1071–1081, 2009.
- [148] M. Molle, L. Marshall, S. Gais, and J. Born. Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. *Journal of Neuroscience*, 22(24):10941–10947, 2002.
- [149] M. Molle, O. Yeshenko, L. Marshall, S.J. Sara, and J. Born. Hippocampal Sharp Wave-Ripples Linked to Slow Oscillations in Rat Slow-Wave Sleep. *Journal of Neurophysiology*, 96(1):62, 2006.
- [150] C. Mulle, A. Sailer, G.T. Swanson, C. Brana, S. O’Gorman, B. Bettler, and S.F. Heinemann. Subunit composition of kainate receptors in hippocampal interneurons. *Neuron*, 28(2):475–484, 2000.
- [151] V. Nadler et al. Kainic acid as a tool for the study of temporal lobe epilepsy. *Life Sciences*, 29(20):2031–2042, 1981.
- [152] A. Neumaier and T. Schneider. Estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Transactions on Mathematical Software (TOMS)*, 27(1):27–57, 2001.

- [153] V. Nimmrich, N. Maier, D. Schmitz, and A. Draguhn. Induced sharp wave-ripple complexes in the absence of synaptic inhibition in mouse hippocampal slices. *The Journal of Physiology*, 563(3):663–670, 2005.
- [154] K. Nishikawa and M.B. MacIver. Membrane and synaptic actions of halothane on rat hippocampal pyramidal neurons and inhibitory interneurons. *Journal of Neuroscience*, 20(16):5915–5923, 2000.
- [155] J. O’Neill, T. Senior, and J. Csicsvari. Place-selective firing of CA1 pyramidal cells during sharp wave/ripple network patterns in exploratory behavior. *Neuron*, 49(1):143–155, 2006.
- [156] N.A. Otmakhova, N. Otmakhov, and J.E. Lisman. Pathway-specific properties of AMPA and NMDA-mediated transmission in CA1 hippocampal pyramidal cells. *Journal of Neuroscience*, 22(4):1199–1207, 2002.
- [157] M.A. Parent, L. Wang, J. Su, T. Netoff, and L.L. Yuan. Identification of the hippocampal input to medial prefrontal cortex in vitro. *Cerebral Cortex*, 20(2):393–403, 2010.
- [158] N. Parga and L.F. Abbott. Network model of spontaneous activity exhibiting synchronous transitions between up and down states. *Frontiers in Neuroscience*, 1(1):57–66, 2007.
- [159] R.J. Paterka, A.C. Sanderson, and D.P. O’leary. Practical Considerations in the Implementation of the French-Holden Algorithm for Sampling of Neuronal Spike Trains. *Biomedical Engineering, IEEE Transactions on*, (2):192–195, 1978.
- [160] G. Paxinos and C. Watson. *The rat brain in stereotaxic coordinates*. San Diego, CA: Academic Press 4th edn, 1998.
- [161] E. Pereda, R.Q. Quiroga, and J. Bhattacharya. Nonlinear multivariate analysis of neurophysiological signals. *Progress in Neurobiology*, 77(1-2):1–37, 2005.
- [162] M. Perouansky, D. Baranov, M. Salman, and Y. Yaari. Effects of halothane on glutamate receptor-mediated excitatory postsynaptic currents: A patch-clamp study in adult mouse hippocampal slices. *Anesthesiology*, 83(1):109–119, 1995.
- [163] M. Perouansky, E.D. Kirson, and Y. Yaari. Halothane blocks synaptic excitation of inhibitory interneurons. *Anesthesiology*, 85(6):1431–1438, 1996.

- [164] A. Peyrache, M. Khamassi, K. Benchenane, S.I. Wiener, and F.P. Battaglia. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nature Neuroscience*, 12(7):919–926, 2009.
- [165] P.F. Pinsky and J. Rinzel. Intrinsic and network rhythmogenesis in a reduced Traub model for CA3 neurons. *Journal of Computational Neuroscience*, 1(1):39–60, 1994.
- [166] J.G. Proakis, C.M. Rader, F. Ling, and C.L. Nikias. *Advanced digital signal processing*, volume 46. Macmillan New York, 1992.
- [167] M. Remondes and E.M. Schuman. Direct cortical input modulates plasticity and spiking in CA1 pyramidal neurons. *Nature*, 416(6882):736–740, 2002.
- [168] M. Remondes and E.M. Schuman. Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory. *Nature*, 431(7009):699–703, 2004.
- [169] A. Rodríguez-Moreno, O. Herreras, and J. Lerma. Kainate receptors presynaptically downregulate GABAergic inhibition in the rat hippocampus. *Neuron*, 19(4):893–902, 1997.
- [170] P. Room, F.T. Russchen, H.J. Groenewegen, and A.H.M. Lohman. Efferent connections of the prelimbic (area 32) and the infralimbic (area 25) cortices: an anterograde tracing study in the cat. *The Journal of Comparative Neurology*, 242(1):40–55, 1985.
- [171] U. Rudolph and B. Antkowiak. Molecular and neuronal substrates for general anaesthetics. *Nature Reviews Neuroscience*, 5(9):709–720, 2004.
- [172] Y. Saito and H. Harashima. Tracking of Information within Multichannel {EEG} record Causal analysis in {EEG}. *Recent advances in {EEG} and {EMG} data processing*, pages 133–146, 1981.
- [173] P.A. Salin and D.A. Prince. Electrophysiological mapping of GABAA receptor-mediated inhibition in adult rat somatosensory cortex. *Journal of Neurophysiology*, 75(4):1589–1600, 1996.
- [174] K. Sameshima and L.A. Baccalá. Using partial directed coherence to describe neuronal ensemble interactions. *Journal of Neuroscience Methods*, 94(1):93–103, 1999.

- [175] M.V. Sanchez-Vives and D.A. McCormick. Cellular and network mechanisms of rhythmic recurrent activity in neocortex. *Nature neuroscience*, 3(10):1027–1034, 2000.
- [176] J.R. Sato, D.Y. Takahashi, S.M. Arcuri, K. Sameshima, P.A. Morettin, and L.A. Baccalá. Frequency domain connectivity identification: An application of partial directed coherence in fMRI. *Human Brain Mapping*, 30(2):452–461, 2009.
- [177] R.J. Sayer, M.J. Friedlander, and S.J. Redman. The time course and amplitude of EPSPs evoked at synapses between pairs of CA3/CA1 neurons in the hippocampal slice. *Journal of Neuroscience*, 10(3):826–836, 1990.
- [178] B. Schelter, M. Winterhalder, M. Eichler, M. Peifer, B. Hellwig, B. Guschlbauer, C.H. Lücking, R. Dahlhaus, and J. Timmer. Testing for directed influences among neural signals using partial directed coherence. *Journal of Neuroscience Methods*, 152(1-2):210–219, 2005.
- [179] B. Schelter, M. Winterhalder, B. Hellwig, B. Guschlbauer, C.H. Lücking, and J. Timmer. Direct or indirect? Graphical models for neural oscillators. *Journal of Physiology-Paris*, 99(1):37–46, 2006.
- [180] D. Schmitz, S. Schuchmann, A. Fisahn, A. Draguhn, E.H. Buhl, E. Petrasch-Parwez, R. Dermietzel, U. Heinemann, and R.D. Traub. Axo-axonal coupling a novel mechanism for ultrafast neuronal communication. *Neuron*, 31(5):831–840, 2001.
- [181] T. Schneider and A. Neumaier. Algorithm 808: ARfitA Matlab package for the estimation of parameters and eigenmodes of multivariate autoregressive models. *ACM Transactions on Mathematical Software (TOMS)*, 27(1):58–65, 2001.
- [182] S.M. Schnider, R.H. Kwong, F.A. Lenz, and H.C. Kwan. Detection of feedback in the central nervous system using system identification techniques. *Biological Cybernetics*, 60(3):203–212, 1989.
- [183] G. Schwarz. Estimating the dimension of a model. *Annals of Statistics*, 6(2):461–464, 1978.
- [184] S.R. Sesack, A.Y. Deutch, R.H. Roth, and B.S. Bunney. Topographical organization of the efferent projections of the medial prefrontal cortex in the rat:

- an anterograde tract-tracing study with Phaseolus vulgaris leucoagglutinin. *The Journal of Comparative Neurology*, 290(2):213–242, 1989.
- [185] A.G. Siapas and M.A. Wilson. Coordinated interactions between hippocampal ripples and cortical spindles during slow-wave sleep. *Neuron*, 21(5):1123–1128, 1998.
- [186] A. Sik, M. Penttonen, A. Ylinen, and G. Buzsáki. Hippocampal CA1 interneurons: an in vivo intracellular labeling study. *The Journal of neuroscience*, 15(10):6651–6665, 1995.
- [187] A. Sik, N. Tamamaki, and T.F. Freund. Complete Axon Arborization of a Single CA3 Pyramidal Cell in the Rat Hippocampus, and its Relationship With Postsynaptic Parvalbumin-containing Interneurons. *European Journal of Neuroscience*, 5(12):1719–1728, 1993.
- [188] A. Sirota, J. Csicsvari, D. Buhl, and G. Buzsaki. Communication between neocortex and hippocampus during sleep in rodents. *Proceedings of the National Academy of Sciences*, 100(4):2065–2069, 2003.
- [189] W.E. Skaggs and B.L. McNaughton. Replay of neuronal firing sequences in rat hippocampus during sleep following spatial experience. *Science*, 271(5257):1870–1873, 1996.
- [190] K.E. Sorra and K.M. Harris. Occurrence and three-dimensional structure of multiple synapses between individual radiatum axons and their target pyramidal cells in hippocampal area CA1. *Journal of Neuroscience*, 13(9):3736–3748, 1993.
- [191] M. Steriade. *Neuronal substrates of sleep and epilepsy*. Cambridge University Press, 2003.
- [192] M. Steriade. Grouping of brain rhythms in corticothalamic systems. *Neuroscience*, 137(4):1087–1106, 2006.
- [193] M. Steriade, D. Contreras, R. Curro Dossi, and A. Nunez. The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. *Journal of Neuroscience*, 13(8):3284–3299, 1993.
- [194] M. Steriade, DA McCormick, and TJ Sejnowski. Thalamocortical oscillations in the sleeping and aroused brain. *Science*, 262(5134):679, 1993.

- [195] M. Steriade, A. Nunez, and F. Amzica. A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. *Journal of Neuroscience*, 13(8):3252–3265, 1993.
- [196] M. Steriade, A. Nunez, and F. Amzica. Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. *Journal of Neuroscience*, 13(8):3266–3283, 1993.
- [197] P. Stoica and R.L. Moses. *Introduction to Spectral Analysis*. Prentice Hall Upper Saddle River, NJ, 1997.
- [198] G.G. Supp, A. Schlögl, N. Trujillo-Barreto, M.M. Müller, and T. Gruber. Directed cortical information flow during human object recognition: analyzing induced EEG gamma-band responses in brain’s source space. *PLoS One*, 2(8):1169–1177, 2007.
- [199] L.W. Swanson. A direct projection from Ammon’s horn to prefrontal cortex in the rat. *Brain Research*, 217(1):150–154, 1981.
- [200] M. Takagishi and T. Chiba. Efferent projections of the infralimbic (area 25) region of the medial prefrontal cortex in the rat: an anterograde tracer PHA-L study. *Brain Research*, 566(1-2):26–39, 1991.
- [201] D.Y. Takahashi, L.A. Baccala, and K. Sameshima. Inference between neural structures via partial directed coherence. *Journal of Applied Statistics*, 34(10):1255–1269, 2007.
- [202] J. Taxidis, B. Coomber, R. Mason, and M. Owen. Assessing cortico-hippocampal functional connectivity under anesthesia and kainic acid using generalized partial directed coherence. *Biological Cybernetics*, 102(4):327–340, 2010.
- [203] J. Taxidis, S. Coombes, R. Mason, and M.R. Owen. Modeling Sharp Wave-Ripple Complexes Through a CA3-CA1 Network Model with Chemical Synapses. *Hippocampus*, doi: 10.1002/hipo.20930.
- [204] P.L. Tierney, E. Dégenétais, A.M. Thierry, J. Glowinski, and Y. Gioanni. Influence of the hippocampus on interneurons of the rat prefrontal cortex. *European Journal of Neuroscience*, 20(2):514–524, 2004.

- [205] I. Timofeev, F. Grenier, M. Bazhenov, T.J. Sejnowski, and M. Steriade. Origin of slow cortical oscillations in deafferented cortical slabs. *Cerebral Cortex*, 10(12):1185–1199, 2000.
- [206] I. Timofeev, F. Grenier, and M. Steriade. Impact of intrinsic properties and synaptic factors on the activity of neocortical networks in vivo. *Journal of Physiology-Paris*, 94(5-6):343–355, 2000.
- [207] I. Timofeev, F. Grenier, and M. Steriade. Disfacilitation and active inhibition in the neocortex during the natural sleep-wake cycle: an intracellular study. *Proceedings of the National Academy of Sciences of the United States of America*, 98(4):1924–1929, 2001.
- [208] I. Timofeev and M. Steriade. Low-frequency rhythms in the thalamus of intact-cortex and decorticated cats. *Journal of Neurophysiology*, 76(6):4152–4168, 1996.
- [209] G. Tononi, M. Massimini, and B.A. Riedner. Sleepy Dialogues between Cortex and Hippocampus: Who Talks to Whom? *Neuron*, 52(5):748–749, 2006.
- [210] R.D. Traub. Simulation of intrinsic bursting in CA3 hippocampal neurons. *Neuroscience*, 7(5):1233–1242, 1982.
- [211] R.D. Traub and A. Bibbig. A model of high-frequency ripples in the hippocampus based on synaptic coupling plus axon-axon gap junctions between pyramidal neurons. *Journal of Neuroscience*, 20(6):2086–2093, 2000.
- [212] R.D. Traub, J.G. Jefferys, R. Miles, M.A. Whittington, and K. Toth. A branching dendritic model of a rodent CA3 pyramidal neurone. *The Journal of Physiology*, 481(Pt 1):79–95, 1994.
- [213] R.D. Traub, J.G.R. Jefferys, and M.A. Whittington. *Fast oscillations in cortical circuits*. The MIT Press, 1999.
- [214] R.D. Traub and R. Miles. *Neuronal networks of the hippocampus*. Cambridge Univ Pr, 1991.
- [215] R.D. Traub, R. Miles, and G. Buzsaki. Computer simulation of carbachol-driven rhythmic population oscillations in the CA3 region of the in vitro rat hippocampus. *The Journal of Physiology*, 451(1):653–672, 1992.

- [216] R.D. Traub, R. Miles, and R.K. Wong. Model of the origin of rhythmic population oscillations in the hippocampal slice. *Science*, 243(4896):1319–1325, 1989.
- [217] R.D. Traub, D. Schmitz, J.G.R. Jefferys, and A. Draguhn. High-frequency population oscillations are predicted to occur in hippocampal pyramidal neuronal networks interconnected by axoaxonal gap junctions. *Neuroscience*, 92(2):407–426, 1999.
- [218] R.D. Traub, M.A. Whittington, E.H. Buhl, J.G.R. Jefferys, and H.J. Faulkner. On the mechanism of the gamma right-arrow beta frequency shift in neuronal oscillations induced in rat hippocampal slices by tetanic stimulation. *Journal of Neuroscience*, 19(3):1088–1105, 1999.
- [219] R.D. Traub, M.A. Whittington, S.B. Colling, G. Buzsaki, and J.G. Jefferys. Analysis of gamma rhythms in the rat hippocampus in vitro and in vivo. *The Journal of Physiology*, 493(2):471–484, 1996.
- [220] R.D. Traub, M.A. Whittington, I.M. Stanford, and J.G. Jefferys. A mechanism for generation of long-range synchronous fast oscillations in the cortex. *Nature*, 383(6601):621–624, 1996.
- [221] R.D. Traub, R.K. Wong, R. Miles, and H. Michelson. A model of a CA3 hippocampal pyramidal neuron incorporating voltage-clamp data on intrinsic conductances. *Journal of Neurophysiology*, 66(2):635–650, 1991.
- [222] R.P. Vertes. Interactions among the medial prefrontal cortex, hippocampus and midline thalamus in emotional and cognitive processing in the rat. *Neuroscience*, 142(1):1–20, 2006.
- [223] R.P. Vertes, W.B. Hoover, K. Szigeti-Buck, and C. Leranth. Nucleus reuniens of the midline thalamus: Link between the medial prefrontal cortex and the hippocampus. *Brain Research Bulletin*, 71(6):601–609, 2007.
- [224] J. Villablanca and M.E. Salinas-Zeballos. Sleep-wakefulness, EEG and behavioral studies of chronic cats without the thalamus: the 'athalamic' cat. *Archives italiennes de biologie*, 110(3):383–411, 1972.
- [225] C. Vreeswijk and D. Hansel. Patterns of synchrony in neural networks with spike adaptation. *Neural Computation*, 13(5):959–992, 2001.

- [226] J.Y. Wang, J.Y. Chang, D.J. Woodward, L.A. Baccala, J.S. Han, and F. Luo. Corticofugal influences on thalamic neurons during nociceptive transmission in awake rats. *Synapse*, 61(5):335–42, 2007.
- [227] J.Y. Wang, H.T. Zhang, J.Y. Chang, D.J. Woodward, L.A. Baccalá, and F. Luo. Anticipation of pain enhances the nociceptive transmission and functional connectivity within pain network in rats. *Molecular Pain*, 4(1):34, 2008.
- [228] X.J. Wang and G. Buzsaki. Gamma oscillation by synaptic inhibition in a hippocampal interneuronal network model. *Journal of Neuroscience*, 16(20):6402–6413, 1996.
- [229] P. Welch. The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *Audio and Electroacoustics, IEEE Transactions on*, 15(2):70–73, 1967.
- [230] J.A. White, C.C. Chow, J. Rit, C. Soto-Trevino, and N. Kopell. Synchronization and oscillatory dynamics in heterogeneous, mutually inhibited neurons. *Journal of Computational Neuroscience*, 5(1):5–16, 1998.
- [231] M.A. Whittington, I.M. Stanford, S.B. Colling, J.G. Jefferys, and R.D. Traub. Spatiotemporal patterns of gamma frequency oscillations tetanically induced in the rat hippocampal slice. *The Journal of Physiology*, 502(3):591–607, 1997.
- [232] M.A. Whittington, R.D. Traub, H.J. Faulkner, J.G.R. Jefferys, and K. Chettiar. Morphine disrupts long-range synchrony of gamma oscillations in hippocampal slices. *Proceedings of the National Academy of Sciences of the United States of America*, 95(10):5807–5811, 1998.
- [233] M.A. Whittington, R.D. Traub, and J.G.R. Jefferys. Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. *Nature*, 373(6515):612–615, 1995.
- [234] C.M. Wierzynski, E.V. Lubenov, M. Gu, and A.G. Siapas. State-Dependent Spike-Timing Relationships between Hippocampal and Prefrontal Circuits during Sleep. *Neuron*, 61(4):587–596, 2009.
- [235] M.A. Wilson and B.L. McNaughton. Reactivation of hippocampal ensemble memories during sleep. *Science*, 265(5172):676–679, 1994.

- [236] M. Winterhalder, B. Schelter, W. Hesse, K. Schwab, L. Leistritz, D. Klan, R. Bauer, J. Timmer, and H. Witte. Comparison of linear signal processing techniques to infer directed interactions in multivariate neural systems. *Signal Processing*, 85(11):2137–2160, 2005.
- [237] T. Wolansky, E.A. Clement, S.R. Peters, M.A. Palczak, and C.T. Dickson. Hippocampal slow oscillation: a novel EEG state and its coordination with ongoing neocortical activity. *Journal of Neuroscience*, 26(23):6213–6229, 2006.
- [238] C. Wu, M.N. Asl, J. Gillis, F.K. Skinner, and L. Zhang. An in vitro model of hippocampal sharp waves: regional initiation and intracellular correlates. *Journal of Neurophysiology*, 94(1):741–753, 2005.
- [239] C. Yamamoto, M. Higashima, and S. Sawada. Quantal analysis of potentiating action of phorbol ester on synaptic transmission in the hippocampus. *Neuroscience Research*, 5(1):28–38, 1987.
- [240] H. Yang, J.Y. Chang, D.J. Woodward, L.A. Baccalá, J.S. Han, and F. Luo. Coding of peripheral electrical stimulation frequency in thalamocortical pathways. *Experimental Neurology*, 196(1):138–152, 2005.
- [241] A. Ylinen, A. Bragin, Z. Nadasdy, G. Jando, I. Szabo, A. Sik, and G. Buzsaki. Sharp wave-associated high-frequency oscillation (200 Hz) in the intact hippocampus: network and intracellular mechanisms. *Journal of Neuroscience*, 15(1):30–46, 1995.