TU Graz

Graz University of Technology

DISSERTATION

to obtain the title of

Doctor of Technical Sciences

Defended by

Gregor Michael Hörzer

Analysis of Neural Data and Models of Neural Networks related to Working Memory

Thesis Advisor: Univ.-Prof. Dr. Wolfgang Maass

defended on May 27, 2011

Jury :

Advisor:	Univ.-Prof. Dr. Wolfgang Maass	-	TU Graz
Reviewer:	Doz. Dr. Alois Schlögl	-	IST Austria
Dean of Studies:	Univ.-Prof. Dr. Horst Bischof	-	TU Graz

Eidesstattliche Erklärung

Ich erkläre an Eides statt, dass ich die vorliegende Arbeit selbständig verfasst, andere als die angegebenen Quellen / Hilfsmittel nicht benutzt, und die den benutzten Quellen wörtlich und inhaltlich entnommene Stellen als solche kenntlich gemacht habe.

Statutory Declaration

I declare that I have authored this thesis independently, that I have not used other than the declared sources / resources, and that I have explicitly marked all material which has been quoted either literally or by content from the used sources.

Graz, May 2011   ....
                                                              (signature)

Abstract

The ability to store information and to retain this information at a later point in time is probably one of the most essential functions of one of the most fascinating systems evolution has developed - the brain. Working memory is used to maintain information for relatively brief time intervals, usually on the time scale of up to several seconds. This thesis investigates several aspects of working memory and provides new insights to the formation of memory traces in the brain. Two approaches to this topic are covered by the present work. The first approach studies the phenomenon of working memory formation from the perspective of electrophysiological data analysis and signal processing, while in the second approach working memory traces and memory-dependent computations are generated through reward-modulated learning in a computational model of a recurrent neural network.

The first part, comprised of two chapters, reveals local interaction patterns between neural populations within extrastriate visual area V4 as well as long-range interactions between two distant cortical areas, V4 and the lateral prefrontal cortex (lPF) of macaque monkeys that perform a visual short-term memory task. The analysis is based on simultaneous recordings of local field potentials (LFP) and spiking activity of single units obtained from extracellular recordings in the awake and behaving monkey, provided by my collaborators from the MPI for Biological Cybernetics in Tübingen. Within V4, analysis with multivariate autoregressive models reveals new insights into the patterns of directed information flow between neural populations on the level of the local field potential which are most prominent in the theta frequency band, and moreover shows that these interaction patterns are a rather local phenomenon. Between V4 and lPF, results from Wavelet-based methods for phase synchronization analysis suggest that the synchronization of oscillatory activity in the theta range between these distant cortical sites is likely to provide the basis for the coordination of spiking activity in both areas during the memory phase of the task.

The second part tries to extend and modify previous results from the field of reservoir computing and provides experimental evidence for the ability of a rate-based recurrent neural network with trained readout units to learn to produce coherent patterns of activity, memory traces and to carry out memory-dependent computations by employing a purely local reward-modulated Hebbian learning rule. In contrast to the traditionally used fully supervised methods, learning in the proposed model is solely based on correlations between the presynaptic activity and postsynaptic noise perturbations, modulated by a global binary signal that provides the system with information whether the overall system performance has recently increased. In this way, the present results provide a new perspective for the emergence of complex computations through learning in biological neural systems.

Keywords: Neural data analysis, oscillatory synchrony, autoregressive models, computational neuroscience, reservoir computing, reward-modulated learning

Zusammenfassung

Die Fähigkeit, Informationen zu speichern und diese zu einem späteren Zeitpunkt wieder abzurufen, ist wahrscheinlich eine der entscheidendsten Funktionen eines der wohl faszinierendsten Systeme, die die Evolution hervorgebracht hat - des Gehirns. Das Arbeitsgedächtnis wird benötigt, um Informationen für relativ kurze Zeitintervalle in Erinnerung zu behalten, üblicherweise auf einer Zeitskala von einigen Sekunden. Diese Dissertation untersucht verschiedene Aspekte des Arbeitsgedächtnisses und liefert neue Einsichten in die Entstehung von Gedächtnisspuren im Gehirn. Zwei Ansätze zu diesem Thema werden in der vorliegenden Arbeit behandelt. Der erste Ansatz untersucht das Phänomen der Entstehung von Arbeitsgedächtnis aus der Perspektive der Analyse elektrophysiologischer Daten und der Signalverarbeitung, während im zweiten Ansatz Gedächtnisspuren und gedächtnisabhängige Berechnungen durch belohnungsmoduliertes Lernen in einem Computermodell eines rekurrenten neuralen Netzes erzeugt werden.

Der erste Teil, bestehend aus zwei Kapiteln, deckt lokale Interaktionsmuster zwischen neuralen Populationen im extrastriären visuellen Areal V4 sowie Interaktionen zwischen zwei voneinander weit entfernten kortikalen Arealen, V4 und dem lateralen präfrontalen Kortex (lPF) von Rhesusaffen, die eine visuelle Kurzzeitgedächtnis-Aufgabe durchführen, auf. Die Analyse basiert auf simultanen Aufzeichnungen von lokalen Feldpotentialen (LFP) und Spike-Aktivität einzelner neuraler Einheiten, die mittels extrazellulärer Messungen im wachen Affen erzeugt wurden. Diese Daten wurden von meinen Kollegen vom MPI für Biologische Kybernetik in Tübingen zur Verfügung gestellt. Innerhalb von V4 liefert die Analyse mittels multivariaten autoregressiven Modellen neue Einsichten in die Muster gerichteten Informationsflusses zwischen neuralen Populationen auf der Ebene des LFP, die am stärksten im Theta-Frequenzbereich ausgeprägt sind. Darüber hinaus wird gezeigt, dass es sich dabei um ein eher lokales Phänomen handelt. Zwischen V4 und lPF legen die Resultate aus Analysen mittels Wavelet-basierten Methoden nahe, dass die Synchronisation von oszillatorischer Aktivität im Theta-Band zwischen diesen voneinander weit entfernt liegenden kortikalen Arealen vermutlich die Basis für die Koordination von Spike-Aktivität in den beiden Arealen während der Gedächtnisphase der Aufgabe bietet.

Der zweite Teil versucht, frühere Resultate aus dem Bereich des Reservoir Computing zu erweitern und zu modifizieren und experimentelle Evidenz dafür zu liefern, dass ein ratenbasiertes rekurrentes neuronales Netzwerk mit trainierten Readout-Einheiten unter Zuhilfenahme einer rein lokalen belohnungsmodulierten Hebb'schen Lernregel fähig ist, kohärente Aktivitätsmuster und Gedächtnisspuren zu erzeugen und gedächtnisabhängige Berechnungen durchzuführen. Im Unterschied zu traditionell benutzen vollständig überwachten Verfahren basiert Lernen im vorgeschlagenen Modell rein auf Korrelationen zwischen präsynaptischer Aktivität und postsynaptischem Rauschen, moduliert durch ein globales binäres Signal, welches das System mit Information darüber versorgt, ob die Gesamtleistung des Systems kürzlich zugenommen hat. Auf diese Weise liefern die vorliegenden Resultate eine neue

Perspektive für die Entstehung von komplexen Berechnungen durch Lernen in biologischen neuronalen Systemen.

Schlüsselwörter: Neurale Datenanalyse, oszillatorische Synchronizität, autoregressive Modelle, computerorientierte Neurowissenschaften, Reservoir Computing, belohnungsmoduliertes Lernen

Acknowledgements

First of all, I want to thank my advisor Wolfgang Maass for his support, for offering me the great opportunity to join the neuroscience community and for enabling me to work on very interesting projects. Moreover, I am very grateful to my other co-authors, Stefanie Liebe, Gregor Rainer and Nikos Logothetis from the Max Planck Institute for Biological Cybernetics in Tübingen, as well as to Robert Legenstein and Alois Schlögl from TU Graz and IST Austria for very fruitful collaborations and inspiring discussions. I would also like to thank Alois Schlögl for agreeing to review my thesis and to attend my defense.

I would like to show my gratitude to my colleagues at the Institute for Theoretical Computer Science, who always used their best endeavors to support me where necessary and made my stay here much more exciting. This includes colleagues from both the scientific and the administrative staff. I would also like to acknowledge that this thesis would not have been possible without the financial support by the research programs of the European Union and the Austrian Science Fund FWF.

I would like to thank my family and friends, who make life much more enjoyable. At last, my deepest gratitude goes to my parents for their unconditional love and for the endless encouragement they have provided to me throughout my entire life.

Contents

Introduction

> Memory is the mother of all wisdom.
> — *Aeschylus, 5th century BC*

As Aeschylus already expressed with the above quote in ancient Greece, the ability to retain information and to reconstruct this information at a later point in time is a major component of our personhood, even though we usually do not notice its substantial influence in everyday life unless we explicitly reflect on it. Memory is involved in virtually every act that we perform. For example, we use memory of the temporal patterns of muscle activations in order to perform specific motions, memory of the consequences of our previous actions in particular contexts in order to choose future actions, memory of consistent observations that we made previously to infer general principles, and we construct entire worlds based on our memories while dreaming. Since memory is so fundamental for us, the question of how the nervous system is able to store information in a way that allows us to selectively access it at a later point in time has attracted a large number of scientists in various disciplines, from philosophers via biologists and psychologists through to computer scientists.

Some fundamental principles that form the basis of our contemporary understanding of the mechanisms that underlie learning and memory have been proposed by Donald Hebb. He concluded from his studies of the impact of brain tissue removal on memory and intelligence that the neural substrate of memory has to be widely distributed in the brain rather than strictly localized because intelligence was only influenced little even if large amounts of brain tissue were removed (Klein, 1999). The same conclusion had been drawn earlier based on animal lesion studies by his former PhD advisor, Karl Lashley (Lashley, 1929).

A question of particular interest within the scope of this thesis that is raised by the above conclusion is how spatially distant neural populations that are involved in memory processes are able to communicate with each other. Since the German researcher Hans Berger provided the first evidence for oscillatory processes in human electroencephalogram recordings (Berger, 1929), researchers have been interested in the function of oscillatory processes in the brain. Oscillatory processes have been related to memory in many studies, both in animals and humans. Analysis of electrophysiological data of intracortically recorded local field potentials (LFP) as well as noninvasive electroencephalogram (EEG) and magnetoencephalogam (MEG) recordings revealed two classes of memory-related oscillations,

low frequency theta oscillations and high frequency gamma oscillations. For example, it has been found that spiking activity in prefrontal cortex of rodents is synchronized with theta oscillatory activity in the Hippocampus during spatial memory (Siapas et al., 2005; Jones and Wilson, 2005; Buzsaki, 2006), suggesting that these low frequency oscillations are used to coordinate the mediation of information between distant cortical sites. A recent review of the relation of brain oscillations and memory is provided in Duezel et al. (2010). Also, the introductory sections of chapters 2 and 3, which investigate interaction patterns within extrastriate visual area V4 and between V4 and lateral prefrontal cortex during visual short-term memory, respectively, provide an overview on the relations of theta oscillatory activity and memory in the recent literature.

Besides his and Lashley's conclusion about the distributed nature of intelligence and memory, Donald Hebb also postulated a widely influential concept in his book "The organization of behavior" in 1949, known as "Hebbian learning". More specifically, Hebb postulated the ability of synaptic connections to increase their efficacy if the activity of a pre- and a postsynaptic neuron is correlated. In other words, Hebb's postulate states that the strength of a synaptic connection between two neurons is increased if the pre- and the postsynaptic neuron repeatedly fire together (Hebb, 1949).

In computational neuroscience, the idea of Hebbian learning is reflected in a large set of learning rules that have been developed for rate-based as well as for spiking neuron models and modify and extend the basic concept. One prominent example of such an extension that is a topic of current research is the concept of spike timing dependent plasticity (STDP). Assuming that an asymmetric STDP function like the one in Bi and Poo (1998) is used, the synaptic connection between a pre- and a postsynaptic neuron is strengthened if the postsynaptic neuron repeatedly fires shortly after the presynaptic neuron, resulting in long term potentiation (LTP). On the other hand, if the postsynaptic neuron repeatedly fires shortly before the presynaptic neuron, the synaptic connection is weakened, resulting in long term depression (LTD). A recent review of STDP and the experimental evidence for its implementation in biological nervous systems can be found in Caporale and Dan (2008), and a theory linking STDP and working memory has been provided in Szatmary and Izhikevich (2010).

In addition to unsupervised learning rules, extensions have been developed in which the learning process is modulated by external reward signals. For STDP, such a learning rule has been proposed for example by Izhikevich (2007) and further investigated in subsequent studies (Legenstein et al., 2008; Fremaux et al., 2010). Another reward modulated learning mechanism that is based on Hebbian learning has been developed in the lab of Sebastian Seung (Xie and Seung, 2004; Fiete and Seung, 2006) and successfully applied in a model of birdsong learning (Fiete et al., 2007). In their approach, instead of pairing pre- and postsynaptic spikes that occur with a small temporal delay to each other, presynaptic spikes are paired with simultaneously occuring spikes from a noise source that spikes at random points in time. These noise spikes are also projected to the postsynaptic neuron, and

whenever a positive reward is given in response to a noise spike, the projections from the presynaptic neurons that spiked together with the noise input are strengthened. Similarly, the projections from presynaptic neurons that spiked in absence of a simultaneous spike in the noise input are weakened if a negative reward is given in response. The learning rule I employ in chapter 4 is similar to this kind of learning rule, but instead of employing a spiking neural network, I focus on reward-modulated Hebbian learning of memory-related computational functions in a rate-based recurrent neural network model which is based on the reservoir computing paradigm (Maass et al., 2002; Jaeger, 2003) and its extensions that incorporate feedback projections from trained readout units into the recurrent neural network and let them contribute substantially to the recurrent network dynamics (Jaeger and Haas, 2004; Maass et al., 2007; Sussillo and Abbott, 2009).

While we have substantial evidence nowadays that the brain stores information in a highly distributed way, involving plasticity mechanisms that regulate the efficacy of the synaptic connections, and a vast amount of both experimental and computational neuroscience studies that try to resolve some aspects of the puzzle, the exact mechnisms that underlie the storage and retrieval of memory contents still remain a mystery.

In this thesis, a few pieces are added to the puzzle. The work presented in the subsequent chapters considers some aspects of memory formation, both from the perspective of electrophysiological data analysis (chapters 2 and 3) and computational modelling (chapter 4). In this work, I focus on working memory, i.e. memory that maintains information for relatively brief time intervals up to a few seconds (Purves et al., 2008). In the literature, three types of memory are distinguished from each other: long-term memory, short-term memory and working memory. While all the definitions of these three types of memory are still controversial, the distinction between short-term and working memory is probably the most subtle one. A recent attempt to clarify the differences has been conducted by Cowan (2008), who characterizes short-term memory by temporal decay and chunk capacity limits, while subsuming short-term memory and the mechanisms to process short-term memory in order to make use of it in a particular context under the more general term working memory. However, in the work presented here, I do not rigorously distinguish between the terms "working memory" and "short-term memory".

1.1 Organization of the Thesis

This thesis is comprised of this introduction and three chapters which are based on papers to which I contributed during my PhD studies. Chapters 2 and 3 focus on results which I obtained by analyzing electrophysiological data and were carried out in collaboration with the Max Planck Institute for Biological Cybernetics in Tübingen, Germany, where the monkey experiments and electrophysiological recordings were designed and performed by my collaborators, Gregor Rainer and Stefanie Liebe. In contrast to these chapters, chapter 4 focuses on a novel mech-

anism for reward-modulated learning in computational models of recurrent neural networks and its application to computational tasks that involve working memory. Note that since the papers on which the subsequent chapters are based were conducted in collaboration with others, my contributions to each paper are indicated at the beginning of each chapter.

In chapter 2, I explore functional coupling patterns in the extrastriate visual cortical area V4 of macaque monkeys performing a delayed matching to sample (DMS) task in which they had to discriminate whether two visual stimuli presented briefly with a delay of approximately 1.5s between them were identical or not. During the delay period, the monkeys therefore had to retain a memory trace of the first stimulus in order to be able to compare it to the second one. In the analysis, I compared the functional coupling between the activities of neural populations, reflected in the local field potential (LFP) recorded simultaneously at several parts of V4, during this delay period to the period before the presentation of the sample stimulus. In order to gain insights into the specific functional connectivity patterns that arise during short-term memory, I fitted multivariate autoregressive models to the simultaneously recorded data. Based on these models, I analyzed the interactions between different LFP traces using coupling measures that are based on the model parameters and build upon the concept of "Granger causality" (Granger, 1969). These coupling measures reveal transient interactions between neural populations by means of oscillatory activity in the frequency domain. I found that these interactions during visual short-term memory predominantly arise within the theta frequency band and are strongly directed rather than symmetric. Moreover, the strength of these interactions decreases with increasing distance between the recording sites. Taken together, the presented results give some important insights into the spatiotemporal connectivity patterns that arise between local populations of neurons in V4 during visual short-term memory.

In chapter 3, I further analyze interactions in electrophysiological data. Specifically, I explore the synchronization of local field potentials ("LFP phase locking") between two distant brain areas, namely the extrastriate visual area V4, which I also investigated in chapter 2, and the lateral prefrontal cortex (lPF) of macaque monkeys during a visual memory task. The task the monkey performed was the same DMS task as described in the previous paragraph. Moreover, I investigate coupling of the spiking activity from single units within both areas to the phase of the local field potentials in the respective other area ("spike phase locking"). Other than in the previous chapter, I use Wavelet-based methods to extract the instantaneous phase of the LFP oscillations, and employ several analysis techniques based on this phase information to measure long range interactions. The analysis reveals that the local field potential oscillations within the theta range exhibit an increase in phase locking between V4 and lPF during the delay period of the task that is accompanied by an increase in spike phase locking between these areas. Interestingly, the findings show that the degree of inter-cortical interactions in the theta range between both areas is correlated with task performance, suggesting that oscillatory synchrony on the level of the LFP provides the basis for the coordination of the

timing of spiking activity in both areas during visual short-term memory.

In contrast to chapters 2 and 3, I explore some aspects of working memory formation from a different perspective in chapter 4. In this chapter, I focus on computational models of recurrent networks of rate-based neurons with feedback from trained readouts. Modifying and extending previous results in the field of reservoir computing, I provide a novel approach for the training of recurrent networks in which the weights of the synaptic connections from the network to the readouts are adapted using a reward-modulated Hebbian learning rule. More precisely, the weight adaptations are based on the correlations of the presynaptic activity with random perturbations of the readout's output in the presence of a global modulatory signal that indicates in a binary fashion whether or not the performance of the system has recently improved due to these perturbations. In the simulations, I show that after an appropriate training time, one or several readouts are able to compute specific oscillatory patterns of activity, persistent memory of previously observed stimuli, and other memory-dependent computations. Therefore, reward-modulated Hebbian learning is a viable alternative to fully supervised learning methods which are traditionally employed in the reservoir computing approach that is based on biologically more plausible mechanisms.

Directed coupling in LFPs of macaque V4 during visual short-term memory

Contents

This chapter is based on the paper "Directed coupling in local field potentials of macaque V4 during visual short-term memory revealed by multivariate autoregressive models", which was published in the journal "Frontiers of Computational Neuroscience" in May 2010. The work was conducted in collaboration with Stefanie Liebe (SL), Gregor Rainer (GR) and Nikos Logothetis (NL), affiliated with the Max Planck Institute for Biological Cybernetics in Tübingen, Germany, and Alois Schlögl (AS), affiliated with the Institute for Human-Computer Interfaces at TU Graz, Austria, at the time the presented work was conducted. The monkey experiments and electrophysiological recordings were designed by SL and GR and conducted by SL. The data analysis was performed by myself and SL. Specifically, I provided the multivariate autoregressive modelling work including the coupling analysis and contributed to the statistical evaluation and the interpretation of the results. The paper was written by myself and SL, with additional input from AS, NL and GR. The results of this work, in which for the first time these analysis techniques are applied to simultaneously recorded V4 LFP data from the macaque monkey performing a visual memory task, reveal new insights into the interaction patterns of populations of neurons within extrastriate visual area V4 during visual short-term memory.

Processing and storage of sensory information is based on the interaction between different neural populations rather than the isolated activity of single neurons. In order to characterize the dynamic interaction and transient cooperation of subcircuits within a neural network, multivariate autoregressive (MVAR) models have proven to be an important analysis tool. In this study, I apply directed functional coupling based on MVAR models and describe the temporal and spatial changes of functional coupling between simultaneously recorded local field potentials (LFP) in extrastriate area V4 during visual memory. Specifically, I compare the strength and directional relations of coupling based on Generalized Partial Directed Coherence (GPDC) measures while two rhesus monkeys perform a visual short-term memory task. In both monkeys I find increases in theta power during the memory period that are accompanied by changes in directed coupling. These interactions are most prominent in the low frequency range encompassing the theta band (3-12 Hz) and, more importantly, are asymmetric between pairs of recording sites. Furthermore, it is shown that the degree of interaction decreases as a function of distance between electrode positions, suggesting that these interactions are a predominantly local phenomenon. Taken together, the results show that directed coupling measures based on MVAR models are able to provide important insights into the spatial and temporal formation of local functionally coupled ensembles during visual memory in V4. Moreover, the findings suggest that visual memory is accompanied not only by a temporary increase of oscillatory activity in the theta band, but by a direction-dependent change in theta coupling, which ultimately represents a change in functional connectivity within the neural circuit.

2.1   Introduction

Cortical oscillatory activity measured from local field potential recordings (LFP) or electroencephalogram (EEG) is a wide-spread neuronal phenomenon and is considered to underlie the communication of local and distant neural populations throughout the brain (Fries, 2005). Different parameters of oscillations in distinct frequency bands often show correlations with various aspects of sensory information processing (Buzsaki and Draguhn, 2004). A prominent example is the modulation of gamma synchrony in visual cognition, for example in tasks involving the manipulation of visual attention (Fries et al., 2001), binocular rivalry (Gail et al., 2004) or object recognition (Supp et al., 2007).

In contrast to visual processing, several studies revealed a specific role of theta oscillations (3-12 Hz) in mnemonic processing, for example in spatial memory in rodents (O'Keefe, 1993; Buzsaki, 2005), working memory in humans (Klimesch, 1999; Raghavachari et al., 2001, 2006) and visual short term memory in non-human primates (Rainer et al., 2004; Lee et al., 2005). In the latter study, neuronal oscillations in the theta band in extrastriate area V4 have been shown to mediate the coding and maintenance of relevant visual information within short-term memory. Thus, theta oscillations in V4 could provide a possible mechanism for supporting

and coordinating cross-neuronal interactions within neuronal ensembles during visual memory. However, physiological evidence for directed oscillatory interactions in the theta frequency range during short-term memory has not been obtained yet.

A description of the interaction patterns of oscillatory processes is provided by different measures that quantify various aspects of functional coupling. For example, some measures such as the phase locking value (Lachaux et al., 1999, 2000) provide insights into the instantaneous phase-relationship between two oscillatory processes and are derived from Wavelet- or Hilbert transform-based methods. In contrast, coupling measures derived from multivariate autoregressive (MVAR) models are becoming increasingly important as they capture not only instantaneous interactions between neural signals, but can give insights into the causal relationship between oscillations as well as the direction of their interaction. Thus, MVAR models are powerful in capturing the complex nature of oscillatory interactions and their role in neural processing.

MVAR models are a generalization of univariate autoregressive (AR) models, which were among the first methods that were applied to EEG data to reveal the spectral properties of brain signals already in the late 1960s (Zetterberg, 1969). MVAR models are able to take the interactions of multiple simultaneously recorded brain signals into account. A large set of coupling measures in the frequency domain such as coherency (Nunez et al., 1997, 1999), Directed Transfer Function (DTF, (Kaminski and Blinowska, 1991)) or Partial Directed Coherence (PDC, (Baccala and Sameshima, 2001)) as well as variants of these and similar measures can be derived using the MVAR model parameters (Schlögl and Supp, 2006; Porcaro et al., 2009) and the implementation of coupling analyses is readily achieved by various toolboxes (Schlögl and Brunner, 2008; Cui et al., 2008).

Importantly, DTF and PDC, unlike coherency, assess the directionality of couplings between signals, i.e. they measure the direction of information flow between different channels. Both measures are based on the concept of Granger causality (Granger, 1969), which can be informally stated as follows: If the observation of a time series $x(t)$ significantly improves the prediction of a time series $y(t)$, $x(t)$ "Granger-causes" $y(t)$. PDC differs from DTF by having the ability to reveal exclusively direct couplings, which means that it does not assess indirect couplings via intermediate sites. For example, if the model incorporates three observed channels, with a connection structure $A \rightarrow B \rightarrow C$, PDC is not expected to show a connection from A to C. It is important to note that Granger causality is not identical to physical causality, but is a statistical measure reflecting the improvement of predictability of one signal based on the information of another.

Previously, autoregressive models have been applied to EEG data and LFP data for various brain areas and frequency bands of interest and have revealed important insights into the functional relations between neuronal assemblies involved in sensorimotor behavior, sensory integration and visual attention (Supp et al., 2007; Bressler et al., 1999; Liang et al., 2000, 2001; Brovelli et al., 2004; Bressler et al., 2007; Chen et al., 2006; Liang et al., 2003; Kayser and Logothetis, 2009; Anderson et al., 2009).

In the present study, I applied multivariate autoregressive modelling to simultaneous LFP recordings from multiple electrodes in V4 while monkeys performed a visual identification task. MVAR models have been used previously to examine causal influences in area V4 in order to elucidate physiological mechanisms underlying neuronal oscillations in the alpha frequency range (i.e. 10-15Hz) (Bollimunta et al., 2008).

In this study, my goal was to exploit the advantages of MVAR models in order to investigate the directed functional relationship between multiple sources underlying theta oscillations during visual memory in V4. In order to gain insights into the direct interaction between multiple oscillatory components (i.e. bypassing coupling due to indirect influences) our MVAR models incorporated LFP activity of more than two simultaneously recorded channels. In addition, I evaluated the temporal and spatial dynamics of these direct interactions and provide a first description of causal and directed oscillatory coupling in the theta frequency range during visual memory.

2.2   Materials and Methods

In the following, I describe the procedure that was used for the analysis. Afterwards, the experimental procedures for the data aquisition are described.

2.2.1   Preprocessing

LFP data was preprocessed using standard techniques, as described for example in (Ding et al., 2000). First, I resampled the data to a frequency f_s of 200 Hz. This sampling rate is low enough to be able to use a sufficiently low MVAR model order while being high enough for an adequate representation of the frequency bands of interest. Then, I used a 50 Hz notch filter to suppress the electrical supply line noise. Afterwards, the data was normalized by subtracting the mean waveform across trials (grand-averaged mean waveform) from each single trial and subsequently dividing the result by the standard deviation across trials. This is necessary to remove first order instationarities from the data and to set the ensemble mean of the resulting data set to zero. I did not apply the same normalization procedure using the temporal mean and standard deviation for each separate trial, which is also frequently proposed, because this can lead to an underestimation of the low frequency components in which we were particularly interested.

2.2.2   Multivariate Autoregressive Modelling

To assess coupling between different LFP channels, I separately generated linear multivariate autoregressive (MVAR) models of the data for each recording session and each time interval of interest. The MVAR model can be expressed as

$$\underline{y}(t) = \sum_{p=1}^{P} \mathbf{A}_p \underline{y}(t-p) + \underline{x}(t).$$

The model tries to predict the data at sample t from a linear combination of the P previous samples of all M channels. Here, $\underline{y}(t)$ is the vector of M simultaneously observed LFP recordings, P is the model order stating the number of preceding samples that are used to predict the data at sample t, and the innovation process $\underline{x}(t)$ (sometimes addressed as the "residual error" or "prediction error", see (Supp et al., 2007; Schlögl, 2000) for comments) is assumed to be a multivariate white noise process and is equal to the difference between the model prediction and the actual data. In order to estimate the model parameter matrices $\mathbf{A}_p$ that weight the previous samples of the time series to predict the current one such that the mean quadratic error is minimized, I use the Burg-type method of Vieira-Morf (Marple, 1987) which, according to (Schlögl, 2006), is expected to provide the most accurate estimates of the model parameters. I used 250 ms windows for the time-frequency analysis, with an overlap of 200 ms for subsequent time intervals ((Ding et al., 2000) called this procedure an Adaptive MVAR or AMVAR approach), and 1 s windows for the assessment of statistical significance of coupling and change in coupling between the two investigated task conditions (cf. section on the experimental task). Note that the model assumes the data to be stationary, which is usually not the case for longer time segments of electrophysiological data, but for the short time intervals that are investigated in this study, the data is assumed to be quasi-stationary. I used the freely available open source Matlab implementation of the `BioSig` Toolbox for biomedical signal processing (Schlögl and Brunner, 2008) for the analysis, which can be found at `http://biosig.sf.net/`.

There exists a number of criteria for estimating the optimal model order for each data set such as the Akaike Information Criterion (AIC, (Akaike, 1974)) or Schwarz's Bayesian Information Criterion (BIC, (Schwarz, 1978)) which try to estimate the optimal model order for the MVAR model.

Both criteria take the goodness of fit to the empirical data into account, but also penalize for increasing numbers of free parameters to avoid overfitting to the data. Note that smaller values indicate better model orders. Unfortunately, the optimal model order is usually not consistent for different criteria and different data sets. I tried to estimate the optimal model order (in the range between 1 and 50, which reflects the length of the 250 ms windows we used for the time-frequency analysis) by using these measures, but the results did not show consistent local minima and qualitatively decreased with increasing model order instead (see Figure 2.1).

I compared models of order 20 and 40 for the 250 ms windows and found the resulting average power spectra and couplings to be qualitatively consistent. Therefore, I used a model order P of 20 for every data set, which corresponds to a time window of 100 ms given the sampling frequency of 200 Hz. This model order reflects a tradeoff between spectral resolution (specifically, I make clear that the model order does not determine the spectral resolution, which is in fact infinite, but instead
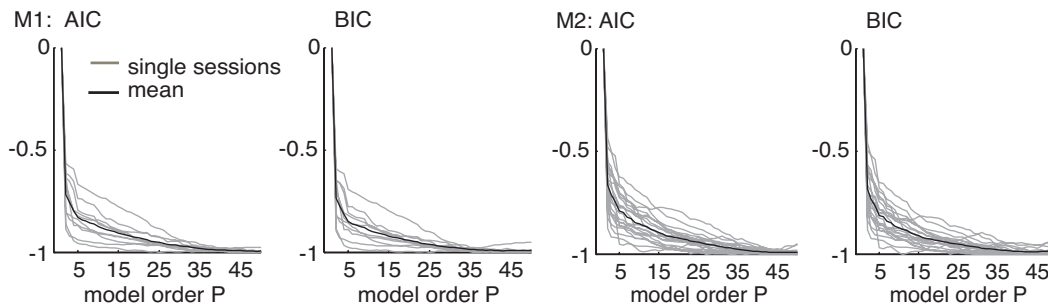
Figure 2.1: Evaluation of model orders using Akaike Information Criterion (AIC) and Schwarz's Bayesian Information Criterion (BIC), normalized between maxima and minima of each session for the 1s data from the delay condition. Gray lines indicate single sessions, black lines correspond to the average over all sessions. Criteria did not show consistent local minima, but qualitatively decreased with increasing model order up to P=50. Smaller values indicate better model orders. For the data from both animals and all sessions, the model order P=20 was chosen as a tradeoff between frequency resolution and overparametrization.

it determines the number of observed frequency components for each pair of channels, which is $P/2$, and relates to the "frequency resolution" in this sense (Schlögl and Supp, 2006)) and overparametrization and approximately corresponds to the model orders used in similar approaches. For example, (Brovelli et al., 2004) used a model order of 10 (corresponding to a 50 ms window) for analyzing beta oscillations, (Supp et al., 2007) revealed couplings in the gamma frequency range using a model order of 15 (30 ms), and (Kayser and Logothetis, 2009) and (Anderson et al., 2009) studied oscillations including the theta range using model orders of 6 (60 ms) and 17 (85 ms), respectively. Additionally, this model order fulfills all the requirements stated in (Schlögl and Supp, 2006) to obtain a sufficient model of the data. Furthermore, one should note that slight changes in the model order do not lead to arbitrarily large changes in the prediction error, but it is still an important parameter for the correct estimation of the couplings (Schlögl et al., 2000).

2.2.3   Generalized Partial Directed Coherence (GPDC)

As mentioned earlier, I used Generalized Partial Directed Coherence (GPDC, (Baccala et al., 2007)) for the analysis, which is a slightly adapted version of PDC with better variance stabilization properties. Analysis of the validity of this coupling measure using simulated and real data for which the ground truth is known as well as a comparison to DTF and other measures can be found elsewhere (Baccala and Sameshima, 2001; Kus et al., 2004; Pereda et al., 2005; Gourevitch et al., 2006; Porcaro et al., 2009). Moreover, (Porcaro et al., 2009) indicated that PDC is the most suitable method for this kind of analysis based on their results on MEG data.

GPDC is derived by first transforming the MVAR model from the time domain into the frequency domain to obtain the frequency representation of the model

parameters

$$\mathbf{A}(f) = I - \sum_{p=1}^{P} \mathbf{A}_p e^{-2\pi i p(f/f_s)},$$

where I refers to the M-dimensional identity matrix and f_s is the sampling frequency. Note that in this equation, $i^2 = -1$.

Then, $GPDC_{ij}$ (which reflects the coupling from channel j to channel i) is calculated to be

$$GPDC_{ij}(f) = \frac{\frac{1}{\sigma_i}|A_{ij}(f)|}{\sqrt{\sum_{k=1}^{M} \frac{1}{\sigma_k^2}|A_{kj}(f)|^2}},$$

where σ_i^2 refers to the variance of the innovation process $x_i(t)$. $GPDC_{ij}$ is normalized in the interval $[0, 1]$, with increasing values for stronger interactions at particular frequencies, and sums up to one for each frequency component over all destination channels including the channel itself. The idea is to calculate the degree of influence of channel j to channel i with respect to the total influence of j on all channels. Note that this normalization procedure of (G)PDC was recently critizised (Schelter et al., 2009) because of some difficulties in comparing interaction strengths for different frequencies. As the values $A_{ij}(f)$ and $A_{ji}(f)$ are not necessarily identical, directionality of coupling is obtained. As $GPDC_{jj}$ has to be interpreted as the remaining amount of coupling that can not be assigned to the influence on other channels, I excluded self-coupling of channel j to itself for the subsequent analysis.

2.2.4 Experimental Task

Note that I did not contribute to the monkey experiments, which were designed by SL and GR and conducted by SL. However, the subsequent paragraph is stated here for the sake of completeness.

Two adult male rhesus monkeys (*Macacca mulatta*) participated in the experiments. All studies were approved by local authorities and were in full compliance with applicable guidelines (EUVD 86/609/EEC) for the care and use of laboratory animals. The behavioral task of the monkeys was a delayed matching to sample task. The monkey was seated in front of a screen at a distance of approximately 110 cm. An initial tone indicated the potential start of a trial. The monkey initiated a trial-start by grasping a lever and fixating on a small fixation spot on the center of the screen (baseline period). After 1500 ms, a first stimulus appeared on the screen for 250 ms, the so-called sample stimulus. As sample stimuli we used different natural images. The stimuli that were used in all of the experiments were chosen from the Corel-Photo-CD "Corel Professional Photos" comprising a collection of natural images showing birds, flowers, monkeys and butterflies in their natural surroundings. The images used in this study were randomly selected. All images were manipulated by Fourier techniques that have been described in detail

elsewhere (Liebe et al., 2009). The sample stimulus was followed by a delay period of 1500 ms during which the monkey held fixation. After the delay, a second stimulus, the so-called test stimulus, was presented. The monkeys were rewarded for a lever release whenever the test stimulus matched the sample stimulus. Whenever the test stimulus did not match the sample, the monkeys' task was to withhold the lever release until, after a brief delay of 200 ms, a second test stimulus appeared, that always matched the sample. This procedure ensured that the monkey had to initiate a behavioral response on every trial. The monkeys were rewarded with juice for every correct trial. Within one session, the different trial types were randomly interleaved. Stimuli were 7°x7° in size, at 24-bit color depth, and presented at the center of gaze on a 21 inch monitor (ViewSonic P810) with linear luminance response as well as linear response at separate color channels (gamma corrected).

2.2.5   Electrophysiology

Note that I did not contribute to the electrophysiological recordings, which were conducted by SL. However, the subsequent paragraph is stated here for the sake of completeness.

Local field potentials (LFP) were recorded from recording chambers placed on the surface of the skull based on stereotaxic coordinates allowing vertical access to the dorsal region of extrastriate area V4. The Hoarsley-Clark coordinates for the center of the recording chambers for monkey 1 were AP: -6.5, ML: -29.7. For monkey 2 the chamber coordinates were AP: -5.2, ML: -29.9. The implantation as well as surgical procedures used are described in detail in (Lee et al., 2005). Neural signals were measured using two custom made micro drives mounted on a plastic grid (Crist Instruments, Hagerstown, MD, USA). In each recording session 4-6 tungsten microelectrodes (UEWLGDSMNN1E, FHC Inc., Bowdoinham, ME, USA) were manually lowered down into the cortex in pairs with a minimal separation between electrodes of 0.5 mm. The impedance of the microelectrodes was approximately $1M\Omega$. The signal from each electrode was preamplified (factor 20, Thomas Recording, Giessen, Germany) using the recording chamber as the external reference. The analog signal was then filtered and amplified (BAK electronics, Germantown, MD, USA) to extract the local field potential (LFP) responses. After an additional waiting period of at least 1 hour the recordings were started. The LFP was obtained by band-pass filtering the signal between 0.1 and 300 Hz and digitizing with a sampling rate of 4464 Hz. One unit of the analog-to-digital converter corresponds to $5\mu V$.

LFP activity was recorded from 44 channels in 10 sessions from monkey 1 and 86 channels in 20 sessions from monkey 2. This resulted in 202 channel pairs for monkey 1 and 398 channel pairs for monkey 2. For each monkey, the minimum number of channels per session was 3, the maximum number of channels was 6. The spatial distribution of all recorded channels for monkey 1 and 2 can be found in Figure 2.2. For each recording site, its location is defined by two dimensions (anterior to posterior, and medial to lateral) based on the recording grid placed within the
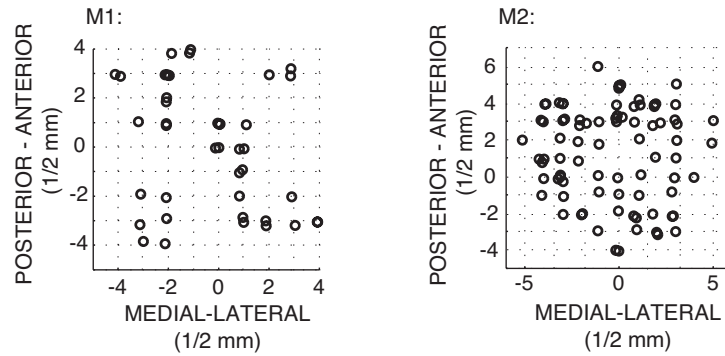
Figure 2.2: Recording locations along the medial - lateral and anterior posterior direction for monkey 1 and 2. 1 unit corresponds to 1/2 mm. Note that symbols are are slightly jittered at their recording locations for better visualization.

recording chamber. In order to measure coupling as a function of distance between recording sites, the Euclidian distance between two sites was calculated based on their respective locations along the two dimenions. The minimal distance between sites was 0.5 mm (i.e. sites directly neighbouring each other within the grid), the maximal distance we obtained was 4 mm.

2.2.6 Statistical Analysis

In order to be able to calculate confidence intervals that can be used to evaluate the significance of differences in coupling between different time intervals, I used a bootstrapping procedure that samples with replacement from the original trial set in order to generate bootstrap samples of the same size as the original data, but with different subsets of trials in them (Efron and Tibshirani, 1993). For each regarded data set of 1 second (last second of baseline and delay period), a set of 1000 bootstrap samples was generated. These bootstrap samples were then independently used to calculate the MVAR models as stated above and to estimate the couplings between the simultaneously recorded LFP channels with their respective confidence intervals. Change in coupling was considered significant if both the 0.01st and 99.9th percentile of the bootstrap distribution was above (increase) or below (decrease) the average baseline level.

Significance of coupling strength compared to the hypothesis that there was no coupling at all was assessed using a shuffling procedure. For each recording channel, trials were independently permuted repeatedly to obtain 1000 shuffled samples. MVAR model estimation was then also applied to these data sets.

For assessing the statistical significance of the effects of coupling as a function of distance between recording sites, I used a shuffling procedure that randomly shuffles the coupling values over distances to obtain 10^4 shuffled samples. Statistical significance of the real rank correlation was then calculated with respect to this
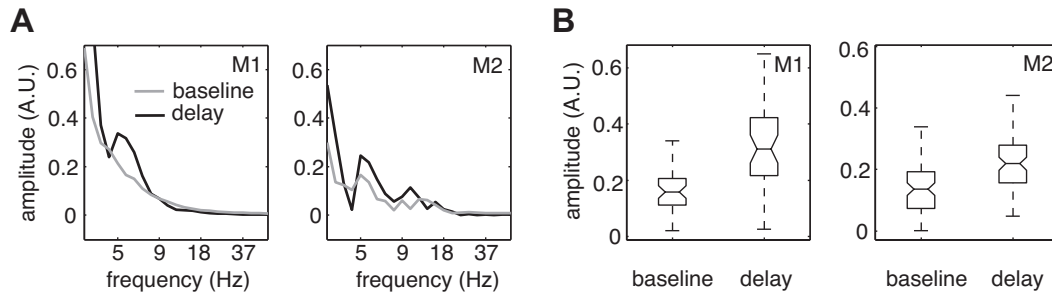
Figure 2.3: Enhanced power in the theta band during the delay period of the task. **(A)** Amplitude spectrum of LFP activity (median across all recorded channels in V4, $N = 44/86$ for monkeys 1/2, respectively) during the pre-stimulus baseline period (gray) and the delay period (black) of the delayed matching to sample task for frequencies from 1 to 40 Hz. Enhanced power in the theta band (4-10Hz) during the delay period is present for both monkeys. **(B)** Boxplots showing the distribution of power at the peak frequency within the theta range for the baseline vs. delay period. Both monkeys show a significant increase in theta power from baseline to delay (Wilcoxon signed rank test $Z = 6.95/4.47$, $p < 0.01$ for monkey 1/2, respectively).

distribution.

2.3   Results

2.3.1   Power Spectra

First, I examined the frequency content of induced oscillations during different periods of the visual memory task. Previously it had been found that there is enhanced power in the theta band during the delay period of the task in V4 (Rainer et al., 2004; Lee et al., 2005). I first sought to confirm these findings and compared the power spectrum for the delay period (i.e. across last 1000 ms before the onset of the test stimulus) to the power spectrum obtained from the 1000 ms time interval preceding the onset of the sample stimulus ("baseline"). Figure 2.3 shows the median amplitude spectra of LFP activity across all recorded channels (panel (A)) derived using a Morlet wavelet based approach (Tallon-Baudry and Bertrand, 1999; Graimann and Pfurtscheller, 2006). In both monkeys, the power spectra showed a local peak in the theta frequency range during the delay period (black) which is absent during the pre-stimulus baseline period (gray). Figure 2.3 (B) shows the distribution of power at the peak frequency within the theta range for the baseline vs. delay period and illustrates a significant change in theta power during the delay compared to the baseline.
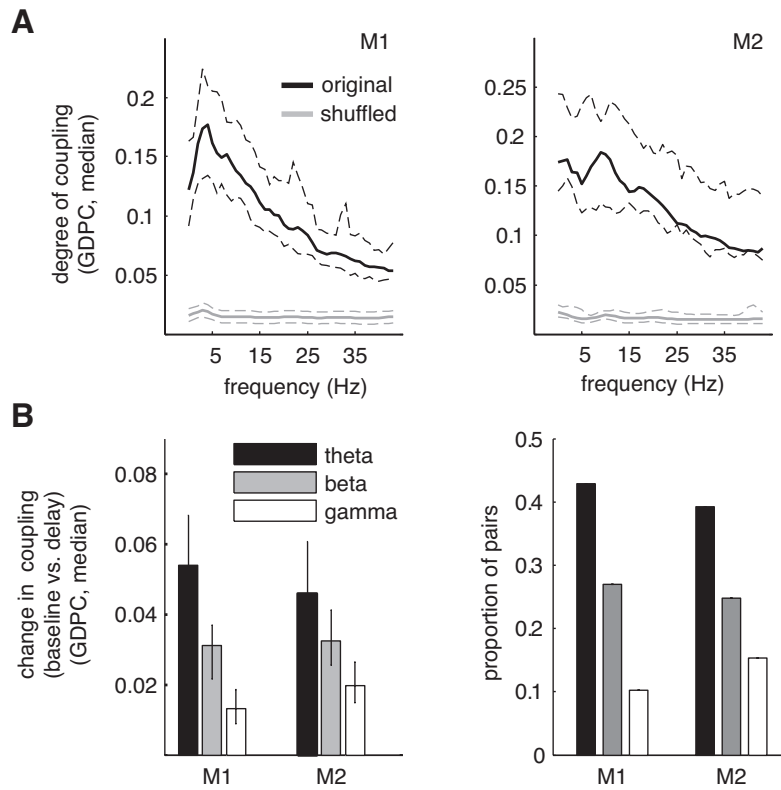
Figure 2.4: GPDC coupling in theta band during the delay period of the task. **(A)** Median degree of coupling (GPDC) as a function of frequency during the delay period (black) across all pairs in V4 for monkeys 1 and 2 for non-shuffled (black) and shuffled pairs. For both monkeys the results show a peak in coupling in the theta frequency range (3-12 Hz), although peak coupling occurred at slightly different frequencies (around 5 Hz for monkey 1, around 9 Hz for monkey 2). Dashed lines correspond to $\pm 34th$ percentiles of values around the median. Note that for the shuffled data, overall coupling was found to be around 0.02 for the whole frequency range. **(B)** Median absolute difference in degree of coupling (GPDC) (left) and proportion of pairs showing a significant change ($p < 0.001$) in coupling between pre-stimulus baseline and delay period for different frequency bands (right). Error bars correspond to $\pm 34th$ percentiles around the median. In both monkeys, the median change in coupling between baseline and delay is highest in the theta band (non-parametric ANOVA Kruskal Wallis Test, $\chi^2 = 11.4/4.23$, $p < 0.01$). Likewise, the proportion of pairs showing significant changes in coupling is also highest in the theta band compared to the other frequency bands (χ^2 test for comparison of proportions, $\chi^2 > 8.2$, $p < 0.001$ for all comparisons).

2.3.2   General Coupling Analysis

Based on the occurence of enhanced theta power during the delay period, I analyzed coupling strength between the different recording sites using generalized partial directed coherence (GPDC) obtained from multivariate autoregressive modelling. I

was interested in whether the enhanced theta power we observed during the delay period of the task coincides with directed coupling in the theta band. Thus, I first examined GPDC coupling as a function of frequency during delay (see Figure 2.4 (A)). Similar to the power spectra, I observed local peaks in GPDC coupling within the theta range (3-12 Hz) for both monkeys, albeit at slightly different frequencies. For monkey 1 the average peak frequency for highest coupling within the theta range was 4.33 ± 3.6 Hz (mean across sessions, ± 1 SD) and was located well within the range of the maximum power peak frequency at 5.86 ± 0.81 Hz (see Figure 2.3). For the second monkey the average peak frequency for highest coupling was larger (8.35 ± 3.8 Hz), and also higher than the average peak of power (5.51 ± 1.77 Hz), but not significantly higher ($p > 0.05$). Similarly, inspection of Figure 2.3 shows that although the peak in power for monkey 2 is around 5 Hz, power is elevated during the delay up to 10 Hz. Thus, the peak frequencies at theta power and theta coupling were overall similar. Note that these and subsequent results are based on the models that were fitted to 1 s time intervals in baseline and delay conditions (equivalent time intervals as for power spectra).

Subsequently, I assessed changes in coupling between the baseline and the delay period for several frequency bands that have been traditionally implicated in the interaction of oscillatory components during sensation and cognition and also follow conventional definitions of theta and beta bands (theta (3-12 Hz), beta (20-35 Hz), gamma (40-80 Hz), (Buzsaki, 2006)). Figure 2.4 (B) shows the median absolute difference in coupling between baseline and delay period (left). The graph shows that the degree of change significantly decreases with increasing frequencies with the largest coupling change occuring in the theta band. Likewise, the proportion of pairs that show significant changes in coupling is highest in the theta range compared to the other frequency bands (right). In the theta band, 116 of 202 pairs showed significant changes in coupling (57%, $p < 0.001$) in monkey 1, in monkey 2 235 pairs showed significant changes in coupling (59%, $p < 0.001$). In both monkeys, I found significant increases as well as decreases in theta coupling during the delay when compared to the baseline. Specifically, in monkey 1 74 pairs showed significant increases, and 42 pairs showed significant decreases in coupling. In contrast, in monkey 2 89 pairs showed increases and 146 pairs decreases. Thus, monkey 1 shows significantly more increases than monkey 2 and vice versa ($\chi^2 = 18.5$, $p < 0.01$). One factor that might contribute to this difference is the different distribution of electrode spacing between the animals, with monkey 1 showing significantly larger distances between electrodes than monkey 2 (mean [median] distance 4.6 [4] / 3.5 [3] for monkey 1/2, respectively; ranksum-test, $Z = 5.23$, $p < 0.01$). This is supported by several facts. First, for the smallest distance between electrodes, i.e. the distance that is identical and therefore comparable between the animals (unit 1, or 0.5 mm), the proportion of increases vs decreases is similar between the monkeys, i.e. statistically identical (50%/27% increases, Z=3.6, $p > 0.05$). Second, for the smallest distance we find an identical proportion of increases and decreases (i.e. 50/50) in monkey 1. Third, the proportion of significant decreases is slightly enhanced for smaller distances (50% at distance 1 vs. 20% at distance 4
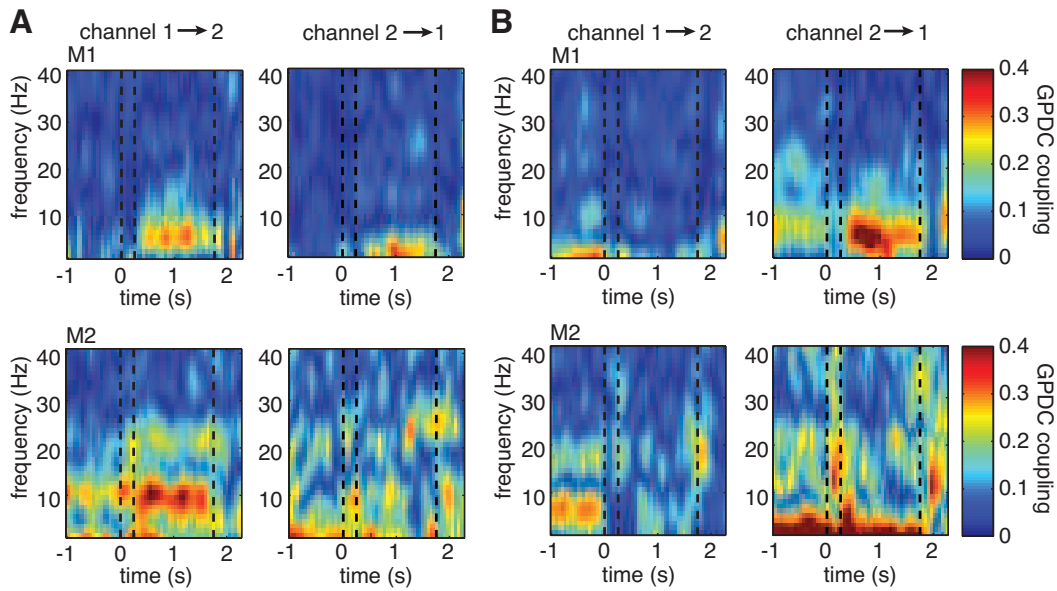
Figure 2.5: Single example LFP channel pairs showing significant changes in coupling during the delay period compared to baseline. Panel **(A)** shows one example for each monkey in which there is significant increase in coupling in the direction from channel 1 to 2 (left). Panel **(B)** plots two examples in which there is significant decrease in theta in one direction (left). The opposite directions show less strong or even opposing trends, indicating that coupling is not symmetric between sites. Dashed lines represent the on- and offset of the sample stimulus, as well as the onset of the first test stimulus during the trials, from left to right, respectively. This convention is also used in the subsequent figures.

for monkey 1, and 72% vs. 53% for distances 1 and 3 for monkey 2) and likewise the proportion of increases reduced at smaller distances. As the distances between electrodes are significantly lower in monkey 2 compared to monkey 1, the percentage of decreases should be higher in monkey 2, and vice versa. Ultimately, due to the limitations in spatial sampling, the differences in spatial configuration can only give an indication of why we find differences in the proportion of significantly increased vs. decreased coupling between the monkeys. In summary, our findings demonstrate that significant directed interactions between local field potentials within V4 during visual memory predominantly occur in the theta frequency range and the frequencies at which highest coupling occurs are comparable to the freqency range of power increases during the delay period. Based on these results we further investigated the time course and directionality of theta coupling during the delay period.

2.3.3   Time course and directionality of coupling

To illustrate the time course of theta coupling during the task, I used moving windows comprising time intervals of 250 ms (with an overlap of 200 ms) and fitted
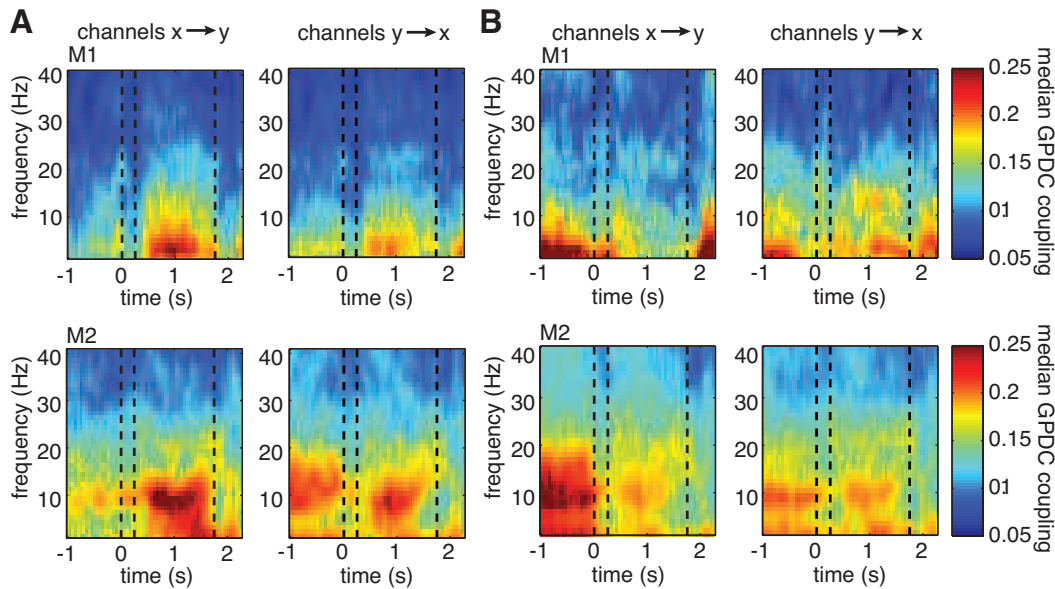
Figure 2.6: Grand-Median directed coupling across all pairs showing significant increase (**A**) or decrease (**B**) during delay compared to baseline ($p < 0.001$, left column) as well as the median coupling for the opposite direction of channel pairs (right column). Note that if there is significant increase (decrease) in both directions (i.e. from x to y and vice versa), both channel pair directions will contribute to the median for both directions of interaction. Otherwise, if all pairs would show significant increase (decrease) in both directions, left and right plots would be identical.

MVAR models to these individual windows. Figure 2.5 shows representative time courses of theta coupling in single recording pairs as well as the time course of coupling in the opposite direction (left/right graphs, respectively). These examples represent channel pairs with a significant ($p < 0.001$) increase or decrease in GPDC during the delay period compared to the baseline period and were chosen based on the previous analyses using coupling measures obtained from 1 s windows (see also Methods).

In all examples, theta increases and decreases occur shortly after the offset of the sample stimulus and are sustained throughout the entire 1500 ms long delay period. Interestingly, in all selected pairs I find differences in coupling strength and even opposing effects between pair directions, for example a significant increase in theta coupling in one, and a significant decrease in theta coupling in the opposite direction (see graph B example for monkey 1). To investigate this assymmetry across all channel pairs in more detail, I first computed the median coupling across all pairs showing significant increases or decreases in the theta band. I subsequently selected all pairs with the respective opposite direction and computed the median coupling across these pairs. The reasoning behind the procedure is as follows: If all channels show significant changes in theta coupling in both directions (i.e. from

channel X to channel Y and from channel Y to channel X), each channel pair will be represented in both groups. Consequently the median coupling across the channel pairs would be the same.

However, this is not the case. Figure 2.6 displays the resulting median coupling strength over all site pairs showing significant couplings within the theta range (left) and their respective opposite direction (right): For both monkeys we find assymmetrical, i.e. more unidirectional increases and decreases in theta coupling during the delay. This result is further illustrated in Figure 2.7 that shows the ratio of median coupling between pairs of channels and their opposite directions separately for each monkey (panels (A) and (B)). Furthermore, panel (C) shows coupling strength in the delay for sites with significant changes from baseline to delay in the theta band versus the couplings in the opposite direction. Similarly to the observed assymmetries in coupling values, significant proportions of pairs (64/48% monkey 1/2, $Z > 15.4$, $p < 0.001$) show significant *increases* in one direction only and significant proportions of pairs (57/58% monkey 1/2, $Z > 18.6$, $p < 0.001$) of the pairs show significant *decreases* in only one direction. Overall, in 3 out of 4 cases, the majority of pairs showed significant changes of coupling in one channel pair direction, but not the other.

Taken together, these findings illustrate that theta coupling during the delay is not symmetric between channel pairs and provide evidence for a complex interaction involving both directionally dependent increases and decreases in coupling during visual memory. In the following, I examine a different aspect of these coupling phenomena, namely their dependence on the spatial layout of the different oscillatory components. Our recording setup allowed us to simultaneously measure the activity of up to 6 LFP electrodes that were spatially distributed across a cortical surface area of approximately 6x6 mm. Therefore, within one session, electrode locations varied in spatial position and distance to each other.

2.3.4 Relation of Coupling Strength and Distance between Recording Sites

Figure 2.8 illustrates the dependence of absolute directed coupling and changes of coupling on the distance between electrodes, with higher direct coupling occuring at lower distances (both monkeys: $\rho_S = -0.52/-0.32, p < 0.0001$). Similar effects were found for the changes in coupling (i.e. decrease M1: rank correlation coefficient $\rho_S = -0.24, p < 0.05$, M2: $\rho_S = -0.31, p \ll 0.01$ and increase M1: $\rho_S = -0.12, p = 0.1$, M2: $\rho_S = -0.2, p < 0.05$) during delay with respect to the baseline). Note that the decrease of change in coupling with higher distance in monkey 1 does not reach a significance level of $p < 0.05$ for increases in coupling, but is at trend level.

The results indicate that both the strength of coupling and the change in coupling from the baseline to the delay condition are stronger for smaller distances between site pairs. This dependence could be found despite the differences in electrode spacing between the animals (see also section 3.2). Thus, not only absolute coupling but also the dynamic changes in coupling are a local phenomenon within
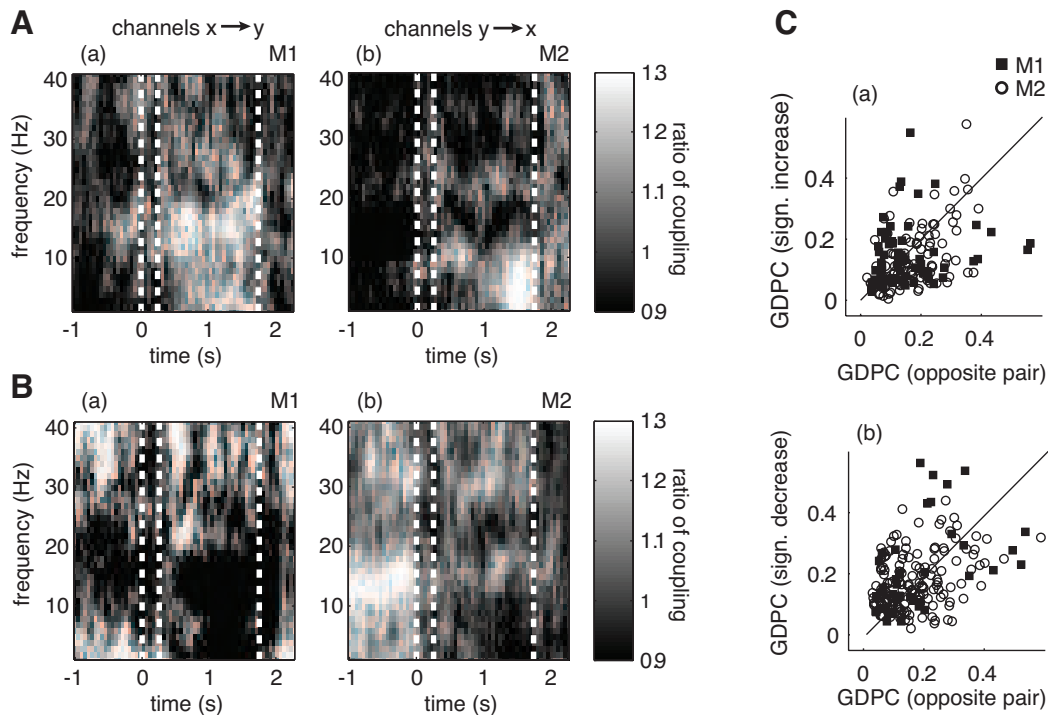
Figure 2.7: Ratio of median coupling between pairs showing significantly enhanced (**A**) and reduced (**B**) directed coupling during baseline and their opposite pairs (i.e. coupling in opposite direction). Plots a-b show the ratio of median coupling across the time course of the trial for different frequencies for monkey 1 and 2, respectively. A ratio larger (or smaller) than one indicates a stronger median coupling in one than the opposite direction, thus the ratio is an indicator of how "directed" or "symmetric" the coupling is between pairs of channels. Graphs (**C**) plot the degree of coupling of significant pairs and their opposite directions (y-axis pairs chosen based on sign difference between delay and baseline, x-axis shows opposite direction) in the theta band. If the degree of coupling is not symmetric between directions, data points deviate from the diagonal. Also note that if there was a significant increase (decrease, respectively) in both directions for a single pair, two data points are shown (one below, one above the diagonal if the respective increase/decrease is not fully symmetric).

the neural network. These findings are consistent with earlier reports for example from V1 recordings of in the macaque showing that pairwise spectral coherence in LFP activity between electrodes decreases as a function of receptive field distance (which is related to spatial distance, (Frien and Eckhorn, 2000)) and from recordings of several sites of the human cortex (Raghavachari et al., 2006) and extend these previous results using directed coupling measures.
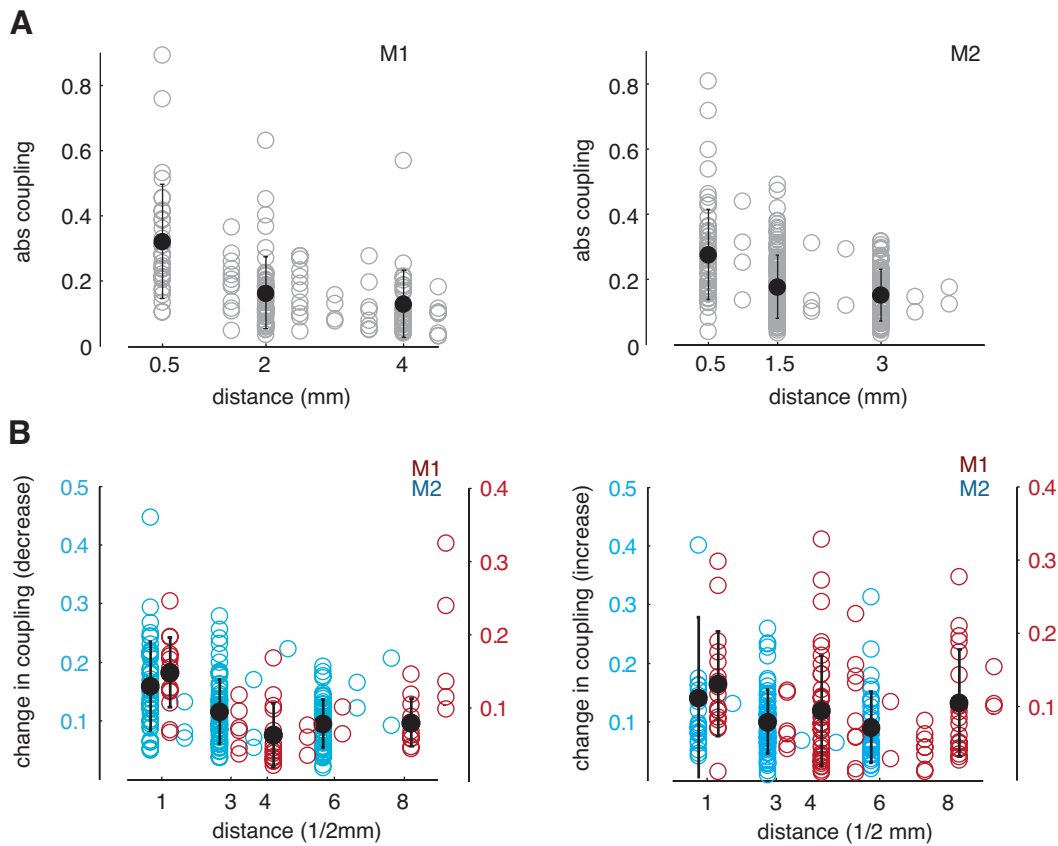
Figure 2.8: Dependence of GPDC in theta band on distance between simultaneously recorded channels. (A) absolute coupling as a function of Euclidian distance across all recorded pairs. Open symbols denote coupling values for single pairs, closed symbols represent the mean ±1 standard deviation across pairs within three bins. For both monkeys, absolute coupling decreases with increasing distance. (B) Dependence of increases and decreases in GPDC during delay relative to baseline in theta band on distance between simultaneously recorded channels. Absolute decrease (left) and increase (right) in coupling as a function of Euclidian distance across all recorded pairs showing decreases (left) or increases (right) in coupling. In both cases, changes in coupling during delay are higher for smaller distances. Note that we shifted the data points for M1 (blue) and M2 (red) slightly to the left and to the right, respectively for better visual discriminability.

2.4   Discussion

Oscillatory activity in neural networks as measured by EEG or LFP recordings is a widespread phenomenon of neural behavior and is thought to arise from the synchronous activity of neuronal populations at various spatial and temporal scales (Buzsaki and Draguhn, 2004; Fries, 2005; Salinas and Sejnowski, 2001). In many studies it has been shown that oscillations in different frequency bands are important for neural computations. Synchronous activity can, for example, establish

and support temporal relationships between different elements within a neural network depending on context, stimulus or behavioral state (Uhlhaas et al., 2009; Tallon-Baudry, 2009) or represent information about sensory events that can not be inferred from spiking activity alone (Montemurro et al., 2008). Thus, oscillatory activity serves the precisely timed cooperation between neural ensembles and could also provide temporal windows that allow for the selective routing and gating of information in an efficient manner (Fries et al., 2001; Salinas and Sejnowski, 2001; Mizuseki et al., 2000). Another important characteristic of neuronal oscillations is that oscillations at different frequencies are thought to subserve different behavioral and cognitive functions. Prominent examples are the involvement of gamma oscillations (>40 Hz) in visual and attention-related processes (Fries et al., 2001; Keil et al., 2001), the role of beta oscillations (15-35 Hz) in sensorimotor tasks (Murthy and Fetz, 1992) and the importance of theta oscillations in memory-related processing (O'Keefe, 1993; Rainer et al., 2004; Raghavachari et al., 2006; Lee et al., 2005).

The quantification of neural synchrony has traditionally been carried out using measures that assess the pairwise and instantaneous correlation in either amplitude or phase between two neural signals, for example using cross-correlation analysis between spike trains of multiple neurons (Aertsen and Arndt, 1989), spike-field coherence between the spiking activity of neurons and LFP activity (Pesaran et al., 2002; Fries et al., 2001) or phase-locking analysis of simultaneously recorded LFP or EEG data (Lachaux et al., 1999, 2000). However, despite the fact that these measures assess the strength with which two neural processes are coupled, they fail to provide information on several aspects of synchronization that can be important to fully describe their interaction, for example the direction of coupling between neural elements.

Here, multivariate autoregressive (MVAR) models have proven to be efficient tools for assessing the direction of coupling and can be more appropriate to capture the complexity of oscillatory dynamics as synchrony between neuronal ensembles changes across time or behavioral conditions. However, it is important to note that unipolar signals are vulnerable to volume conducted far-field effects and issues related to the usage of a common reference against which all differences in electrical potential are measured. Both factors might lead to adversely affected measures of coupling strengths, and elaborate methods for resolving these issues completely (besides using bipolar signals) need still be found (Schlögl and Supp, 2006; Bollimunta et al., 2009). Nevertheless, the application of coupling measures based on MVAR models has revealed important insights into neural interactions in many studies (Bressler et al., 1999; Brovelli et al., 2004; Kayser and Logothetis, 2009; Supp et al., 2007). For example, (Brovelli et al., 2004) analyzed interaction patterns in monkey sensorimotor cortex and found unidirectional couplings from somatosensory areas to motor areas within the beta frequency range that might be used to control motor output. In a different study (Kayser and Logothetis, 2009) investigated interactions of monkey auditory and superior temporal cortices related to sensory integration and found that while interactions from auditory cortex to

superior temporal regions prevail below 20 Hz, interactions in the other direction are more pronounced at frequencies above 20 Hz. A third example are the findings of the study by (Supp et al., 2007) in which the authors demonstrated that visual processing of familiar and unfamiliar objects engages different cortical networks at different degrees of directionality via interactions in the gamma frequency range. In all these studies, MVAR models revealed insights into the directed spatio-temporal dynamics of multiple cortical areas during cognitive processing that went beyond the description of synchrony between these areas.

In the present study I used MVAR models to provide a description of the directed coupling of theta oscillations during short-term memory. This oscillatory phenomenon has been described in a number of studies in relation to short term memory processes both in humans and animals. There are mainly two lines of research that focus on the role of these oscillations for memory-related processes. A large set of studies provides strong evidence on a connection between theta oscillations and (especially spatial) memory for the rat hippocampus, revealing that the timing of spikes both within hippocampus and within regions like prefrontal cortex is strongly connected to the hippocampal theta rhythm (O'Keefe, 1993; Buzsaki and Draguhn, 2004; Mizuseki et al., 2000; Siapas et al., 2005). A second line of research has concentrated on the importance of theta oscillations for memory performance in primates with the focus on EEG and LFP recordings in various cortical areas. Using EEG in human subjects, multiple studies have shown that there are increases as well as decreases in theta power that can depend on the specific nature of the memory task demand (Klimesch, 1999; Raghavachari et al., 2001, 2006). For example, (Raghavachari et al., 2001, 2006) showed an increase in theta power during the delay period of a memory task that co-varied with delay period length in multiple cortical regions, including frontal, temporal and occipital areas. Very recently, one study also incorporated MVAR modelling to reveal directional influences between different cortical regions (Anderson et al., 2009), providing evidence for memory-related theta-frequency interactions between prefrontal and medial temporal sites in the human brain. In contrast to human studies, only few studies have investigated the relation between theta oscillations and visual memory in the non-human primate. Here, research has focused on the extrastriate visual area V4, in which theta oscillations occur during the memory phase of the task, are modulated by task difficulty (Rainer et al., 2004) and are involved in the coding of visual stimuli during visual memory (Lee et al., 2005). In summary, research on theta synchrony during short term memory has thus far provided evidence for the hypothesis that memory processing is accompanied by increased theta power and synchrony.

However, as previously mentioned, measures of directed coupling based on MVAR modelling have shown to be useful for investigating the complex interaction patterns in LFP during cognitive processing. Thus I applied these coupling measures to analyze neural interactions in the theta band during visual short term memory within V4. The analyses firstly confirmed earlier results, showing enhanced theta power during the delay period. Using the coupling measures based on MVAR models, I additionally found increases as well as decreases in coupling

between recording sites in the delay period with respect to the baseline period that were most prominent in the theta band. This was evident in the coupling value estimates as well as in the proportion of site pairs showing significant changes in coupling. More importantly, however, I showed that these changes in coupling tend to be asymmetric between sites, i.e. they depend on the considered direction between site pairs. This finding suggests that not the mere occurrence of oscillatory activity or coherence in the theta band correlates with memory processing. Instead, the selective and direction-dependent change in theta coupling, which ultimately represents a change in functional connectivity within the neural circuit, plays an important role in this process. My results on the asymmetrical nature of directed interaction during memory also favors the hypothesis that theta oscillations and therefore coupling arise locally within the V4 network. In contrast, if coupling would be a phenomenon due to common input, one would expect bidirectional coupling with similar strength in both directions if the data from the neural elements producing the common input is not incorporated by the model. When interpreting the results from MVAR models, one should keep in mind that the model is based on data from only a small subset of the whole neural system (Stevenson et al., 2008). In addition, I was able to confirm earlier work on the relations of interaction strength and spatial distance that showed decreasing coherence of signals with increasing distance between sites and extended their results by measuring direct causal interactions instead of coherence (Frien and Eckhorn, 2000; Raghavachari et al., 2006).

Taken together, these effects would not have been revealed using more traditional methods that incorporate only phase synchronization or coherence. Therefore, the present work clearly shows the advantage of using directed coupling measures based on MVAR models for studying functional connectivity patterns within the brain and highlights the importance of direction-dependent modulations of local interactions between neural populations for studying sensory and cognitive processing.

Finally, I would like to point out that while the methods that I applied provide important insights into the functional connectivity patterns within the brain, their power is still limited because they can only assess linear interactions. While some extensions to nonlinear MVAR models (together with all the issues of nonlinear optimization) have been proposed (Pereda et al., 2005; Sun, 2008; Jachan et al., 2009), there is still work to be done to further improve these methods. In addition, the findings provide only the elementary description of the pattern of interaction between different oscillatory processes during visual memory. Further investigation will be needed to assess the specific role of directed coupling in relation to various cognitive parameters, for example task difficulty, performance or memory load. In addition, MVAR analysis can also be used to assess the interactions of LFP and spiking activity if the spike trains are properly preprocessed for this purpose. Here, it seems to be interesting to see how oscillations at the level of the LFP exert an influence on neuronal firing directly measured from the spiking activity of single neurons.

Acknowledgements

Supplementary Material
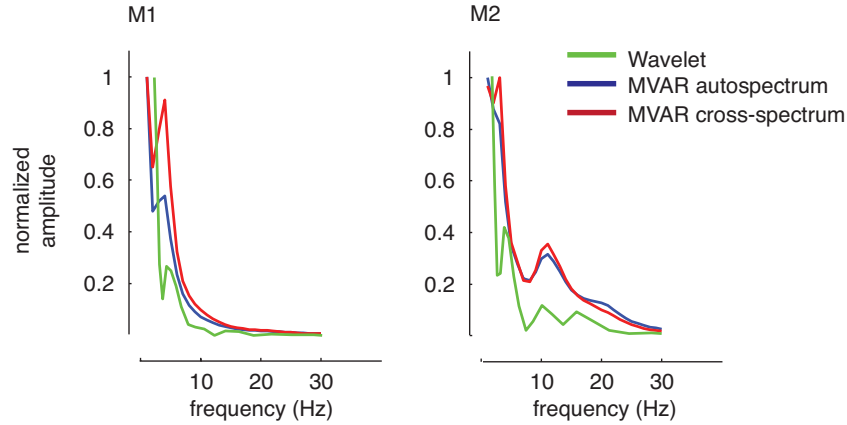
Figure 2.9: The average normalized amplitude spectra obtained from both analysis methods (MVAR and Wavelet Analysis) per monkey during the delay period in which the coupling effects were observed. For the MVAR method, both auto- and cross-spectra are shown. The Figure shows that the spectra are comparable between both approaches and that both analyses yield similar peak frequencies within the theta band. Note that the spectra were normalized between their respective minimum and maximum values within the shown frequency range after averaging.

Long range theta-coupling between V4 and lPF predicts visual short-term memory performance

Contents

This chapter is based on the manuscript "Long range theta coupling between V4 and lateral prefrontal cortex predics visual short-term memory performance" which has been submitted to the journal "Nature Neuroscience" and is currently under review. The work was conducted in collaboration with Stefanie Liebe (SL), Gregor Rainer (GR) and Nikos Logothetis (NL), affiliated with the Max Planck Institute for Biological Cybernetics in Tübingen, Germany, at the time the presented work was conducted. The monkey experiments and electrophysiological recordings were designed by SL and GR and conducted by SL. The data analysis was performed by myself and SL. Specifically, I contributed to the Wavelet-based analysis of LFP oscillations and subsequent estimation of phase locking and spike-phase locking as well as to the statistical evaluation by providing versions of the analysis that were then usually extended by SL. Moreover, I contributed to the interpretation of the results. The paper was written by myself and SL, with additional input from NL and GR. The results of this work, in which for the first time these analysis techniques are applied to simultaneously recorded data from V4 and the lateral prefrontal cortex of the macaque monkey performing a visual memory task, reveal new insights into the synchronization and long-range interactions of populations of neurons between extrastriate visual area V4 and the lateral prefrontal cortex during visual short-term memory.

Short-term memory requires the communication between multiple brain regions that collectively mediate encoding and maintenance of sensory information. Oscillatory synchronization has been suggested to underlie inter-cortical communication between regions mediating these processes. Yet, whether and how distant cortical areas cooperate during visual memory remains elusive. Here, neural interactions between visual area V4 and the lateral prefrontal cortex (lPF) based on simultaneous recordings of local field potentials (LFP) and single unit activity (SUA) in monkeys performing a visual short-term memory task are reported. During the memory period enhanced between-area phase synchronization in the theta band (3-9Hz) of LFPs together with elevated phase locking of SUA to theta oscillations across regions could be observed. Most importantly, however, it is shown for the first time that the strength of inter-cortical locking is predictive of the animals' memory performance. This suggests that theta-band synchronization may provide the basis for the timely coordination of spiking output between V4 and prefrontal cortex and may well reflect facilitated communication of visual information during visual memory.

3.1 Introduction

Every cognitive act entails the participation of multiple brain regions. In visual short-term memory, for example, visual information is initially encoded in sensory brain areas and then communicated to regions that mediate the retention, manipulation and retrieval of information. One prominent hypothesis that addresses the question of how communication between neural ensembles is achieved claims that neuronal oscillations support the timely coordination of neural activity between different brain regions (Buzsaki and Draguhn, 2004; Fries, 2005; Womelsdorf et al., 2007; Hipp et al., 2011). Specifically, neuronal oscillations in the theta frequency band (3-9Hz) have been suggested to underlie the interaction between neural ensembles during mnemonic processing (Buzsaki and Draguhn, 2004; Fries, 2005; Fell and Axmacher, 2011).

One line of evidence supporting this hypothesis stems from studies investigating the role of hippocampal theta in memory formation in rodents (O'Keefe, 1993; Buzsaki, 2002; Jones and Wilson, 2005; Siapas et al., 2005; Lisman and Buzsaki, 2008). For example, phase locking of prefrontal spiking to the hippocampal theta rhythm is enhanced during task periods with enhanced memory demands (Jones and Wilson, 2005). Additional evidence stems from studies measuring surface-based or intracortical electroencephalogram in human subjects. Here, memory performance correlates with an increase in theta power (Klimesch, 1996; Kahana et al., 1999; Tesche and Karhu, 2000; Raghavachari et al., 2001; Sederberg et al., 2003) or enhanced theta synchrony between electric potentials recorded from memory-related areas, for example within and between prefrontal cortex and medial temporal lobe regions (Anderson et al., 2009; Rutishauser et al., 2010), as well as between frontal and occipital regions (Sarnthein et al., 1998; Stam et al., 2002).

These findings raise the question whether theta synchrony measured at the mesoscopic level of electric potentials provide a basis for the timely coordination of spiking output between distant cortical areas that have been traditionally associated with the sensory encoding of visual information on the one hand and mnemonic processing on the other. Moreover, is the precision of coordination between these regions associated with changes in memory performance? To answer these questions, I studied neuronal interactions between the extrastriate visual area V4 and the lateral prefrontal cortex (lPF) while monkeys performed a visual memory task.

While neural activity in V4 has been mostly related to color and shape processing of visual objects (Zeki, 1980; Schein and Desimone, 1990; Pasupathy, 2006; Orban, 2008; Tanigawa et al., 2010) and attentional modulation of visual activity (Fries et al., 2001; Gregoriou et al., 2009), neural responses in lPF have been traditionally associated with "working memory", i.e. the short-term maintenance and manipulation of sensory information in memory tasks (Fuster and Alexander, 1971; Baddeley, 1986; Petrides, 1996; Miller and Cohen, 2001). More recently, however, an increasing number of studies have shown that the neural circuitry underlying short-term retention of sensory information likely entails earlier sensory cortical areas as well; for a review see (Pasternak and Greenlee, 2005). Consistent with this, both prefrontal regions (Pesaran et al., 2002; Pipa et al., 2009) and V4 have been linked to memory-related oscillatory synchrony. Specifically, theta oscillations are enhanced during the delay period of memory tasks in both lateral prefrontal cortex and V4 (Rainer et al., 2004; Hoerzer et al., 2010) and increased oscillatory theta synchrony is accompanied by a phase dependent coding of visual stimuli retained in short-term memory (Lee et al., 2005; Siegel et al., 2009).

Here, enhanced phase locking between local field potentials recorded in V4 and lPF ("inter-area LFP-phase locking") that occurs in the theta range (∼3-9Hz) during the memory period of a visual short-term memory task is reported. Increased LFP phase locking is associated with greater locking of spikes to the phase of the theta oscillations in the respective other area ("inter-area spike-phase locking"). Both inter-area LFP- and spike-phase locking was significantly higher for subsequent correctly remembered stimuli and was also predictive of session-to-session variations in task performance.

3.2 Materials and Methods

3.2.1 Animals and Recordings

Note that I did not contribute to the electrophysiological recordings, which were conducted by SL. However, the subsequent paragraphs are stated here for the sake of completeness.

Two adult male monkeys (macaca mulatta) participated in the experiments. All studies were approved by the local authorities (Regierungspräsidium, Tübingen, Germany) and were in full compliance with the guidelines of the European Community for the care and use of laboratory animals (European Union Directive

86/609/EEC). Single-unit activity and the local field potentials (LFP) were recorded from two recording chambers placed on the surface of the skull based on stereotaxic coordinates allowing access to the subjacent brain areas on the left hemisphere. The Horsley-Clarke coordinates for the centers of the recording chambers for monkey 1 were: V4: AP: -6.5, ML: -29.7, lPF: AP: -33, ML: -23.7. For monkey 2 the coordinates for the chambers were: V4: AP: -5.2, ML: -29.9, lPF: AP: -34.5, ML: -22.6. The implantation as well as surgical procedures used are described in detail in (Lee et al., 2005).

Neural signals were measured using two custom made micro drives mounted on a plastic grid (Crist Instruments, Hagerstown, MD, USA). In each recording session 4-6 tungsten microelectrodes (UEWLGDSMNN1E, FHC Inc., Bowdoinham, ME, USA) were manually lowered in pairs with a minimal separation of 0.5mm between the electrodes. The impedance of the microelectrodes was approximately 1MΩ at 1kHz. The signal from each electrode was preamplified (factor 20, Thomas Recording, Giessen, Germany) using the recording chamber as the external reference. The analog signal was split into two signals and filtered and amplified separately (BAK electronics, Germantown, MD, USA) to extract single unit activity (SUA) as well as the local field potential (LFP) responses.

The spiking activity was obtained by bandpass filtering the signal between 300Hz and 4kHz and digitizing with a sampling rate of 22.231kHz. Single units were extracted from the spiking activity using standard spike sorting routines (Offline Sorter, Plexon, Dallas, TX, USA). In each recording session the electrodes were advanced until one or more single neurons could be reliably isolated. Neurons were not selected based on task selectivity. After an additional waiting period of at least 1 hour the recordings were started. The LFP was obtained by band-pass filtering the signal between 0.1Hz and 300Hz and digitizing with a sampling rate of 4464Hz. One unit of the analog-to-digital converter corresponds to $5\mu V$.

3.2.2   Behavioral Paradigm

Note that I did not contribute to the monkey experiments, which were designed by GR and SL and conducted by SL. However, the subsequent paragraphs are stated here for the sake of completeness.

3.2.2.1   Task

The behavioral task of the monkeys was a delayed matching to sample task. The monkey initiated a trial-start by grasping a lever and fixating on a small fixation spot on the center of the screen. After successful fixation of 1000ms a first stimulus appeared on the screen for 250ms, the so-called sample stimulus. The sample stimulus was followed by a delay period of 1500ms during which the monkey held fixation. After the delay, a second stimulus, the so-called test stimulus, was presented. The monkeys were rewarded for a lever release whenever the test stimulus matched the sample stimulus. Whenever the test stimulus did not match the sample, the mon-

keys' task was to withhold the lever release until, after an additional brief delay of 200ms, a second test stimulus appeared that always matched the sample. This procedure ensured that the monkey had to initiate a behavioral response on every trial. The monkeys were rewarded with juice for every correct trial. In each experiment, 50% of the trials were "match" trials, 50% were "non-match" trials.

3.2.2.2 Stimuli

A detailed description of the behavioral paradigm and the stimuli used is given in a previous study (Liebe et al., 2009). In brief, for each experiment a set of three or four natural images was presented as visual stimuli. The stimuli were 7°x7° in size, with 24-bit color depth and presented at the center of gaze on a monitor (Intergraph 21sd107) with linear luminance response (gamma corrected) at a distance of approximately 110cm from the monkeys. Prior to the recording sessions, the monkeys had been familiarized with the images and it had been ensured that monkeys did not show performance changes due to learning anymore, for details see (Liebe et al., 2009). The images were chosen from the Corel-Photo-CD "Corel Professional Photos" comprising a collection of natural images showing birds, flowers, monkeys and butterflies in their natural surroundings and were degraded with visual noise to various degrees, including showing stimuli that only contained "visual noise" on approximately only 10% of the trials, to ensure motivation of the animals. The trials were presented in a pseudo-randomized fashion to ensure that the animal would perform a minimum number of trials for every condition.

3.2.3 Data Analysis

All of the data analyses were performed with custom software written in Matlab (The MathWorks Inc., 2007). Circular statistics were computed using the Circular Statistics Toolbox written for Matlab (Berens, 2009).

3.2.3.1 Derivation of the analytic signal using Wavelets

All of the spectral analyses were based on the same time-frequency decomposition using complex Morlet wavelets. The LFP was first resampled from 4464Hz to 200Hz and normalized by subtracting the mean waveform over trials and subsequently dividing the result by the standard deviation over trials. Afterwards I extracted the instantaneous amplitude and analytical phase as a function of time and frequency by convolving the raw real-valued time series $x(t)$ with the complex Morlet wavelet $w(t, f_0)$ to obtain the complex output signal $y(t, f_0)$, also denoted as the "analytic signal":

$$y(t, f_0) = x(t) * w(t, f_0) = \int_{-\infty}^{+\infty} x(\tau)w(\tau, f_0)d\tau \qquad (3.1)$$

Here, f_0 denotes the desired center frequency of the wavelet function

$$w(t, f_0) = A e^{-t^2/(2\sigma_t{}^2)} e^{i2\pi f_0 t} \tag{3.2}$$

composed of a Gaussian and a sinusoidal component, where $A = (\sigma_t \sqrt{2\pi})^{-1}$ is a scaling factor, $\sigma_t = (2\pi\sigma_{f_0})^{-1}$ defines the width of the Gaussian component in the time domain, and $\sigma_{f_0} = f_0/c$ defines the position and width of the peak in the frequency domain. Note that the higher the value of c, the better the frequency resolution, but the poorer the time resolution. I used a value of $c=7$ as proposed in the literature (Tallon-Baudry and Bertrand, 1999; Graimann and Pfurtscheller, 2006). Note that the Hilbert transform is another approach that is commonly used to extract the instantaneous phase of a time series, see (Le Van Quyen et al., 2001) for a comparison of the two approaches, that leads to similar results. I chose to use the Wavelet transform because it already provides the necessary band-pass characteristics while the Hilbert transform has to be preceded by filtering the data using a narrow band-pass filter to obtain physically meaningful results. I did not want to restrict our analysis to specific frequency bands and thus I used a logarithmic frequency scale and obtained the instantaneous amplitude and phase of 27 frequency components ranging from 2Hz to 97Hz.

I first examined the frequency content of induced oscillations during different periods of the task (see Figure 3.1A). Here, in order to discard stimulus-evoked components of the oscillations, the data was normalized by subtracting the (evoked) mean waveform across trials from every single trial before applying Wavelet transform filters to obtain a time-frequency representation of the data. Subsequently the amplitude spectra were computed by averaging the spectra across all individual sites recorded in V4 and lPF for each monkey separately. To compare the frequency content of the LFPs between baseline and delay period, the amplitude spectra were averaged within a 1000ms window preceding the sample stimulus ("baseline") and across the last 1000ms preceding the test stimulus, leaving out the first 500ms after sample onset. The latter time window was chosen to ensure that the frequency content was not contaminated by stimulus-evoked activity. For all analyses, unless otherwise stated, the same time windows were chosen for "baseline" and "delay".

3.2.3.2   Phase locking analysis between LFP sites

Phase locking was quantified between individual pairs of LFP sites (m, n) by computing the phase locking value (Lachaux et al., 1999) across the entire trial period. The PLV for two recording channels m and n at a particular center frequency f_0 and time t is defined as

$$\text{PLV}_{mn}(t, f_0) = \frac{1}{K} \left| \sum_{k=1}^{K} e^{i\left(\varphi_k^m(t,f_0) - \varphi_k^n(t,f_0)\right)} \right| \tag{3.3}$$

where K denotes the number of trials and $\varphi_k^m(t, f_0)$ and $\varphi_k^n(t, f_0)$ denote the instantaneous phases of the k-th trial of the two channels that were computed using the Wavelet transform with center frequency f_0. Therefore, the PLV measures

the degree of similarity of the phase difference $\Phi_k^{mn}(t, f_0) = (\varphi_k^m(t, f_0) - \varphi_k^n(t, f_0))$ across different trials k and ranges in the interval $[0, 1]$. The higher the similarity of phase differences across trials, the higher is the PLV. Note also that the PLV is symmetric for pairs of channels, i.e. $\text{PLV}_{mn} = \text{PLV}_{nm}$.

Since the phase locking measure yields a single value, which is obtained across multiple trials, significance between baseline and delay in single electrode pairs was tested using a bootstrapping procedure. Specifically, in order to obtain confidence intervals that can be used to evaluate the significance of differences in phase locking between the baseline and delay period, I sampled (with replacement) from the original trial set and generated 1001 bootstrap samples of the same size as the original data, but containing different subsets of trials (Efron and Tibshirani, 1993). For each subset, the PLV was calculated separately. The distribution of PLVs for the baseline period was derived using the average PLV across the baseline interval (see above) for each bootstrap sample. A PLV was considered significantly different from baseline only if the average PLV over all bootstrap samples was above or below the 99.9th percentile of baseline distribution of the PLV. To obtain the proportion of significantly phase locked LFP pairs in each frequency band, the proportions of significant pairs were averaged across frequencies within the band of interest, i.e. 3Hz to 9Hz for theta, 16Hz to 36Hz for beta, and 47Hz to 97Hz for gamma.

In order to compare phase locking between delay and baseline across the population of LFP pairs, the mean phase locking during baseline was subtracted from every phase locking value across the entire trial period per pair, divided by the standard deviation across pairs and subsequently averaged across all pairs. This procedure was done separately for each monkey and resulted in the normalized Z-score transformed phase-locking values.

3.2.3.3 Locking of spikes to the LFP

To investigate spike-phase locking within each area, LFP-unit pairs recorded at the same electrode site were excluded. Thereby, 601 spike-LFP pairs within V4 (M1: 426, M2: 175), and 397 pairs within lPF (M1: 216, M2: 181) were obtained. In order to examine inter-area locking 110 lPF LFP channels (M1: 67, M2: 43) were paired with the spiking activity of 167 simultaneously recorded V4 units (M1: 114, M2: 53) and 124 V4 LFP channels (M1: 79, M2: 45) with the spiking activity of 130 simultaneously recorded lPF units (M1: 69, M2: 61). This resulted in 660 lPF LFP-V4 unit pairs (M1: 458, M2: 202) and 593 V4 LFP-lPF unit pairs (M1: 335, M2: 258). No channels or units were selected based on any criteria, for example task-related activity or visual responsiveness in any of our analyses. Thus, the data represent a completely unbiased sample of LFP and single-unit activity in both regions.

Circular non-uniformity of the distribution of spikes with respect to the phase of the theta oscillation, i.e. for frequencies between 3Hz and 9Hz, was tested for each pair of simultaneously recorded LFP signal and single unit activity. Specifically, each spike was assigned its respective phase value of the simultaneous LFP oscilla-

tion (the phase was obtained per frequency using the same Wavelet decomposition as used for calculating the frequency spectra and PLV). Due to the fact that the window usually is no integer multiple of the oscillation length of $1/f_0$s, some phase values can occur more often than others, which results in non-uniformity "per se". In order to account for possible non-uniform distribution of phases, the spike counts were normalized for each bin using the number of occurrences of each particular phase bin during the considered window for each trial. If spikes occur more often at particular phases, this procedure results in a non-uniform distribution of phase values across phase bins. Then, the Rayleigh Test for circular uniformity (Siapas et al., 2005; Zar, 2008) was applied to test whether the spikes were significantly locked to the theta oscillation. Phase locking was defined to be present for pairs for which a value of at least $ln(Z) = 1.09$ was obtained, which approximately corresponds to a value of $p = 0.05$ for distributions containing at least 50 spikes (Siapas et al., 2005).

In order to estimate concentration parameter κ and mean direction μ during the baseline, spike trains were cut out within a 500ms window preceding the sample stimulus ("baseline") and a 500ms window during the delay period that were centered on the maximal troughs of the filtered theta oscillations to characterize the locking properties of individual units (frequencies ranging from 3Hz to 9Hz). The central 500ms of the delay period were chosen in order to not contaminate spiking activity by either the sensory response after sample stimulus onset or spiking activity related to test stimulus, thus keeping spike rates between baseline and delay as comparable as possible. Within these windows, spikes were binned according to the phase of a sine wave of frequency f_0 mapped onto the extracted cycle. This procedure was repeated for each trial and the resulting windows were subsequently stacked up to form theta triggered spike train rasters (as for example shown in Figure 3.5). Spikes were then summed across trials in order to obtain histograms containing the spike counts as a function of phase bin. Subsequently, von Mises density functions with mean direction μ and concentration parameter κ were fitted to the spike distribution across phase bins. The von Mises density function is a circular analog of the one-dimensional normal distribution and has two parameters μ and κ, where μ describes the mean direction or preferred phase of firing and κ indicates the concentration of firing as a function of phase, i.e. how peaked the firing distribution is around the preferred phase. Both procedures (Rayleigh Test, von Mises function) have been extensively used to describe phase locking characteristics of individual neurons (Siapas et al., 2005; Zar, 2008).

The median concentration parameter κ was compared between groups of neurons that could potentially vary with respect to mean spike rate. In order to account for these differences, κ was also compared between groups of units whose mean firing rate was matched. To create two sets of prefrontal and V4 neurons with (approximately) the same mean firing rate, the most active cells from V4 and the same number of least active cells from lPF were selected. The number of cells selected was chosen such that the mean firing rates in the two selections were as similar as possible. For monkey 1 the mean rates did not differ significantly

between V4 and lPF during delay (mean rate V4: 18.3Hz, mean rate lPF: 18.9Hz, independent sample T-test $t = 0.43$, $p > 0.05$). Thus, it was not necessary to correct for mean rates for units from this monkey. In contrast, for monkey 2 the mean spike rate during delay was significantly larger for prefrontal than for V4 units (mean rate V4=23.4Hz, mean rate lPF=33.7Hz, independent sample t-test $t = 4.9$, $p < 0.01$). After the selection, the mean rates for units from V4 were 25.3Hz and for lPF 25.4Hz (independent sample t-test $t = 0.05$, $p > 0.05$). The same procedure was applied when to compare the median κ values for the baseline vs. the delay period within the group of prefrontal and V4 units. Prefrontal unit activity tended to be higher during the delay than during the baseline (mean rate delay: 20.8Hz, mean rate baseline: 19.6Hz), but was only significant in monkey 1 (t-test, M1: $t = 5.8$, M2: $t = 1.6$, M1: $p < 0.01$, M2: $p > 0.05$). In V4 mean spike rates tended to be lower during the delay than the baseline (mean rate delay: 16.6Hz, mean rate baseline: 21.5Hz). The difference was significant in both animals (t-test, M1: $t = 4.0$, M2: $t = 5.6$, $p < 0.01$). In order to create two data sets for baseline and delay periods with (approximately) the same mean firing rate, the most active cases from the baseline group were selected and paired with the least active cases from the delay group. Again, the most active units from the baseline group were chosen and paired with the least active units from the delay group in order to obtain data sets with approximately the same (i.e. not significantly different) firing rate.

For all analyses involving single unit activity and LFP signals from the same area (cf. Figure 3.4), only signals recorded simultaneously from neighboring electrodes were analyzed, i.e. at least 0.5mm away, in order to avoid artifacts do to spillover of signals from single unit activity into the LFP signal recorded at the same electrode and in order to have comparable settings for interactions within and between the areas (Siapas et al., 2005; Siegel et al., 2009).

3.2.3.4 Analysis of coupling for correct and incorrect trials

For this analysis only trials were chosen for which stimuli were shown at a high degradation level (45% coherence) as for this condition animals made enough errors to reliably estimate phase locking from both correct and incorrect trials. Correct trials were defined as trials in which the animals correctly identified the previously shown stimulus ("hit"), and non-correct trials were defined as trials in which the animals did not ("miss"). As the performance of animals for this stimulus condition was above chance level, there were typically more correct than non-correct trials. To avoid confounding the phase locking measurements with differences in trial numbers, phase locking between LFPs as well as locking of spiking activity to LFP phase was computed only across the minimum number of trials in each condition.

Whether phase locking significantly differed between correct and incorrect trials was assessed by computing the d-prime value between correct and incorrect sets of trials per pair. To estimate how variable PLV values would be for correct vs. incorrect sets of trials, first the phase-locking for different sub-samples of trials for both conditions (199 repetitions) was computed and then the mean across subsampled

PLVs from correct trials was subtracted from the mean across PLVs from incorrect trials. Subsequently, the difference was divided by the standard deviation across all PLVs. D-prime values were computed for original data sets and datasets in which the trials were randomly exchanged between the correct and incorrect condition (shuffled). The p-values represent the average proportion of cases in which the shuffled d-prime averaged across pairs during delay was higher than the original d-prime, see Figure 3.7, panel D.

To compare spike-phase locking between correct and incorrect trials, the same analyses were performed as described above, again taking only the subset of correct and incorrect trials. In order to correlate the difference in phase locking between correct and incorrect responses (d-prime) with session-to-session variations in performance, each V4 LFP pair was assigned its respective proportion of correct performance value of the session in which the pair was recorded. Subsequently Spearman's rank correlation coefficient was computed across all sessions/pairs from bootstrapped samples drawn with replacement from the original data set (see also Figure 3.9 panels A and C). The mean and 95th percentile confidence intervals of correlations across the bootstrapped samples ($N = 999$) are reported. To illustrate the dependency of d-prime and performance further,d-prime data points were binned using six equally spaced bins for the proportion of correct responses and averaged across d-primes per bin (Figure 3.9, panel B).

3.3   Results

Local field potentials (LFP) and spiking activity in V4 and lPF were simultaneously recorded while monkeys performed a delayed matching-to-sample (DMS) task. During the task, the animals had to retain information about a briefly presented visual stimulus (250ms) over the course of a delay period (1500ms). The sample stimuli consisted of a set of natural images that were shown at different levels of image degradation. The stimulus conditions that were used have been described in detail in a previous study (Liebe et al., 2009). The task sequence is depicted in Figure 3.1A (see also Methods).

The task design allowed (1) to evaluate neural synchrony between the areas as measured by phase locking of simultaneously recorded LFP signals during the memory period of the task, (2) to determine whether changes in neural synchrony between both areas at the level of the LFP are accompanied by changes in the locking of spiking activity to LFP oscillations in the respective other area and (3) to examine whether differences in memory performance are associated with changes in phase locking between LFP sites and between LFP and spiking activity across the areas.

First, I and my collaborators sought to confirm previous results on increased theta oscillations during visual memory by analyzing LFPs from a total of 131 sites recorded in V4 (M1: 86, M2: 45) and 117 sites in lPF cortex (M1: 74, M2: 43) and comparing the amplitude spectra of LFPs during the delay period of the task to a
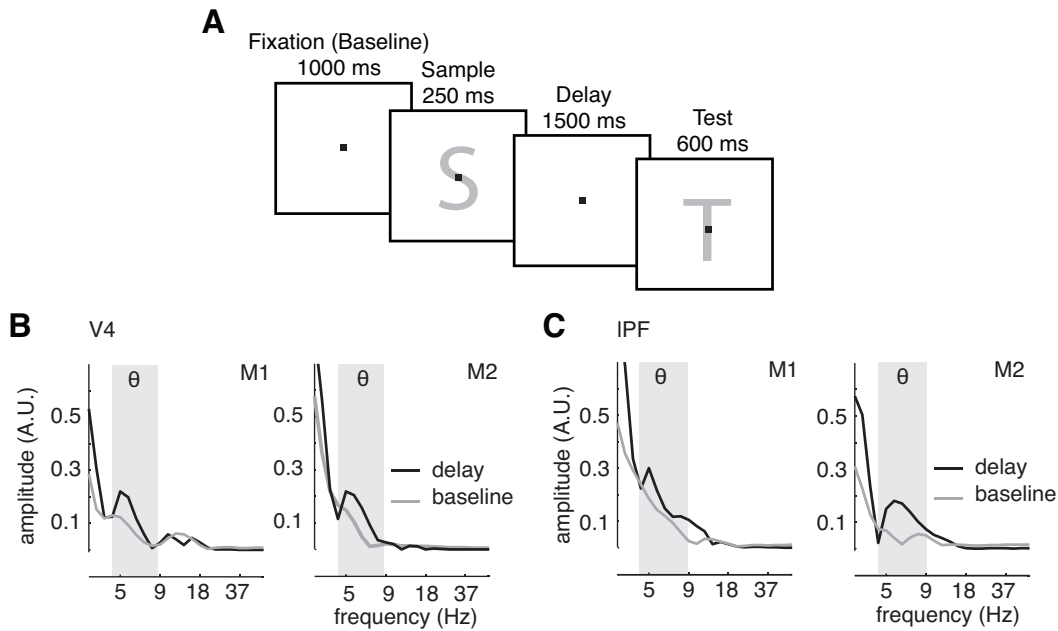
Figure 3.1: Behavioral paradigm used during experiments and amplitude spectra in areas V4 and lPF during baseline and delay. **A:** Events and their respective duration during the delayed matching-to-sample task (DMS task). Monkeys had to fixate a small fixation spot (2°) on the screen, subsequently a sample stimulus was shown for 250ms, which was followed by a delay period during which the animals had to retain fixation. Afterwards, a test stimulus was shown that either matched or did not match the sample stimulus (50% match, 50% non-match trials). **B,C:** Enhanced amplitude of theta oscillations during the delay period of the task both in V4 (B) and lPF (C). Both plots show mean amplitude spectra of LFP activity during the pre-stimulus baseline period (gray) and the delay period (black) for frequencies from 1 to 40Hz (log scale).

baseline period prior to the sample stimulus onset ("fixation"). Figures 3.1B and C show the mean amplitude spectra of LFP activity across all recorded channels per monkey in V4 and lPF, respectively. In both animals, the frequency spectra showed local peaks in the theta band (3-9Hz) in V4 as well as in lPF cortex during the delay period that were significantly higher than during the baseline period (Wilcoxon signed rank test, V4: $Z > 4.3$, $p < 0.01$, lPF: $Z > 5.4$, $p < 0.01$ for both animals). These findings are consistent with previous reports on increases in the amplitude of theta band oscillations during visual memory tasks in both humans and non-human primates (Klimesch, 1996; Raghavachari et al., 2001; Sederberg et al., 2003; Rainer et al., 2004; Lee et al., 2005; Raghavachari et al., 2006; Anderson et al., 2009).

3.3.1 Phase synchrony in the theta band between V4 and lPF sites is enhanced during the delay period of the memory task

Based on the observation of increased theta power during the memory period, the next question was whether there is increased theta-band synchronization between V4 and lPF LFP sites during the delay period. To address this question, all simultaneously recorded V4 and lPF channels ($N = 332$ for M1 and $N = 175$ for M2) were paired and the phase locking value (PLV) was computed between all these pairs for frequencies ranging between 1 and 100Hz (for details see Method section). The phase locking value provides a description of the consistency in the phase relationship between two oscillatory signals and has been used in a wide variety of studies to get insights into the strength and timing of synchronized interactions of electric potentials (Lachaux et al., 1999; Varela et al., 2001).

Figure 3.2A shows an example V4-lPF LFP channel pair showing increased phase locking during the delay period in theta frequency (illustrated at 6.8Hz). Specifically, panel A shows that theta oscillations in V4 and lPF exhibit a consistent phase relationship across the different trials shown during the delay period that is absent during the baseline period of the task. Figure 3.2B displays the PLV of several representative example pairs that exhibit significantly enhanced PLV during the delay period in the theta range, including the pair illustrated in Figure 3.2A (upper left).

The results show increases in theta phase locking during the delay not only for single pairs, but generally across all recorded LFP pairs. This effect is illustrated in Figure 3.3A, which depicts the average change in phase locking (Z-Score, normalized to pre-stimulus baseline) as a function of time and frequency for all recorded pairs (grand average across means per animal). In general, an increase in phase locking during the delay period could be observed that was most prominent between 3Hz and 9Hz, which corresponds well to both the frequency range in which we observed amplitude increases during the delay and the frequency range traditionally associated with the theta band. As can be seen in Figure 3.3A, phase locking in theta starts to increase towards the end of the sample stimulus presentation and is maintained throughout the entire delay period for 1.5s.

In addition to the population analysis I and my collaborators also assessed phase locking for individual LFP pairs during the baseline and the delay for several frequency bands. Across all frequencies tested, 168 of 507 pairs showed significantly elevated phase locking during the delay phase compared to baseline ($p < 0.001$ based on permutation tests, see Methods for details; M1: 135/332, M2: 33/175 pairs, Z-Test for significant proportion, $p < 0.001$). The proportion of pairs showing elevated phase locking during the delay was highest in the theta band (3-9Hz) with 29% and 13% for monkey 1 and 2, respectively (Z-Test for significant proportion, $p < 0.01$). Although theta phase locking was generally increased during the delay, there were also significant decreases in theta locking in a small number of pairs (M1: 13/332, M2: 12/175 pairs). However, this proportion was not significant (Z-Test, $Z < 1.21$, $p > 0.05$) and thus the further analysis focuses on pairs showing elevated theta

synchrony.

During the delay, no other frequency bands seemed to show systematic changes in phase locking. Consistent with this, only a small proportion of pairs showed a significant increase in PLV during the delay in the beta (16-36Hz) and gamma (42-97Hz) frequency range (cf. Figure 3.3B). These proportions were significantly smaller than for theta in both monkeys (Proportions M1: $P_\theta = 0.29$ vs. $P_\beta = 0.07$, $P_\gamma = 0.12$, M2: $P_\theta = 0.13$ vs. $P_\beta = 0.04$, $P_\gamma = 0.06$, χ^2 Test, $\chi^2 > 12.3$, $p < 0.001$ for all comparisons).

Interestingly, there was also a transient increase in phase locking around 16-20Hz between V4 and lPF during the sample stimulus presentation (see Figure 3.3A), which was significant in 29% of the pairs in monkey 1 and 20% of the pairs in monkey 2 ($p < 0.001$, Z-Test). The transient nature of this effect suggests that it is related to stimulus-locked visual evoked potentials that typically arise during the presentation of a visual stimulus. Since this increased phase locking, however, did not extend into the memory period, the subsequent analyses focus on memory-related theta phase locking.

In addition to assessing the magnitude of phase locking, also the timing relationship of theta synchrony between V4 and lPF was examined as it provides insights into the nature of the interaction between different regions: while oscillatory synchrony through common input is often associated with a zero-degree phase lag between oscillatory signals, non-zero degree phase shifts may imply more direct interactions between oscillatory processes (Fries, 2005). Nevertheless, the possibility of synchrony through common input with non-zero phase lags due to different latencies from the source to the different target regions cannot be completely ruled out.

When assessing the distribution of time shifts during the delay for theta frequencies (3-9Hz), I and my collaborators found a prominent peak of time differences around 15ms (Figure 3.3C) which could imply that synchrony between V4 and lPF is mediated via interactions between both areas rather than common input to these regions. Moreover, the time shift estimates are similar to a previous study measuring phase synchrony between V4 and a different prefrontal region, the frontal eye fields (FEF) (Gregoriou et al., 2009). In addition, these findings fit well with estimates of response latencies between different anatomically connected visual areas (Nowak and Bullier, 1997). Thus, time shifts between many V4-lPF pairs could potentially be explained by conduction times and synaptic delays via direct interactions between our recording sites. Nevertheless, it should be noted that there was also a high variance in time shifts across LFP pairs. This is reflected in a higher median time shift (25ms) and a relatively large inter-quartile range of 36ms (time shifts were computed per frequency and then averaged across frequencies). Thus, while conduction- and synaptic related delays could mediate direct interactions between pairs showing small shifts in phase locking (10-15ms), it might be possible that larger time shifts in phase locking are mediated by multi-synaptic connections involving intermediate cortical regions connected to both V4 and lPF, for example inferior temporal cortex.
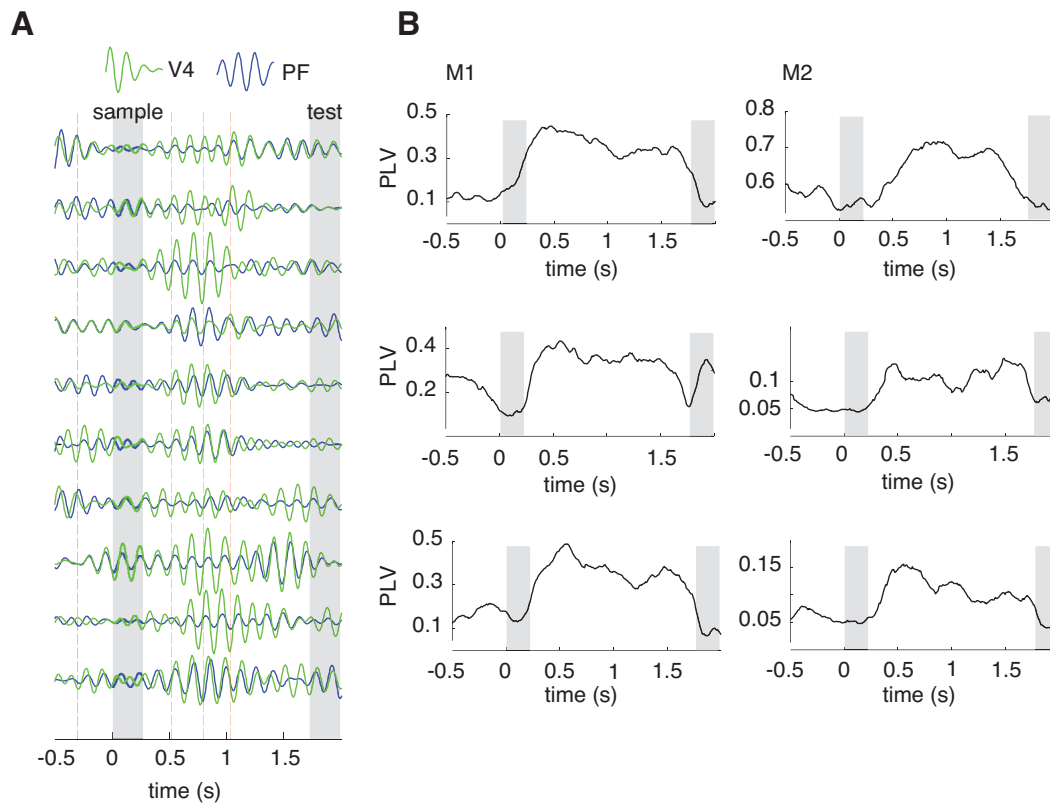
Figure 3.2: Phase-locking between V4 and lPF in theta band. **A:** Example of theta-filtered (at 6.8Hz) single-trial LFP traces from two simultaneously recorded V4 (green) and lPF (blue) channels. Both channels show a consistent phase relationship during the delay period but not during the baseline (red vertical lines are for illustration purposes only). Grey regions indicate the presentation intervals of the sample and test stimulus, respectively. **B:** Representative examples of theta phase locking as measured by the phase locking value (PLV) as a function of time for monkey 1 (left column) and monkey 2 (right column). In all shown pairs there is a significant increase in phase locking in the theta range. Note that the upper left trace corresponds to the example shown in A.

Taken together, the results show that local field potentials recorded in extrastriate area V4 and the lateral prefrontal cortex show enhanced oscillatory phase locking in the theta frequency range during the memory period of a visual memory task. These findings provide new evidence for the emergence of a reliable timing relationship between distant cortical areas during visual memory and support previous reports on the importance of theta oscillations in mediating long-range interactions between different cortical and sub-cortical areas during memory processing (Klimesch, 1996; Sarnthein et al., 1998; Stam et al., 2002; Fell et al., 2003; Jones and Wilson, 2005; Pasternak and Greenlee, 2005; Cashdollar et al., 2009).

Next I and my colleagues analyzed whether increases in theta-band phase locking
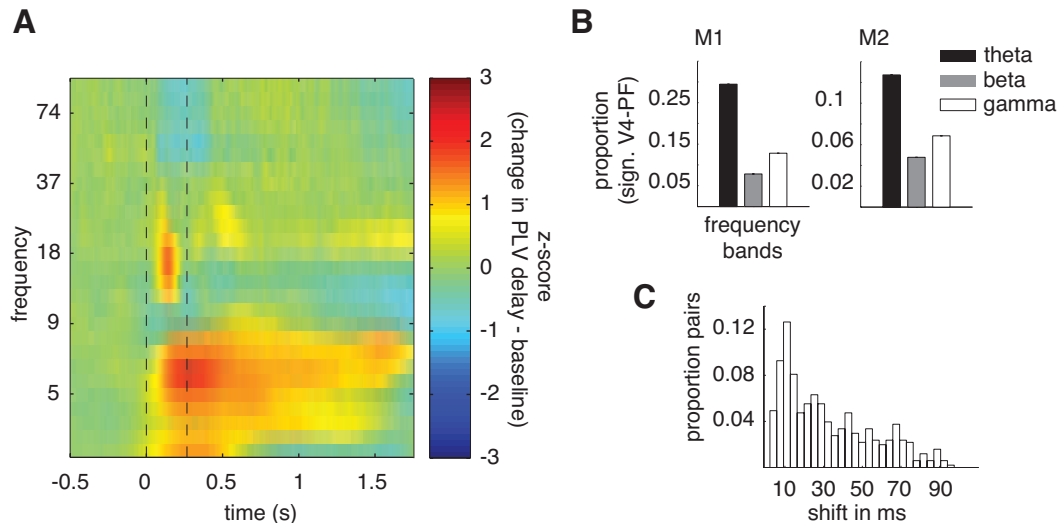
A

B

C

Figure 3.3: Phase-locking between V4 and lPF in theta band. **A:** Time-frequency representation of the averaged normalized PLV difference (Z-score) between baseline and delay across all recorded pairs. Dashed vertical lines indicate the on- and offset of the sample stimulus. **B:** Proportion of significantly locked pairs during delay as a function of frequency band for both animals. Significance thresholds are based on a bootstrapping procedure comparing baseline and delay using $p < 0.001$. **C:** Distribution of time shifts between all V4-PF pairs during the delay period in the frequency range of 3Hz to 9Hz.

between oscillatory LFP activity was associated with increases in spike-phase locking. To assess spike-phase locking, each spike was assigned its respective analytic phase value of the simultaneous LFP oscillation during the baseline and delay period. Subsequently Rayleigh's Z score was used as a test statistic to assess whether spiking of individual units was significantly locked to particular theta phases. In addition, histograms were generated containing the spiking probability as function of phase bins and fitted von Mises density functions to yield estimates of the mean direction μ (i.e. preferred phase) and the concentration κ of spiking (i.e. how peaked the spiking distribution is around the preferred phase). After normalizing for differences in spike rates between baseline and delay (see Methods for details), the proportion of significantly locked units as well as median κ between these task periods was compared. Both approaches have been extensively used to describe phase locking characteristics of individual neurons (Siapas et al., 2005; Zar, 2008).

3.3.2   Spike phase locking in theta increases during the delay within and between areas

Within-area spike phase locking of V4 and prefrontal neurons was first examined during baseline and delay (number of simultaneously recorded unit-channel pairs in V4 M1: 426, M2: 175 and lPF M1: 216, M2: 181). Note that channels or units

were not selected based on any criteria, for example task-related activity or visual responsiveness. Thus, the data represent a completely unbiased sample of LFP and single-unit activity in both regions.

It can be observed that across all recorded units in V4 and lPF, theta-band spike-phase locking was significantly higher during the delay than during baseline (Ranksum Test comparing median κ, $Z > 2.7$, $p < 0.01$ in both areas). Also, V4 units showed a greater enhancement in theta locking than prefrontal units (V4: 23% vs. lPF:14% increase in locking). While this effect was observable across all theta frequencies (3Hz to 9Hz) for one animal ($Z > 9.4$, $p < 0.001$ for both regions), in the other animal the increase in spike-phase locking was mainly present below 6Hz ($Z > 2.5$, $p < 0.01$ for both regions). In addition, similar to the general increase in theta locking at the population level, there was a significant increase in the proportion of significantly coupled units from baseline to delay in V4 (17% increase) and lPF (30% increase, $Z > 2.1$, $p < 0.05$ for both regions and animals). Figure 3.4 shows the median concentration parameters for all V4 and prefrontal units during baseline and delay as well as the phase of firing probability histograms averaged across all significantly locked units in both regions ($p < 0.01$).

Taken together, the results on enhanced spike-phase synchrony during visual memory within each region are similar to previous studies investigating spike-phase locking in V4 (Lee et al., 2005) and prefrontal regions (Siegel et al., 2009) during memory tasks. The more interesting question, however, is whether spiking in V4 and lPF is phase locked to theta oscillations in the respective other area and whether this effect is enhanced during the memory period of the task. If phase synchrony at the level of LFPs provides the timely basis to coordinate spiking output between both regions, the present findings on increased LFP-phase locking between V4 and lPF during visual memory suggest that this could be the case.

To address spike-phase locking between V4 and lPF, the activity of 660 simultaneously recorded lPF LFP-V4 unit pairs (M1: 458, M2: 202) and 593 V4 LFP-lPF unit pairs (M1: 335, M2: 258) was analyzed. Panels A and B of Figure 3.5 show several example units from V4 and lPF whose spike rate is significantly modulated as a function of theta phase during the delay in the respective other area ($p < 0.01$, see Figure caption for estimates of preferred phase, κ and Z scores of all example units). Similarly, at the population level a significant enhancement in inter-area spike-phase locking in theta frequencies during the delay compared to baseline across all V4 and prefrontal units could be observed (Wilcoxon signed rank test, V4 units: $Z > 3.5$, $p < 0.01$, lPF units: $Z > 3.4$, $p < 0.01$ for both animals). Moreover, the increase in inter-area coupling was larger than the effect we observed within the regions. During visual memory, spike-phase locking between V4 units and prefrontal LFPs was enhanced by 31% and between prefrontal units and V4 LFPs by 19%. Consistent with this effect at the population level, the proportion of significantly locked V4 and prefrontal units was higher during delay than during baseline with an average increase of 32% for V4 units (χ^2 test, $p < 0.01$) and an average increase of 29% for lPF units (χ^2 test, $p < 0.01$). Figure 3.6A shows the distribution of log transformed Rayleigh's Z scores as well as the median κ parameters for baseline and delay.
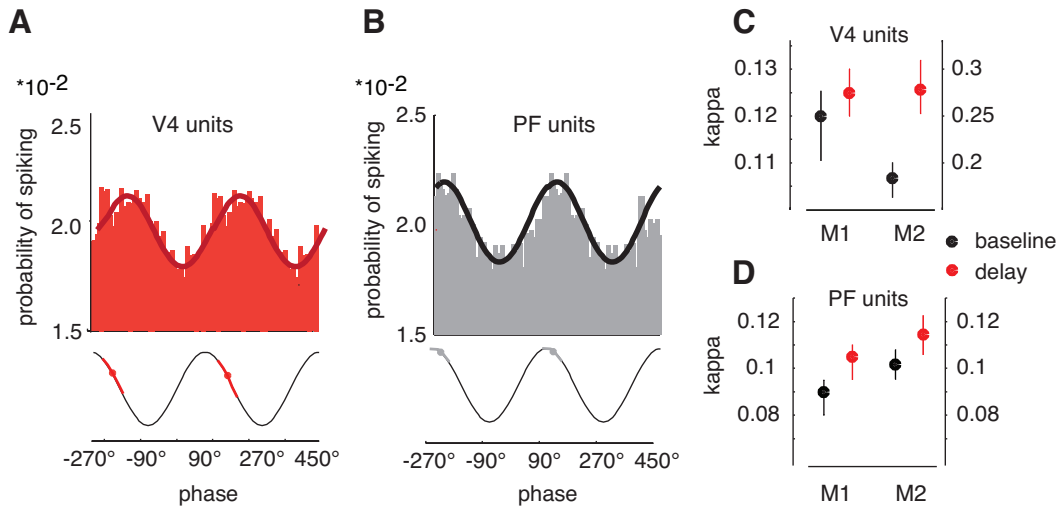
Figure 3.4: Within the theta frequency range inter-area spike-phase locking is enhanced during the delay. **A,B:** Phase of firing probability averaged across significantly locked V4 and prefrontal neurons ($p < 0.01$). Below the histograms the reference LFP cycle is shown. Red and gray symbols mark the average preferred phase across the units and lines show the circular 95% confidence intervals for V4 and prefrontal units, respectively. Plots depict two cycles for illustration purposes. **C,D:** Median kappa values for baseline and delay for V4 and prefrontal units.

I and my collaborators next compared the strength and timing of phase locking between V4 units locked to lPF theta and lPF units locked to V4 theta. Figure 3.6B shows the average phase of firing probability histogram across all significantly locked V4 and lPF units along with the von Mises fits and the referenced oscillation cycle ($p < 0.01$, V4: $N = 125$, upper plot, lPF: $N = 86$, lower plot). While V4 units tended to fire towards the beginning of the theta peak, lPF units had preferred phases towards the second half of the peak and the falling edge of the theta cycle. This difference was well reflected in the distribution of preferred phases of V4 and lPF units, which were significantly different (Watson-Williams test, mean preferred phase $72°$ for V4 units and $108°$ for lPF units, $F = 8.41$, $p < 0.01$, see also Figure 3.6C).

Finally, the distribution of the concentration parameter κ obtained from von Mises fits was compared for each unit-LFP pair along with the median estimates of κ across all significantly coupled pairs ($p < 0.01$) between V4 and lPF units (panel D). Interestingly, significantly higher concentration parameters for V4 units locked on lPF theta could be observed than the other way around (both: $Z = 5.9$, M1: $Z = 4.01$, M2: $Z = 1.99$, $p < 0.05$).

In summary, the analyses revealed that V4 neurons exhibit significant phase locking to theta rhythms in lPF cortex and likewise, lPF neurons exhibit phase locking to V4 theta oscillations. More importantly, the degree of phase locking was
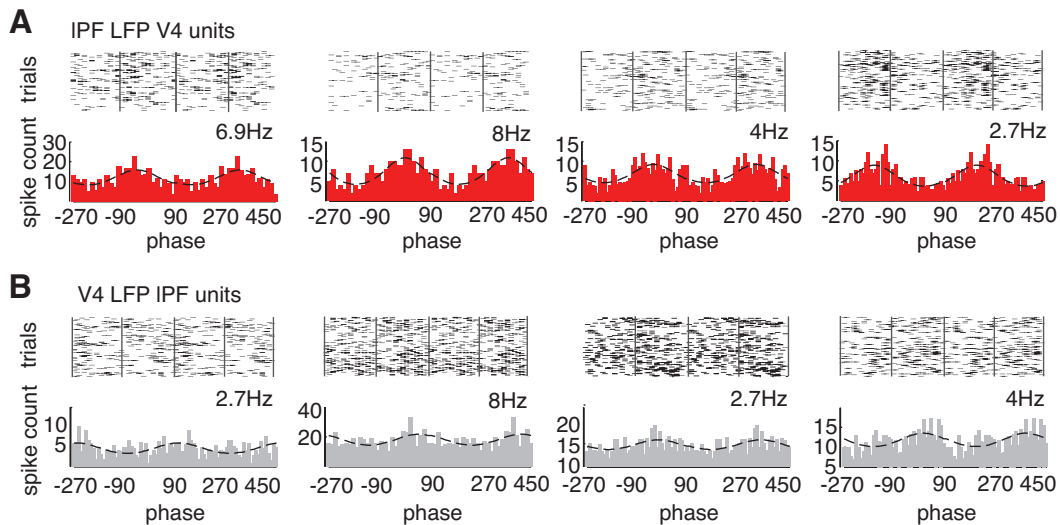
Figure 3.5: Inter-area theta phase locking of V4 and prefrontal neurons during visual memory. **A:** Spiking activity of single units in V4 significantly locked to the theta rhythm in lPF ($p < 0.01$). Raster plots show spike trains centered at the trough of the theta oscillation during the delay period. The resulting segments were stacked up across individual trials. Also shown are the von Mises fits for each example unit parameterized by mean direction μ and concentration parameter κ. The parameters of the von Mises fits (from left to right) are as follows (μ, κ): (-0.68, 0.39), (0.09, 0.22), (-2.6, 0.46) and (-0.39, 0.34). The first two plots correspond to examples from monkey 1, the latter two from monkey 2. The corresponding log transformed Rayleigh's Z scores are: 2.17, 2.08, 2.76 and 2.71. Spike trains and histograms are plotted for two cycles for visual clarity. Phase values in degrees on the horizontal axes refer to the mapping onto a sine wave. **B:** Spiking activity of single units in lPF significantly locked to the theta rhythm in V4 ($p < 0.01$). Comparable plots as in (a). The parameters of the von Mises fits (from left to right) are as follows (μ, κ): (1.78, 0.29), (0.96, 0.21), (-0.24, 0.27) and (0.28, 0.25). The corresponding log transformed Rayleigh's Z Scores are 1.99, 1.72, 2.06 and 1.64.

dependent on the task period: during the memory period of the task spikes in V4 as well as in prefrontal cortex tended to occur more concentrated around particular theta phases in the respective other area as compared to the pre-stimulus fixation period. The increase in spike-phase locking during the delay period is in good agreement with enhanced inter-cortical locking at the level of LFPs between the regions. Thus, theta synchrony at the mesoscopic level of electric potentials may provide a mechanism by which spiking activity becomes more reliably coordinated in time across two distant cortical areas that are both involved in visual working memory. These findings support the notion that visual memory comprises multiple memory systems that not only include regions that have been traditionally associated with the storage of information, such as the prefrontal cortex, but also include sensory cortical areas, such as V4 (Pasternak and Greenlee, 2005). These findings further suggest that theta oscillations play an important role in mediating memory-related
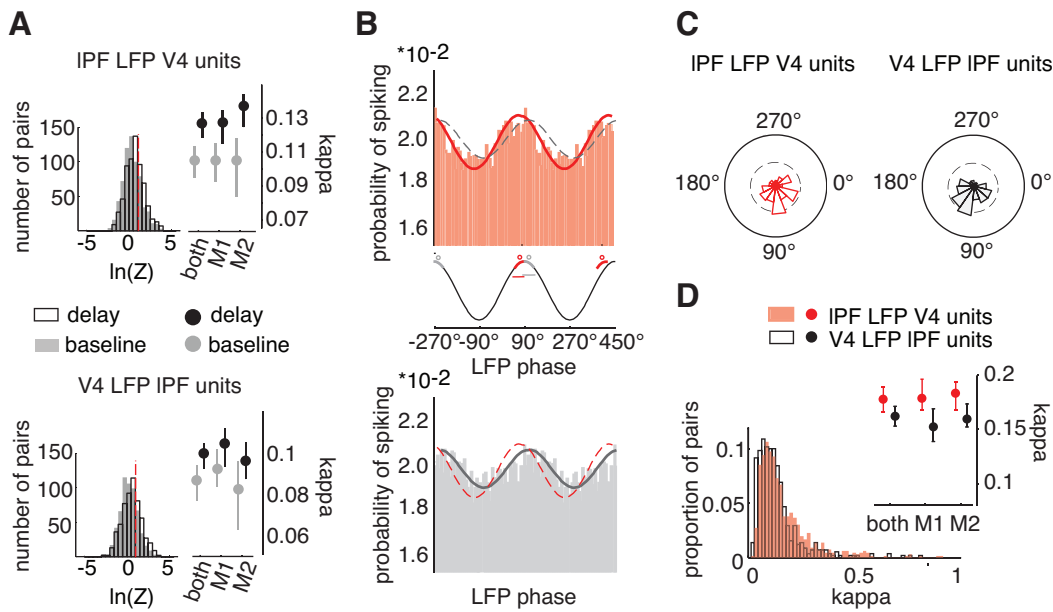
Figure 3.6: (see next page for Figure caption)

long-range interaction and communication between these systems (Fries, 2005; Fell and Axmacher, 2011).

3.3.3 Inter-cortical theta coupling is significantly higher for correctly remembered stimuli and correlates well with session-to-session variations in memory performance

Thus far the analyses revealed an increase in inter-cortical theta coupling both at the level of LFPs and of spiking activity that is specific for the memory period of the task. Based on these results, I and my collaborators were interested in the functional significance of this effect and thus we asked whether changes in phase locking are also associated with changes in memory performance. The task design allowed us to investigate this particular question, as the animals did not only have to remember non-degraded stimuli, for which task performance is almost perfect, but also natural images that were degraded with visual noise. Thus, phase locking between trials in which the animals committed a substantial amount of errors could be compared.

First, phase locking between LFP pairs and spike-LFP pairs was computed separately for correct and incorrect trials. The findings are summarized in Figures 3.7 and 3.8. Panel A of Figure 3.7 shows the average raw phase-locking value across all V4-lPF LFP pairs obtained from incorrect (left) and correct (right) trials as a function of frequency and time. Phase locking between V4 and lPF showed to be increased for correct compared to incorrect responses during the delay and that this

Figure 3.6: Inter-area theta phase locking of V4 and prefrontal neurons during visual memory. **A:** Histograms showing the distribution of log transformed Rayleigh's Z scores during the baseline as well as during the delay. The upper histogram corresponds to V4 units, the lower histogram to prefrontal units. The red line indicates the significance associated with the corresponding log Z-score $ln(Z) > ln(-ln(p))$ at $p = 0.05$. To the right of each histogram median κ values from von Mises fits of spike-phase distributions during baseline (gray) and delay (black) are shown. Error bars indicate the $\pm 34\%$ confidence intervals around the median obtained using a bootstrapping procedure. Prefrontal phase locking on V4 theta as well as V4 phase locking on prefrontal theta is significantly elevated during the delay compared to the baseline period when comparing κ. **B:** Phase of firing probability averaged across significantly locked V4 and prefrontal neurons along with von Mises fits (red V4, gray lPF, with mean direction μ and concentration parameter κ for V4 units $\mu=1.26$ and $\kappa=0.08$, prefrontal units $\mu=1.86$ and $\kappa=0.06$). Between histograms the reference LFP waveform is shown. Red and gray symbols mark the average preferred phase across the units and thick red/gray lines show the circular 95% confidence intervals for V4 and prefrontal units, respectively. **C:** Circular distribution of preferred phases of V4 (red) and prefrontal (black) neurons. The mean directions were significantly different (Watson-Williams test, $F = 8.4$, $p < 0.01$). **D:** Distribution of concentration parameter κ for all pairs of simultaneously recorded V4 units and prefrontal LFPs as well as prefrontal units and V4 LFPs and median κ values for V4 units locked to lPF theta and vice versa plotted across units from both animals and per animal (for pairs statistically significant at $p < 0.05$). Median κ values were significantly higher for V4 units locked on prefrontal theta than the other way around.

effect occurred again predominantly in the theta band. This is further illustrated in panel B, which shows the normalized difference between the two conditions (d-prime) averaged across all recorded pairs as a function of frequency during the delay period. Interestingly, the difference in phase locking between correct and incorrect trial conditions varied across time, with larger differences occurring during the early period of the delay and again during the later phase.

To test whether the increase in phase locking for correct trials is significant, d-prime values using the original data sets were compared with d-prime values obtained from shuffled data sets, i.e. in which phase locking was computed between groups of trials for which the "correct" or "incorrect" conditions had been randomly assigned. Panel C shows the normalized difference in phase locking between correct and incorrect trial groups from the original (green) and shuffled (gray) data sets per frequency averaged across the delay period. In both animals, phase locking was significantly elevated in theta frequencies ($p < 0.01$ for frequencies averaged within 3-9Hz, for both animals, see also panel D for p-values of individual frequencies) during the delay phase of trials in which the animals would later correctly identify the previously shown stimulus compared to trials in which the animals would fail to do so.

The increase in LFP phase locking was also associated with an increase in spike-phase locking for correct compared to incorrect trials. Interestingly, this effect seemed to be stronger for V4 units than for prefrontal units. While median κ
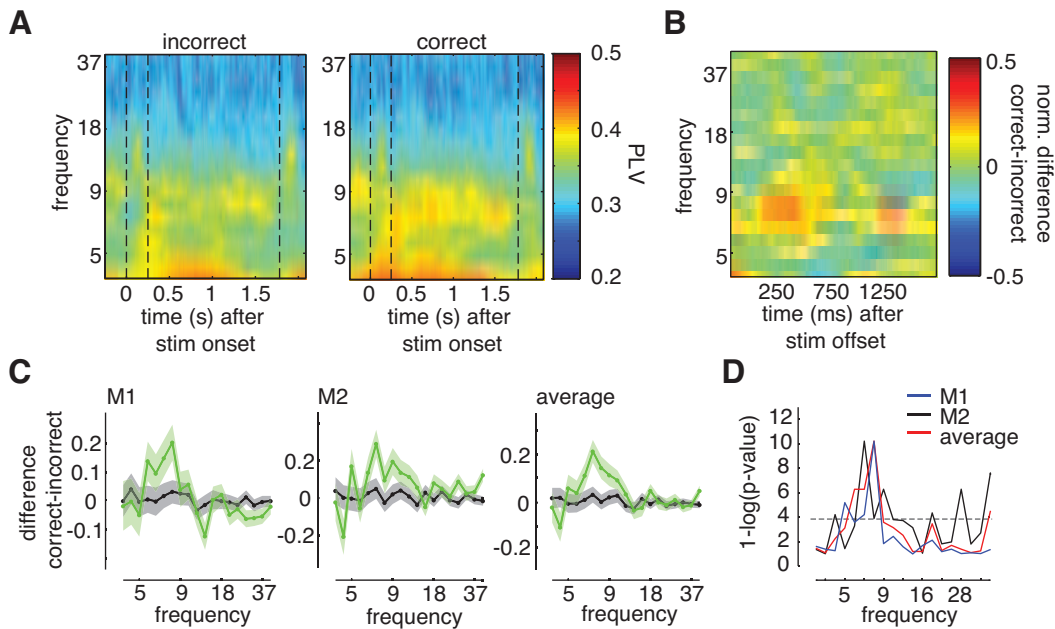
Figure 3.7: Phase locking during the delay is higher for correct compared to incorrect trials in theta. **A:** Phase locking (average across all pairs) per frequency during the course of the task. Phase locking is increased during correct (right) compared to incorrect trials (left) for theta frequencies, peaking between 5 to 9Hz. **B:** Normalized difference (d-prime) of correct minus incorrect trials during the delay period shows increases during the early delay phase (around 250ms after stimulus offset) and again during the late delay phase (around 250ms prior to test stimulus onset). **C:** Normalized difference (d-prime) in phase locking between correct and incorrect trials averaged across the delay per frequency for pairs from monkey 1 (left), monkey 2 (middle). In both animals, there was a systematic increase in phase locking for theta frequencies. Error bars correspond to ± 1SD. **D:** Level of significance (1-log(p-value)) for comparison between d-primes from original and shuffled data sets per frequency.

values were significantly higher for correct compared to incorrect trials in V4 in both animals (Signed rank test, $Z > 2.95$, $p < 0.001$ for both animals), this was only true in one animal, but not in the other in lPF (Signed rank test M1: $Z = 0.37$, $p > 0.05$, M2: $Z = 2.1$, $p < 0.05$). To further investigate whether the difference in spike phase locking between correct and incorrect responses is larger in V4 neurons than in prefrontal neurons, d-prime values (normalized difference in κ between correct and incorrect trials) were compared between reciprocally matched pairs: d-prime values of V4 neurons recorded at V4 site X that were locked on prefrontal LFP site Y were paired with d-prime values of prefrontal neurons recorded at site Y that were locked at V4 site X. In a significant majority of pairs, V4 neurons exhibited higher d-prime values, i.e. larger differences in locking between correct and incorrect responses than their matched prefrontal neurons recorded at the respective other
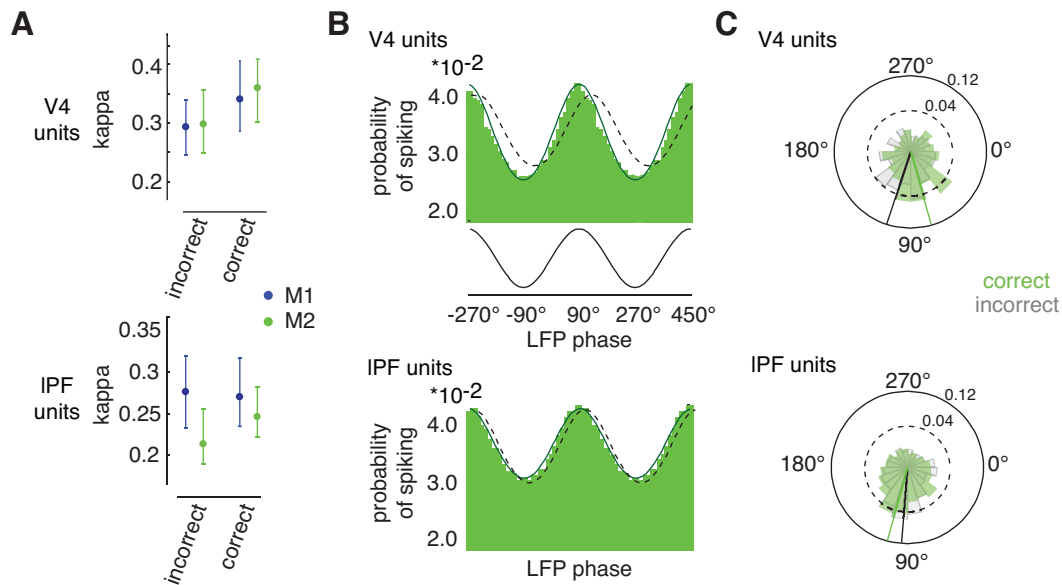
Figure 3.8: Phase locking during the delay is higher for correct compared to incorrect trials in theta. **A:** Median κ estimates from phase-of-firing probabilities of correct and incorrect trials for V4 neurons (left) and prefrontal neurons (right, locking was calculated for significantly locked pairs, $p < 0.01$). Spike concentration κ was significantly enhanced in correct trials for both animals in V4, whereas there was no systematic effect in prefrontal neurons. Error bars correspond to the ± 34th percentile around the median. **B:** Phase of firing probability histogram of V4 neurons (upper) and prefrontal neurons for correct trials aligned to preferred phase per neuron and then averaged across neurons. Lines correspond to von Mises fits to correct (green) and incorrect (gray) spike histograms. **C:** Distribution of preferred phase of V4 neurons and prefrontal neurons in correct and incorrect trials. Preferred phases were significantly different in correct vs. incorrect trials for V4 neurons, whereas we found no difference for prefrontal neurons.

LFP site (Binomial Test, $p < 0.02$, for both animals).

In addition to an increase in spike phase locking during correct trials, there were also differences in the timing relationship between spiking and theta oscillations between correct and incorrect trial types. To illustrate this effect, panel B of Figure 3.8 shows the phase of firing probability histograms across V4 and prefrontal units aligned to their preferred phase and panel C depicts the circular distribution of preferred phases for both conditions and areas. In both animals, differences in preferred phase between correct and incorrect trials were significantly larger for V4 units than for prefrontal units (Watson-Williams Test, $F > 5.3$, $p < 0.05$ for both animals). In other words, preferred phases were significantly different in correct vs. incorrect trials for V4 neurons (mean phase correct $71.5°$, incorrect $103.5°$, $F = 16.69$, $p < 0.001$), but not prefrontal neurons (mean phase correct: $105.1°$, incorrect $94.4°$, $F = 2.1$, $p > 0.05$). Interestingly, preferred phases of V4 and prefrontal neurons showed a higher similarity during correct compared to incorrect
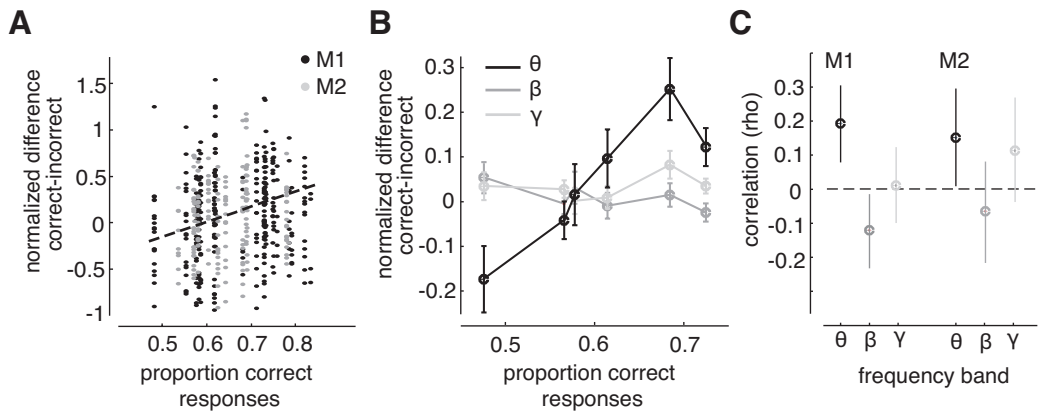
Figure 3.9: Correlation of phase locking with performance across single sessions. **A:** Scatter plot showing normalized difference between correct and incorrect responses (y-axis) as a function of the proportion of correct responses per session. For every session, the performance measure was pairs with d-primes of all pairs recorded during that session. Black symbols correspond to monkey 1, gray symbols to monkey 2. **B:** d-prime values binned into equally spaced percentiles of proportion of correct responses for three frequency bands (theta: black, beta: dark gray and gamma: light gray). Error bars correspond to ±1SE. **C:** Rank correlation coefficients estimated from raw data as shown in A for three frequency bands and each monkey separately. Both monkeys show a significant positive correlation between performance and d-prime values of phase locking in theta, whereas no significant effect was found for the other frequency bands. Error bars correspond to 95th percentile confidence intervals.

trials.

For the previous analyses, correct and incorrect responses averaged across different sessions were compared. On average, correct identification was shown to be associated with increases in theta-band phase locking for correct compared to incorrect responses. Does this difference in phase locking also co-vary with changes in performance across single sessions? To address this question, d-prime values (correct minus incorrect) were correlated per pair with the proportion of correct responses across sessions. The results are summarized in Figure 3.9. Indeed, the d-prime values were positively correlated with the proportion of correct responses, i.e. larger (positive) d-primes occurred during sessions with higher performance (see panel A). This effect was again most dominant in the theta frequency range and can be seen more clearly in panel B, which presents pooled d-prime values across different percentiles of the proportion of correct responses for the three tested frequency bands. Consistent with this, significant positive rank correlation coefficients for theta frequencies, but not in beta or gamma frequency bands could be observed (panel C, average rank correlations: M1 and M2: ρ_θ=0.18 and 0.14, $p < 0.05$, ρ_β=0.11 and -0.06, $p > 0.05$ and ρ_θ=0.01 and 0.11, $p > 0.05$).

3.4   Discussion

The goal of this study was to investigate the interaction between two distant cortical brain regions that have both been previously linked to memory processing. Specifically, simultaneously recorded LFPs as well as single unit activity (SUA) from ensembles of neurons in the extra-striate visual area V4 and the lateral prefrontal cortex were analyzed and it was examined whether neural synchrony based on oscillatory components of the LFP and concurrently measured spiking activity is associated with the retention of visual information during the delay period of a delayed-match to sample task.

To assess neural interaction, first the phase relationship between LFP oscillations in both areas was examined. LFP signals from V4 and lPF exhibited a consistent phase relationship during the delay period of the memory task that was significantly reduced during the baseline period. The consistent phase relationship was reflected in a significant increase in the phase locking value (Lachaux et al., 1999) and occurred predominantly in the theta frequency range (3-9Hz). In addition, the increase in phase locking during the delay was accompanied by an overall increase in theta power in both areas during the delay interval compared to the baseline condition.

These results suggest that neural interactions based on oscillatory synchrony between two distant cortical regions are related to the maintenance and possible communication of visual information during visual short-term memory. Our findings are consistent with previous studies recording human EEG or intra-cortical local field potentials that also showed increased theta power in either frontal or occipital regions during memory processing (Kahana et al., 1999; Raghavachari et al., 2001; Schack et al., 2002; Sederberg et al., 2003; Rainer et al., 2004; Lee et al., 2005; Siegel et al., 2009) as well as previous observations of increased coherence between EEG signals from frontal and occipital regions in human subjects (Sarnthein et al., 1998; Weiss et al., 2000; Stam et al., 2002). Taken together these findings support the notion that prefrontal and visual cortical areas are part of a large scale network involved in visual memory processing in primates (Klimesch, 1996; Pasternak and Greenlee, 2005).

For many V4-lPF pairs the synchronous theta oscillations in V4 and lPF were shifted in time by about 10-15ms, which is roughly consistent with previously reported time shifts between LFPs recorded in V4 and the frontal eye fields, a different prefrontal region (Gregoriou et al., 2009). In addition, the observed time shift is about an order of magnitude shorter than the cycle lengths of the underlying theta oscillations (120-250 ms) which is similar to what has been reported for example for gamma or beta coupling (Fries, 2005; Gregoriou et al., 2009). It has been suggested that phase shifts occur in this time range in order to account for the conduction delays and integration times between different areas, such that presynaptic firing affects postsynaptic neurons at their peak depolarization phase, i.e. during their most excitable state (Nowak and Bullier, 1997; Volgushev et al., 1998; Fries, 2005; Tiesinga et al., 2008; Gregoriou et al., 2009). This could also be the case in the

present study, where for example V4 spiking is structured to affect downstream areas including the lPF during their most excitable period or vice versa.

To determine whether changes in neural synchrony at the level of the LFP are indeed accompanied by changes in the locking of spiking activity to oscillations in the respective other area, I and my collaborators measured spike-phase locking of V4 units to prefrontal theta and vice versa. An increase in locking of V4 single unit activity to the phase of the prefrontal theta rhythm ("spike-phase locking") across areas could be observed: prefrontal and V4 spikes occurred around distinct phases within the theta cycle of the respective other area. Again, this effect was significantly stronger during the memory phase of the task. This finding, however, might not be that surprising given that spikes within each area appeared to be phase-locked to their "own" theta rhythm. Thus, if theta oscillations between the areas become synchronized during the delay, spike phase locking across areas is likely to occur. Nevertheless, our findings suggest that the increase in LFP-based theta synchrony mediates this increased spike-phase locking between the regions, and thus influences the timely coordination of neuronal activity from distant brain areas during memory processing. This is further supported by the fact that memory-related enhancement in theta spike-phase locking was stronger between areas than within areas.

In addition to the general increase in inter-area spike phase locking during the delay, spiking activity in V4 was more strongly locked to prefrontal theta than vice versa. This could imply that prefrontal regions are more involved in generating and sustaining theta oscillations during memory processing as V4 spiking appears to be more sensitive to prefrontal theta than the other way around. However, given the idea that phase alignment of presynaptic spikes to postsynaptic membrane potential oscillations is a mechanism that establishes synchronous input into an area and therefore may increase the probability of postsynaptic firing (Klimesch, 1996; Volgushev et al., 1998; Singer, 1999; Salinas and Sejnowski, 2001; Fries, 2005; Womelsdorf et al., 2007; Haider and McCormick, 2009), the observations in this study could indicate that V4 spiking may be consequently more effective in driving prefrontal activity than vice versa. Taken together, the asymmetry in spike-phase locking between the regions implies a possible directedness in the interaction and communication pattern between the regions, of which the exact nature and timing relationship remains to be established.

Another important observation was that theta synchrony of LFPs between V4 and lPF predicted subsequent recognition of visual stimuli within each session and, in addition, correlated well with session-to session variations in memory performance. These findings are consistent with recent observations relating strength of theta coupling to memory performance in humans using magneto- or electroencephalography (Sederberg et al., 2003; Guderian and Duzel, 2005; Sato and Yamaguchi, 2007; Fell et al., 2008; Cashdollar et al., 2009; Fuentemilla et al., 2010). More importantly, successful recognition was associated with tighter coordination of spike timing with theta oscillations in the respective other region. This effect is similar to a previous report of locally enhanced spike-phase locking within the medial

temporal lobe during successful retention in humans (Rutishauser et al., 2010). Our findings thus not only confirm a direct functional link between spike timing relative to theta oscillation and memory performance but also - at least to my knowledge - for the first time extend this principle to long-range interactions between distant cortical areas. In addition, the findings are consistent with the hypothesis that phase alignment between rhythmically active neural ensembles promotes effective communication between these ensembles (Fries, 2005; Fell and Axmacher, 2011). Thus, theta phase synchronization may lead to enhanced memory performance by facilitating communication between V4 and lPF.

Interestingly, both V4 and prefrontal units had their preferred phase shortly before and after the peak of the theta oscillations, respectively, i.e. near the depolarized phase of the oscillation. These conditionsare ideal to induce long-term potentiation (LTP) (Pavlides et al., 1988), which is a synaptic mechanism that promotes the encoding of sensory information into memory, see also Martin et al. (2000); Fell and Axmacher (2011). Thus, theta band across-area synchronization between V4 and lPF coordinates spiking output in a way that is ideal for the stimulus to undergo memory encoding. This effect is enhanced during the retention period of correctly remembered stimuli since inter-area coupling is greater for correct compared to incorrect trials. Thus, higher memory performance could reflect facilitated encoding of visual information into memory.

A remaining question is how theta-based synchrony between V4 and prefrontal cortex is generated. One hypothesis is that V4 and lPF directly interact and entrain each other's rhythms through cortico-cortical connections. The non-zero phase shift between synchronous LFP oscillations we found across areas supports this view. Direct entrainment could, for example, be easily established through monosynaptic and reciprocal connections between V4 and the lateral prefrontal cortex. However, to my knowledge, only one monosynaptic, reciprocal connection between V4 and a very localized area of the prefrontal cortex, the frontal eye field (FEF), has been reported thus far (Ungerleider et al., 2008). Since the used recordings were not located in this area but in the lateral part of the prefrontal cortex (area 46) anterior to the FEF (Walker, 1940; Barbas and Pandya, 1989) and given that it is unclear whether direct anatomical connections exist between our recording regions, it is questionable whether direct entrainment of theta synchrony via direct reciprocal connections between V4 and lPF lead to the increase in phase locking between the two regions.

A different way by which phase synchronization between V4 and lPF could be established is through indirect polysynaptic connections involving intermediate areas. For example, synchronized theta activity in V4 could entrain and/or synchronize intermediate areas such as the inferior temporal cortex (IT) and subsequently synchronize with local rhythms in prefrontal regions or vice versa. This hypothesis is consistent with reports on working memory-related theta frequency oscillations that have been observed in human infero-temporal areas (Sederberg et al., 2003; Raghavachari et al., 2006; Rizzuto et al., 2006) and is also supported by the anatomical connectivity between V4 and lPF: the major feedforward output from V4 goes

to IT (Ungerleider et al., 2008) and the major feedforward output from IT reaches lateral prefrontal cortex, especially its ventral portion (Ungerleider et al., 1989).

Another possibility is that neural ensembles in V4 and prefrontal cortex become synchronized by receiving common input from a third region (Basar and Bullock, 1992; Buzsaki, 2006). Although the non-zero degree phase shift argues against this hypothesis, differences in conduction delays between V4 and lPF and a third region could account for the lags in phase shift. A prominent candidate in synchronizing neuronal ensembles via oscillations in the theta band is the hippocampal formation including the surrounding parahippocampal region (Sirota et al., 2008). This compound structure is located in the medial temporal lobe of the primate brain and its major substructures are the hippocampus and the perirhinal, enthorhinal and parahippocampal cortices (Andersen et al., 2007).

Theta rhythms are a well-known characteristic in the hippocampal LFP of non-human primates, rats and humans (Green and Arduini, 1954; Winson, 1972; Arnolds et al., 1980; Kahana et al., 1999; Tesche and Karhu, 2000; Cashdollar et al., 2009) and the functional role of theta oscillations observed in the hippocampal formation is closely related to memory processing. For example, in the hippocampus, theta oscillations are associated with the generation and enhancement of long-term potentiation (LTP) and depression (LTD), a synaptic mechanism relevant for the encoding and retrieval of new information into and from memory (Buzsaki, 2002, 2006). In addition, hippocampal and cortical theta oscillations show similar modulations in response to cognitive demands in memory tasks (Klimesch, 1996; Tesche and Karhu, 2000). Moreover, it has been shown that theta-phase dependent firing within the hippocampus and several cortical regions (Winson, 1978; Lisman, 1999; Buzsaki, 2002; Lee et al., 2005; Siegel et al., 2009) but also between the hippocampus and cortical areas (Jones and Wilson, 2005; Siapas et al., 2005) is related to memory processes.

Based on the extensive bidirectional projections to many neocortical areas in non-human primates, including the lateral prefrontal cortex and V4 (Felleman and Van Essen, 1991; Miller, 1991; Suzuki and Amaral, 1994; Burwell, 2000; Ungerleider et al., 2008) and the functional similarity between hippocampal and cortical theta it has been suggested that theta oscillations synchronize widely distributed cortical areas during memory processing via (para-)hippocampal-cortical feedback loops (Miller, 1991; Buzsaki, 1996; Klimesch, 1996; Sarnthein et al., 1998; Kahana et al., 2001; Anderson et al., 2009). This could also be the case in the present study, where V4 and the lateral prefrontal cortex become synchronized in the theta band via cortico-hippocampal connections during memory processing.

Finally, the presented findings are also compatible with computational models proposing phase-dependent coding as a means to store and retrieve information in and from memory (Jensen and Lisman, 2005; Lisman and Buzsaki, 2008) as well as with more recent models on synaptic mechanisms that underlie working memory (Mehta et al., 2000; Mongillo et al., 2008; Verduzco-Flores et al., 2009). According to these models, visual information about an item held in memory is not exclusively stored by selective and sustained spiking activity during the delay period, a type of

response pattern that has been traditionally associated with maintaining information during memory (Fuster and Alexander, 1971; Miller and Cohen, 2001). Instead, information is thought to be maintained by calcium-mediated short-term synaptic facilitation and read out or reactivated by population spikes, i.e. synchronously spiking neural ensembles, at a rate corresponding to theta oscillations (Mongillo et al., 2008). Along these lines, oscillatory synchrony has been recently linked to the ability to flexibly route information between neural populations (Akam and Kullmann, 2010).

Taken together, the findings suggest that theta synchrony mediates the timed cooperation between neural ensembles in distant cortical areas. Increased coupling during successful retention may thus reflect more efficient and selective routing and gating of information during short-term memory and also promote encoding of information into memory. Ultimately, the observations are in agreement with the proposition that oscillatory synchrony, and in particular phase synchronization, is a general principle of integrating information within and between different cortical and subcortical areas during memory processing (Buzsaki and Draguhn, 2004; Fries, 2005; Fell and Axmacher, 2011).

Acknowledgements

Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-modulated Hebbian Learning

Contents

This chapter is based on the manuscript "Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-modulated Hebbian Learning" which is curently in preparation and will be submitted in 2011. The work was conducted together with Robert Legenstein (RL) and Wolfgang Maass (WM). The experiments were designed by myself, RL and WM. All the simulations and statistical evaluations were conducted by myself, and the paper was written by myself, RL and WM. This work proposes a novel framework for online training of recurrent neural networks which builds upon the reservoir computing paradigm and is based on mechanisms that are biologically more plausible than the traditionally applied fully supervised methods.

In spite of substantial research efforts, it has remained a mystery how complex computational structures can emerge in a biological neural system without an intelligent "supervisor" that masterminds connection patterns of networks of neurons, the synaptic weights of their connections, and the resulting network dynamics. Some progress has recently been made by showing that complex computations can be carried out by randomly connected neural circuits with chaotic dynamics, provided that an intelligent supervisor controls the evolution of the synaptic weights to readout neurons. Here I show that this surprising capability can also emerge without a supervisor, just on the basis of global signals that carry information whether the computational performance for a given task has recently improved. More precisely, random perturbations of neuronal activity suffice to steer - in the presence of global reward signals - randomly connected networks of neurons through trial and error to network configurations where they can carry out a diverse set of complex computations with high precision. In particular, the network can learn in this "self-supervised" fashion to maintain selected information in working memory, to adjust the processing of input streams in dependence of the current content of such working memory, and to route incoming information to other networks in a state-dependent manner. This demonstrates that complex computational structures can emerge from initially chaotic neural networks through experimentally supported reward-modulated synaptic plasticity mechanisms.

4.1   Introduction

Biological networks of neurons have to be able to carry out a large set of complex computational functions that involve memory in several ways. One example is the memory - and in turn the capability of production - of specific oscillatory activation patterns, which are a necessity for a large variety of body functions such as breathing and mastication or locomotive actions such as walking or swimming. Those functions are often associated with neural circuits that act as so-called "central pattern generators" (CPGs), which have been found primarily in the spinal cord and the brain stem (Marder and Bucher, 2001), but there is also evidence for movement-related rhythmic activity in primary motor cortex. For example, rhythmic activity related to jaw and tongue movements has been recorded in primary motor cortex of primates (Yao et al., 2002), an area that has been shown to be involved in motor skill learning (Molina-Luna et al., 2009; Hosp et al., 2011). Another example is short-term memory of previously observed input stimuli that is involved in virtually every type of higher cognitive process, e.g. while comparing subsequent visual stimuli or during processing of grammatical information. For example, it has been observed that single neurons in the prefrontal cortex of macaque monkeys respond to specific visual stimuli with a persistent increase or decrease of their firing rates for time intervals in the range of seconds during the delay period of delayed matching-to-sample tasks, suggesting that this behavior is a neuronal correlate of working memory (Fuster and Alexander, 1971; Goldman-Rakic, 1995; Miller et al.,

1996).

Evidence for CPG networks that produce rhythmic activity has been reported in a large variety of species (Wolf and Pearson, 1987; Baessler and Bueschges, 1998; Marder and Bucher, 2001; Prinz et al., 2004; Marder and Goaillard, 2006; Bueschges et al., 2008; Guertin, 2009), subserving body functions such as breathing, chewing and swallowing and locomotive actions such as walking, running or swimming. For the role of CPG networks for locomotion, see for example the reviews by MacKay-Lyons (2002); Grillner (2006) and Ijspeert (2008) as well as the references therein. One well-known example of the function of a CPG network that has been thoroughly studied is the pyloric rhythm of the crustacean stomatogastric nervous system. This network produces a stereotyped rhythmic activity with a cycle length of more than a second that controls the function of the constrictor muscles of the lobster's stomach. Previous results suggest that while the pyloric rhythm itself is similar across different individuals, the network parameters, e.g. synaptic strengths, may be quite diverse. Moreover, electrophysiological recordings in juvenile and adult lobsters have shown that the pyloric rhythm is maintained during growth, indicating that instead of genetically predetermined network parameters, the network activity is regulated by homeostatic mechanisms that are capable of searching for any attainable parameter setting that allows the system to generate the required periodic signal (Marder and Bucher, 2001; Prinz et al., 2004; Marder and Goaillard, 2006).

How, i.e. by means of which mechanisms, these networks can learn to accomplish these functions is one of the major issues being investigated in the field of computational neuroscience, but in spite of substantial research efforts and some major advances in the field, it remains poorly understood how neurons can learn to solve complex computational tasks in the face of ubiquitous plasticity of synapses, changes in neuronal excitability, and a host of other plasticity processes on different time scales. This becomes especially difficult when there is no intelligent "supervisor" that guides the learning process. However, the introduction of such a mechanism poses a problem because it is unclear how such a supervisor could be implemented in a biological system. Since neural networks in biological systems are highly recurrent, a learning mechanism that is capable of finding an appropriate setting for the synaptic weights in recurrently connected neural networks has to be employed for the tasks mentioned above.

The approach of "reservoir computing" has introduced a simple but computationally powerful paradigm for the training of recurrent neural networks. The two major exponents, Liquid State Machines (LSMs, (Maass et al., 2002)) and Echo State Networks (ESNs, (Jaeger, 2003)) were simultaneously and independently developed by Wolfgang Maass and Herbert Jaeger, respectively and share the same concept. Both approaches are based on the idea of using large networks (or "reservoirs") of sparsely interconnected neurons to which one or multiple input streams are provided. The reservoir then hosts a large variety of nonlinear functions of the interaction of the input stream(s) and the internal state of the recurrent network. While recurrent connections are usually generated randomly and stay constant, learning is restricted to the synaptic weights from the circuit to one or multiple

readout neurons. Each readout is trained to produce a particular target function. While ESNs focus more on the engineering approach, including applications in nonlinear signal processing and machine learning, LSMs were developed from the background of computational neuroscience and focus on the investigation of computational properties of neural microcircuits in the brain, usually using biologically realistic spiking neuron and synapse models (see Buonomano and Maass (2009); Lukosevicius and Jaeger (2009) for recent reviews).

More recently, this approach was extended by adding feedback loops that project the output of the readout neuron(s) back into the neural circuit and therefore make it take part in the system dynamics (Jaeger and Haas, 2004; Maass et al., 2007), the latter reference showing that in the idealized case without noise such a system can carry out any conceivable digital or analog computation on time-varying inputs. This huge repertoire of possible computations places even more weight on the open problem of how these networks can be "programmed" by a biological neural system, i.e., how a particular computation can be selected from an infinite repertoire of possible computations.

In these traditional approaches, the feedback is replaced by a noisy version of the target function during training, a method also known as "teacher forcing". This ensures that the network is driven to and stays within the correct regime during training and that the system is robust to imprecisions in the feedback. However, I am not aware of any biologically plausible mechanism that would serve as the basis of such a teacher signal. In fact, this would require a second independent network that is already able to generate the desired function. Moreover, the weights of the readout neurons are traditionally trained using batch regression, which means that the system is simulated for a substantial amount of time without a learning process that takes place online, and after sufficiently many snapshots of the network activity have been taken in order to properly represent the temporal dynamics of the network, the weights are estimated on the basis of these samples. As an alternative to batch regression, online regression mechanisms such as least mean squares (LMS) or recursive least squares (RLS) regession (Haykin, 2001) are also frequently used to train the readout neurons.

Modifying and extending these models, Sussillo and Abbott (2009) provided an online learning procedure for rate-based reservoirs which they called "FORCE learning". This method uses the online regression techniques stated above and feeds back the actual output of the readout neurons to the network also during training, which is more natural than using a teacher signal. The trained feedback is used to drive a chaotic network into a non-chaotic state in order to produce coherent patterns of activity. Chaoticity is an important aspect of their approach. In their experiments, they use spontaneously active networks of neurons that exhibit chaotic dynamics, i.e. irregular activity that is exponentially dependent to the initial conditions, in a way that exploits these dynamics for learning (Sussillo and Abbott, 2009). In contrast to their approach, the strength of the recurrent connections is traditionally scaled in order to avoid chaotic activity while staying close to the "edge of chaos" in order to yield maximal computational performance for a large set of computational

tasks in which the system is driven by time-varying inputs (Legenstein and Maass, 2007). However, due to the fact that Sussillo and Abbott predominantly take tasks into account in which there is no input provided to the system, chaotic activity is crucial in order to have a system that produces rich enough dynamics for generating the desired rhythmic activity. I adopt the idea of using chaotic networks in combination with feeding back the actual readout output during training (i.e., no teacher forcing) in our model. However, since Sussillo and Abbott apply standard online regression methods to train the readouts, their procedure relies on explicit information on the error of each readout unit, making it a fully supervised learning mechanism. This makes their FORCE learning procedure problematic in terms of biological plausibility.

An alternative line of research in computational neuroscience is concerned with learning on the basis of correlations of pre- and postsynaptic activity, better known as "Hebbian learning", named after Donald Hebb, who postulated in 1949 that the synaptic strengh between two interconnected neurons is increased if the pre- and the postsynaptic neuron fire together (Hebb, 1949). Besides unsupervised learning mechanisms which are solely based on pre- and postsynaptic correlations, several learning rules have been proposed that modulate correlation-based learning using a global modulatory signal. Based on this concept, a large family of learning rules has been proposed, both for rate-based (Barto et al., 1988; Mazzoni et al., 1991; Williams, 1992; Legenstein et al., 2010) and spiking (Xie and Seung, 2004; Fiete and Seung, 2006; Pfister et al., 2006; Izhikevich, 2007; Farries and Fairhall, 2007; Legenstein et al., 2008) neural network models in various contexts. Most of these models, however, take only network structures into account that are largely feedforward, or are employed in less complex computational tasks like the enhancement of the rate of a single neuron within a recurrent network while keeping the rates of other neurons approximately stable.

Here, I employ a variant of a reward-modulated Hebbian learning rule proposed by Legenstein et al. (2010) in order to produce oscillatory activity as well as memory-dependent computations from initially chaotic recurrent neural networks through feedback from trained readout neurons without the need for an intelligent supervisor that lets the readout neurons know explicitly about their error. Instead, the synaptic weight changes are based solely on local variables and a global binary modulatory signal. In the proposed model, I do not use teacher forcing of the feedback, but use the actual output of the readout also during the training phase. This approach represents an alternative to the fully supervised learning procedures that are commonly used in reservoir computing and is based on more biologically plausible mechanisms. Instead of knowing explicitly about the error, it suffices to know whether or not the overall performance of the system has recently improved due to random noise perturbations. Moreover, with the specific rule employed here, the readout neuron does not have to know about the noise either, but estimates the noise from its own output (Legenstein et al., 2010).

In the following sections, I show that reward-modulated Hebbian learning applied to the weights of a readout neuron with feedback to the recurrent circuit

is able to induce the production of a specific rhythmic activity pattern with high precision in a recurrent network without additional input. Moreover, I show that using this method, multiple independent traces of persistent memory about the identity of previously observed stimuli can be learned on the basis of a common binary modulatory signal. Similarly, I show that such memory traces and other complex computational functions that rely on these memory traces, such as state-dependent routing of information and other nonlinear computations, can be aquired simultaneously by such a learning mechanism.

In summary, the results of this article provide a new perspective and framework for the emergence of complex computations in biological neural systems.

4.2   Materials and Methods

4.2.1   Recurrent neural network model

For the recurrent network, I use a set of N leaky integrator neurons that are sparsely connected in a recurrent fashion. The output of neuron i at time t, which is interpreted as its current firing rate, is denoted by $r_i(t)$. The neurons within the network receive additional projections from M inputs $u_j(t)$ and also from L linear noisy readouts with output $z_{\xi,j}(t)$.

The network dynamics is given by

$$\tau \dot{x}_i(t) = -x_i(t) + \lambda \sum_{j=1}^{N} W_{ij}^{\mathrm{rec}} r_j(t) + \sum_{j=1}^{M} W_{ij}^{\mathrm{in}} u_j(t) + \sum_{j=1}^{L} W_{ij}^{\mathrm{fb}} z_{\xi,j}(t), \qquad (4.1)$$

where $x_i(t)$ is the state of the i-th neuron and represents its somatic activation potential at time t, and τ is the membrane time constant. The neuronal output is given by $r_i(t) = \tanh(x_j(t)) + \xi_j^{\mathrm{state}}(t)$, where $\xi_j^{\mathrm{state}}(t)$ models zero-mean noise drawn from a uniform distribution in the interval $[-\theta^{\mathrm{state}}, \theta^{\mathrm{state}}]$. Since the sigmoidal $\tanh(\cdot)$ function nonlinearly maps the neuron's state $x_j(t)$ onto the interval $[-1, 1]$, the neuronal output $r_j(t)$ can take values in the interval $[-1 - \theta^{\mathrm{state}}, 1 + \theta^{\mathrm{state}}]$.

Parameters W_{ij}^{rec}, W_{ij}^{in} and W_{ij}^{fb} denote the weights for recurrent connections within the network, connections from inputs to the network, and feedback connections from readout neurons to the network neurons respectively. The parameter λ scales the strength of the recurrent connections. If λ is smaller than one, the network tends to dampen its activity, as usual in a traditional reservoir, while values of λ greater than one lead to networks that tend to produce chaotic activity (Sussillo and Abbott, 2009). Especially in the case where the actual feedback is used during training without any additional input, a value of $\lambda > 1.0$ is crucial for successful learning, because the readout needs sufficiently rich initial dynamics to produce the target function.

Neurons in the recurrent network are connected with a probability p. In other words, each connection strength W_{ij}^{rec} is set and held to zero with probability $(1-p)$. The weights of the network connectivity matrix W^{rec} are drawn from a Gaussian

distribution with zero mean and a variance of $1/(pN)$. The input weights W^{in} and feedback weights W^{fb} are drawn from a uniform distribution in the interval $[-1, 1]$. The output weights $\mathbf{w}_i$, which I adapt during training, are initialized to zero in my experiments, but can also be set to nonzero values. However, these values need to be relatively small in order to avoid that the readout output is too sensitive to noise on the network state because large values need to cancel each other out in order to obtain a small readout output. Note also that while the readouts with feedback in our experiments receive input and also project their output to all neurons within the network, each readout neuron could in principle also be substituted by multiple neurons with sparse connectivity.

4.2.2 Comparison to FORCE Learning

Because I compare the performance of the reward modulated learning rule to the performance of the FORCE learning rules proposed by Sussillo and Abbott (2009), I state them here in short. In contrast to the learning rule used here, the two variants of the FORCE learning mechanism use the exact error $e_i(t) = z_i(t) - f_i(t)$ of the i-th readout to update its weights $\mathbf{w}_i$. In the local least mean squares (LMS) based FORCE rule, readout weights are adapted according to

$$\Delta\mathbf{w}_i(t) = -\eta(t)e_i(t)\mathbf{r}(t). \tag{4.2}$$

The more powerful recursive least squares (RLS) based FORCE rule is defined as

$$\Delta\mathbf{w}_i(t) = -e_i(t)P(t)\mathbf{r}(t), \tag{4.3}$$

the matrix $P(t)$ being a running estimate of the inverse of the correlation matrix of the network output $\mathbf{r}(t)$ plus a regularization term (cf. Sussillo and Abbott (2009); Haykin (2001) for details). Therefore, the RLS-based FORCE rule uses global information about the whole presynaptic state of the readout to modify the learning rate of an individual synapse. This procedure makes the learning rule more powerful than the simple LMS-based rule, but at the expense of locality and simplicity of the rule, which involves some rather complicated computations to generate the matrix $P(t)$.

4.2.3 Teacher forcing versus usage of the actual feedback during training

I consider in this work network architectures where one or several readout neurons project their output back into the neural circuit (cf. panels A of Figures 4.2, 4.6 and 4.7). There are two approaches to deal with this feedback loop during training. In the traditional method called "teacher forcing" (Jaeger and Haas, 2004; Maass et al., 2007) instead of the actual output of the readout a teacher signal $\hat{f}_i(t) = f_i(t) + \xi_i^{tf}(t)$ is fed back into the network during training. Here, $\hat{f}_i(t)$ is a noisy version of the target function $f_i(t)$ with zero-mean noise $\xi_i^{tf}(t)$. This guarantees that the network dynamics stay in the expected regime during training. However, I am not aware of

any biologically plausible mechanism that would provide the basis for the teacher signal.

In this work, I employ the other option, as proposed by Sussillo and Abbott (2009) for their "FORCE learning" procedure, where the actual output $z_{\xi,i}(t)$ of readout i is used as the feedback signal also during training. In this setup, it is necessary that the weight adaptation is fast enough such that the readout reproduces the target function (with slight variations that are necessary for stable performance) quickly after the beginning of training in order to mimic the desired effect of the teacher signal. During training, ongoing weight updates keep the deviation between the readout output and the target output small until a set of time-independent weights is found that lets the readout follow the target function without further weight adaptation (Sussillo and Abbott, 2009).

4.2.4 Performance evaluation

To evaluate the performance of the trained system for different parameter settings in the periodic trajectory production experiment below, I use the following procedure. Since there are no input signals in this experiment (cf. Figure 4.2A), the readout has no reference during the testing period, and small errors in the frequencies of the trained signal components therefore lead to a varying shift between the target and the readout output over time, i.e. the actual cycle of the output of the readout is slightly longer or shorter than the target signal's cycle. To see whether the shape of the target signal is nevertheless accurately reproduced by the readout's output, I cut the readout's output during the testing interval into successive time slices of one second, which is the cycle length of the target function $f(t)$. Then, I calculate the minimum mean squared error (MSE) between each time slice and circularly shifted versions of a one-cycle slice of the target signal instead of just calculating the MSE between the target function $f(t)$ and the readout's output $z(t)$ directly.

In the comparative experiments employing the FORCE learning procedures, I conducted experiments with and without adding exploration noise during training. In any case, the error e_i of the i-th readout that was used for weight adaptation was based on the readout output without exploration noise in these experiments.

To evaluate the performance of the memory units in the "persistent memory" and the "switchable routing" experiment, I calculate the percentage of time steps during the testing interval in which the absolute difference between the target function and the output of the readout neuron exceeds a certain threshold, which I set to 0.5 (half the difference between the two target values for the two different states).

4.2.5 Parameter settings for the simulations

In my simulations, I used the following parameter settings: network size $N = 1000$ units, internal connectivity $p = 0.1$, time constant $\tau = 50$ms for the first and $\tau = 10$ms for the other experiments, chaoticity level $\lambda = 1.5$ for the first experiment (without input) and $\lambda = 1.2$ for the other (partly input driven) experiments,

simulation time step $\Delta t = 1$ms, state noise $\xi^{\text{state}}(t)$ is drawn from a uniform distribution in the interval $[-0.05, 0.05]$, exploration noise $\xi(t)$ is drawn from a uniform distribution in the interval $[-0.5, 0.5]$.

I use a decaying learning rate of $\eta(t) = \eta_{\text{init}}/(1 + t/T)$ with initial learning rate η_{init} and a constant parameter T. The initial learning rate $\eta_{\text{init}} = 0.0005$ for the reward modulated Hebbian learning rule, and the initial learning rate $\eta_{\text{init}} = 0.0001$ for the LMS-based FORCE rule, with learning rate decay constant $T = 20$s. The running average of the noisy readout output is given by $\bar{z}_{\xi,i}(t) = 0.8\bar{z}_{\xi,i}(t - \Delta t) + 0.2z_{\xi,i}(t)$, and the running average of the overall reward is given by $\bar{R}(t) = 0.8\bar{R}(t - \Delta t) + 0.2R(t)$, where Δt is the simulation time step.

The periodic function which has to be produced in the first experiment is

$$f(t) = 1.3/1.5 \sin(2\pi 1 t) + 1.3/3 \sin(2\pi 2 t) + 1.3/9 \sin(2\pi 3 t) + 1.3/3 \sin(2\pi 4 t).$$

Pulses in the input streams of the second and third experiment have an amplitude of 0.4 and a duration of 100 ms (before smoothening), smoothening filter time constant is 50ms, target values for the memory units are -0.5 for the "off" state and 0.5 for the "on" state. Temporally correlated noise inputs $u_3(t)$ and $u_4(t)$ in the third experiment are drawn from a uniform distribution, filtered with a time constant of 0.5s and then scaled to have a standard deviation of 0.25. Moreover, bias values of 0.3 and 0.15 are added to these inputs, respectively. With the nonlinear computation unit, I compute the function

$$f_3(t) = 0.5 \left(u_3(t)^2 + u_4(t)^2 + u_3(t)u_4(t) \right).$$

4.3 Results

In this section, I provide an overview of the proposed learning procedure, and show that three typical computational tasks in the context of rhythmic activity and memory dependent processing can be acquired by recurrent neural networks through reward-modulated Hebbian learning.

I employed in my experiments a network model consisting of N recurrently connected neurons, where the output of each individual neuron is described by its firing rate $r_j(t)$. The recurrent network model is generic in the sense that it is not designed for a particular computational function. Instead, connections within the network are randomly drawn such that neurons are sparsely connected by excitatory and inhibitory synapses (cf. *Materials and Methods*). Similar network models have been used previously to model the dynamics of recurrent biological networks of neurons (Amari, 1972; Hopfield, 1984; Haykin, 1999; Sussillo and Abbott, 2009), including work in the context of CPG networks (Doya and Yoshizawa, 1992). If necessary for the experimental task, the circuit neurons receive projections from additional input steams $u_j(t)$. Specific computational functions are aquired through synaptic modification of the weights from the network to so-called readout neurons (Maass et al., 2002; Jaeger, 2003). These readout neurons can also feedback their
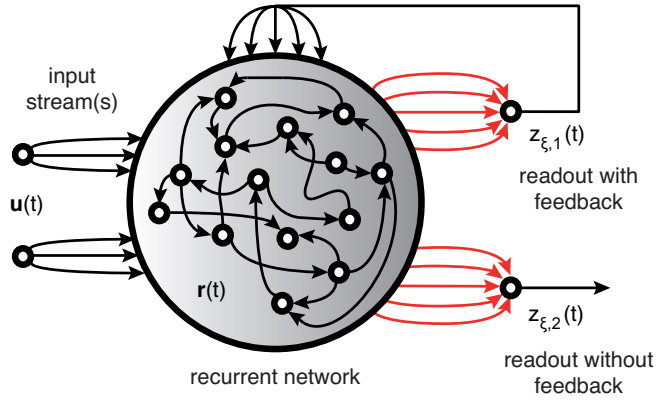
Figure 4.1: Schematic of the basic network topology. A recurrent network of sparsely interconnected neurons receives input from M input streams $u_j(t)$ at time t, assembled into a column vector $\mathbf{u}(t)$. The firing rates of N individual neurons within the network, denoted by $r_j(t)$, are assembled into a column vector $\mathbf{r}(t)$. The network activity $\mathbf{r}(t)$ is then projected to a set of readout neurons. Some readouts feed back their output into the network (see $z_{\xi,1}(t)$), others do not feed back their output (see $z_{\xi,2}(t)$). Only the weights from the network to the readout neurons, denoted by red arrows, are trained.

activity into the recurrent neural network (Jaeger and Haas, 2004; Maass et al., 2007). See Figure 4.1 for a schematic of the basic network topology.

In the presented model, the state $x_j(t)$ of neuron j within the network (cf. *Materials and Methods*) represents its somatic activation potential at time t resulting from excitatory and inhibitory inputs with regard to its resting state (Hopfield, 1984; Haykin, 1999). The output of the j-th neuron in the recurrent network is given by

$$r_j(t) = \tanh(x_j(t)) + \xi_j^{\text{state}}(t), \tag{4.4}$$

where $\xi_j^{\text{state}}(t)$ models zero-mean noise on the firing rate of the neuron. I assemble the activities of the network neurons at time t and a bias term with value one into a column vector $\mathbf{r}(t)$. Assembling the weights connecting these neurons to readout neuron i into a column vector $\mathbf{w}_i$, this readout computes the function

$$z_i(t) = \mathbf{w}_i^T \mathbf{r}(t). \tag{4.5}$$

Reward-modulated Hebbian learning of readout weights demands the readout neurons to be noisy (Legenstein et al., 2010). The output of the i-th readout is therefore modeled by

$$z_{\xi,i}(t) = z_i(t) + \xi_i(t), \tag{4.6}$$

where $\xi_i(t)$ models zero-mean exploration noise on the readout firing rate. This noisy readout signal is also fed back to the circuit neurons if the readout provides feedback into the recurrent circuit. Exploration noise $\xi_i(t)$ is only applied during training, hence $z_{\xi,i}(t) = z_i(t)$ during testing.

Biological networks of neurons operate under substantial amounts of noise and trial-to-trial variability. Therefore, stability of the learned network trajectories is crucial in such systems, because the readouts have to learn to react properly to perturbations. Such stability can be promoted by training in noisy conditions. Two sources of noise in the network play a role in this context. First, the exploration noise $\xi(t)$ in the readout output (cf. Equation 4.6), which is primarily the driving force of learning, also perturbs the network dynamics via the feedback loop. Due to the slow dynamics of the neurons within the network, leading to a lowpass filtering effect, these perturbations are rather small even for high levels of exploration noise. Second, the state noise $\xi^{\text{state}}(t)$ on the output activity of the neurons within the network (cf. Equation 4.4) can have a positive effect on the robustness of the trained system if also applied during the training period. The latter type of perturbations has been used as an alternative to perturbations of the teacher signal by Jaeger and Haas (2004) to stabilize learning, but has turned out to be less effective (Sussillo and Abbott, 2009).

Following the approach of Sussillo and Abbott (2009), and in contrast to traditional reservoirs, the dynamics of the recurrent network used in our simulations tend to be in the chaotic regime prior to training. Different dynamic regimes, from nonchaotic to chaotic, can be accomplished by scaling the recurrent synaptic connections within the network by an appropriate constant λ (cf. *Materials and Methods*). It is important to note that during training, the readout neurons drive the network activities into a nonchaotic state via the feedback pathway, even though the recurrent weights within the network are in principle scaled to support chaotic dynamics (Sussillo and Abbott, 2009).

4.3.1 Reward-modulated learning rule

In contrast to Sussillo and Abbott (2009), who use a fully supervised online learning rule to train the network, I investigate the capabilities of a biologically more plausible reward-modulated online learning rule. Since synaptic plasticity in biological systems can only rely on imperfect error signals, I use a minimal global modulatory signal stemming from an external critic for the training of the synaptic weights of the readout neurons. This signal is minimal in the sense that it communicates in a binary fashion whether the common performance of all readouts has recently increased instead of providing full information on the sign and magnitude of the error. Importantly, this modulatory signal is shared among all readouts and does not contain any information on the amount of improvement or the identity of the readout(s) responsible for the reward increase.

Specifically, the global binary modulatory signal at time t

$$M(t) = \begin{cases} 1 & \text{if } R(t) - \bar{R}(t) > 0, \\ 0 & \text{if } R(t) - \bar{R}(t) \leq 0, \end{cases} \tag{4.7}$$

depends on the the momentary reward $R(t)$ compared to a temporal average of the reward over the recent past $\bar{R}(t)$ (cf. *Materials and Methods*). Therefore, $R(t) - \bar{R}(t)$

indicates the amount of increase or decrease of the reward.

The reward $R(t)$, as calculated by the external critic, depends on the differences between the readout outputs $z_{\xi,i}(t)$ and their respective target functions $f_i(t)$ and is given by

$$R(t) = -\sum_{i=1}^{L}(z_{\xi,i}(t) - f_i(t))^2 \tag{4.8}$$

in my experimental setup. It is important to note that the specific values of $R(t)$ and $\bar{R}(t)$ are only visible to the external critic, while the readout neurons only receive the binary modulatory signal $M(t)$.

In order to produce oscillatory activity as well as memory-dependent computations, this binary modulatory signal guides adaptation of readout weights $\mathbf{w}_i$ via a reward-modulated Hebbian learning rule. I employ a variant of the rule proposed by Legenstein et al. (2010), where the weight change $\Delta\mathbf{w}_i(t)$ of readout i at time t is given by

$$\Delta\mathbf{w}_i(t) = \eta(t)(z_{\xi,i}(t) - \bar{z}_{\xi,i}(t))M(t)\mathbf{r}(t). \tag{4.9}$$

Here, $\eta(t)$ is a small learning rate that can either be constant or decaying over time such that learning is slowed down as training progresses. I use a decaying learning rate in our experiments. Moreover, $\bar{z}_{\xi,i}(t)$ is a running average of the noisy output of the readout during the previous time steps (cf. *Materials and Methods*). If one assumes that $z_i(t)$ changes only slightly within the time scale of the filter and the noise is only weakly correlated over time, then the term $z_{\xi,i}(t) - \bar{z}_{\xi,i}(t)$ approximates the noise $\xi_i(t)$. Therefore, the rule does not need explicit information on the exploration noise part of the output signal, but instead estimates the noise from the output itself (Legenstein et al., 2010).

While traditional Hebbian learning rules use the correlation between the presynaptic and the postsynaptic activity for learning, this learning rule is Hebbian in the sense that it uses the correlation of the postsynaptic noise and the activity of the presynaptic neuron together with a global binary modulatory signal in order to train the local weights. I want to point out that the exploration noise $\xi_i(t)$ is the driving force of learning. Without perturbations of the readouts' output, no learning process would take place.

In all of the following experiemnts, I used $N = 1000$ neurons in the recurrent network. Each pair of these neurons was connected with a probaility $p = 0.1$. These and other basic network parameters were chosen such that they correspond to the values used by Sussillo and Abbott (2009). Other parameters, such as the amplitudes of the exploration noise $\xi(t)$ and state noise $\xi^{\text{state}}(t)$ as well as the initial learning rate η_{init} were manually tuned to produce considerably accurate results. Note that I used the same values for the parameters above for all experiments and did not perform a rigorous parameter search in order to find the setup with optimal performance for each experiment, see *Materials and Methods* for details on the parameter settings used in my experiments.

4.3.2 Emergence of a periodic trajectory through reward-modulated Hebbian learning

Biological neural networks produce many different types of rhythmic activities for various purposes such as muscle activations for breathing, mastication or locomotion. I studied whether reward-modulated Hebbian learning can lead to autonomously generated rhythmic activity in recurrent neural circuits. I simulated a network that receives no inputs besides the feedback projections from a single readout neuron (see also Figure 4.2A). The task of the readout neuron was to produce a specific periodic trajectory and to repeat this periodic trajectory in a stable manner.

In contrast to teacher forcing, which is problematic from the perspective of biological plausibility since the actual output of the readout is replaced by the desired target value, I always used the actual output of the readout neuron to drive the network dynamics via the feedback loop. In order to keep the network dynamics in the expected regime, the readout output has to resemble the target computational function already shortly after the beginning of training. To accomplish this, the synapses of the readout neuron have to adapt rapidly during training, such that the readout approximately follows the target trajectory until a set of time-independent weights is found that lets the readout produce the target function without further weight adaptations (Sussillo and Abbott, 2009). Therefore, the first step was to see whether the proposed reward-modulated Hebbian learning rule is able to adapt the readout weights quickly enough in order to fulfill this requirement. This is a prerequisite for using the actual activity of the readout during training instead of a teacher signal.

Figure 4.2B shows a representative example of the readout activity at the onset of the learning procedure. Within less than 50ms, the readout is able to adapt its activity in order to reach the desired target, and to approximately follow the target signal henceforth. Since the network dynamics are slow compared to the update rate of the learning procedure, the input to the readout as well as the desired target only change slightly on the time scale of the learning procedure. Therefore, the readout is able to adapt its output to move closer to the target output each time the applied noise leads to an increase in the reward. Hence, in spite of the general belief that reward-modulated Hebbian learning is slow as compared to other supervised learning rules, these results indicate that it is fast enough to ensure quick convergence of the readout output to the current target value.

The goal of training is to find a set of time-independent weights such that the system is able to keep producing the target function when the learning mechanism is switched off after an appropriate training time. Figure 4.2C shows that this goal is accomplished by reward-modulated Hebbian learning. After a training time of 400s, which corresponds to 400 cycles of the periodic trajectory, the readout keeps producing the desired trajectory in spite of the lack of any further weight adaptation. The rhythmic activity of the readout, which drives the network via the feedback pathway, has a strong influence on the internal network dynamics. Figures
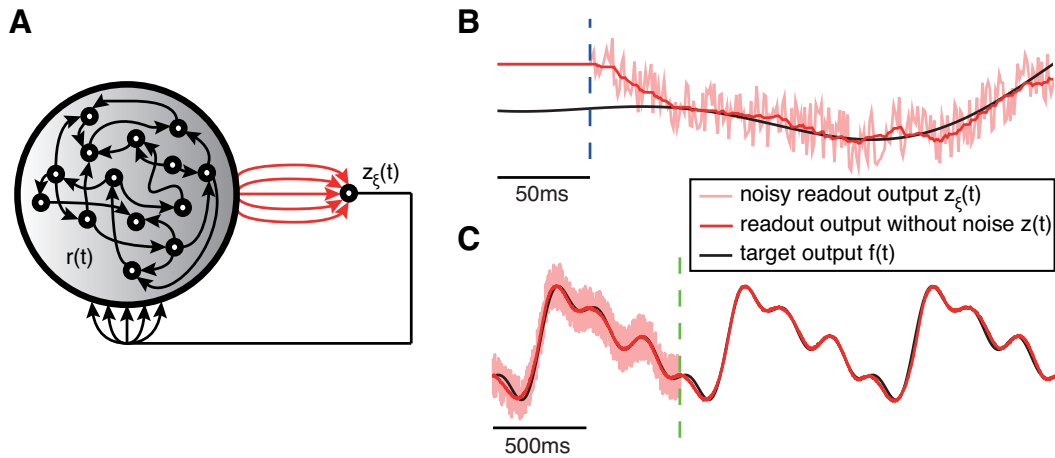
Figure 4.2: Emergence of periodic activity through reward-modulated learning. **A:** A recurrent network receives feedback from a readout neuron, but no other inputs. The readout is trained to produce a periodic trajectory composed of four sinusoids with an overall period of one second. **B:** Beginning of the training interval (first quarter period shown). The dashed blue line indicates the onset of training. After less than 50ms, the readout output without exploration noise $z(t)$ (red) approximately follows the target function $f(t)$ (black), a prerequisite for learning without teacher forcing. The actual feedback signal $z_\xi(t)$ (light red) that is provided to the network includes the exploration noise $\xi(t)$, which is the driving force of learning. **C:** Beginning of the testing interval where synaptic weights remain fixed (last training cycle and first two testing cycles shown). The dashed green line indicates the beginning of the testing period. After an appropriate training time (400s for the presented example), the readout continues to approximately produce the target function without further weight adaptation.

4.3A and 4.3B show a subset of 10 random units from within the network at the onset of training and at the transition from training to testing, respectively. Before the system is being trained, the network exhibits chaotic dynamics and produces rich spontaneous activity. Shortly after the onset of learning, a stable periodic pattern emerges due to the driving force of the feedback loop (panel A). This stable periodic pattern persists during the testing interval when there is no further weight adaptation (panel B).

Based on these encouraging results, I further evaluated the performance of the system. First, I investigated how much training is needed before the network reliably reproduces the oscillatory pattern. To address this question, I conducted 50 independent simulation trials for each of several considered training times. Training was followed by a testing interval of 500s, which also corresponds to 500 cycles of the periodic trajectory. Figure 4.4A shows the trial-averaged performance of the system for varying training times between 10s and 400s. The performance of the system increases with increasing training time. Importantly, for training times of 200s or more, the error stays approximately constant throughout the whole testing interval, with only a slight increase during the first 50 cycles of testing due to a few
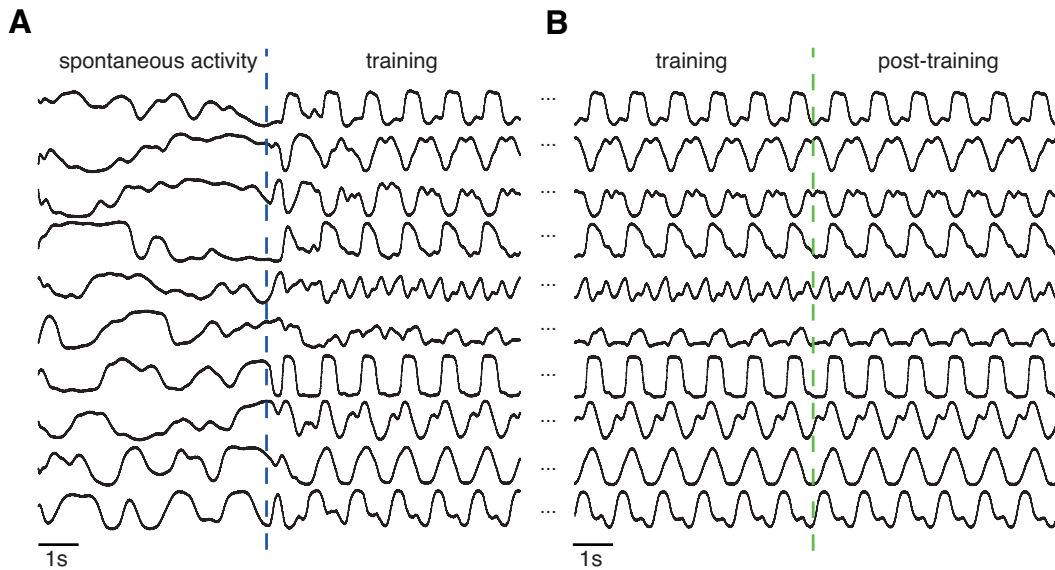
Figure 4.3: Emergence of a stable periodic pattern of the network state through learning. **A:** Beginning of the training interval (first six cycles shown). The dashed blue line indicates the onset of training. While the network produces spontaneous activity before learning, a stable periodic pattern emerges shortly after the onset of training due to the drive of the feedback loop of the readout. **B:** Beginning of the testing interval (last six training cycles and first six testing cycles shown). The dashed green line indicates the beginning of the testing period. After sufficiently long training, the stable periodic pattern keeps being produced during the testing interval.

trials in which the readout diverges from the desired trajectory after some time (cf. also Figure 4.4C, showing that the majority of trials has a very small test error, indicating that the desired trajectory is in fact stably produced until the end of the testing period).

To further evaluate the performance of the system in comparison to FORCE learning (cf. *Materials and Methods* for a brief description of the FORCE learning rules), I compared the performance of the system which employs the reward-modulated Hebbian rule with the performance of systems employing the two different FORCE learning rules (Sussillo and Abbott, 2009), again for training times between 10s and 400s. For the FORCE-trained systems, I did not apply exploration noise during training. However, control simulations with noise showed that this noise does not significantly change the system behavior.

Figure 4.4B shows the result of this comparison. The RLS-based FORCE rule (green) performs best, leading to a good approximation of the target signal after as few as 10 training cycles. This is not surprising since the RLS-based rule is known to be superior to the LMS-based rule with respect to performance, but uses non-local information on the whole presynaptic state of the readout neurons to adapt the learning rates of the single synapses individually, which seems to be problematic
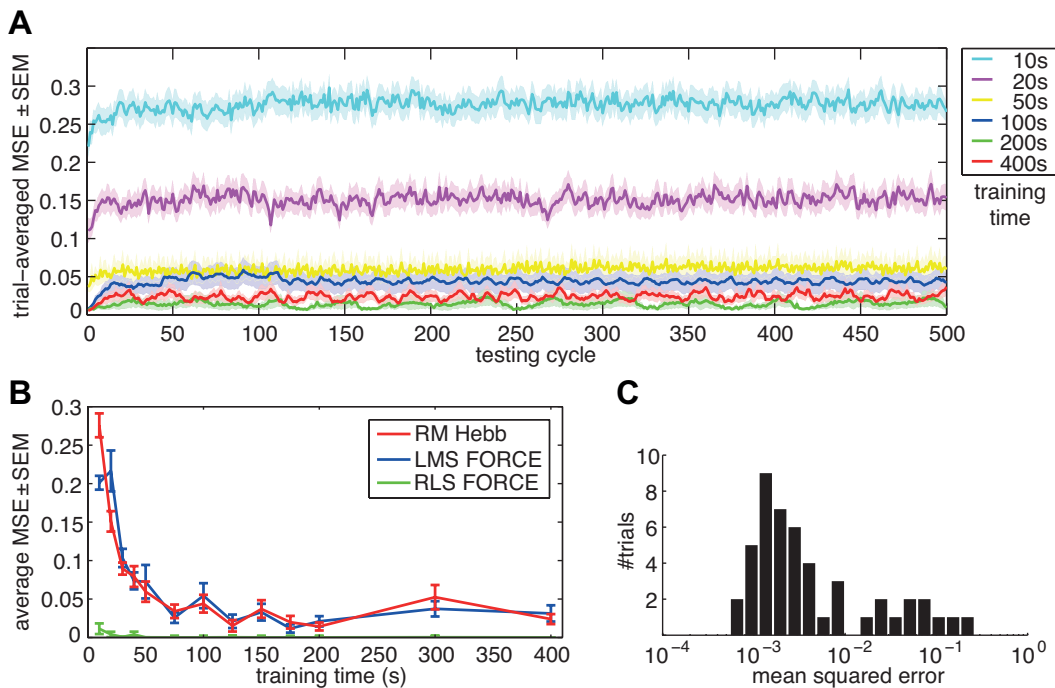
A

B

C

Figure 4.4: Evaluation of the testing performance of the reward modulated Hebbian learning rule for the periodic signal production experiment and comparison to FORCE learning. **A:** Test performance (average mean squared error over 50 independent trials per training time) for varying training times of 10s to 400s. The test error deceases for increasing training time. Importantly, for large training times, the error stays approximately constant during the whole testing interval, with only a slight increase during the first 50 cycles due to a few unstable trials. Shaded areas around the mean squared error (MSE) indicate the standard error of the mean (SEM). **B:** Comparison of the average test error for three different learning rules. While the non-local RLS-based FORCE rule (green) is able to approximate the target function well after only 10 training cycles, the employed reward-modulated Hebbian learning rule (red) and local LMS-based FORCE rule (blue) perform similar to each other, but take approximately 100 cycles to reach a considerable level of performance. **C:** Distribution of the MSE at the end of testing (last testing cycle, log scale) over all 50 trials for the reward modulated Hebbian learning rule and 400s of training. Most of the trials keep producing the target function properly until the end of testing, resulting in a small MSE. Only in a few trials the readout diverges from the desired trajectory after some time, causing the slight increase in the MSE during the first 50 cycles of testing (cf. panel A).

from the point of biological plausibility. However, the reward modulated Hebbian learning rule (red) is able to perform similar to the local LMS-based FORCE rule (blue). With both local rules, the network needs to be trained for approximately 100 cycles before a good performance level is reached.

In order to see whether a certain level of chaoticity - which I regulate by the parameter λ that scales the weights of the recurrent circuit - is necessary for accurate results, I tested the system performance for different values of λ. Figure 4.5A shows
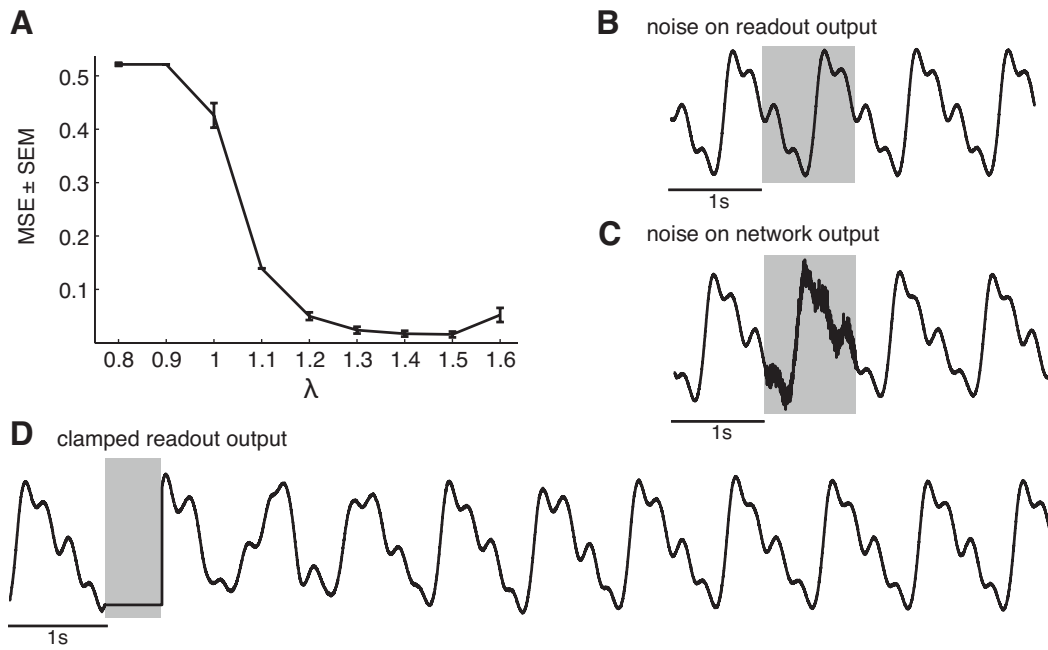
Figure 4.5: (see next page for Figure caption)

the results of this evaluation for training intervals of 400s, and indicates that a certain level of chaoticity is indeed necessary for an accurate performance of the system. On the other hand, if the chaoticity exceeds a certain level, the drive from the feedback loop becomes too weak to drive the network dynamics into a stable regime. This is in line with the results of Sussillo and Abbott (2009). Initial chaotic dynamics are needed because the network has to initially produce sufficiently rich dynamics in order to properly generate the target function if no further input is provided to the system.

Biological networks of neurons have to be able to operate under the influence of substantial amounts of noise and strong trial-to-trial variability. Therefore, another important aspect that needs to be ensured for this type of system is whether it is robust to strong noise perturbations during the testing interval. In order to test this, I perturbed the network in two different ways, for a time interval of one second each. First, the system was perturbed with strong noise in the feedback signal of the same amplitude as the exploration noise applied during training. Afterwards, it was perturbed with state noise on the network output that was ten times stronger than the state noise usually used in the simulations. Figures 4.5B and C show that the trained system is robust against perturbations of the feedback signal (panel B) that affect the readout only via the network, and perturbations of the network output (panel C), which directly affect the input to the readout neuron itself. While perturbations of the feedback signal have no visible effect on the output of the readout since they are filtered by the slow dynamics of the network, the

Figure 4.5: Relation between network performance and chaoticity of the network and robustness of the system against perturbations. **A:** Performance of the network after 400s of training for different values of λ. The network needs a certain level of chaoticity, leading to rich enough dynamics, in order to generate the target function without further input. The system performs best for λ values between 1.3 and 1.5. **B, C:** Noise robustness of the trained trajectory. In order to test the robustness of the system to strong noise perturbations applied during testing, we added uniformly distributed noise to the output of the readout $z(t)$ (panel B, noise added in the shaded 1s interval, we show the output before adding noise) as well as to the network output $\mathbf{r}(t)$, which is in turn the input to the readout unit (panel C, noise added in the shaded 1s interval). Noise on the readout output does not produce any visible deviations of the readout output, because the noise influence is small due to the lowpass filtering properties of the leaky integrator units within the network. On the other hand, noise on the network output directly affects the input to the readout and produces deviations that fluctuate around the desired trajectory. However, the desired trajectory is quickly restored after the noise is removed again. **D:** To further evaluate the robustness to perturbations, I clamped the readout output to a constant value for a time interval of 500ms. After unclamping the output of the readout, the system is able to restore the desired trajectory within a few cycles.

strong perturbations of the network output do have a visible effect on the readout output. However, while this type of noise lets the readout fluctuate around the desired trajectory during the time interval in which the system is perturbed, the desired trajectory is restored shortly after the noise is removed again. Interestingly, the network is able to restore the desired trajectory even after extremely severe perturbations of the readout output. For example, I clamped the readout output such that it remained constant for 500ms. After unclamping of the readout output, the system is able to recover the desired trajectory within a few cycles (see panel D of Figure 4.5).

4.3.3   Simultaneous learning of multiple persistent memory traces

High-level cognitive processing in the brain demands working memory in the sense that neuronal circuits are able to temporarily store contextual information such as the identity of previously observed stimuli. Neuronal correlates of working memory have been observed for example in single neuron recordings from the prefrontal cortex of macaque monkeys during visual working memory in delayed matching-to-sample tasks (Fuster and Alexander, 1971; Goldman-Rakic, 1995; Miller et al., 1996). In these experiments, it was observed that prefrontal cortex neurons hold information of previously observed stimuli by a persistent increase or decrease of their firing rates for a time period in the range of seconds.

I tested whether such memory-dependent processing can emerge in my model through reward-modulated learning in a generic working memory task. In this task, the network should hold information about two independent contextual variables in its network state and make this state visible through the activity of two readout neurons, one for each variable. The current state of each of these variables is
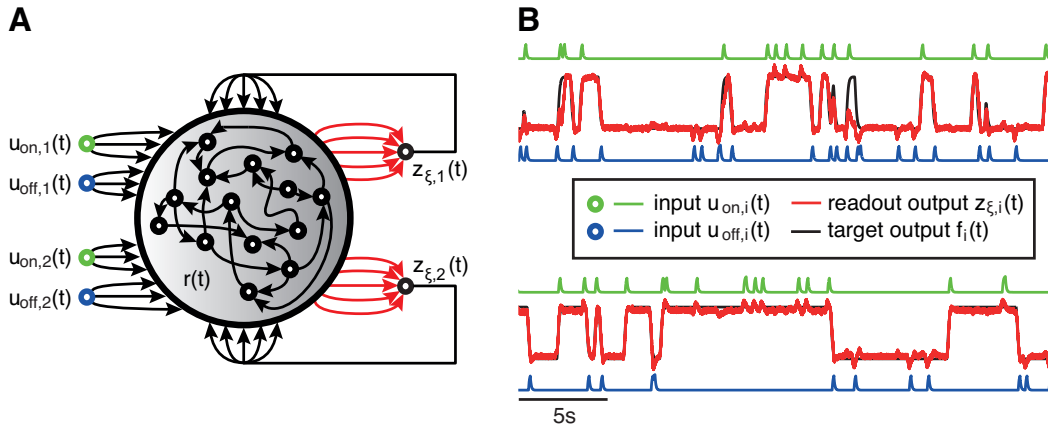
A

B

Figure 4.6: Learning of independent persistent memory traces based on a common modulatory signal. **A:** In this task, two readouts are trained by adapting their weights (red arrows in panel A) using a common binary modulatory signal. Each readout $z_{\xi,i}(t)$ is trained to produce a memory trace (red trace in panel B, the black trace represents the target function $f_i(t)$) by changing its firing rates depending on the activity of the two input streams it is associated with. If the associated "on" input ($u_{\mathrm{on},i}(t)$, green inputs in panel A, and green traces above the readout output $z_i(t)$ and target function $f_i(t)$ in panel B) is briefly activated, the readout is trained to switch to a high firing rate. If the associated "off" input ($u_{\mathrm{off},i}(t)$, blue inputs in panel A, and blue traces below $z_i(t)$ and $f_i(t)$ in panel B) is activated, it is trained to switch to a low firing rate. **B:** After a training time of 500s, the system is tested for another 500s. The last 30s of the testing period are shown for both readouts, aligned with the respective readout neurons in the schematic in panel A. Both readouts nicely produce the desired memory traces.

defined by the past activity in two input streams, an "on"-stream and an "off"-stream (see Figure 4.6A). The task of each readout neuron is to learn to change its firing rate depending on the pulses in the two associated input streams. A brief (100ms) activation of the "on" input demands the readout to switch to a high firing rate, while a brief activation of the "off" input demands the readout to switch to a low firing rate. If pulses occur subsequently in the same input stream without an intermediate pulse in the other input stream, the readout is expected to keep the same rate. Moreover, the readout has to remain independent of the inputs that are associated with the respective other readout and also of the feedback of the state of the other readout itself. Figure 4.6A shows the network setup for this persistent memory experiment.

Figure 4.6B shows a representative example of the result of training the network for 500s and subsequent 500s without further weight adaptation. The last 30s of the testing interval are presented. The firing rates of the two readouts correctly change depending on the associated inputs. Note that the system learned the correct behavior for both readouts based on a single modulatory signal that indicated only whether the combined performance of both readouts recently improved.

In order to evaluate how much time the system spends in the correct state during a 500s testing interval, I performed 50 independent simulations and calculated the percentage of time in which the readouts are in the wrong rate. The average fraction of time in the wrong state is 4.57±0.73 percent of the testing time, with an even lower median fraction of 2.9 percent (based on 50 trials, average proportion over both readouts per trial).

In summary, this experiment shows that such a circuit is able to hold information about recently observed inputs for at least several seconds and both readouts are able to learn to stay independent of each other despite the fact that both readouts strongly modulate the network dynamics and do only learn on the basis of a common binary reward signal.

4.3.4   Simultaneous learning of memory traces and state-dependent computations such as switchable routing

Solely having a working memory trace does not suffice for complex cognitive processes, for example during the task of processing grammatical information or during action selection based on contextual information sensed previously. We therefore address the question whether the proposed model can simultaneously generate memory traces of preceding inputs, like in the previous experiment, and perform different complex computational functions on the basis of these memory traces by employing reward-modulated Hebbian learning.

In my experimental setup (see Figure 4.7A), I simulated a network that receives input from four input streams. Two inputs, $u_1(t)$ and $u_2(t)$, are series of short pulses occurring at random points in time, like in the persistent memory experiment above. The other two inputs, $u_3(t)$ and $u_4(t)$, are temporally correlated noise signals. Moreover, the network projects to three readout units that are expected to solve three different computational tasks using the common modulatory signal. The first readout, the "memory unit", feeds its output signal $z_{\xi,1}(t)$ back into the network. Its function is the same as the function of each readout in the "persistent memory" experiment. The other two readouts, which do not feed back to the network, have different target functions. The second readout, the "routing unit" with output $z_{\xi,2}(t)$, is trained to output the first noise input $u_3(t)$ if the memory unit is in the "on" state, and to output the second noise input $u_4(t)$ if the memory unit is in the "off" state. With respect to this readout, the system can thus be viewed as a switchable routing network that routes information from one of the two input channels depending on its memory about a contextual variable. The third readout, the "nonlinear computation unit" with output $z_{\xi,3}(t)$, is trained to compute a nonlinear function of the noise inputs $u_3(t)$ and $u_4(t)$ which is independent of the state of the memory unit.

Figures 4.7B and C show the readout outputs of the final system after 500s of training. Figure 4.7B shows an example in which the transition of the memory unit from the "on" state to the "off" state is correctly executed. At the time of this state transition, the routing unit also changes its output from approximately representing
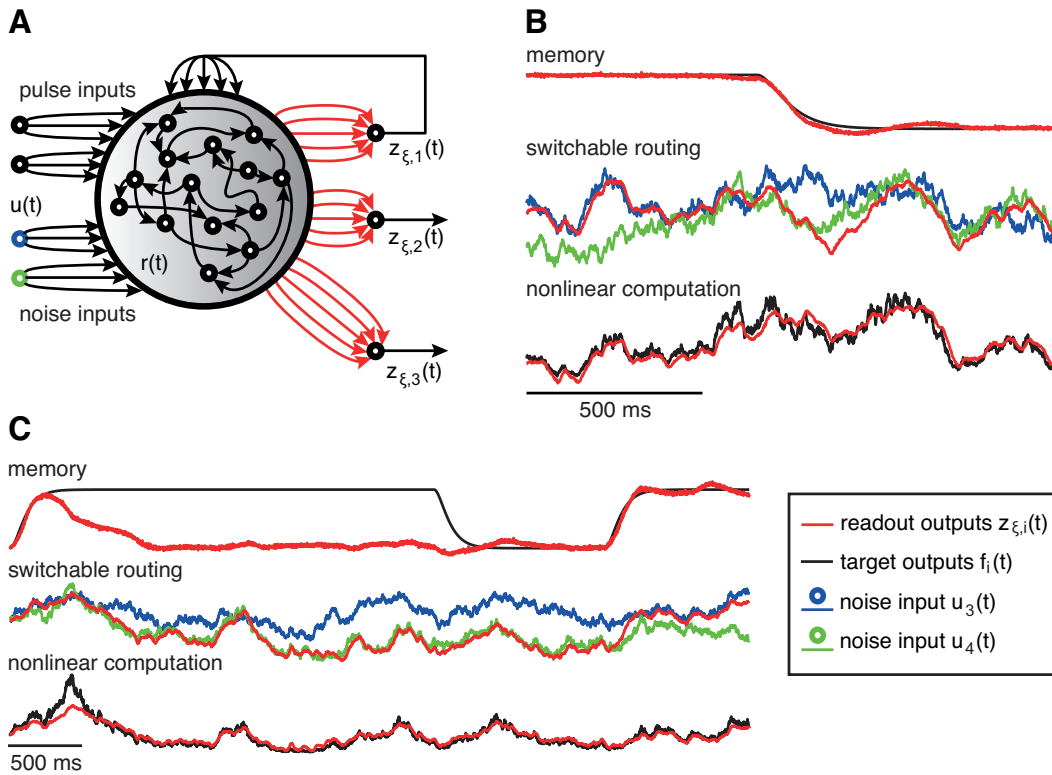
Figure 4.7: Simultaneous learning of memory traces and state-dependent computations.
A: Three readouts are trained using a common modulatory signal. The "memory unit" $z_{\xi,1}(t)$ provides feedback to the circuit and switches between two firing rates based on the pulses in the two "pulse inputs" (cf. persistent memory experiment). The "switchable routing unit" $z_{\xi,2}(t)$ routes one of two temporally correlated "noise inputs" through the network, depending on the state of the memory unit. The third readout $z_{\xi,3}(t)$ computes a nonlinear function of the two noise inputs independent of the state of the memory unit. **B:** Output traces of the three readout units at the end of a 500s testing period after 500s of training, aligned with the corresponding readout units in the schematic in panel A. The memory unit $z_{\xi,1}(t)$ (red, upper trace) switches properly from the "on" state to the "off" state. The routing unit's output $z_{\xi,2}(t)$ (red, middle trace) also switches from approximately reproducing the input $u_3(t)$ (blue) to reproducing the input $u_4(t)$ (green). The third readout $z_{\xi,3}(t)$ properly computes a nonlinear function of the two correlated noise inputs $u_3(t)$ and $u_4(t)$. **C:** Improper switch of the memory unit's state. The memory unit's output $z_{\xi,1}(t)$ (upper trace) switches back to the "off" state while its target state $f_1(t)$ is the "on" state. As expected, the routing unit's output $z_{\xi,2}(t)$ (middle trace) also reproduces the wrong input, in this case $u_4(t)$ (green) instead of $u_3(t)$ (blue), since its output depends on the memory unit's state. As soon as the memory unit is again in the correct state, also the routing unit reproduces the correct input. The nonlinear computation unit $z_{\xi,3}(t)$ remains largely unaffected by the wrong switch state.

the first noise input $u_3(t)$ (blue) to representing the second noise input $u_4(t)$ (green). The third readout is not affected by the state switch and correctly computes the

nonlinear function of both noise inputs throughout. This shows that the same circuit is able to perform concurrently complex memory-dependent operations and memory-independent nonlinear computations. Figure 4.7C shows an example from the same trial approximately 50s earlier in which a switch of the memory unit to the "on" state fails. Similar to the previous experiment, the memory unit is in the wrong state for as few as 5.75 ± 0.38 percent of the whole training time (average over 50 trials). The routing unit behaves as expected, i.e. since the memory unit is in the wrong state, the routing unit also represents the "wrong" input $u_4(t)$ (green) instead of $u_3(t)$ (blue), which is in turn the correct behavior of the routing unit because its activity depends on the memory unit's state. As soon as the memory unit switches to the "on" state accurately, also the routing unit switches to representing $u_3(t)$ (blue) again. The additional readout unit that computes the nonlinear function of the noise inputs $u_3(t)$ and $u_4(t)$ remains largely unaffected by the wrong state of the memory unit.

4.4   Discussion

In extension and modification of previous results (Jaeger and Haas, 2004; Maass et al., 2007; Sussillo and Abbott, 2009; Legenstein et al., 2010), I have shown that a simple reward-modulated online learning rule that is used for the training of a set of readout neurons with feedback projections is able to steer the dynamics of an initially chaotic recurrent neural network in order to produce patterns of oscillatory activity as well as persistent memory and memory-dependent computations.

In contrast to previous results, a system that is trained with the presented method is able to learn to perform complex computations without the need for an intelligent "supervisor". Instead, it suffices that a global modulatory signal is present which indicates whether or not the overall performance of the system has recently improved due to random noise perturbations. The learning procedure does only incorporate local synapse-specific information (i.e., the output of the readout itself and the output of the neuron that is presynaptic to the readout via that specific synapse) as well as a global binary modulatory signal. Since it does not depend on exact information on the error, the learning rule proposed here is more biologically plausible than the fully supervised learning rules that are traditionally used to train such networks.

I want to stress the importance of the globality of this signal to biological plausibility of the system. In traditional error-based learning approaches involving multiple readouts, each readout has to be supervised individually by providing it with its exact error. In the proposed system, it is not necessary that an individual modulatory signal is present for each readout that indicates whether the performance has increased or not. Instead, I use the same modulatory signal for all readouts, indicating only whether the common performance of all readouts has increased due to random noise perturbations. I therefore propose to experimentally test this by simultaneously measuring responses of multiple neural populations that are involved

in coding for the persistent memory of multiple independent cues, probably within prefrontal cortex, together with the local concentrations of neuromodulatory signals during the entrainment of a memory task. If the local concentration of these signals would show to be only correlated with changes in the performance of the local neural population at the same location, but not with changes in the performance of the other populations at different locations, this would suggest multiple individual performance-related signals. However, the presented model predicts that the concentration of neuromodulatory signals would be correlated not only with changes in the local population's performance, but with changes in the performance of several populations with individual tasks each.

Moreover, since I always use the actual feedback of the readout neurons, the model does not need a teacher signal that already resembles the target activity of the readout and thereby drives the network activity into the expected regime during training. This approach is adopted from Sussillo and Abbott (2009) and increases the biological plausibility of such a system in comparison to the traditional teacher-forced systems Jaeger and Haas (2004); Maass et al. (2007), since the teacher signal would require the existence of another circuit that is already able to perform the desired function, reducing the task of the network to merely reproducing the behavior of the other circuit. Moreover, assuming that such an additional circuit would already exist, the network would still have to switch from the "open loop" topology using the teacher signal to a "closed loop" topology using the actual readout output as soon as the desired target function is obtained by the readout, either gradually or instantaneously. This would require another involved rewiring mechanism that detects that the readout has acquired a considerable level of performance and induces the transition from the "open loop" to the "closed loop" topology.

CPG networks, recurrent neural structures which autonomously produce rhythmic activity, have been found to be located in the spinal cord, associated with reflexes and locomotion patterns, and in the brain stem, associated with breathing, chewing, swallowing and eye movements. Those regions receive projections from the basal ganglia, a structure that is known to be important for the control and coordination of CPGs, including tonic inhibition of CPG networks during resting conditions. Also, the basal ganglia are thought to be highly involved in reward-based learning. I therefore suggest that synaptic modification within CPG networks that are involved in locomotion might be modulated by signals stemming from the basal ganglia in order to steer the activity towards entrainment of specific rhythms. After training, the modulatory signals are not required to produce the entrained rhythms, which is consistent with results from lesion studies showing that the adaptatability of CPG networks is impaired after lesions of the basal ganglia, while individual CPGs can still be activated through stimulation. Also, the responsiveness of these systems to enhanced or reduced dopaminergic input may point towards this hypothesis (Grillner, 2006). The question whether our hypothesis holds has to be addressed by experimental studies that focus on the availability of modulatory signals during the entrainment of activation patterns by CPGs. Specifically, it has to be tested whether the concentration of these signals is correlated with changes in

network performance.

Moreover, movement-related rhythmic activity patterns related to jaw and tongue movements have been found in the primary motor cortex of primates (Yao et al., 2002), which has been shown to be involved in learning of fine motor skills (Molina-Luna et al., 2009). Also, traces of persistent memory reflected in sustained firing activity of single neurons in response to specific visual stimuli have been recorded in prefrontal cortex (Fuster and Alexander, 1971; Goldman-Rakic, 1995; Miller et al., 1996). Primary motor cortex and prefrontal cortex both receive input projections from midbrain dopaminergic neurons, and dopamine has been related to the expression of synaptic plasticity in these areas (Descarries et al., 1987; Lewis et al., 1987; Molina-Luna et al., 2009; Hosp et al., 2011). I therefore suggest that the entrainment of such movement-related activity patterns in primary motor cortex and persistent memory traces in prefrontal cortex may also be guided by modulatory input from midbrain neuromodulatory signals.

Assuming that the proposed tasks are performed using such a learning mechanism, the presented model predicts that the synaptic adaptations that keep the desired trajectories stable during learning depend on the availability of global signals, such as specific neuromodulators. Without the presence of such signals, adaptation would not be possible. This is consistent with studies showing that both working memory performance and motor skill learning are impaired if the dopaminergic system of the brain is degenerated, as in patients with Parkinson's disease (Durstewitz and Seamans, 2002; Doyon, 2008; Molina-Luna et al., 2009), and also with studies showing that working memory performance is is impaired if dopaminergic input to the prefrontal cortex is blocked (Durstewitz and Seamans, 2002). Moreover, our simulation results are also consistent with results indicating that dopaminergic signaling in primary motor cortex is involved in learning new motor skills, but not in executing a previously learned skill (Molina-Luna et al., 2009; Hosp et al., 2011).

However, further experiments remain to be conducted to explore how the performance of such learning rules scales with larger sets of readouts. Moreover, it is an open question to which extent variants of this mechanism can be applied to spiking neural networks to solve such tasks. In summary, reward-modulated learning rules provide an interesting alternative to the fully supervised rules that are commonly used in reservoir computing approaches. My results give an indication how stable rhythmic patterns as well as persistent memory and memory-dependent computatations can be formed from initially chaotic recurrent networks and also are in line with the hypothesis that noise plays a crucial role for learning instead of only being a disturbing factor.

Acknowledgements

List of Publications

---

1. G. M. Hörzer. *Extraction of Information about the Behavioral State of Monkeys from Neuronal Recordings with Methods from Machine Learning.* Master's thesis, Institute for Theoretical Computer Science, Graz University of Technology. 2008.

2. S. Liebe, G. M. Hörzer, N. K. Logothetis, W. Maass, and G. Rainer. *Long Range Coupling between V4 and PF in Theta Band during Visual Short-Term Memory.* Poster at the Annual Meeting of the Society for Neuroscience, Chicago, USA. 2009.

3. G. M. Hörzer, S. Liebe, A. Schlögl, 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(14). 2010.

4. G. M. Hörzer. *Reward-modulated Hebbian Learning is able to induce Coherent Patters of Activity and Simple Memory Functions in initially Chaotic Recurrent Neural Networks.* Workshop on Cognitive and Neural Models for Automated Processing of Speech and Text (CONAS), Ghent, Belgium. 2010.

5. S. Liebe, G. M. Hörzer, N. K. Logothetis, and G. Rainer. *Oscillatory Neuronal Synchronization between Prefrontal and Extrastriate Visual Cortex during Visual Memory.* Poster at the Annual Meeting of the Society for Neuroscience, San Diego, USA. 2010.

6. S. Liebe, G. M. Hörzer, N. K. Logothetis, and G. Rainer. *Long range coupling in theta between V4 and prefrontal cortex predicts visual memory performance.* Poster at the Computational and Systems Neuroscience Conference (COSYNE), Salt Lake City, USA. 2011.

7. S. Liebe, G. M. Hörzer, N. K. Logothetis, and G. Rainer. *Long range theta coupling between V4 and lateral prefrontal cortex predics visual short-term memory performance.* Submitted. 2011.

8. G. M. Hörzer, R. Legenstein and W. Maass. *Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-modulated Hebbian Learning.* In preparation. 2011.

A.1   Comments and Contributions to Publications

The first publication, *Extraction of Information about the Behavioral State of Monkeys from Neuronal Recordings with Methods from Machine Learning*, is my master's thesis, which was jointly supervised by Wolfgang Maass and Gregor Rainer and provided the basis for the collaboration during my doctoral studies which resulted in most of the publications above (except for publications 4 and 8). The results from my master's thesis are not contained in this thesis.

The publication *Directed coupling in local field potentials of macaque V4 during visual short-term memory revealed by multivariate autoregressive models* was a joint paper with Stefanie Liebe (SL), Gregor Rainer (GR) and Nikos Logothetis (NL) from the Max Planck Institute for Biological Cybernetics in Tübingen, Germany[1], and Alois Schlögl (AS) from the Institute for Human-Computer Interfaces at TU Graz, Austria[2]. The monkey experiments and electrophysiological recordings were designed by SL and GR and performed by SL, the data analysis was performed by myself and SL, and the paper was written by myself and SL with input from AS, NL and GR. This paper is the basis for Chapter 2 of this thesis.

The publication *Reward-modulated Hebbian Learning is able to induce Coherent Patters of Activity and Simple Memory Functions in initially Chaotic Recurrent Neural Networks* is a contribution to the Workshop on Cognitive and Neural Models for Automated Processing of Speech and Text (CONAS) in Ghent, Belgium, in July 2010 and was accepted for an oral presentation. The experiments were performed and the paper was written by myself. The results from this publication are not included in this thesis, but provided the basis for further work that resulted in the paper *Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-modulated Hebbian Learning*.

The publications *Long Range Coupling between V4 and PF in Theta Band during Visual Short-Term Memory* as well as *Oscillatory Neuronal Synchronization between Prefrontal and Extrastriate Visual Cortex during Visual Memory* and *Long range coupling in theta between V4 and prefrontal cortex predicts visual memory performance* are posters which present results from joint work with Stefanie Liebe (SL), Nikos Logothetis (NL) and Gregor Rainer (GR). The posters were prepared by SL and myself with input from NL and GR, and presented by SL at the Annual Meeting of the Society for Neuroscience (SfN) in 2009 and 2010 as well as at the Computational and Systems Neuroscience Conference (COSYNE) in 2011, respectively.

The paper *Long range theta coupling between V4 and lateral prefrontal cortex predics visual short-term memory performance* is based on joint work also pre-

---

[1]Stefanie Liebe is currently affiliated with the Department of Physiology, Pharmacology and Neuroscience at the University College London, United Kingdom, and Gregor Rainer is currently the head of the Visual Cognition Laboratory of the Department of Medicine at the University of Fribourg, Switzerland.

[2]Alois Schlögl is currently affiliated with the Institute of Science and Technology Austria in Maria Gugging, Austria.

sented in the posters at SfN and COSYNE (see previous paragraph). The monkey experiments and electrophysiological recordings were designed by SL and GR and performed by SL, the data analysis was performed by myself and SL, and the paper was written by myself and SL with input from NL and GR. This paper has been submitted to "Nature Neuroscience" in February 2011 and is currently under review. This paper is the basis for Chapter 3 of this thesis.

The paper *Emergence of Complex Computational Structures from Chaotic Neural Networks through Reward-modulated Hebbian Learning* is a joint paper with Robert Legenstein (RL) and Wolfgang Maass (WM) and extends the work presented in the workshop contribution *Reward-modulated Hebbian Learning is able to induce Coherent Patters of Activity and Simple Memory Functions in initially Chaotic Recurrent Neural Networks*. The experiments were designed by myself, RL and WM, performed by myself, and the paper was written by myself, RL and WM. This paper is in preparation and will be submitted in 2011. This paper is the basis for Chapter 4 of this thesis.

Bibliography

Aertsen, A. M. H. J. and Arndt, M. (1989). Response synchronization in the visual cortex. *Current Opinion in Neurobiology*, 3(4):586–594. 24

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19(6):716–723. 11

Akam, T. and Kullmann, D. M. (2010). Oscillations and filtering networks support flexible routing of information. *Neuron*, 67(2):308–320. 56

Amari, S.-I. (1972). Characteristics of random nets of analog neuron-like elements. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-2(5):643–657. 65

Andersen, P., Morris, R., Amaral, D., Bliss, T., and O'Keefe, J. (2007). The hippocampal formation. In Andersen, P., Morris, R., Amaral, D., Bliss, T., and O'Keefe, J., editors, *The hippocampus book*. Oxford University Press. 55

Anderson, K. L., Rajagovindan, R., Ghacibeh, G. A., Meador, K. J., and Ding, M. (2009). Theta oscillations mediate interaction between prefrontal cortex and medial temporal lobe in human memory. *Cereb Cortex*, 20(7):1604–1612. 9, 12, 25, 30, 39, 55

Arnolds, D. E., Lopes da Silva, F. H., Aitink, J. W., Kamp, A., and Boeijinga, P. (1980). The spectral properties of hippocampal EEG related to behaviour in man. *Electroencephalogr Clin Neurophysiol*, 50(3-4):324–328. 55

Baccala, L. A. and Sameshima, K. (2001). Partial Directed Coherence: A new concept in neural structure determination. *Biological Cybernetics*, 84:463–474. 9, 12

Baccala, L. A., Sameshima, K., and Takahashi, D. Y. (2007). Generalized Partial Directed Coherence. *15th International Conference on Digital Signal Processing*, pages 163–166. 12

Baddeley, A. (1986). *Working Memory*. Oxford University Press. 31

Baessler, U. and Bueschges, A. (1998). Pattern generation for stick insect walking movements - multisensory control of a locomotor program. *Brain Research Reviews*, 27:65–88. 59

Barbas, H. and Pandya, D. N. (1989). Architecture and intrinsic connections of the prefrontal cortex in the rhesus monkey. *J Comp Neurol*, 286(3):353–375. 54

Barto, A. G., Sutton, R. S., and Anderson, C. W. (1988). Neuronlike adaptive elements that can solve difficult learning control problems. pages 535–549. MIT Press. 61

Basar, E. and Bullock, T. H. (1992). *Induced rhythms in the brain*. Birkhaeuser. 55

Berens, P. (2009). Circstat: A Matlab toolbox for circular statistics. *Journal of Statistical Software*, 31(10). 33

Berger, H. (1929). Ueber das Elektroenkephalogramm des Menschen. *Arch Psychiatr Nervenkrankh*, 87:527–570. 1

Bi, G. Q. and Poo, M. M. (1998). Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type. *J Neuroscience*, 18(24):10464–10472. 2

Bollimunta, A., Chen, Y., Schroeder, C. E., and Ding, M. (2008). Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. *Journal of Neuroscience*. 10

Bollimunta, A., Chen, Y., Schroeder, C. E., and Ding, M. (2009). Characterizing oscillatory cortical networks with granger causality. *Coherent Behavior in Neuronal Networks*, pages 169–189. 24

Bressler, S. L., Ding, M., and Yang, W. (1999). Investigation of cooperative cortical dynamics by multivariate autoregressive modeling of event-related local field potentials. *Neurocomputing*, 26–27:625–631. 9, 24

Bressler, S. L., Richter, C. G., Chen, Y., and Ding, M. (2007). Cortical functional network organization from autoregressive modeling of local field potential oscillations. *Statistics in Medicine*, 26:3875–3885. 9

Brovelli, A., Ding, M., Ledberg, A., Chen, Y., Nakamura, R., and Bressler, S. L. (2004). Beta oscillations in a large-scale sensorimotor cortical network: Directional influences revealed by Granger causality. *Proceedings of the National Academy of Sciences of the United States of America*, 101(26):9849–9854. 9, 12, 24

Bueschges, A., Akay, T., Gabriel, J. P., and Schmidt, J. (2008). Organizing network action for locomotion: Insights from studying insect walking. *Brain Research Reviews*, 57:162–171. 59

Buonomano, D. and Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews in Neuroscience*, 10(2):113–125. 60

Burwell, R. D. (2000). The parahippocampal region: corticocortical connectivity. *Ann N Y Acad Sci*, 911:25–42. 55

Buzsaki, G. (1996). The hippocampo-neocortical dialogue. *Cereb Cortex*, 6(2):81–92. 55

Buzsaki, G. (2002). Theta oscillations in the hippocampus. *Neuron*, 33(3):325–340. 30, 55

Buzsaki, G. (2005). Theta rhythm of navigation: link between path integration and landmark navigation, episodic and semantic memory. *Hippocampus*, 15(7):827–840. 8

Buzsaki, G. (2006). *Rythms of the Brain*. Oxford University Press. 2, 18, 55

Buzsaki, G. and Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, 304(5679):1926–1929. 8, 23, 25, 30, 56

Caporale, N. and Dan, Y. (2008). Spike timing-dependent plasticity: A hebbian learning rule. *Annual Reviews in Neuroscience*, 31:25–46. 2

Cashdollar, N., Malecki, U., Rugg-Gunn, F. J., Duncan, J. S., Lavie, N., and Duzel, E. (2009). Hippocampus-dependent and -independent theta-networks of active maintenance. *Proc Natl Acad Sci USA*, 106(48):20493–20498. 42, 53, 55

Chen, Y., Bressler, S. L., and Ding, M. (2006). Frequency decomposition of conditional Granger causality and application to multivariate neural field potential data. *Journal of Neuroscience Methods*, 150:228–237. 9

Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? *Progress in Brain Research*, 169:323–338. 3

Cui, J., Xu, L., Bressler, S. L., M., D., and Liang, H. (2008). Bsmart: a matlab/c toolbox for analysis of multichannel neural time series. *Neural Networks*. 9

Descarries, L., Lemay, B., Doucet, G., and Berger, B. (1987). Regional and laminar density of the dopamine innervation in adult rat cerebral cortex. *Neuroscience*, 21(3):807–824. 80

Ding, M., Bressler, S. L., Yang, W., and Liang, H. (2000). Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modelling: data preprocessing, model validation, and variablilty assessment. *Biological Cybernetics*, 83:35–45. 10, 11

Doya, K. and Yoshizawa, S. (1992). Adaptive synchronization of neural and physical oscillators. In Moody, J. E. nd Hanson, S. J. and Lippmann, R. P., editors, *Advances in Neural Information Processing Systems*, volume 4, pages 109–116. 65

Doyon, J. (2008). Motor sequence learning and movement disorders. *Current Opinion in Neurology*, 21:478–483. 80

Duezel, E., Penny, W. D., and Burgess, N. (2010). Brain oscillations and memory. *Current Opinion in Neurobiology*, 20:143–149. 2

Durstewitz, D. and Seamans, J. K. (2002). The computational role of dopamine d1 receptors in working memory. *Neural Networks*, 15:561–572. 80

Efron, B. and Tibshirani, R. S. (1993). *An introduction to the Bootstrap*, volume 57. Chapman Hall. 15, 35

Farries, M. A. and Fairhall, A. L. (2007). Reinforcement learning with modulated spike timing dependent synaptic plasticity. *Journal of Neurophysiology*, 98(6):3648–3665. 61

Fell, J. and Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature Reviews Neuroscience*, 12(2):105–118. 30, 47, 54, 56

Fell, J., Klaver, P., Elfadil, H., Schaller, C., Elger, C. E., and Fernandez, G. (2003). Rhinal-hippocampal theta coherence during declarative memory formation: interaction with gamma synchronization? *Eur J Neurosci*, 17(5):1082–1088. 42

Fell, J., Ludowig, E., Rosburg, T., Axmacher, N., and Elger, C. E. (2008). Phase-locking within human mediotemporal lobe predicts memory formation. *Neuroimage*, 43(2):410–419. 53

Felleman, D. J. and Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cereb Cortex*, 1(1):1–47. 55

Fiete, I. R., Fee, M. S., and Seung, H. S. (2007). Model of birdsong learning based on gradient estimation by dynamic perturbation of neural conductances. *Journal of Neurophysiology*, 98(4):2038–2057. 2

Fiete, I. R. and Seung, H. S. (2006). Gradient learning in spiking neural networks by dynamic perturbation of conductances. *Physical Review Letters*, 97:48104. 2, 61

Fremaux, N., Sprekeler, H., and Gerstner, W. (2010). Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience*, 30(40):13326–13337. 2

Frien, A. and Eckhorn, R. (2000). Functional coupling shows stronger stimulus dependency for fast oscillations than for low frequency components in striate cortex of awake monkey. *European Journal of Neuroscience*, 12:1466–1478. 22, 26

Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci*, 9(10):474–480. 8, 23, 30, 41, 47, 52, 53, 54, 56

Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508):1560–1563. 8, 24, 31

Fuentemilla, L., Penny, W. D., Cashdollar, N., Bunzeck, N., and Duzel, E. (2010). Theta-coupled periodic replay in working memory. *Curr Biol*, 20(7):606–612. 53

Fuster, J. M. and Alexander, G. E. (1971). Neuron activity related to short-term memory. *Science*, 173:652–654. 31, 56, 58, 74, 80

Gail, A., Brinksmeyer, H. J., and Eckhorn, R. (2004). Perception-related modulations of local field potential power and coherence in primary visual cortex of the awake monkey during binocular rivalry. *Cerebral Cortex*, 14:300–313. 8

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron*, 14:477–485. 58, 74, 80

Gourevitch, B., Le Bouquin-Jeannes, R., and Faucon, G. (2006). Linear and non-linear causality between signals: methods, examples and neurophysiological applications. *Biological Cybernetics*, 95(4):349–369. 12

Graimann, B. and Pfurtscheller, G. (2006). Quantification and visualization of event-related changes in oscillatory brain activity in the time-frequency domain. *Prog Brain Res*, 159:79–97. 16, 34

Granger, C. W. J. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438. 4, 9

Green, J. D. and Arduini, A. A. (1954). Hippocampal electrical activity in arousal. *J Neurophysiol*, 17(6):533–557. 55

Gregoriou, G. G., Gotts, S. J., Zhou, H., and Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*, 324(5931):1207–1210. 31, 41, 52

Grillner, S. (2006). Biological pattern generation: The cellular and computational logic of networks in motion. *Neuron*, 52:751–766. 59, 79

Guderian, S. and Duzel, E. (2005). Induced theta oscillations mediate large-scale synchrony with mediotemporal areas during recollection in humans. *Hippocampus*, 15(7):901–912. 53

Guertin, P. A. (2009). The mammalian central pattern generator for locomotion. *Brain Research Reviews*, 62(1):45–56. 59

Haider, B. and McCormick, D. A. (2009). Rapid neocortical dynamics: cellular and network mechanisms. *Neuron*, 62(2):171–189. 53

Haykin, S. (1999). *Neural Networks: A Comprehensive Foundation*. Prentice Hall, New Jersey, 2nd edition. 65, 66

Haykin, S. (2001). *Adaptive Filter Theory (4th Edition)*. Prentice Hall. 60, 63

Hebb, D. O. (1949). *The Organization of Behavior*. Wiley, New York. 2, 61

Hipp, J. F., Engel, A. K., and Siegel, M. (2011). Oscillatory synchronization in large-scale cortical networks predicts perception. *Neuron*, 69(2):387–396. 30

Hoerzer, G. M., Liebe, S., Schloegl, A., Logothetis, N. K., and Rainer, G. (2010). Directed coupling in local field potentials of macaque V4 during visual short-term memory revealed by multivariate autoregressive models. *Front Comput Neurosci*, 4(14):1–13. 31

Hopfield, J. J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *PNAS*, 81:3088–3092. 65, 66

Hosp, J. A., Pekanovic, A., Rioult-Pedotti, M. S., and Luft, A. R. (2011). Dopaminergic projections from midbrain to primary motor cortex mediate motor skill learning. *Journal of Neuroscience*, 31(7):2481–2487. 58, 80

Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: A review. *Neural Networks*, 21:642–653. 59

Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of STDP and dopamine signaling. *Cerebral Cortex*, 17:2443–2452. 2, 61

Jachan, M., Henschel, K., Nawrath, J., Schad, A., Timmer, J., and Schelter, B. (2009). Inferring direct directed-information flow from multivariate nonlinear time series. *Physical Review E (Statistical, Nonlinear, and Soft Matter Physics)*, 80(1):011138. 26

Jaeger, H. (2003). Adaptive nonlinear system identification with echo state networks. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press. 3, 59, 65

Jaeger, H. and Haas, H. (2004). Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless communication. *Science*, 304:78–80. 3, 60, 63, 66, 67, 78, 79

Jensen, O. and Lisman, J. E. (2005). Hippocampal sequence-encoding driven by a cortical multi-item working memory buffer. *Trends Neurosci*, 28(2):67–72. 55

Jones, M. W. and Wilson, M. A. (2005). Theta rhythms coordinate hippocampal-prefrontal interactions in a spatial memory task. *PLoS Biol*, 3(12):e402. 2, 30, 42, 55

Kahana, M. J., Seelig, D., and Madsen, J. R. (2001). Theta returns. *Curr Opin Neurobiol*, 11(6):739–744. 55

Kahana, M. J., Sekuler, R., Caplan, J. B., Kirschen, M., and Madsen, J. R. (1999). Human theta oscillations exhibit task dependence during virtual maze navigation. *Nature*, 399(6738):781–784. 30, 52, 55

Kaminski, M. and Blinowska, K. J. (1991). A new method of the description of the information flow in brain structures. *Biological Cybernetics*, 65:203–210. 9

Kayser, C. and Logothetis, N. K. (2009). Directed interactions betwen auditory and superior temporal cortices and their role in sensory integration. *Frontiers in Integrative Neuroscience*, 3. 9, 12, 24

Keil, A., Gruber, T., and Müller, M. M. (2001). Functional correlates of macroscopic high-frequency brain activity in the human visual system. *Neuroscience and Biobehavioral Reviews*, 25:527–534. 24

Klein, R. (1999). Donald O. Hebb. In Wilson, R. A. and Keil, F. C., editors, *MIT Encyclopedia of the Cognitive Sciences*, pages 366–367. MIT Press. 1

Klimesch, W. (1996). Memory processes, brain oscillations and EEG synchronization. *Int J Psychophysiol*, 24(1-2):61–100. 30, 39, 42, 52, 53, 55

Klimesch, W. (1999). EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis. *Brain Research Reviews*, 29:169–195. 8, 25

Kus, R., Kaminski, M., and Blinowska, K. J. (2004). Determination of EEG activity propagation: pair-wise versus multichannel estimate. *Biomedical Engineering, IEEE Transactions on*, 51(9):1501–1510. 12

Lachaux, J. P., Rodriguez, E., Le Van Quyen, M., Lutz, A., Martinerie, J., and Varela, F. J. (2000). Studying single-trials of phase synchronous activity in the brain. *International Journal of Bifurcation and Chaos*, 10(10):2429–2439. 9, 24

Lachaux, J. P., Rodriguez, E., Martinerie, J., and Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4):194–208. 9, 24, 34, 40, 52

Lashley, K. S. (1929). *Brain mechanisms and intelligence: A quantitative study of injuries to the brain.* University of Chicago Press. 1

Le Van Quyen, M., Foucher, J., Lachaux, J. P., Rodriguez, E., Lutz, A., Martinerie, J., and Varela, F. J. (2001). Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J Neurosci Methods*, 111(2):83–98. 34

Lee, H., Simpson, G. V., Logothetis, N. K., and Rainer, G. (2005). Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron*, 45(1):147–156. 8, 14, 16, 24, 25, 31, 32, 39, 44, 52, 55

Legenstein, R., Chase, S. M., Schwartz, A. B., and Maass, W. (2010). A reward-modulated Hebbian learning rule can explain experimentally observed network reorganization in a brain control task. *The Journal of Neuroscience*, 30(25):8400–8410. 61, 66, 68, 78

Legenstein, R. and Maass, W. (2007). Edge of chaos and prediction of computational performance for neural circuit models. *Neural Networks*, 20(3):323–334. 61

Legenstein, R., Pecevski, D., and Maass, W. (2008). A learning theory for reward-modulated spike-timing-dependent plasticity with application to biofeedback. *PLoS Computational Biology*, 4(10):1–27. 2, 61

Lewis, D. A., Campbell, M. J., Foote, S. L., Goldstein, M., and Morrison, J. H. (1987). The distribution of tyrosine hydroxylase-immunoreactive fibers in primate neocortex is widespread but regionally specific. *Journal of Neuroscience*, 7(1):279–290. 80

Liang, H., Bressler, S. L., Ding, M., Desimone, R., and Fries, P. (2003). Temporal dynamics of attention-modulated neuronal synchronization in macaque V4. *Neurocomputing*, 52-54:481–487. 9

Liang, H., Ding, M., and Bressler, S. L. (2000). On the tracking of dynamic functional relations in monkey cerebral cortex. *Neurocomputing*, 32-33:891–896. 9

Liang, H., Ding, M., and Bressler, S. L. (2001). Temporal dynamics of information flow in the cerebral cortex. *Neurocomputing*, 38-40:1429–1435. 9

Liebe, S., Fischer, E., Logothetis, N. K., and Rainer, G. (2009). Color and shape interactions in the recognition of natural scenes by human and monkey observers. *J Vis*, 9(5):14.1–16. 14, 33, 38

Lisman, J. E. (1999). Relating hippocampal circuitry to function: recall of memory sequences by reciprocal dentate-CA3 interactions. *Neuron*, 22(2):233–242. 55

Lisman, J. E. and Buzsaki, G. (2008). A neural coding scheme formed by the combined function of gamma and theta oscillations. *Schizophr Bull*, 34(5):974–980. 30, 55

Lukosevicius, M. and Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149. 60

Maass, W., Joshi, P., and Sontag, E. D. (2007). Computational aspects of feedback in neural circuits. *PLoS Computational Biology*, 3(1):e165, 1–20. 3, 60, 63, 66, 78, 79

Maass, W., Natschlaeger, T., and Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560. 3, 59, 65

MacKay-Lyons, M. (2002). Central pattern generation of locomotion: A review of the evidence. *Physical Therapy*, 82(1):69–83. 59

Marder, E. and Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Current Biology*, 11:R986–R996. 58, 59

Marder, E. and Goaillard, J.-M. (2006). Variability, compensation and homeostasis in neuron and network function. *Nature Reviews Neuroscience*, 7:563–574. 59

Marple, S. L. (1987). *Digital Spectral Analysis with Applications*. Englewood Cliffs: Prentice Hall. 11

Martin, S. J., Grimwood, P. D., and Norris, R. G. M. (2000). Synaptic plasticity and memory: An evaluation of the hypothesis. *Annual Reviews in Neuroscience*, 23:649–711. 54

Mazzoni, P., Andersen, R. A., and Jordan, M. I. (1991). A more biologically plausible learning rule for neural networks. *Proceedings of the National Academy of Sciences*, 88:4433–4437. 61

Mehta, M. R., Quirk, M. C., and Wilson, M. A. (2000). Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron*, 25(3):707–715. 55

Miller, E. K. and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu Rev Neurosci*, 24:167–202. 31, 56

Miller, E. K., Erickson, C. A., and Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *Journal of Neuroscience*, 16:5154–5167. 58, 74, 80

Miller, R. (1991). *Cortico-hippocampal Interplay and the representation of contexts in the brain*, volume 17 of *Studies of Brain Function*. Springer Verlag. 55

Mizuseki, K., Sirota, A., Pastalkova, E., and Buzsaki, G. (2000). Theta oscillations provide temporal windows for local circuit computation in the entorhinal-hippocampal loop. *Neuron*, 64:267–280. 24, 25

Molina-Luna, K., Pekanovic, A., Röhrich, S., Hertler, B., Schubring-Giese, M., Rioult-Pedotti, M.-S., and Luft, A. R. (2009). Dopamine in motor cortex is necessary for skill learning and synaptic plasticity. *PLoS ONE*, 4(9):e7082. 58, 80

Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, 319(5869):1543–1546. 55, 56

Montemurro, M., Rasch, M. J., Murayama, Y., Logothetis, N. K., and Panzeri, S. (2008). Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Current Biology*, 18:375–380. 24

Murthy, V. N. and Fetz, E. E. (1992). Coherent 25- to 35-hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proceedings of the National Academy of Sciences*, 15:5670–5674. 24

Nowak, J. and Bullier, L. G. (1997). The timing of information transfer in the visual system. In Kaas, K. S., Peters, J. H., and Rockland, A., editors, *Cerebral Cortex: Extrastriate Cortex in Primates*, volume 12, page 36. Plenum. 41, 52

Nunez, P. L., Silberstein, R. B., Shi, Z., Carpenter, M. R., Srinivasan, R., Tucker, D. M., Doran, S. M., Cadusch, P. J., and Wijesinghe, R. S. (1999). EEG coherency II: experimental comparisons of multiple measures. *Clinical Neurophysiology*, 110(3):469–486. 9

Nunez, P. L., Srinivasan, R., Westdorp, A. F., Wijesinghe, R. S., Tucker, D. M., Silberstein, R. B., and Cadusch, P. J. (1997). EEG coherency I: statistics, reference electrode, volume conduction, laplacians, cortical imaging, and interpretation at multiple scales. *Electroencephalography and Clinical Neurophysiology*, 103(5):499–515. 9

O'Keefe, J. (1993). Hippocampus, theta, and spatial memory. *Curr Opin Neurobiol*, 3(6):917–924. 8, 24, 25, 30

Orban, G. A. (2008). Higher order visual processing in macaque extrastriate cortex. *Physiol Rev*, 88(1):59–89. 31

Pasternak, T. and Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nat Rev Neurosci*, 6(2):97–107. 31, 42, 46, 52

Pasupathy, A. (2006). Neural basis of shape representation in the primate brain. *Prog Brain Res*, 154:293–313. 31

Pavlides, C., Greenstein, Y. J., Grudman, M., and Winson, J. (1988). Long-term potentiation in the dentate gyrus is induced preferentially on the positive phase of theta-rhythm. *Brain Res*, 439(1-2):383–387. 54

Pereda, E., Quian Quiroga, R., and Bhattacharya, J. (2005). Nonlinear multivariate analysis of neurophysiological signals. *Progress in Neurobiology*, 77:1–37. 12, 26

Pesaran, B., Pezaris, J. S., Sahani, M., Mitra, P. P., and Andersen, R. A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci*, 5(8):805–811. 24, 31

Petrides, M. (1996). Specialized systems for the processing of mnemonic information within the primate frontal cortex. *Philos Trans R Soc Lond B Biol Sci*, 351(1346):1455–1462. 31

Pfister, J.-P., Toyoizumi, T., Barber, D., and Gerstner, W. (2006). Optimal spike-timing dependent plasticity for precise action potential firing in supervised learning. *Neural Computation*, 18:1309–1339. 61

Pipa, G., Stadtler, E. S., Rodriguez, E. F., Waltz, J. A., Muckli, L. F., Singer, W., Goebel, R., and Munk, M. H. (2009). Performance- and stimulus-dependent oscillations in monkey prefrontal cortex during short-term memory. *Front Integr Neurosci*, 3(25):1–13. 31

Porcaro, C., Zappasodi, F., Rossini, P. M., and Tecchio, F. (2009). Choice of multi-variate autoregressive model order affecting real network functional connectivity estimate. *Clinical Neurophysiology*, 120(2):436–448. 9, 12

Prinz, A. A., Bucher, D., and Marder, E. (2004). Similar network activity from disparate circuit parameters. *Nature Neuroscience*, 7(12):1345–1352. 59

Purves, D., Brannon, E., Cabeza, R., Huettel, S. A., Labar, K., Platt, M., and Woldorff, M. (2008). *Principles of Cognitive Neuroscience*. Palgrave Macmillan. 3

Raghavachari, S., Kahana, M. J., Rizzuto, D. S., Caplan, J. B., Kirschen, M. P., Bourgeois, B., Madsen, J. R., and Lisman, J. E. (2001). Gating of human theta oscillations by a working memory task. *J Neurosci*, 21(9):3175–3183. 8, 25, 30, 39, 52

Raghavachari, S., Lisman, J. E., Tully, M., Madsen, J. R., Bromfield, E. B., and Kahana, M. J. (2006). Theta oscillations in human cortex during a working-memory task: evidence for local generators. *J Neurophysiol*, 95(3):1630–1638. 8, 22, 24, 25, 26, 39, 54

Rainer, G., Lee, H., Simpson, G. V., and Logothetis, N. K. (2004). Working-memory related theta (4-7Hz) frequency oscillations observed in monkey extrastriate visual cortex. *Neurocomputing*, 58–60. 8, 16, 24, 25, 31, 39, 52

Rizzuto, D. S., Madsen, J. R., Bromfield, E. B., Schulze-Bonhage, A., and Kahana, M. J. (2006). Human neocortical oscillations exhibit theta phase differences between encoding and retrieval. *Neuroimage*, 31(3):1352–1358. 54

Rutishauser, U., Ross, I. B., Mamelak, A. N., and Schuman, E. M. (2010). Human memory strength is predicted by theta-frequency phase-locking of single neurons. *Nature*, 464(7290):903–907. 30, 54

Salinas, E. and Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nat Rev Neurosci*, 2(8):539–550. 23, 24, 53

Sarnthein, J., Petsche, H., Rappelsberger, P., Shaw, G. L., and von Stein, A. (1998). Synchronization between prefrontal and posterior association cortex during human working memory. *Proc Natl Acad Sci U S A*, 95(12):7092–7096. 30, 42, 52, 55

Sato, N. and Yamaguchi, Y. (2007). Theta synchronization networks emerge during human object-place memory encoding. *Neuroreport*, 18(5):419–424. 53

Schack, B., Vath, N., Petsche, H., Geissler, H. G., and Moller, E. (2002). Phase-coupling of theta-gamma EEG rhythms during short-term memory processing. *Int J Psychophysiol*, 44(2):143–163. 52

Schein, S. J. and Desimone, R. (1990). Spectral properties of V4 neurons in the macaque. *Journal of Neuroscience*, 10(10):3369–3389. 31

Schelter, B., Timmer, J., and Eichler, M. (2009). Assessing the strength of directed influences among neural signals using renormalized partial directed coherence. *Journal of Neuroscience Methods*, 179(1):121–130. 13

Schlögl, A. (2000). *The electroencephalogram and the adaptive autoregressive model: theory and applications.* Shaker Verlag. 11

Schlögl, A. (2006). A comparison of multivariate autoregressive estimators. *Signal Processing*, 86(9):2426–2429. 11

Schlögl, A. and Brunner, C. (2008). BioSig: A free and open source software library for BCI research. *Computer*, 41(10):44–50. 9, 11

Schlögl, A., Roberts, S. J., and Pfurtscheller, G. (2000). A criterion for adaptive autoregressive models. In *Proceedings of the 22nd annual international conference of the IEEE engineering in medicine and biology society*, volume 2, pages 1581–1582. 12

Schlögl, A. and Supp, G. (2006). Analyzing event-related EEG data with multivariate autoregressive parameters. *Progress in Brain Research*, (159):135–147. 9, 12, 24

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6:461–464. 11

Sederberg, P. B., Kahana, M. J., Howard, M. W., Donner, E. J., and Madsen, J. R. (2003). Theta and gamma oscillations during encoding predict subsequent recall. *J Neurosci*, 23(34):10809–10814. 30, 39, 52, 53, 54

Siapas, A. G., Lubenov, E. V., and Wilson, M. A. (2005). Prefrontal phase locking to hippocampal theta oscillations. *Neuron*, 46(1):141–151. 2, 25, 30, 36, 37, 43, 55

Siegel, M., Warden, M. R., and Miller, E. K. (2009). Phase-dependent neuronal coding of objects in short-term memory. *Proc Natl Acad Sci U S A*, 106(50):21341–21346. 31, 37, 44, 52, 55

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron*, 24(1):49–65. 53

Sirota, A., Montgomery, S., Fujisawa, S., Isomura, Y., Zugaro, M., and Buzsaki, G. (2008). Entrainment of neocortical neurons and gamma oscillations by hippocampal theta rhythm. *Neuron*, 60(4):683–697. 55

Stam, C. J., van Cappellen van Walsum, A. M., and Micheloyannis, S. (2002). Variability of EEG synchronization during a working memory task in healthy subjects. *Int J Psychophysiol*, 46(1):53–66. 30, 42, 52

Stevenson, I. H., Rebesco, J. M., Miller, L. E., and Körding, K. P. (2008). Inferring functional connections between neurons. *Current Opinion in Neurobiology*, 18:582–588. 26

Sun, X. (2008). Assessing nonlinear granger causality from multivariate time series. In *ECML PKDD '08: Proceedings of the European conference on Machine Learning and Knowledge Discovery in Databases - Part II*, pages 440–455. Springer. 26

Supp, G. G., Schlögl, A., Trujillo-Barreto, N., Müller, M. M., and Gruber, T. (2007). Directed cortical information flow during human object recognition: analyzing induced EEG gamma-band responses in brain's source space. *PLoS ONE*, 2(1). 8, 9, 11, 12, 24, 25

Sussillo, D. and Abbott, L. F. (2009). Generating coherent patterns of activity from chaotic neural networks. *Neuron*, 63(4):544–557. 3, 60, 62, 63, 64, 65, 67, 68, 69, 71, 73, 78, 79

Suzuki, W. A. and Amaral, D. G. (1994). Perirhinal and parahippocampal cortices of the macaque monkey: cortical afferents. *J Comp Neurol*, 350(4):497–533. 55

Szatmary, B. and Izhikevich, E. M. (2010). Spike-timing theory of working memory. *PLoS Comput Biol*, 6(8):e1000879. 2

Tallon-Baudry, C. (2009). The roles of gamma band oscillatory synchrony in human visual cognition. *Frontiers in Biosciences*, 14:321–332. 24

Tallon-Baudry, C. and Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. *Trends in Cognitive Sciences*, 3(4):151–162. 16, 34

Tanigawa, H., Lu, H. D., and Roe, A. W. (2010). Functional organization for color and orientation in macaque V4. *Nature Neuroscience*, 13(12):1542–1548. 31

Tesche, C. D. and Karhu, J. (2000). Theta oscillations index human hippocampal activation during a working memory task. *Proc Natl Acad Sci USA*, 97(2):919–924. 30, 55

Tiesinga, P., Fellous, J. M., and Sejnowski, T. J. (2008). Regulation of spike timing in visual cortical circuits. *Nat Rev Neurosci*, 9(2):97–107. 52

Uhlhaas, P., Pipa, G., Lima, B., Melloni, L., Neuenschwander, S., Nikolic, D., and Singer, W. (2009). Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience*, 3(17):1–19. 24

Ungerleider, L. G., Gaffan, D., and Pelak, V. S. (1989). Projections from inferior temporal cortex to prefrontal cortex via the uncinate fascicle in rhesus monkeys. *Exp Brain Res*, 76(3):473–484. 55

Ungerleider, L. G., Galkin, T. W., Desimone, R., and Gattass, R. (2008). Cortical connections of area V4 in the macaque. *Cereb Cortex*, 18(3):477–499. 54, 55

Varela, F., Lachaux, J. P., Rodriguez, E., and Martinerie, J. (2001). The brainweb: phase synchronization and large-scale integration. *Nat Rev Neurosci*, 2(4):229–239. 40

Verduzco-Flores, S., Bodner, M., Ermentrout, B., Fuster, J. M., and Zhou, Y. (2009). Working memory cells' behavior may be explained by cross-regional networks with synaptic facilitation. *PLoS One*, 4(8):e6399. 55

Volgushev, M., Chistiakova, M., and Singer, W. (1998). Modification of discharge patterns of neocortical neurons by induced oscillations of the membrane potential. *Neuroscience*, 83(1):15–25. 52, 53

Walker, E. (1940). A cytoarchitectural study of the prefrontal area of the macaque monkey. *Journal of Comparative Neurology*, 98:59–86. 54

Weiss, S., Muller, H. M., and Rappelsberger, P. (2000). Theta synchronization predicts efficient memory encoding of concrete and abstract nouns. *Neuroreport*, 11(11):2357–2361. 52

Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8(3):229–256. 61

Winson, J. (1972). Interspecies differences in the occurrence of theta. *Behav Biol*, 7(4):479–487. 55

Winson, J. (1978). Loss of hippocampal theta rhythm results in spatial memory deficit in the rat. *Science*, 201(4351):160–163. 55

Wolf, H. and Pearson, K. G. (1987). Comparison of motor patterns in the intact and deafferented flight system of the locust. ii. intracellular recordings from flight motoneurons. *Journal of Comparative Physiology A*, 160:269–279. 59

Womelsdorf, T., Schoffelen, J. M., Oostenveld, R., Singer, W., Desimone, R., Engel, A. K., and Fries, P. (2007). Modulation of neuronal interactions through neuronal synchronization. *Science*, 316(5831):1609–1612. 30, 53

Xie, X. and Seung, H. S. (2004). Learning in neural networks by reinforcement of irregular spiking. *Physical Review E*, 69(4):41909. 2, 61

Yao, D., Yamamura, K., Narita, N., Martin, R. E., Murray, G. M., and Sessle, B. J. (2002). Neuronal activity patterns in primate motor cortex related to trained or semiautomatic jaw and tongue movements. *Journal of Neurophysiology*, 87:2531–2541. 58, 80

Zar, J. (2008). *Biostatistical Analysis*. Prentice Hall, 5 edition. 36, 43

Zeki, S. (1980). The representation of colours in the cerebral cortex. *Nature*, 284(5755):412–418. 31

Zetterberg, L. (1969). Estimation of parameters for a linear difference equation with application to EEG analysis. *Mathematical Biosciences*, 5:227–275. 9