

Interpreting the Local Field Potential as a
Reflection of Cooperative Neuronal Spiking
Dynamics

Dissertation

zur Erlangung des akademischen Grades des
Doktors der Naturwissenschaften (Dr. rer. nat.)

eingereicht im Fachbereich Biologie, Chemie, Pharmazie
der Freien Universität Berlin

vorgelegt von

Michael Denker

aus Göttingen

August 2009

Die in dieser Dissertation vorgestellte Forschung entstand in der Zeit vom Februar 2004 bis Juli 2009 am Institut für Neurobiologie, Freie Universität Berlin, und am RIKEN Brain Science Institute, Wako-shi, Japan. (The research presented in this thesis was carried out from February 2004 until July 2009 at the Institute for Neurobiology, Free University Berlin, Germany, and the RIKEN Brain Science Institute, Wako-shi, Japan.)

- 1. Gutachter (1st referee): Herr Prof. Dr. Dr. h.c. Randolph Menzel, Institut für Neurobiologie, Freie Universität Berlin, Königin-Luise-Str. 28/30, 14915 Berlin, Germany.
- 2. Gutachter (2nd referee): Frau PD Dr. Sonja Grün, RIKEN Brain Science Institute, 2-1 Hirosawa, Wako-shi 351-0198, Japan.

Disputation am 6.10.2009

List of Published Work and Author Contributions

The work presented in [Sec. 2.2](#) and [Sec. 3.1](#) has been published in peer-reviewed international journals. The work presented in [Sec. 2.1](#), [Sec. 3.2](#), and [Sec. 3.3](#) is currently under review in peer-reviewed international journals.

[Sec. 2.1](#) is submitted to the European Journal of Neuroscience as:

M. Denker, R. Finke, F. Schaupp, S. Grün, and R. Menzel. Neural Correlates of Odor Learning in the Honeybee Antennal Lobe. Publisher homepage: <http://www.blackwell-synergy.com>

Michael Denker designed and performed the analysis, and wrote the manuscript. Robert Finke designed and performed the experiments, designed and performed a preliminary analysis, and was involved in writing an early version of the manuscript. Frank Schaupp designed and set up the experiments, was involved in performing the experiments, and helped in various technical issues. Sonja Grün designed the analysis and wrote the manuscript. Randolph Menzel designed the experiments, was involved in the design of the analysis, and wrote the manuscript.

[Sec. 2.2](#) has been published as:

A. Sharott, C. K. E. Moll, G. Engler, M. Denker, S. Grün, and A. K. Engel. Different subtypes of striatal neurons are selectively modulated by cortical oscillations. *J Neurosci* 29, 4571-4585 (2009).

The published article is available at <http://www.jneurosci.org>
doi: [10.1523/JNEUROSCI.5097-08.2009](https://doi.org/10.1523/JNEUROSCI.5097-08.2009)

Andrew Sharott wrote the manuscript, and designed and performed the following parts of the analysis: neuron type classification, power analysis, and cross correlation analysis. Andrew Sharott, Christian Moll, and Gerhard Engler designed and performed the experiments. Michael Denker designed and performed the analysis relating spikes and LFP, performed a preliminary power analysis, and wrote the manuscript. Sonja Grün was involved in the design of the analysis. Andreas Engel was involved in the design of the experiments and in the writing of the manuscript.

Sec. 3.1 has been published as:

M. Denker, S. Roux, M. Timme, A. Riehle, and S. Grün. Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomputing* 70, 2096-2101 (2007).

The published article is available at <http://www.elsevier.com>
doi: [10.1016/j.neucom.2006.10.088](https://doi.org/10.1016/j.neucom.2006.10.088)

Michael Denker designed and performed the analysis, and wrote the manuscript. Sébastien Roux designed and performed the experiments. Marc Timme contributed to the theoretical aspects of the analysis design. Alexa Riehle designed the experiments. Sonja Grün was involved in the design of the analysis and in writing the manuscript.

Sec. 3.2 will be submitted as:

M. Denker, S. Roux, H. Lindén, M. Diesmann, A. Riehle, and S. Grün. Local field potentials reflect neuronal assemblies.

Michael Denker designed and performed the analysis, and wrote the manuscript. Sébastien Roux designed and performed the experiments. Henrik Lindén performed a preliminary analysis that is not contained in the present form of the manuscript. Markus Diesmann was involved in the design of the analysis and in writing the manuscript. Alexa Riehle designed the experiments and contributed to writing the manuscript. Sonja Grün was involved in the design of the analysis and in writing the manuscript.

Sec. 3.3 is submitted to the Journal of Computational Neuroscience as:

M. Denker, A. Riehle, M. Diesmann, and S. Grün. Estimating the contribution of assembly activity to cortical dynamics from spike and population measures.

Publisher homepage: <http://www.springer.com>

Michael Denker designed and performed the analysis, and wrote the manuscript. Alexa Riehle designed the original experiments. Markus Diesmann designed the analysis and wrote the manuscript. Sonja Grün designed the analysis and wrote the manuscript.

Summary of Main Findings

Brain activity is observed using various methods of measurement that operate on a large range of spatial and temporal scales. In the case of extracellular recordings, the electrode signal is commonly interpreted on two such levels: while its high-frequency components reveal the precise timing of locally generated action potentials, the low-frequency contributions constitute a mesoscopic population signal –referred to as the local field potential (LFP)– that is dominated by the superposition of activities of neurons in a larger volume. Although it is commonly assumed that the oscillatory features of the LFP result from the synchronous subthreshold activity in the neuronal population, it is a question of ongoing debate how these oscillations are related to the spiking discharge, and to synchrony on the level of spikes in particular. This thesis contains five reports that describe different aspects of how the relationship between the concerted spiking activity and the LFP is informative of the features and the dynamical organization of the underlying neuronal system. The focus of this thesis rests on three of these reports that uncover the missing link between precise synchrony of action potentials and the LFP.

In order to reliably quantify the typically weak entrainment of neurons to the LFP a novel technique was developed that originates in the phase synchronization analysis of continuous signals. In contrast to commonly used measures based on signal averaging, such as the spike-triggered average, the method directly analyzes the spike-LFP phase relationship on the basis of the single spike. All the reports in this thesis demonstrate how this methodological tool sensitively identifies and characterizes the neuronal ensembles that are coupled to the observed population oscillations.

The first study demonstrates learning-related response changes to odor stimulation in the output region of the antennal lobe of the honeybee during olfactory conditioning. We show that learning induces a restructuring of the ensemble representation of odors (based on the rate responses of individual neurons) that is strongest for the rewarded odor. Moreover, observed response changes in LFP power in a given frequency band for the rewarded odor correlate with changes in the size and composition of the neuronal subpopulation phase-locked to the same band. Therefore, the analysis in this report reveals strong support for the hypothesis that the LFP acts as monitor

of the learning-induced reorganization of the temporal ensemble representation of an external stimulus.

In the second study, we analyze separately the role of four identified neuronal subtypes in the striatum of anesthetized rats in relaying oscillations of cortical origin to the basal ganglia network. The methods based on spike-LFP and spike-ECOG (electrocorticogram) phase coupling, augmented by the analysis of interneuronal cross-correlations, allows us to identify the neuronal subtype, fast spiking interneurons, that is linked to the generation of the oscillatory LFP component in the high-frequency (gamma) range. Here, the LFP reflects the synchronized action of a specific physiological neuron type.

The third and largest study of this thesis consists of three successive reports that reveal how spike synchrony on a millisecond time-scale is related to the synchronous mass action visible in the LFP. The experimental data is recorded from motor cortex of monkey in a delayed pointing task. We first show that the amplitude (or envelope) of the LFP is indicative of the proportion of spikes that phase-lock to the LFP. In the following report we directly prove the long-standing hypothesis that synchronous spikes are reflected in the field potential by analyzing the LFP coupling separately for single spikes and precise spike coincidences. Contrary to intuition, we show that this relation holds only for synchronous spikes that occur in time windows where the number of observed coincidences significantly exceeds the expectation. These excess coincidences are signatures of the coordinated discharge patterns of specific neuronal subgroups (assemblies). A conceptual model consistently explains the experimental data in the context of the assembly framework. In the final study we show how the combination of measures of synchrony on the spike and population levels reveals a macroscopic parameter that estimates the network-wide degree to which active cell assemblies contribute to neuronal spike data. In summary, these three studies provide first-time evidence that LFP oscillations are an image of the orchestrated activity of neuronal ensembles, as predicted by one of today's most intriguing theories on neuronal computation.

Contents

1	Introduction	10
1.1	The Local Field Potential (LFP) as a Probe for Population Activity	11
1.1.1	Origin and Correlates of the LFP	11
1.1.2	Interpretation of Oscillatory LFP	14
1.1.3	LFPs in the Framework of the Assembly Hypothesis of Neuronal Coding	16
1.2	Phase Locking between Spiking Activity and LFP	18
1.2.1	Introducing the Phase Analysis	18
1.2.2	Sensitivity and Reliability of the Phase Analysis	21
1.3	Overview of the Thesis	22
2	LFP as a Monitor of Neuronal Activity	25
2.1	Neural Correlates of Odor Learning in the Honeybee Antennal Lobe	25
2.2	Different Subtypes of Striatal Neurons Are Selectively Modulated by Cortical Oscillations	58
3	Relating LFPs to Neuronal Assemblies	83
3.1	Phase Synchronization between LFP and Spiking Activity in Motor Cortex During Movement Preparation	83
3.2	Local Field Potentials Reflect Neuronal Assemblies	90
3.3	Estimating the Contribution of Assembly Activity to Cortical Dynamics from Spike and Population Measures	122
4	Conclusion	150
4.1	Summary and Interpretation	150
4.2	General comments	159

<i>CONTENTS</i>	9
4.3 Outlook	161
A Zusammenfassung	165
Bibliography	168

Chapter 1

Introduction

Even after more than a century of intense research in the neurosciences, it is fascinating to realize that we are still struggling to understand the basic principles by which our brains work. The stereotypical, all-or-none pulse emitted by neurons, the action potential or spike, is commonly identified as the most likely carrier of information in the nervous system. At least two fundamental problems obstruct our attempts in trying to interpret the role of individual spikes in the light of cortical processing: undersampling and complexity. First, if the detailed knowledge of the emission times of spikes is necessary to understand the workings of our brain, then how should we deal with the fact that even the most sophisticated recording technologies capture only a small proportion of spikes? Second, even if all action potentials were known to us, we are in need of methods that deal with the complexity of the collected data. The problem is more than a mere technicality: it requires knowledge of the subtle signatures that a hypothesized dynamical process would imprint on a large set of observations.

Virtually parallel to the discovery of the microscopic potentials across membranes in the single cell, large-scale electric potentials exhibiting wave-like properties were identified from recordings of scalp electrodes, the EEG. These signals are easy to record and sample across a large part of the neuronal population, but are unspecific in origin. From the large size of such population signals compared to the simple transmembrane potentials, they are believed to arise from coherence within the neuronal population. This dynamical interpretation gives hope that the combination of large-scale brain signals and the detailed information on spiking level may provide insights

on the underlying mechanisms that determine the spiking activity of the network. However, how the population signals emerge from the activity of individual neurons is still a matter of ongoing research. In this study, we investigate the relationship of coherence phenomena on the single cell level to the oscillatory activity of the local field potential (LFP), a mesoscopic population measure operating on a much smaller spatial scale than the EEG. In the Introduction, we first give a brief overview of current research on the biophysical origins of the LFP and the common interpretation of its oscillatory features. In this context, we introduce one of the major theories on higher brain function developed in the last decades, and state its hypothesized relation to the LFP. In the second part we explain the analytic tools we developed to evaluate the synchronization between the spiking activity and population signals. Finally, we briefly guide the reader through the sequence of reports that make up the main part of this thesis.

1.1 The Local Field Potential (LFP) as a Probe for Population Activity

1.1.1 Origin and Correlates of the LFP

Neuronal data is typically recorded on various spatial and temporal scales (Varela et al., 2001), ranging from large-scale measurements of brain activity, such as functional magnetic resonance imaging (fMRI) or electroencephalography (EEG), to observations on the single cell level in an intracellular recording. On the level of extracellular recordings, the rich electrode signal is typically evaluated on two distinct time scales. To this end, the raw electrode signal is filtered to extract the frequency components of interest before further processing is performed. High frequencies, where a typical range of values might be 500 Hz to 5 kHz, yield a signal from which the extracellular spike wave forms emitted by neurons in a close vicinity ($\approx 140 \mu\text{m}$) of the electrode may be extracted (Buzsáki, 2004). In addition to determining the exact spike times from the filtered signal, the magnitude of the rectified electrode signal, sometimes referred to as multiple-unit activity (MUA), is regarded as a measure of the average spiking intensity of neurons near the recording site (Eckhorn et al., 1988). In contrast to these signals, the slow components of the electrode signal, typically obtained below frequencies of

200 Hz (sometimes higher, see [Buzsáki and Draguhn, 2004](#)), are collectively termed the local field potential, or LFP.

The origin of the LFP is commonly attributed to the weighted average of transmembrane currents of neurons in the vicinity of the recording electrode: Currents flowing into one cell (current sink) at one location (e.g., an influx of positive ions) will cause a counteracting outward current (source) at another location some distance apart ([Nunez and Shrinivasan, 2006](#)). The resulting dipole will evoke an electric field visible in the conducting extracellular space. Its major contributions are assumed to originate from currents associated with excitatory and inhibitory synaptic activity ([Mitzdorf, 1985](#); [Viswanathan and Freeman, 2007](#); [Monosov et al., 2008](#)), though other contributors to the subthreshold dynamics are also likely to influence LFP dynamics. In particular, various voltage-gated membrane oscillations ([Kamondi et al., 1998](#)) and the afterpolarizations following soma-dendritic spikes ([Buzsaki et al., 1988](#)) are candidates for secondary factors that could influence LFP dynamics (for a current summary see, [Goense and Logothetis, 2008](#)). Nevertheless, the LFP is commonly considered to be a reflection of the mean input into a specific region of cortex, as opposed to the generated output spikes (see, e.g., [Eckhorn and Obermueller, 1993](#); and for a direct demonstration: [Viswanathan and Freeman, 2007](#)). Highlighting this view, it was recently shown in the visual system that measured LFPs at a certain location are predictable from the timing of spikes of distant neurons, given knowledge of their respective average postsynaptic effect on the LFP at that location ([Nauhaus et al., 2009](#)). Although it is assumed that the fast action potential dynamics of individual neurons is not visible in the recorded LFP ([Elul, 1971](#); [Logothetis, 2003](#)), low-frequency components of the spike may still enter the frequency-range of the LFP ([Ray et al., 2008](#)). Therefore, care must be taken when comparing LFP and spike signals from the same electrode to avoid spurious correlations between the two signals, especially at high LFP frequencies (personal observation). In cortex, the elongated, large shape of (excitatory) pyramidal neurons and the aligned, layered arrangement of their apical dendrites make them primary candidates for LFP generators ([Logothetis and Wandell, 2004](#)), since this geometry supports an additive overlap of fields (open field), in contrast to less orderly configurations where fields would cancel each other out (closed field).

A question under current debate is the radius of the region over which

the LFP is integrated. In general, although the extracellular medium may be well approximated by an Ohmic resistor, neither is its conductivity easy to measure nor would the complex structure of the extracellular space suggest it to be uniform and direction-independent (Elul, 1971; Mitzdorf, 1985). Moreover, evidence indicates that oscillations in lower frequency bands tend to exhibit increased volume conductance (Buzsáki and Draguhn, 2004). Nevertheless, current-source-density methods (Mitzdorf, 1985; Pettersen et al., 2006) successfully employ a far-field approximation to estimate the current sinks and sources underlying the generation of the LFP. Fortunately, a recent study suggests that the particulars of the extracellular medium are well approximated by an isotropic, frequency-independent resistor, as deviations from this assumption influence the LFP to a lesser degree than the magnitudes of the underlying generators of the oscillation (Logothetis et al., 2007). Data from experimental and modeling studies indicate that indeed in a columnar architecture the low-pass LFP signal may be assumed to decay less rapidly than the high-pass filtered MUA (Legatt et al., 1980; Somogyvári et al., 2005; Gold et al., 2007; Logothetis et al., 2007; Pettersen et al., 2008). Early estimates of the LFP radius based on estimates of the lateral accuracy of current-source density techniques range between 1 mm and 3 mm (Mitzdorf, 1987). Other studies (Murthy and Fetz, 1996a; Destexhe et al., 1999; Juergens et al., 1999) that analyze the correlation between population activity at different recording sites observe that these correlations decrease over distances up to around 3 – 6 mm depending on the underlying system and behavioral state. Therefore, these studies cannot disentangle what contribution to this estimate is due to strong correlations in the input patterns at both recorded patches, in contrast to volume conductance of the LFP. In fact, in areas such motor cortex, LFPs display long-range correlations and exhibit a wave-like propagation (Rubino et al., 2006) that is unlikely to be due to conductance effects. Under some conditions, LFPs may even become phase-coupled across brain areas (Courtemanche and Lamarre, 2005). A recent modeling study (Lindén et al., 2009) estimates that the LFP radius is on the order of 500 μm in the absence of interneuronal correlations, and finds that nearby correlated inputs may drastically increase the power content of the LFP. An inventive study by Katzner et al. (2009) yields an experimental estimate of the LFP radius as a Gaussian distribution with a half-width corresponding to an integration area of $\approx 250 \mu\text{m}$, by combining information

from visual orientation maps and a linear superposition model of LFPs.

While LFPs are typically considered as an epiphenomenon of the synaptic network activity, it is still an open question in how far the field potential oscillations in turn influence the spike timing of individual neurons by ephaptic effects. A recent study by [Radman et al. \(2007\)](#) demonstrates theoretically and experimentally that indeed spike timing can be affected by oscillatory fields of even small amplitudes, comparable to the field strength of LFPs. Other studies in hippocampus confirm a possible role of field potentials in enhancing spiking precision ([Anastassiou et al., 2008](#)). A different modeling study ([Holt and Koch, 1999](#)) suggests that ephaptic effects of extracellular potentials due to spiking activity on pyramidal cells are small, and should have an effect only on neurons close to threshold. As current evidence on the relevance of ephaptic effects is inconclusive for the neuronal systems investigated in this thesis, we discuss this topic further in [Sec. 4.2](#).

1.1.2 Interpretation of Oscillatory LFP

LFP signals typically exhibit an oscillatory structure, and can often be decomposed into a number of oscillatory modes. These modes are historically classified into different bands ([Buzsáki and Draguhn, 2004](#)) based on their approximate frequency composition and behavioral correlates. In this study we focus on two of the fastest bands, namely the β - (approx. 15 – 45 Hz) and γ -bands (approx. 30 – 100 Hz), which are sometimes collectively described ([Donoghue et al., 1998](#)). For example definitions of both bands in a single study see, e.g., [Kopell et al. \(2000\)](#). Due to its origin as an average over many neurons, the LFP should reflect a highly complex signal. Indeed, features of the LFP oscillation are informative not only of external stimuli or behavioral aspects ([Laurent, 2002](#); [Mehring et al., 2003](#); [O’Leary and Hatsopoulos, 2006](#); [Mazzoni et al., 2008](#)), but also correlate with internal processes, such as movement planning ([Murthy and Fetz, 1996a](#); [Donoghue et al., 1998](#); [Roux et al., 2006](#)) or attentional modulation ([Fries et al., 2001](#); [Lakatos et al., 2008](#); [Rotermund et al., 2009](#)).

Functionally, it was hypothesized early on that the oscillatory nature of the LFP is likely to result from non-cancelling potential contributions of a subpopulation of synchronized neurons, which is coarsely estimated to be as low as 10% of the whole network ([Elul, 1971](#)). Evidence from sensorimotor areas shows that the intrinsic rhythmicity of neurons is not necessary for

the generation of oscillations (Witham and Baker, 2007), corroborating the view that oscillations emerge as a network phenomenon. While in general the LFP is only weakly - if at all - predictive of spike times (Donoghue et al., 1998; Rasch et al., 2008), it has been established that single spikes can become transiently entrained to the LFP in a rhythmic or non-oscillatory fashion (Eckhorn and Obermueller, 1993; Murthy and Fetz, 1996a). In the case of γ oscillations, a growing body of evidence supports the view that they involve the entrainment of excitatory cells to an underlying rhythm created by a synchronized inhibitory cell network (Fries et al., 2007; Berens et al., 2008). Indeed, simple models of an interacting network of excitatory and inhibitory neurons are able to generate rhythms compatible with the γ -range (Mazzoni et al., 2008). The proposal of Fries et al. (2007) that the oscillations act as a mechanism to transform stimulus intensity to a phase-of-firing code is currently under debate (Ray et al., 2008). Another hypothesis (Womelsdorf et al., 2007) states that rhythms act as a gate to channel the information carried by excitatory neurons (see also Sec. 1.1.3). The degree of phase locking between neurons and the LFP depends in general on the strength of LFP oscillations (see Murthy and Fetz, 1996b and Sec. 3.1), and the phase relationships are behaviorally and functionally informative (Harris et al., 2002; Friedrich et al., 2004; Montemurro et al., 2008). The population rate of groups of neurons tends to show correlations with the LFP (Mukamel et al., 2005) most likely as a result of the underlying network dynamics (Viswanathan and Freeman, 2007). Furthermore, in LFP spindles (brief epochs of strong oscillations, about 5 – 10 cycles) spiking activity exhibits oscillatory features (Murthy and Fetz, 1996a). Combining the latter two observations, Nir et al. (2007) show that the firing rate profiles correlate with γ -band LFP power only if the level of interneuronal rate correlation (correlations on slow time scales) is high. On a finer temporal scale, distinct spike patterns across neurons and their phase relationship to LFP oscillations encode a substantial amount of surplus of information about the stimulus compared to information contained in the firing rate alone (Kayser et al., 2009). A recent review summarizing the main interpretations of oscillatory LFP is given by Singer et al. (2009).

Despite these findings, it remains unclear what mechanisms are responsible for the typically weak correlations between the two biophysically distinguished (Sec. 1.1.1) levels of observation, single spikes and the LFP, in the

absence of global synchronized network state. Contrary to intuition, correlations in membrane potentials do not in general predict spike synchrony, in particular in the awake, active state of cortical networks. In recordings in the anesthetized animal (Lampl et al., 1999) that exhibit strong correlations between the membrane potentials of neurons, the action potentials evoked by an artificial excitatory drive become synchronized on a fine time scale. However, in the absence of strong current injections the synchronizing effect of subthreshold oscillations diminishes. A recent study by Poulet and Petersen (2008) demonstrates that although LFPs are highly correlated with the synchronized membrane potentials of nearby neurons, synchronous action potentials occur largely independently of the field potential. Supporting this view, theoretical studies agree that the correlation transfer from synaptic input to spikes is weak within neurons (Tetzlaff et al., 2008). Therefore, a key question is to identify how synchronized spiking activity on a fine temporal scale (millisecond scale) is related to the LFP.

1.1.3 LFPs in the Framework of the Assembly Hypothesis of Neuronal Coding

An intriguing hypothesis of information coding in the brain originates from a proposal by Hebb (1949) that has since then been steadily extended. According to this theory, the identities of groups of synaptically coupled and synchronously activated neurons, termed assemblies, are the primary carriers of information. In contrast to models of rate coding (Hubel and Wiesel, 1959; Georgopoulos et al., 1986), the precise timing of spikes is attributed a decisive role. Theoretically, the concept offers elegant solutions for a number of conceptual problems associated with rate-based approaches, such as feature-binding of related and complementary percepts (von der Malsburg and Schneider, 1986; Singer, 1999; Yu et al., 2008). The concept of the synfire chain (Abeles, 1991) has provided the framework to study the properties of stable propagation of information via the successive activation of groups of synchronized neurons (Diesmann et al., 1999; see, e.g., Abeles et al., 2004; Hayon et al., 2005 for recent studies that probe computational capabilities of synfire models). Experimentally, despite recent technological advances on the experimental (Nicolelis et al., 1997; Csicsvari et al., 2003; Euston et al., 2007; Fujisawa et al., 2008) and technical levels (Brown et al., 2004; Gerstein, 2004; Schrader et al., 2008), the undersampling of the neuronal

population makes it difficult to detect assemblies from the spiking activity. Nevertheless, a hallmark signature of the assembly is the behavior-related, repeated coactivation of two or more neurons belonging to the same assembly (Gerstein et al., 1989). The Unitary Events method (for a summary, see Grün, 2009) detects such imprints of assembly activations by comparing the number of coincidences between two neurons to the prediction based on the firing rates. A number of studies have found evidence for precise spike synchrony in a behavioral context (Vaadia et al., 1995; Riehle et al., 1997; Kohn and Smith, 2005; Fujisawa et al., 2008; Maldonado et al., 2008), some interpreting observed spatio-temporal patterns of activation as indications of synfire propagation (e.g., Ikegaya et al., 2004), while some studies questioned the functional relevance of observed correlations (Baker and Lemon, 2000).

Early on, it was suggested that the correlations between spiking signals and the fast β - or γ -band LFP rhythms exhibited by roughly one third to one half of the recorded population may reflect those neurons that participate in the assembly dynamics (Eckhorn et al., 1988; Murthy and Fetz, 1996b; Donoghue et al., 1998). In this interpretation, it is not necessarily a single population of neurons that is related to the LFP rhythm, but neurons may engage transiently in different active assemblies. This idea of 'multiplexing' assemblies on a grid given by the oscillation is illustrated in Singer (1999). A recent study has given some preliminary evidence for such a picture: Womelsdorf et al. (2007) showed that interactions between multi-unit activities, defined by the magnitude of their power correlation, are most prominent in those trials where their oscillatory activity becomes locked to a specific phase of the LFP. It is commonly assumed that pyramidal neurons, whose synaptic inputs mainly contribute to the LFP (cf., Sec. 1.1.1), constitute the assemblies. Therefore, one reasonable interpretation to reconcile this picture with the findings of the previous section is that fast oscillations, generated by a recurrent inhibitory network (cf., Sec. 1.1.2), generate a time grid by which the assembly activations are structured (Buzsáki and Draguhn, 2004). Notably, the hypothesis of assemblies being reflected in the LFP rhythms is in good agreement with recent findings of Poulet and Petersen (2008): strong and specific imprints of synaptic activity are observed in membrane potentials (which themselves correlate strongly with the LFP) preceding synchronized spikes, such as required for the propagation of synchronous activity in the assembly framework. However, there exists no direct proof for

any hypothesized relation between synchrony on the spike-level and the LFP in the active cortical network.

1.2 Phase Locking between Spiking Activity and LFP

1.2.1 Introducing the Phase Analysis

Central to the analysis performed in each of the studies presented in this thesis is the evaluation of occurrences of spikes with respect to the oscillation cycle, in other words their locking. In this section we present a short overview of the analysis method used throughout the remainder of this thesis. A classical, heuristic approach to relate spikes to the LFP is the spike-triggered average (STA) of the LFP, which is obtained by calculating the cross-correlation function $c(\tau) = \langle S(t) L(t - \tau) \rangle_t$ of the binned spike train $S(t)$ and the LFP signal $L(t)$. The result corresponds to the average LFP centered at the time points of spike occurrences. Intrinsic rhythms in the signals and correlations between the spike train and the LFP induce an oscillatory component in the resulting STA. However, phase lag, amplitude, and frequency composition of the LFP all influence the resulting shape of the STA. Thus, the STA (or its related Fourier transform, termed spike-field coherence; see [Rosenberg et al., 1989](#); [Fries et al., 2002](#)) cannot quantitatively distinguish to what extent its features result from contributions of spikes that are well entrained to the rhythm of the LFP or those that occur at a higher amplitude (i.e., envelope) of the oscillation.

A more direct way to assess the degree of synchrony between two signals is to analyze the phase synchronization between them directly ([Rosenblum et al., 1996](#)). In systems of interacting oscillators, this measure is informative of the true coupling of individual oscillators, as opposed to their mere entrainment to a common rhythm ([Pikovski et al., 2001](#)). The underlying idea is to extract the instantaneous phase $\phi_i(t)$ and $\phi_j(t)$ from two continuous signals as the relevant observables. [Quyen et al. \(2001\)](#) shows that comparable results are obtained by two commonly used methods to compute these phases, based on the complex-valued analytic signal (often inaccurately termed Hilbert transform) and the convolution of the signal by a complex wavelet function, respectively. In contrast to the STA, both methods are

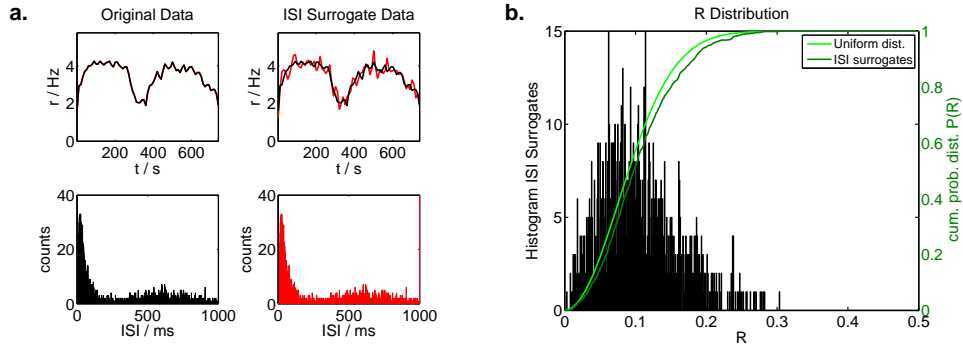


Figure 1.1: Generation of surrogate data for the phase synchronization analysis. **a.** Original data data (black) compared to one surrogate data trace (red, right). The time-dependent rate profile of the original data is approximated by the local shuffling procedure (top graphs), the interspike interval distributions (bottom graphs) are exactly retained. **b.** Distribution of R (histogram) from one thousand surrogates constructed as in panel a. Due to the clear structure in the first-order interval statistics, the cumulative distribution of the surrogate data (dark green curve) differs from the Poisson assumption (i.e., Rayleigh criterion) for R (light green curve). All data and analysis parameters taken from the study presented in [Sec. 2.2](#).

independent of the instantaneous amplitudes (envelopes) of the underlying signals. In fact, the time-dependent amplitude may be conveniently obtained as a completely decoupled measurement in a similar fashion. As time series, such as the LFP, typically contain a superposition of several frequency components, it becomes necessary to pre-filter the signal with a band pass around the frequency range of interest in order to interpret the phase in a meaningful manner ([Boashash, 1992](#)). As a further consequence, transient non-oscillatory periods where the instantaneous phase is undefined should be excluded from the analysis. In a subsequent step, the constancy of the absolute phase difference $|\phi_i(t) - \phi_j(t)|$ is evaluated over the length of the time series ([Quiroga et al., 2002](#)) as a measure of the coupling between the two signals. Phase synchronization analysis has become a standard tool in the analysis of continuous neuronal time series (e.g., [Lachaux et al., 1999](#); [Rodriguez et al., 1999](#); for reviews, see [Varela et al., 2001](#); [Quyén and Bragin, 2007](#)).

In contrast to the original definition of phase synchronization, spike data is only characterized by single points in time. Besides the possibility of converting the spike data to a time-dependent continuous rate function before performing the analysis described above ([Hurtado et al., 2004](#)), one alternative is to analyze the phase of a continuous signal at the points in time

indicated by the discrete process (Schäfer et al., 1999). This approach has been successfully used in a number of areas where a clear phase relationship between spiking activity and the continuous signal is observed, most prominently in studies of phase-precession in hippocampus (e.g., Harris et al., 2002) or auditory processing (e.g., Liu et al., 2006). In the work contained in this thesis we extended the use of phase synchronization analysis to the irregular spiking activity typically encountered in cortical spike data, an approach that is becoming increasingly popular (e.g., Kayser et al., 2009).

To this end, we first compute the instantaneous phases $\phi(t_i)$ of the LFP at the time of the spike occurrences t_i , $i = 1, \dots, N$. The modulation of the resulting phase distribution, also termed locking strength or phase locking value (Lachaux et al., 1999), can be determined by the magnitude R of the vector sum $R \exp(2\pi i \mu_0) = \sum_{i=1}^N \exp(2\pi i \phi(t_i))$. For better interpretation, this value can be transformed to the circular standard deviation $\sigma = \sqrt{-2 \ln R}$, which corresponds to the familiar standard deviation σ of a Gaussian distribution wrapped around the unit circle in the limit of small R (Mardia and Jupp, 2000). The value of μ_0 reveals the mean phase of the circular distribution. To test whether a given phase distribution indicates spike-LFP locking that exceeds chance level, we test R against the null hypothesis of a uniform distribution (Rayleigh test). This test assumes a stationary oscillation and a Poisson spiking statistics.

Clearly, a regular spiking process would intrinsically show locking to a well defined oscillation of comparable frequency. Therefore, the Rayleigh test is not conclusive of whether the generators of the two signals are indeed coupled, i.e., phase-locked better than expected assuming their independence. To account for spiking statistics that is more regular than Poisson, as found in many neuronal systems (Shinomoto et al., 2009), we resort to a surrogate technique. Random spike trains are generated by shuffling the original order of interspike intervals, thus retaining to first order the regularity inherent in the spike train (see Fig. 1.1 for an illustration). From the distribution of R obtained from these surrogates we determine whether the precise timing of measured spikes leads to a stronger phase-locking than expected. Note that both tests evaluate the strength of entrainment only, independent of the preferred phase μ_0 . A concise treatment of circular statistics, further testing methods, and the use of boot-strap techniques for circular data, is given by Mardia and Jupp (2000) and Stark and Abeles (2005).

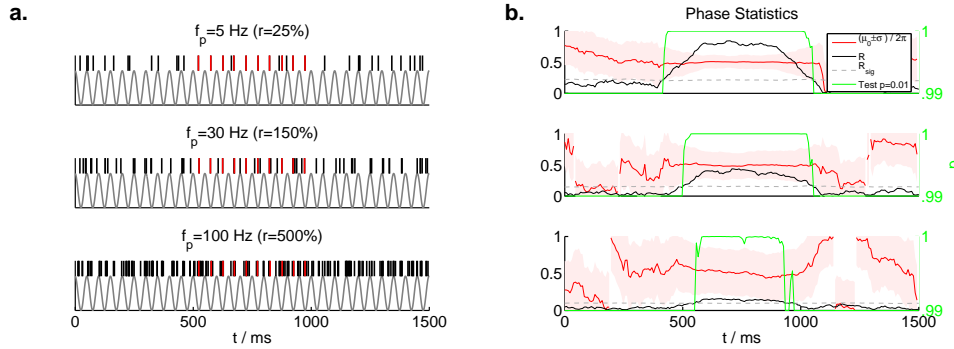


Figure 1.2: Reliability in detecting locked spikes embedded in an otherwise unlocked background process. **a.** Example spike train (ticks) and LFP (simulated as a sinusoid) of one trial generated for 3 different choices of the rate f_p of the unlocked Poisson spike train in the central locking region ($t = 500 - 1000$ ms). The rate of the locked spike train (red) is kept fixed at $f = 20$ Hz (one spike per period). The total spike train outside the center region has a fixed rate of $f + f_p$. The percentage r is defined by the ratio f_p/f . **b.** Results of the phase locking analysis for the three choices of the parameter r (calculated using $n = 20$ trials per parameter; each trial is similar to the respective trial in panel a). Analysis is carried out across trials and in sliding windows of length $T = 200$ ms (4 periods of the LFP oscillation). The black line indicates the raw value of the vector strength R . The red curve shows the mean phase μ_0 ; the circular standard deviation σ is visualized by the red shaded area. The green curve indicates the p -value of the Rayleigh test, plotted for the last percentile (right axis). For $p > 0.99$, the window is significantly non-uniform at the 1% level.

1.2.2 Sensitivity and Reliability of the Phase Analysis

In the following, we illustrate the sensitivity and reliability of the measures introduced in Sec. 1.2.1 in detecting subtle, transient changes in the locking characteristics of neurons. This section highlights findings from a larger body of work presented in Denker et al. (2005). As a basis of each analysis that follows, for given parameters we construct $n = 20$ simulated trials of spike trains and LFPs (examples of a single trial for different parameters are shown in Fig. 1.2a and Fig. 1.3a). Each trial consists of a phase-locked region in the middle part ($t = 500 - 1000$ ms), where one spike per LFP cycle (colored red) is randomly placed within a locking region of width $0 < \Delta\phi < 2\pi$. These locked spikes are framed by Poisson spike trains ($t = 0 - 500$ ms and $t = 1000 - 1500$ ms) such that the rate of the entire process is constant. The LFP is modeled as a pure sine function with a fixed frequency. Since we choose to model the unlocked spikes by a Poisson process, we are able to use the simple Rayleigh test to detect the presence of locked spikes, instead of the more involved surrogate technique. We first test whether the procedure is able to detect a central region of perfect locking (i.e., $\Delta\phi = 0$) in the

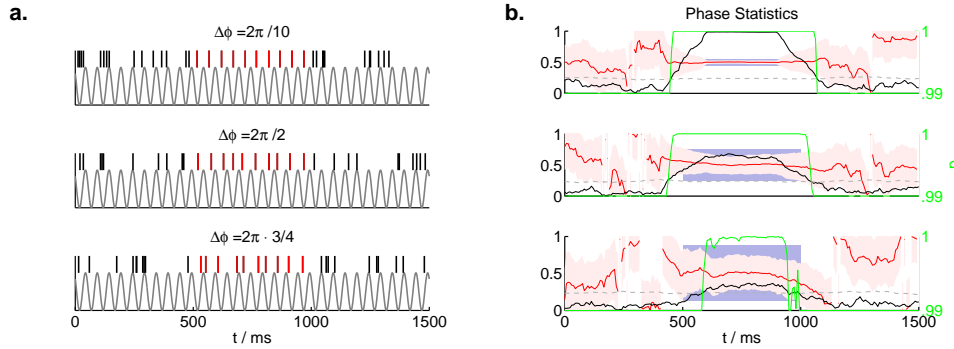


Figure 1.3: Reliability in detecting weakly locked spikes. **a.** Compared to Fig. 1.2, the detectability of phase-locked spikes is probed for three different choices the phase jitter $\Delta\phi$ and in the absence of non-locked spikes ($f_p = 0$). Locked spikes (red) are placed randomly in an interval $[\pi - \Delta\phi/2; \pi + \Delta\phi/2]$ (indicated in panel b). The total rate of the spike process is fixed to $f = 20$ Hz. **b.** Results of locking analysis for the three choices of the parameter $\Delta\phi$. Curves correspond to their equivalents in Fig. 1.2: value of vector strength R (black), circular mean μ_0 (red; shading indicates the circular standard deviation σ) and p -value of the Rayleigh test (green). The dark shading indicates the original locking interval defined by $\Delta\phi$. As before, analysis is carried out across trials and in sliding windows of length $T = 200$ ms.

presence of additional non-locked spikes characterized by the rate f_p of their generating Poisson process (Fig. 1.2). Second, we test whether locking is reliably detected in the absence of unlocked spikes (i.e., $f_p = 0$), but with a non-zero phase jitter $\Delta\phi$ of the locked spikes (Fig. 1.3).

We observe that for variation of the respective parameters f_p and $\Delta\phi$, the method is reliably able to detect the central region of locked spikes (indicated by the green curves in Fig. 1.2b and Fig. 1.3b). Furthermore, probing the sensitivity reveals that this holds even for high values of these parameters. Denker et al. (2005) show that the estimation remains surprisingly good given the simultaneous variation of both parameters. Therefore, given reasonable model parameters that one might encounter in electrophysiological recordings, the phase synchronization method sensitively and reliably detects transient periods where spike trains become significantly coupled to the underlying oscillation.

1.3 Overview of the Thesis

Equipped with a review of relevant state-of-the-art hypotheses on the generation and interpretation of the LFP signal (Sec. 1.1) and the detailed explanation of the common analysis method used to analyze spike-LFP re-

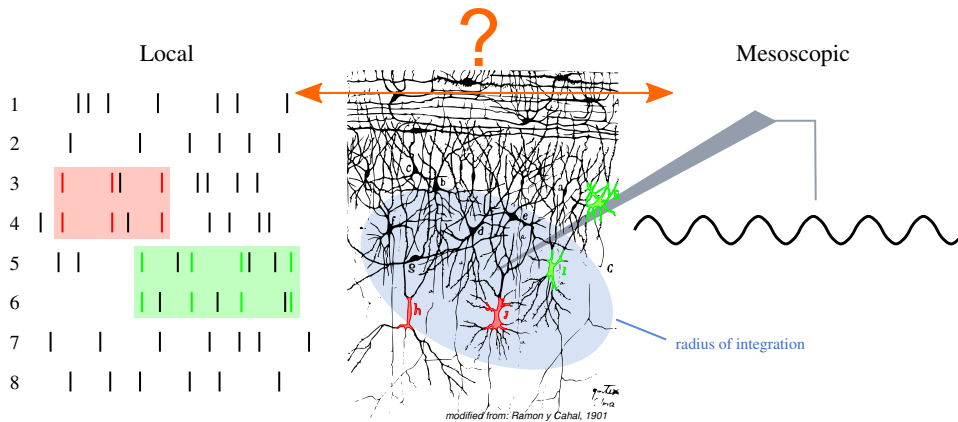


Figure 1.4: Schematic representation of the central idea linking the studies presented in this thesis. Neuronal activity is recorded on two spatial scales: spiking activity in parallel recordings of neurons (local, left) and LFPs integrating over the potentials arising from transmembrane currents in a large area surrounding the electrode (mesoscopic, right). In the absence of a global synchronized state governing the spiking activity, only a subset of synchronized neurons is expected to couple to (or generate) the observed LFP oscillations. The composition of this subset may change over time. Shown schematically are spikes of two groups of neurons within the network (red and green), which become transiently (shaded areas) synchronized (colored spike pairs represented by ticks). The grouping may be induced by different mechanisms, such as: neurons are part of an ensemble coding for a specific behavioral correlate or stimulus (Sec. 2.1), neurons belong to a neuronal subtype specifically involved in the propagation of the oscillation (Sec. 2.2), or neurons are part of an internal network process that involves their entrainment to the LFP (Chap. 3). Background staining modified from Ramon y Cahal, 1901.

relationships (Sec. 1.2), we now turn to introduce the individual studies that are combined in this thesis. The central theme linking all the reports is the search for neuronal populations that are reflected in the oscillations exhibited by the LFP (compare sketch in Fig. 1.4). The different studies are grouped into two chapters: Chap. 2 highlights information gained from the analysis of the coupling of single spikes to the population signal, while Chap. 3 investigates how coincident spiking activity is related to the LFP.

In Chap. 2, we investigate the relationship of single spikes to the LFP in order to identify neuronal subpopulations reflected in the LFP based on behavioral correlates or the physiological characterization of neurons. The study in Sec. 2.1 analyzes the plastic changes induced by olfactory conditioning in the output region of the antennal lobe of the honeybee. The correlations between learning-induced changes in the LFP power spectrum and changes in the synchronization of neurons to the LFP leads to the conclusion that the LFP likely reflects the temporal representation of the odor stimulus by means of a synchronized ensemble. The second study presented in Sec. 2.2

shows that in recordings from the basal ganglia of the anesthetized rat, the strong interneuronal correlations of a specific subtype of striatal neurons (the fast-spiking interneurons) are related to a prominent synchronization of this neuron type to fast population oscillations. Thus, in this case we identify a specific neuron type as the group of neurons coupled to the LFP oscillations.

In [Chap. 3](#) we present the main finding of this thesis in a series of three reports that investigate the link between synchronous activity in spike data to oscillations in the LFP. The experimental data used throughout this chapter is recorded from the primary motor cortex of monkeys in a time-delay pointing task. In [Sec. 3.1](#) we present an analysis that identifies spikes that occur during strong oscillations of the LFP as those that are most strongly phase-locked to the LFP cycle. In [Sec. 3.2](#) we provide the central experimental study that explains how neuronal cell assemblies are reflected in the oscillatory structure of the LFP (cf., [Sec. 1.1.3](#)). In a theoretical companion study given in [Sec. 3.3](#), we not only provide increased detail on how the experimental findings are to be reconciled with the theoretical framework of assembly coding, but show how the combination of measurements of synchrony on the spike and population levels provide access to a network-wide parameter that describes the extent to which assemblies are represented in the spiking activity of the brain.

Finally, in [Chap. 4](#) we provide a summary that recapitulates, links, and interprets the aspects revealed by the individual reports. Furthermore, we discuss issues that are touched by all presented studies, and provide an outlook to questions that follow from the results presented in [Chap. 3](#).

Chapter 2

The LFP as a Monitor of Neuronal Populations

2.1 Neural Correlates of Odor Learning in the Honeybee Antennal Lobe

The following manuscript is submitted to the European Journal of Neuroscience as:

M. Denker, R. Finke, F. Schaupp, S. Grün, and R. Menzel. Neural Correlates of Odor Learning in the Honeybee Antennal Lobe. Publisher homepage: <http://www.blackwell-synergy.com>

Neural Correlates of Odor Learning in the Honeybee Antennal Lobe

Michael Denker^{1,*}, Robert Finke^{2,*}, Frank Schaupp²,
Sonja Grün^{1,3}, Randolph Menzel^{2,3}

June 5, 2009

¹ RIKEN Brain Science Institute, Wako-Shi, 351-0198 Saitama, Japan

² Freie Universität Berlin, Institut für Biologie - Neurobiologie, 14195 Berlin,
Germany

³ Bernstein Center for Computational Neuroscience, Berlin, Germany

*The first two authors contributed equally to this work.

Abstract

Extracellular spiking activity and local field potentials (LFP) were recorded via tetrodes at the output of the antennal lobe in the honeybee brain during olfactory conditioning. Odors induce reliable rate responses which consist of either odor-unspecific phasic-tonic responses, or complex responses with odor-specific temporal profiles. In addition, odors evoke consistent responses of LFP oscillations in the 50 Hz band during the phasic ON response to odor stimulation, and variable LFP responses at other frequency bands during the sustained response. A principle component analysis of the ensemble activity during differential conditioning consistently indicates the largest changes in response to the learned odor (CS+). Relative LFP power increases for CS+ in the 15-40 Hz frequency band during the sustained response, and decreases for frequencies above 45 Hz. To quantify the relationship between these population responses given by the ensemble spiking activity and LFP, we show that for CS+ the learning-related changes in the degree of the phase-locked spiking activity correlate with the power changes in the corresponding frequency bands. Our results indicate associative plasticity in the antennal lobe of the bee leading to both enhancement and decrease of neuronal response rates. LFP power changes and the related changes in the correlation between spike and LFP serve as further evidence for a learning-induced restructuring of temporal ensemble representations. The observed changes appear to decorrelate

the responses to the learned and non-learned odors and may provide a learning-related neural signature to the learned odor.

Introduction

Learning leads to changes of neural connectivity, giving the network the power to store previous experience and to retrieve it later for behavioral control (e.g. hippocampus: Sutherland & McNaughton, 2000; prefrontal cortex: Goldman-Rakic, 1995; orbito-frontal cortex: Rolls *et al.*, 1996; reward system in the ventral tegmentum: Schultz, 1998). The processes involved in learning-related adaptive changes have been intensively studied in the olfactory system of both vertebrates (e.g. Wilson *et al.*, 1987; Keverne, 1995; Ravel *et al.*, 2003) and invertebrates - in particular, of insects (Faber *et al.*, 1999; Daly *et al.*, 2004; Yu *et al.*, 2004). Odors are represented at the level of the first stage of neural integration (the antennal lobe, AL, in insects) by spatially distributed activity patterns (Joerges *et al.*, 1997; Wang *et al.*, 2003; Hallem & Carlson, 2006) and the timing of spike activity both at the level of the single neuron and in neuronal populations (for a review, see Laurent, 2002). Most second-order neurons (projection neurons, PNs, in the insect AL) connecting the first-order neuropil with downstream integration areas receive input from a single glomerulus. The across-fiber activity patterns appear to reflect the spatial glomerular coding together with the result of lateral neural interactions via local inhibitory and excitatory interneurons. These local inhibitory networks shape the chemical response profiles of these PNs (Sachse & Galizia, 2002; Krofczik *et al.*, 2009) and induce temporal structures in the odor responses that code odors also in the time domain (Laurent *et al.*, 2001). Stopfer *et al.* (1997) observed that pharmacological interference with the GABA inhibitory circuit in the honeybee antennal lobe changed the oscillatory field potential and reduced learned odor discrimination.

Learning was also found to alter the spatial coding of odors in the antennal lobe of *Drosophila* (Yu *et al.*, 2004) and the honeybee (Faber *et al.*, 1999; Sandoz *et al.*, 2003), and thus one would expect that the activity patterns across different fibers in the PNs might be changed by learning. In rats, odor learning leads to a shift of the dominant frequency from the gamma band (60-90 Hz) to the beta band (15-40 Hz) in local field potentials of the olfactory bulb (Ravel *et al.*, 2003 ; Martin *et al.*, 2006), indicating changes in the odor-induced timing of spike activity at the population level. In the terrestrial mollusc *Limax* olfactory conditioning also changes the oscillatory activity in the primary olfactory neuropil (Gelperin, 1990; Tank *et al.*, 1994).

Here we ask whether activity patterns of PNs in the honeybee brain change

in the course of associative learning. We monitor neural activity by recordings of multi-unit spiking activities composed of single to few units, as well as by simultaneous recordings of local field potentials (LFP). We find both enhancement and decrease of spike rates to the odors learned, which result in a specific restructuring of the ensemble network response to conditioned odors. The power of LFP induced by the odor stimuli also changes in a characteristic manner indicating that the correlation between spike and LFP may serve as a signature for a learning-induced restructuring of the temporal ensemble code.

Materials and Methods

Electrophysiology

Data were obtained from 21 foraging honeybees (*Apis mellifera carnica*) caught at the hive entrance, immobilized by cooling, and mounted in a metal tube. Bees were fed with sucrose solution and kept overnight in the dark at 20° C. The next day the head was fixed with dental wax and dissected by removing a piece of cuticle between the antennae, the ocelli and the compound eyes. A hole for the reference electrode was cut into one compound eye. Small droplets of bee saline (in mmol/l: 130 NaCl, 6 KCl, 2 MgCl₂, 10 HEPES, 17 glucose, 6 fructose, 160 sucrose, pH 6.7) were applied to prevent the brain from drying out (Mauelshagen, 1993).

Two shanks of tetrodes of a silicon probe (supplied by the Center for Neural Communication Technology, University of Michigan, Daryl R. Kipke) were inserted at the inner border of the AL into a depth of 60 – 90 μm , thus close to the median and lateral antenno-calycal tract (m and lACT) exiting the AL here. The distance between the shanks was 60 μm , and they were inserted such that the two ACTs were expected to run between the two shafts. The input resistances of the electrodes varied between 2 and 4 M Ω . Signals from tetrodes were amplified by two eight-channel amplifiers (Lynx-8, Neuralynx, Tucson, AZ). An A/D converter (DT3010 Data Translation A/D board, Data Translation Inc., Marlboro, MA) digitalized the extracellular signals with a sampling rate of 32 kHz. Gain and filter settings were adjusted via software (Cheetah Data Acquisition Software, Neuralynx, Tucson, AZ). All channels were filtered between 0.3 and 6 kHz for recording the spike signals. Two probe channels were additionally filtered between 10 and 125 Hz for recording the local field potential (LFP). Line hum was removed from the LFP signal with a Hum-Bug filter (50 Hz, Quest Scientific, Vancouver, BC, Canada).

A thin silver wire was used as an electromyographic electrode for recording the electrical activity of the M17 muscle. Activity of the M17 indicates the

extension of the proboscis as a monitor for the unconditioned and conditioned response (a threshold of 5 M17 muscle action potentials was applied to detect a response).

Differential conditioning

In differential conditioning animals learn to discriminate between a rewarded (CS+) and an unrewarded (CS-) stimulus (Bitterman *et al.*, 1983). The sequence of trials is shown in Fig. 1a. In a pre-test phase, three different odors were presented within three trials each with duration of 1 s and an inter-stimulus interval of 60 s. Odor A then became the CS+ in the next phase, odor B became the CS-. Odor C, the control odor (Ctrl), was not presented during differential conditioning, and acted as a test for generalization in the third phase, the post-test. In the differential conditioning phase the CS+ was paired with sucrose solution (2 M), which was applied with a toothpick to both antennae and then to the proboscis. The CS+ and CS- conditioning trials followed each other in intervals of 1 minute (massed conditioning). Five minutes after the end of the differential conditioning trials, the post-test phase followed, in which all three odors were presented in the same way as in the pre-test (extinction trials).

An olfactometer was used for odor stimulation as described in Galizia *et al.* (1997). A custom-made software program (Visual Basic) controlled the stimulation conditions. Three different odors out of ten (1-octanol, 1-nonanol, 1-heptanol, 2-hexanone, 2-heptanone, hexanal, isoamylacetat, methylsalicylate, cineol, and spearmint) were used for an individual animal. Across different animals the odors were balanced with respect to their function as CS+, CS- and Ctrl. The odor source consisted in five microliters of pure substance on a filter paper inserted into each container of the olfactometer. These stimulus conditions were found to give the best conditioning results (Pelz *et al.*, 1997). Since no difference was found between the 10 odors with respect to the learning related changes, the responses to the different odors will be pooled and analyzed according to their function during conditioning (CS+, CS-, Ctrl).

Spike extraction and LFP selection

For spike detection of the 8 recording sites those were selected which gave the best signal to noise ratio, and which did not record from the same unit(s). Two forms of spike signals were observed. In the first case the signal-to-noise ratio was high, and the spikes were of the same waveform and amplitude (Fig. 1b). This type of spike signal was observed in 18% of all recorded data. In the other 82%, the signal-to-noise ratio was lower and spikes occurred with different amplitudes (Fig. 1c). In such cases it was not possible to separate spike forms

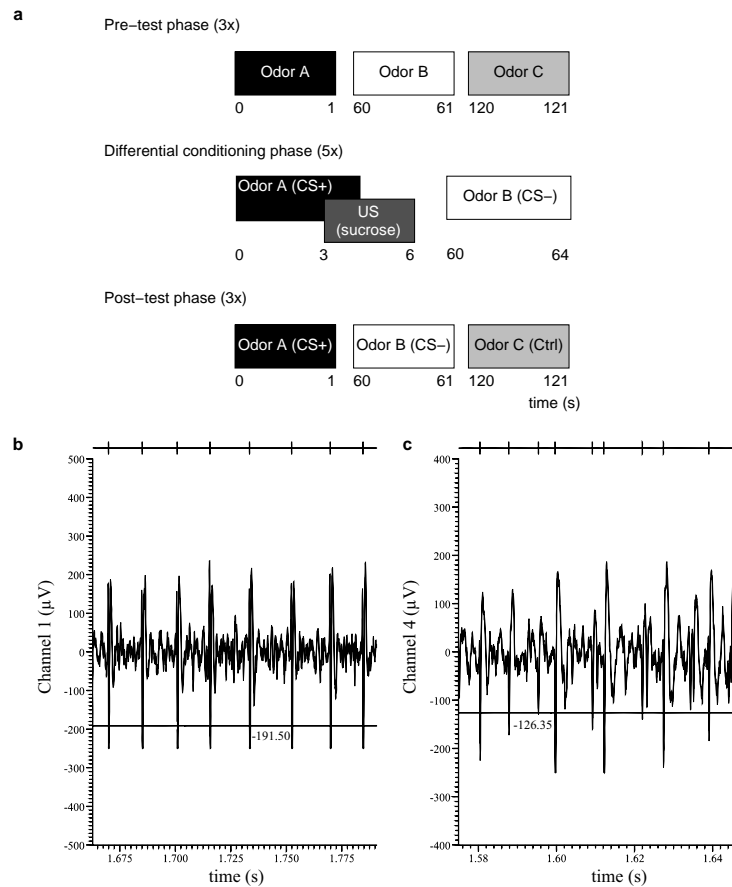


Figure 1: Design of the differential conditioning experiment and spike detection. (a) The odors A, B, and C functioning as CS+, CS-, and Ctrl were selected from 10 different odors. The experiment is composed of three phases: pre-test, differential conditioning, and post-test. Boxes indicate the presentation of odors, attached numbers illustrate the relative duration of application in seconds. The box labeled 'US' marks the time of the sucrose reward. Odors of each phase were applied at intervals of one minute, and each set of three odors was repeated 3 times (5 times for the differential conditioning phase). The interval between the three phases was 5 minutes. (b) Example of a single-unit recording during odor stimulation. The horizontal line at $-191.50 \mu\text{V}$ marks the chosen spike threshold, and extracted spikes are marked as ticks (top line). (c) Same as (b), but for a different recording with a lower signal-to-noise ratio. Here, the choice of the threshold at $-126.35 \mu\text{V}$ yielded a spike train from presumably two contributing neurons.

unambiguously due to the noise of our recordings, and a threshold level was chosen that may have included the spikes of more than one form. Judging by eye we consider that only in very rare cases three forms of spikes might have pooled together. We shall refer to “unit” for those spikes pooled together throughout the text recognizing that such “units” may well receive spikes from more than one neuron, possibly two (and in rare cases even three) very closely attached neurons. We included into our analysis only such recordings which were stable with respect to the spike forms over the whole time of the experiment (>1 h). We also compared the stability of the recordings between clear single neuron recordings and other recordings, and found no difference in stability of the recordings. On average, we simultaneously recorded four units from each animal, and the total number of recorded units was 99.

Both spike signals and LFP signals were recorded during odor presentation as well as for an interval of one second before and one second after each odor presentation. LFP time courses obtained from the two recording channels were typically highly correlated, and therefore we performed all analysis involving LFP on the channel with the highest total power during spontaneous activity of the pre-test phase averaged across all three presented odors. The data were visualized in Spike2 (Cambridge Electronic Design Ltd., Cambridge, UK). Further data analysis was performed using the MATLAB software package (The MathWorks, Inc., Natick, MA, USA).

Response properties

During the pre-test phase most of the units (87 out of 99) showed a characteristic rate response to repeated odor stimuli to at least one of the three odors presented. Here, an excitatory response was defined as a rate increase exceeding two standard deviations above the spontaneous level in at least 2 of the 3 trials; analogously an inhibitory response corresponded to a rate decrease falling below one standard deviation from the spontaneous mean. Across all animals, 51% of units responded to all three odors, whereas 20% responded selectively to only two and 17% to only one of the presented odors. The majority of detected responses to odors (60%) responded with phasic ON-excitation, whereas only 10% responded with phasic inhibition, and 30% did not show an ON response to any odor. Here and in the following, ON is defined as the 500 ms time window starting at odor stimulation onset. The sustained responses during the remaining 500 ms of the odor presentation were more variable (31% excitatory, 16% inhibitory, 53% no response), and complex time courses of excitation and inhibition were seen frequently. Most units (36%) gave excitatory ON bursts (around 100 ms after stimulus onset) often in combination with short non-responsive

to inhibitory sustained responses or OFF-responses (500 ms time window following the end of the odor presentation). Individual units exhibited diverse response patterns: Though typical ON responses were excitatory, a small percentage of units (5%) showed inhibitory responses to one or two of the odors, and excitation to the remaining. The average sustained response rates for the different odors were 1-heptanol: 47 Hz, 1-octanol: 62 Hz, 1-nonanol: 73 Hz, 2-hexanone: 42 Hz, 2-heptanone: 48 Hz, hexanal: 47 Hz, isoamylacetat: 70 Hz, spearmint: 55 Hz, cineol: 50 Hz, and methylsalicilate: 68 Hz. The average latencies (time between stimulus onset and begin of excitatory response) varied between 98.2 ± 78.0 ms (isoamylacetat) and 163.8 ± 54.2 ms (cineol). All of these response patterns were required to be stable during the pre-test phase and showed no stimulus repetition effect.

Spectral analysis

Power spectra were computed using a 256-point Fourier transform tapered by a Hanning-window to reduce spectral leakage. Time-resolved power spectra were calculated in sliding windows of 300 ms in steps of 30 ms. Baseline power was evaluated during 700 ms before odor presentation.

Phase analysis

To investigate a possible relationship between spike rates and LFP, we analyzed the phase relation between the LFP oscillations and individual spikes using methods originating in phase synchronization analysis (Denker *et al.* 2007; Harris *et al.* 2002; Hurtado *et al.* 2004; Sharott *et al.*, 2009; Varela *et al.* 2001). In a first step, the LFP was filtered to a specified frequency band (compare also Fig. 2b). The cutoff frequencies $[f/1.1; f/0.9]$ of the filter (8th order Butterworth) were chosen such that they contained frequencies corresponding to ± 10 percent of the period T of the center frequency $f = T^{-1}$. The determination of the oscillation phase $\phi(t)$ of the filtered LFP signal was based on the analytic signal obtained by a Hilbert transformation $\tilde{x}(t) = \frac{1}{\pi} \text{P.V.} \int \frac{x(\tau)}{t-\tau} d\tau$ of the original signal $x(t)$, where P.V. denotes that the integral is to be taken as the Cauchy principal value (Le Van Quyen *et al.*, 2001; Rosenblum *et al.*, 1996). In a next step, we analyzed the distribution of instantaneous phases $\phi(t_i)$ of the LFP obtained at spike time occurrences. In particular, we obtained the mean phase ϕ of the circular average $R e^{i\phi} = N^{-1} \sum_{i=1}^N e^{i\phi(t_i)}$, where $\phi(t_i)$ indicates the phase of the field potential at time t_i of spike i .

A natural way to quantify whether spikes show a pronounced phase relationship to the LFP is to test for non-uniformity of the resulting distribution

of phases (e.g., Rayleigh test for uniform phase distribution with correction for small sample size; cf. Mardia and Jupp, 2000). However, the inter-spike interval distributions suggest that spiking activity has a tendency for regularity, in particular during strong rate responses of the initial phase of the odor presentation. Therefore, we might expect a non-uniform phase distribution as a mere consequence of the underlying regularity. To investigate locking that is not explained by such underlying regularity we employed surrogate spike trains to quantify the intrinsic degree of locking between spikes and LFP. We constructed these surrogates by shuffling the inter-spike intervals that occurred during each of the 500 ms time windows, while choosing a random time point for the first spike. Thus the spike count and the first-order interval statistics were preserved. Using these surrogates we built our test statistics for of the phase locking from the lengths R of their circular average.

Odor response classification

To prepare for the analysis of learning induced changes in the complete neuronal ensemble, we constructed a representation in a high-dimensional space where individual dimensions represent the activity of the various units of the ensemble. Thus a point in this space represents the ensemble activity contributions of all units for a given stimulus type, experimental phase and trial. We investigated consistent changes in the ensemble representation of the three stimulus types. A common method for dealing with such high-dimensional representations is to employ the principal component analysis (PCA) to reduce the data to a few dimensions.

In practice, we defined for each stimulus type (CS+, CS-, Ctrl) and for each trial of the pre-test, differential conditioning, and post-test phase a 99-dimensional vector, where each element represents the rate of one unit (estimated by its spike count during one of the time windows, typically ON). Thus for each time window we obtain 28 such vectors (3 stimulus types with 3 trials in pre-test and post-test phase, and 2 stimulus types with 5 trials in the differential conditioning phase). The PCA was applied to these vectors to determine the directions of the largest variance in the data, which are then used as a new basis (principle components, PC). The component (or loading) w_i of a PC in the yields the weight with which each unit i contributes to the PCA representation. Reducing the dimensionality by using only the first three PCs, 74% of the variance in our data is explained. This reduction allows us to graphically visualize the structure of the data in an optimal way (cf., Fig. 4).

Results

Odor coding: Relation between LFP and spiking activity

Odor responses are not only reflected in rate changes of individual units, but also in complex LFP responses. Spindle-shaped oscillations in the LFP occurred both during spontaneous activity and during odor stimulations. First, we investigated the characteristics of these LFP responses to odor presentations during the pre-test phase. We analyzed the LFPs for dynamic changes in their frequency composition. The time-resolved and trial-averaged LFP power spectra to different odors in the same animal typically reveal a complex structure of frequency components in time. These patterns differ for different odors in the same animal, and are not odor-specific across animals.

To identify common response components across all odors and animals, we characterized the average frequency content for the 500 ms windows during spontaneous activity (baseline), phasic ON, sustained and OFF responses, respectively (Fig. 2a, top panels). During phasic ON, the total power in all frequency bands is larger than in all other time windows. In addition, a broad peak of enhanced power (compared to base line) between 40 Hz and 60 Hz emerges, which decays and shifts toward lower frequencies during the sustained and OFF responses, while the overall power decays.

We calculated a baseline-corrected and time-resolved power spectrum (Fig. 2a, bottom panels) to investigate the temporal structure of the observed increase in power from base line during odor stimulation (phasic ON, sustained and OFF response). The spectrum reveals that the peak observed during the ON response is a broad odor response across high LFP frequencies about 100 ms after odor onset. This peak starts out with a maximum at about 45 Hz, and sharpens out toward 50 Hz during the latter course of the odor presentation. Therefore, despite the high response variability across the animals, a consistent increase in power above baseline is found in the gamma frequency regime. These frequencies are not necessarily the dominant frequencies in individual animals/odors, i.e. those with the strongest power, but rather the frequencies which are consistently enhanced.

For single odor presentations we found that 73% of the odor applications in all animals evoked a significant change in the total power of the spectrum. Here, enhancement in power was defined as significant, when the power within 300 ms after stimulus onset exceeded the tolerance interval of two standard deviations around the spontaneous mean. These responses to odor stimulations were typically 300-400 ms long, and had an average response latency of 55 ± 63 ms. In 10 (out of 21) animals all three odors evoked a response, in 6 animals

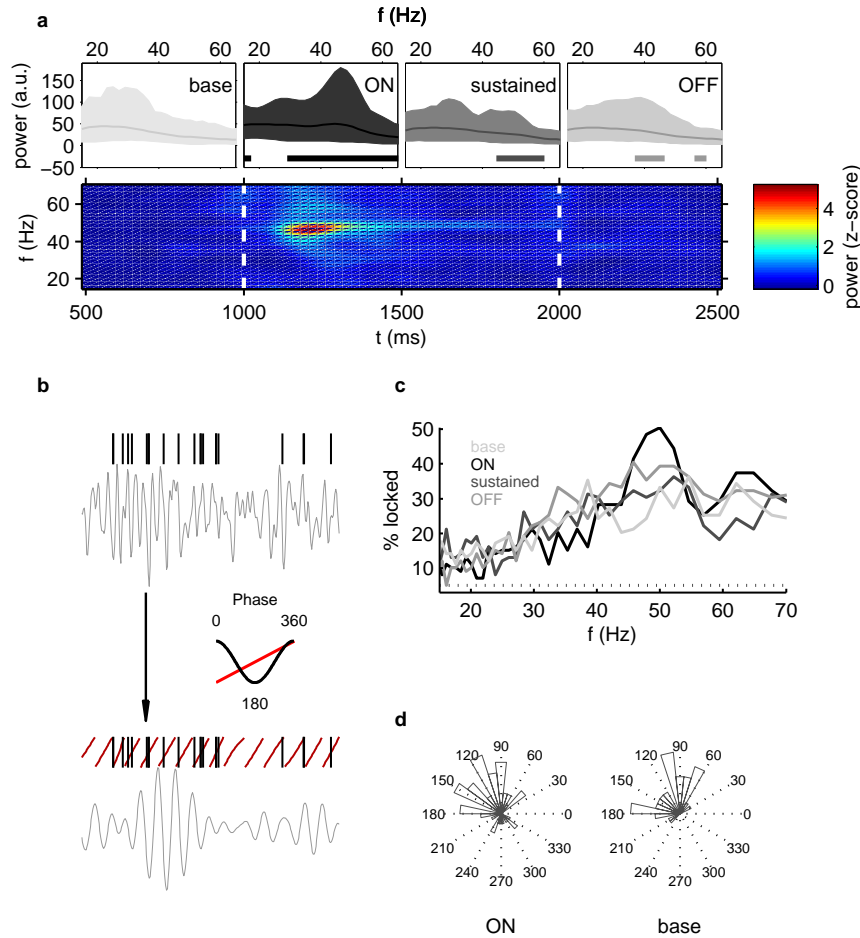


Figure 2: Frequency analysis of LFP and relationship of spikes to LFP. (a) Top: Raw power spectra for the four consecutive 500 ms time windows: spontaneous base line activity, phasic ON response, sustained response, and OFF response. Each box shows the trial-averaged power spectra in pre-test pooled across all animals and odors (curves). Shaded areas enclose at each point 90% of all recorded power spectra. Bars in the three right-most panels indicate frequency bins where the distribution of power (at that frequency) significantly deviates from the one during base (Wilcoxon signed-rank test, $p=0.05$). Bottom: Corresponding pooled, time-resolved, baseline-corrected power spectra of the LFP (300 ms sliding window; averaged over the three trials.). Here, at each time bin the spectra (indicated by color) are normalized to baseline (1000 ms before odor presentation) by subtracting its mean and dividing by its standard deviation. Beginning and end of odor presentation are marked by dashed white lines. (b) To assess the degree of phase locking between units and LFP (top, tick marks and trace, respectively), the LFP is filtered to a specified frequency band (bottom). Its instantaneous phase (implied by red lines, each spanning one oscillation cycle) is computed at the time points of spike occurrences (180 degrees: trough). (c) Percentage of units that show significant locking (as determined by surrogate distribution, $p=0.05$) to the LFP during each of the four time windows (colors as in (a)) as a function of the center frequency of the applied band-pass filter (dotted line: 5% significance threshold). (d) The distribution of mean phases of all significantly locked unit-LFP pairs (at the frequency band centered at $f=50$ Hz, cf. peak in panel (c)) during ON (left) and base (right) shows a consistent tendency for spike occurrence on the falling phase of the LFP.

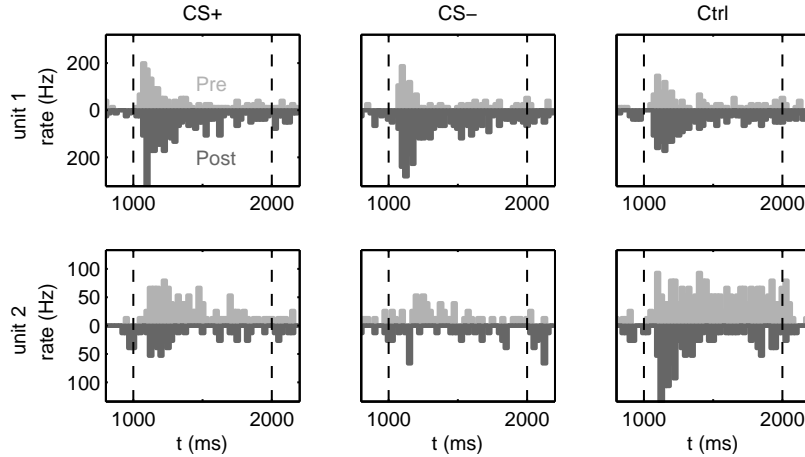


Figure 3: Typical examples of unit rate responses. Rate profiles of two units are shown before and after differential conditioning, separately for CS+ (left), CS- (middle) and Ctrl (right). Two simultaneously recorded units (top and bottom row, respectively) are presented. The rate responses are quantified by the PSTHs (bin size 25 ms) that include the three trials of the pre-test (light gray) and the post-test (dark gray, reflected about the x-axis for comparison). Dashed lines mark the interval of the stimulus presentation.

two, and in the 4 animals for one of the three odors.

Odor coding: Relationship of spikes to LFP

In a subsequent step, we analyzed how the oscillatory LFP responses are related to the responses of the spiking activity (compare Fig. 2b). To this end, we calculated the phase relationship between spiking activity of individual units and the simultaneously recorded LFP during the pre-test phase. Fig. 2c shows for all frequencies the relative number of LFP-unit pairs that are locked more than expected by chance ($p=0.05$) given their spike interval statistics. The analysis reveals that unit activity shows such genuine locking to the LFP predominantly for oscillations around 50 Hz. This effect is most pronounced for the phasic ON response (Fig. 2c, black line) in comparison to the other time windows.

The mean preferred LFP phases of all significantly locked units are not evenly distributed, but rather exhibit a strong tendency to prefer spiking on the falling phase of the oscillation. Although the mean phase shifts with the frequency band, it remains on the falling phase for all LFP frequencies investigated (15-70 Hz). The distributions of mean phases during phasic ON response and baseline (Fig. 2d) have a similar mean and variance, indicating that the preferred phase of the spikes to the LFP oscillation is not stimulus-dependent.

Learning-related changes: Spike rate responses

In the following sections we will investigate how the observed responses to odors change during the course of differential conditioning. We first asked whether the spike rate responses are affected by learning. Most of the rate responses during the pre-test stimuli are excitatory (cf. Methods). The average spike rates across all animals and trials were calculated for the responses to CS+, CS- and Ctrl during the various phases of the experiment (pre-test phase, differential conditioning phase and post-test phase). Average responses for the CS+ and CS- during the differential conditioning phase exhibit a slight but not significant decrease. However, the rate responses showed a very high variability across the population of units for a particular stimulus and experiment phase.

A more detailed analysis of spike rates performed for each of the units separately revealed either an increase or a decrease of the respective spike rates when comparing the rates during the pre-test phase with those during the post-test phase for both CS+ and CS- (see Fig. 3 for typical examples of rate changes). Therefore, even within an individual animal, the simultaneously recorded units may individually change their odor responses by either increasing or decreasing their spike rates. Averaging the changes in responses between pretest phase and post-test phase across the population would thus eliminate these complex response characteristics in individual neurons.

In order to quantify the learning-induced changes in spike rate, we calculated for each unit the average rate of the three trials in the pretest phase and evaluated whether the rates in post-test trials differed significantly (rate change of +/- 1 standard deviation estimated from the variability during spontaneous activity). We classified a change in response if at least two of the three trials of the post-test phase differed significantly. Thus we characterized unit ON responses solely by the fact that they changed, irrespective of increase or decrease of rates. 79% of the 99 units changed their rate response significantly for CS+, 67% for CS- and 74% for Ctrl. These results apply to the phasic responses. A similar effect was found for the sustained responses and the phasic OFF responses.

Besides these quantitative effects we also observed qualitative changes during the course of differential conditioning. A large number of units changed the type of their phasic ON response (33 changed their responses to CS+, 32 to CS- and 31 to Ctrl). Of these, 5 units (2 for CS+, 1 for CS-, 2 for Ctrl) switched from excitatory to inhibitory responses. In the latter cases the switch occurred abruptly in the second or third trials of the differential conditioning phase. In 39 cases an excitatory response disappeared in the post-test phase. For CS- and Ctrl each, 5 units showed an excitatory response exclusively in the post-test phase. Similarly, 9 (CS+), 11 (CS-), and 8 (Ctrl) units responded with

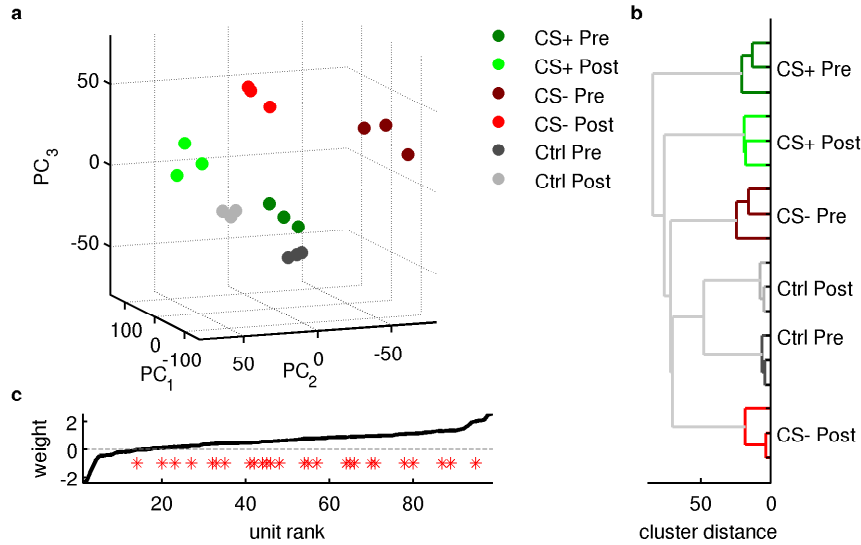


Figure 4: The result of the principle component analysis (PCA) for the ensemble of all 99 unit response rates during the ON response. (a) Projection of the response rate vectors (dots) onto the first three principle components (PCs). Each rate vector represents the response of all units for one trial of one stimulus presentation (green: CS+, red: CS-, gray: Ctrl) during the pre-test phase (dark shades) or post-test phase (light shades). (b) The hierarchical cluster tree shows the Euclidean cluster distances between all data points in panel (a). Trial responses in the same condition are tightly clustered, and cluster distances of pre-test and post-test phases of the same stimulus are closer than those between stimuli. (c) Rank-ordered contributions (weights) of each individual unit to the PC space representation of the data. Units marked by a red asterisk were recorded in animals that were classified as behavioral learners based on their M17 response.

inhibition only in the post-test phase. Finally, 7 (CS+), 3 (CS-), and 3 (Ctrl) units showed inhibitory responses in the pre-test phase and no response in the post-test phase.

In conclusion, both positive and negative changes in spike rates are observed mostly for CS+. This effect is shadowed if rates are simply averaged across the population, since these changes are compensated. To capture the highly selective response properties, we turn to more detailed measure of the learning-induced neuronal population response.

Learning-related changes: Ensemble response changes

Assuming that neural coding is a network property, learning-related changes may express themselves in global adjustments of its properties. Any analysis focusing on single neuron properties would not capture important aspects of the learning-induced network change. Below we consider the complete set of recorded units as a representative sample drawn from the network, irrespective of the fact that not all units were recorded simultaneously in the same animal.

Similar approaches were taken for calculating the direction-selective ensemble responses in the motor cortex (Georgopoulos *et al.*, 1986) or in the analysis of odor coding in the moth *Manduca* (Daly *et al.*, 2004) and in the locust (Laurent, 1996).

In a first step we concentrated on learning-induced changes from the pre-test to the post-test phase during the phasic ON response. To this end we constructed a 99-dimensional feature vector for each odor representation in both phases, whose entries are set to the respective response rates of each of the 99 recorded units (see Methods). The results of a principle component analysis (PCA) for all trials before and after differential conditioning are shown in Fig. 4a. Each resulting vector represents the projected response rates of all units to a stimulus type for one trial of the pre-test or post-test phase. Projections onto the first three principle components (PCs) captured 74% of the variance in the data. Pre-test and post-test data cluster in a stimulus-specific way. During the sustained response, the quality of clustering rapidly decreased and disappeared for the OFF response.

Another way to visualize the result of the PCA is by hierarchical clustering, as revealed by a cluster tree (Fig. 4b). This tree visualizes the clustering of the vectors by means of the Euclidean distances, which are represented by the lengths of the tree’s connecting branches. The distance between the corresponding data points in these two phases exceeds the variance of the distances in the individual trials of each phase. In particular, the distance for CS+ slightly exceeds that for CS-, whereas the distance for Ctrl is substantially lower than for the other two stimuli. We confirmed that this clustering could not be attributed to a dominant influence of only a few odors by systematically excluding single odor identities from the analysis.

In addition, we further tested the possibility that only few recordings (e.g., with a high rate) dominate the clustering. We analyzed the impact of each recording, i.e., the coefficients weighting the contributions w_i of each unit (99 weights) onto individual PCs. Fig 4c shows an ordered representation of these weights summed overall PCs, revealing that nearly all units share a substantial contribution to the ensemble representation. Next, we identified those units with an combined absolute impact $w(i)$ on the first three PCs that differed clearly from zero ($|w| > 1$). We find that all 21 animals contribute at least with one unit thus suggesting that no small group of outliers from particular animals was responsible for the clustering. For additional control, we asked how many animals and therefore, how many units had to be pooled in order to obtain a stimulus-specific clustering. We created random subgroups of increasing size from of the set of 21 animals and performed the PCA on each of these subgroups in an identical manner. A similar qualitative clustering using the first three PCs

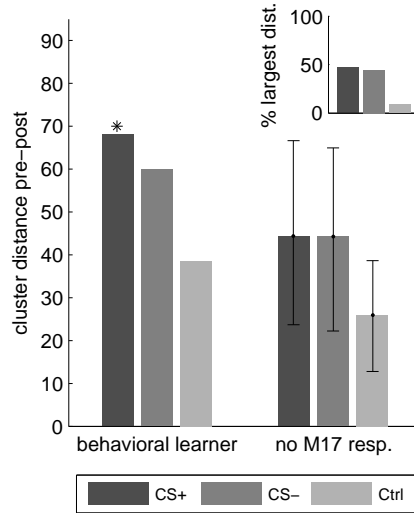


Figure 5: Comparison between animals classified as behavioral learners (left) and non-learners (right) as indicated by the response of the muscle M17 to CS+ and CS- in the post-test phase. The set of three left bars shows for learners the Euclidean cluster distances of the PCA rate vectors between pre-test and post-test phases during the ON response separately for CS+, CS-, and Ctrl. To compare these distances to the group of non-learners, we account for the different sample sizes of the two groups by a boot strap procedure. Non-learner distances were recomputed 1000 times using a random selection of animals, each consisting of an equal number (5) of animals as in the group of learners. The set of three right bars shows the mean of the distributions of distances obtained by this procedure for each of the odors (error bars indicate the 5% quantiles of the distribution). The star indicates a significant increase above the 5% level in the learner group compared to this resampled distribution. For each of the 1000 recomputations the odor exhibiting the largest cluster distance was identified. The inset shows the distribution of outcomes across the three odors (CS+: left, CS-: middle, Ctrl: right).

was obtained for all subgroups consisting of at least 5 animals.

To relate the observed changes in network dynamics to the learning process, we tested whether the observed clustering was specific to behavioral learners. These were identified by their conditioned responses (measured with the M17 recording) to CS+ but not to CS- during the post-test phase. The rather low learning rate (5 out of 21 animals) is most likely due to the impairing dissection. The PCA contributions of units recorded in behavioral learners (units indicated by asterisks in Fig. 4c) are comparable to those of the non-learners (16 animals). Furthermore, we contrasted the Euclidean cluster distances between pre-test and post-test phases of the five behavioral learners with 1000 recomputations considering five animals each that did not show any M17 response (Fig. 5). Pre- and post-test clusters are significantly further apart for CS+ in the group of learners ($p=0.05$, comparison to distribution of distances obtained from groups non-learners), indicating a stronger restructuring for the rewarded odor in animals that showed learning in the behavioral response. Nev-

ertheless, most of these shuffled groups of non-learners consistently showed the largest pre-post cluster distance for CS+ and CS-, compared to Ctrl (Fig. 5, inset). Taken together, these findings reveal that the change of response rates represented in the PCA is more pronounced in the group of behavioral learners, but not restricted to that group. Furthermore, the analysis suggests that the response characteristics are evenly distributed among all recorded animals, and justify our approach of pooling across animals to obtain a network sample.

To further characterize the changes in the network responses, we analyzed the transition between the observed clusters of pre-test and post-test phases during the course of conditioning. The distances between unit rate response vectors of successive trials belonging to the same phase of the experiment were smaller for the pretest and posttest phase as compared to those of the differential conditioning phase (for both, CS+ and CS-). Therefore, we conclude that without conditioning, an odor retains a rather stable representation in unit rate space, while the learning process induces a strong change in its ensemble rate representation. The finding serves as a strong indication that the change in response rates is due to conditioning and not to a gradual change over the whole experiment.

In summary, a PCA of all unit responses across animals revealed a homogeneous stimulus specificity of the rate responses in the pre-test and post-test phase of the experiment, where the distance between the two phases for the CS+ exceeded that for the CS- and Ctrl.

Learning-related changes: LFP power

To further characterize learning-induced changes in the network response, we address the question whether also changes in the LFP responses occurred due to differential conditioning. In particular, we ask if the LFP power shifted with respect to the contributions of certain frequency components. The average time-resolved power spectra for CS+ in the pre-test phase as compared to the post-test phase (Fig. 6a) shows prominent changes in the normalized power mainly during ON response. In particular, we observe a reduction in power at high LFP frequencies around 50 Hz shortly after stimulus onset. To quantify this observation we calculated the average across animals of the difference between pre- and post-test power (Fig. 6b). For each animal we averaged the normalized spectra in the post-test trials and subtracted the average of the normalized spectra of the respective pre-test trials separately for the three stimulus conditions (CS+, CS-, and Ctrl). Only for CS+, the relative power in the 15–40 Hz band increases significantly from the pre-test to the post-test phase, while the power in the frequency components above 40 Hz decrease significantly (vari-

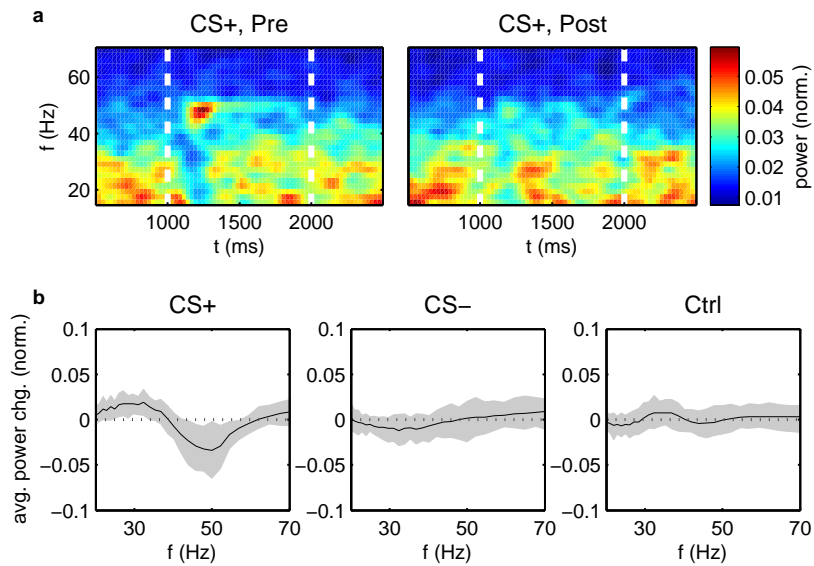


Figure 6: Changes in LFP power from pre-test to post-test phase. (a) Time-resolved power spectra for pre- (left) and post-test (right) averaged across all animals for CS+. Analysis is carried out in sliding windows of 300 ms and the average of the three trials per animal is normalized to unit area per time slice before pooling. White lines indicate stimulus onset and offset, power is indicated by color. (b) Average power change during the ON response from pre- to post-test phase for CS+ (left), CS- (middle), and Ctrl (right), resolved by individual frequency bands. Before averaging across animals, the differences between power (averaged across the three trials and normalized to unit area) in pre-test and post-test are calculated. Error bars ($\pm 95\%$) were obtained using 100 bootstraps.

ability across animals estimated by bootstrapping). In contrast, no significant power changes were observed for the other two stimuli (CS-, Ctrl). In the other time windows, no significant changes were seen for any of the stimuli.

To clarify the nature of the observed change in power for the CS+, we investigated its time course also during the differential conditioning phase. The change in LFP power occurs only with the start of the post-test phase. Thus, we might speculate that the LFP power reflects a consolidation effect, whereas the continuous rate change during the conditioning phase is indicative of immediate and local changes in the network.

The frequency components were also analyzed separately for the behavioral learners. For these five animals, we observed similar changes in power as for the whole population. Likewise, for the rest of the animals that did not show an M17 response we found comparable power changes. Thus we conclude that the observed result is independent of behavioral learning, in agreement with our results of the unit ensemble analysis.

Learning-related changes: Relationship between LFP and spiking activity

As shown in the previous section, we observed a relative change in the power composition of the LFP for the CS+ odor. However, learning-induced changes in unit rates involved a complex restructuring of the ensemble responses, which is not reflected in the averaged network response. Therefore, we did not observe direct correlations between the global changes reflected by the LFP and the average rate changes of units recorded in the same animal.

However, we were able to directly investigate changes in the relation between spikes and LFP oscillations. Learning-related changes were investigated with respect to an emergence or loss of significantly locked units in the post-test phase as compared to the pre-test phase. In a first step we calculated the change in the number of units locked to the LFP more precisely than expected given their regularity during the ON response as a function of the LFP frequency (Fig. 7a, cf. also Fig. 2). We observed a tendency for a decrease in the number of locked units at 45-60 Hz for CS+ and Ctrl, and a slight increase in the 15-40 Hz band for CS+ and CS-.

This frequency-resolved analysis shows a strong similarity between the change in the number of locked units and the change in power from pre-test to post-test phases (compare Fig. 6b). We quantified the similarity by calculating the correlation coefficient between these two measures across frequencies. For CS+, changes in locking and changes in power are highly correlated across frequencies with a correlation coefficient of 0.58 during the ON response. For the other

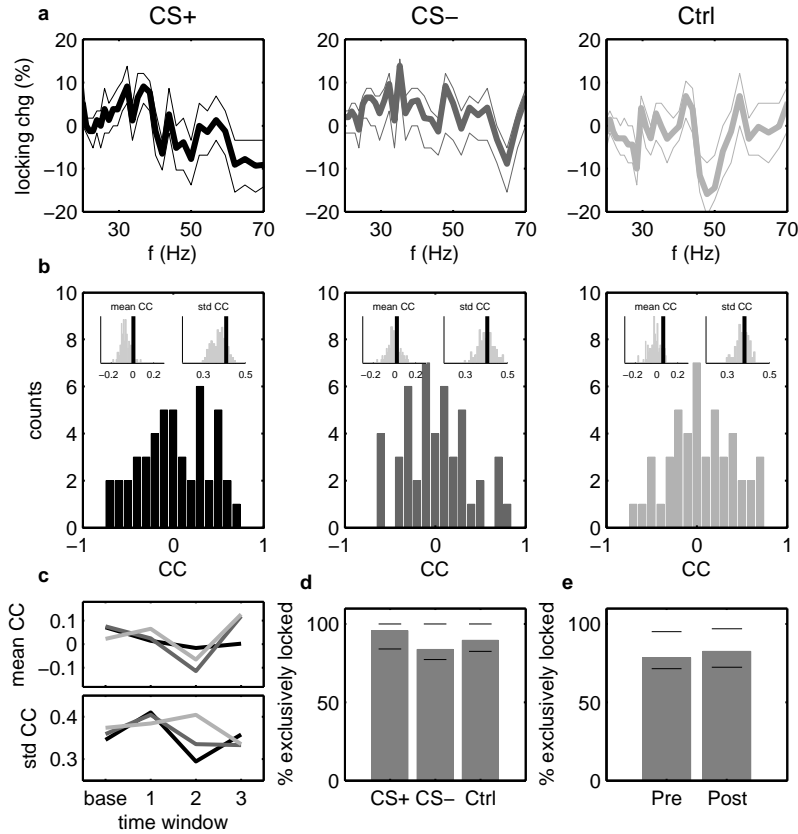


Figure 7: Learning-related changes in the phase relationship of spikes to LFP. a) Difference from pre-test to post-test phases in the relative number of units locked significantly to LFP oscillations for CS+ (left), CS- (middle), and Ctrl (right) for frequencies from 20-70 Hz. The change is expressed in percent of the total number of units available per odor. Error bars ($\pm 95\%$, thin lines) are calculated using a resampled data sets comprising 75% of the data. (b) Distribution of correlation coefficients between the frequency-resolved changes in power and locking obtained for each unit-LFP pair in during ON. Insets show mean and standard deviation of this distribution (black lines), in addition to distributions (gray) of means and standard deviations of the distributions of correlation coefficients obtained from 100 surrogates (randomized trial-by-trial pairing of power and locking). (c) Mean and standard deviation of distribution of correlation coefficients as in panel b, but for all 4 time windows. (d) Number of units during CS+, CS-, and Ctrl, respectively, that are locked exclusively to either only pre-test or only post-test for frequencies between 40 and 60 Hz, relative to the number of units that are locked in any of these two phases. (e) Number of units during pre- and post-test, respectively, that are locked to either only CS+, only CS- or only Ctrl, relative to the number of units locked to any of these three odor stimuli. Error bars in (d) and (e) indicate the range of exclusively locked units expected assuming randomly chosen identities of locked units based on the measured frequency of locking.

two stimuli, correlations are weaker and in fact anti-correlated for CS-, with correlation coefficients of -0.27 and -0.38 for CS- and Ctrl, respectively.

Given the large variability between individual animals, we investigated whether a similar relationship between changes in locking and power can be observed for single recorded units. Fig. 7b shows histograms of the correlation coefficients measured between the change in power and the change in locking (represented as a binary vector for each unit, indicating for each individual frequency locked or unlocked as 1 or 0, respectively) for all animals during the ON response. In addition, the mean and standard deviations of the resulting distributions are shown (insets), along with the means and standard deviations obtained from a set of 100 surrogates. For each surrogate, the power changes corresponding to one unit were correlated with the locking changes obtained from a different, randomly-chosen unit. The distributions appear largely unimodal and centered around zero. Therefore, the relationship between the average locking change and average power change does not indicate a consistent trend for individual units. Similar results are obtained for the remaining time windows (Fig. 7c). To understand this effect, it is helpful to realize that an individual unit is typically locked to only a particular frequency component. Therefore, on the level of individual units, we find no correlation across frequencies. In contrast, on the network level, LFP power for a certain frequency correlates with the probability of neurons becoming entrained to oscillatory network activity at that frequency.

Next, we investigated the stability of locked units from pre-test to post-test (Fig. 7d), and specificity of locked units across different CS+, CS- and Ctrl odors (Fig. 7e). The analysis reveals that the composition of units that show significant locking to the LFP is highly variable. In fact, most units are highly specific in locking to the experienced odor (CS+, CS- or Ctrl). In addition, even though the total number of units locked to the LFP does not change on average, the identities of units that exhibit locking after conditioning have changed to a large extent. Due to the low relative number of units that show significant locking, this specificity may be well explained by a random selection of the units identities (as indicated by the black lines in Fig 7d,e). Therefore, the question of whether this observation is in fact due to a systematic restructuring that enforces this specificity cannot be answered in this analysis. Nonetheless, these results support our previous findings that differential conditioning induces a broad restructuring of responses within the network, not only on the rate level, but more over in the temporal coordination of individual units to the network activity expressed through oscillatory LFP activity.

Discussion

Learning during differential conditioning of the bee leads to spike rate increases and decreases in responses to the reward odor (CS+), the unrewarded odor (CS-), and the control odor (Ctrl) introduced to test generalization. An analysis of the ensemble activity indicated the largest learning related difference for CS+. Furthermore, LFP power increases for CS+ in the 15-40 Hz frequency band during the odor presentation, and decreases for frequencies above 45 Hz. For CS+, learning-related changes in the size of the phase-locked spiking activity correlate with the LFP power changes. These results reflect associative plasticity in the antennal lobe of the bee resulting from a restructured odor coding network. The observed changes appear to decorrelate the responses to the learned and non-learned odors and may provide a learning-related neural signature of odor learning.

Characteristics of unit and LFP responses

Our recordings aimed for PNs by positioning the two tetrodes close to the AL exit of the mACT and lACT. The responses of units recorded resemble those collected with intracellular electrodes in PNs (Sun *et al.*, 1993; Abel *et al.*, 2001; Müller *et al.*, 2002, Krofczik *et al.*, 2009). Odors elicit excitatory phasic-tonic responses, inhibitory tonic or phasic responses, and excitatory and inhibitory OFF-responses. The response patterns of 17% of our recordings, and their odor response profiles are in accordance with the findings for the lACT type of PNs (Müller *et al.*, 2002). 83% of the recordings showed response properties that are similar to those of the mACT neurons (Abel *et al.*, 2001; Müller *et al.*, 2002). They respond more specifically to odors, and their response patterns consist of bursting-like spiking, inhibitory responses or off-responses, properties which were only seen in mACT neurons (Müller *et al.*, 2002). The average response rates in our recordings were between 1.5 and 2 times higher than those observed in intracellular recordings of mACT neurons.

LFP oscillations occurred both spontaneously and odor-induced. It is still under debate what kind of information lies in such oscillations with respect to olfactory information processing. Recordings from the AL of *Manduca sexta* revealed that the frequency spectra of oscillations are more specific for the recording site than for odor identities, and may be restricted to single glomeruli (Christensen *et al.*, 2000, 2003). This could mean that the oscillations reflect the summed activity of neuronal populations processing odor information within a single glomerulus or a small number of glomeruli and thus may depend strongly on the spatial organization of the AL. Recordings in *Schistocerca americana* suggested a more global function of LFP oscillations and are interpreted as a

reference for the transient synchronization of subpopulations of PNs (reviews: Laurent, 2002; Laurent *et al.*, 2001). These latter authors conclude that such a transient synchronization serves as a mechanism for the transfer of olfactory information from the AL to the MB, where MB-intrinsic Kenyon cells may act as coincidence detectors for the synchronous firing of PNs during fixed cycles of the LFP oscillation period. This hypothesis is particularly interesting in the context of odor learning, because changes in the synchronization of PN spike activity may reflect a signature of learned odors (Finelli *et al.*, 2008, see below).

The frequency spectra of LFPs for different odors were specific within the same animal, and different for the same odor in different animals. Since the units picked up may differ between animals, we conclude that odor-evoked LFP oscillations are specific for both the recording site and the odor identity. This supports the idea that odors are coded by a spatio-temporal glomerular activity pattern within the AL of the honeybee, as indicated by optical recordings (for a review, see Galizia & Menzel, 2000). The oscillations would then reflect the summed activity of odor-specific subpopulations of neurons.

Yet the sources of the LFP oscillations remain unclear. The LFP signals were nearly identical at the two shafts (separated by 60 μm). LFPs recorded at two remote sites, the AL and the mushroom body (calyx or alpha lobe) on the same side of the brain are also very similar and synchronous, whereas LFPs recorded on the two sides of the brain are less similar and not synchronous (Szyyszka and Menzel, pers. observation). This indicates that the sources of LFP in the honeybee brain are distributed, or – if localized – are synchronous for the distributed sites in one half of the brain. We conclude that the LFP - which we recorded with the same electrodes as the spike activity - represents not only the low-frequency components of summed spikes at the recording site, but reflects further-ranging low-frequency components indicative of coherence phenomena in a population of neurons. It is, therefore, not surprising that odor-induced unit responses and LFP oscillations were not always related to each other. In 10% of all odor stimulations even strong unit responses were not associated with strong LFP oscillations. Furthermore, strong LFP oscillations were not always connected to unit responses (as observed in 3% of all odor stimulations). Furthermore, the high frequency component was found in our recordings at a considerable higher frequency (around 45-60 Hz) than in recordings from the bee brain by Stopfer *et al.* (1997) around 30 Hz. The reasons for this difference is elusive but may be related to the fact that we ensured temperature at the recording site above 23° C because we found in other experiments that a temperature below 23° C leads to lower spontaneous activity, lower response rates and spikes with smaller amplitude.

The high 50 Hz power occurs during the phasic on-response of PN, a response

component that codes odor identity less reliable than the sustained response component (Sun *et al.*, 1993; Abel *et al.*, 2001; Müller *et al.*, 2002, Krofczik *et al.*, 2009). Thus the phasic synchronization of the fast odor-induced activity may serve other functions than specific odor coding (possibly as a reference point for latency coding, and/or general excitation of mushroom body-intrinsic and -extrinsic neurons). The delayed synchronization in the frequency band of <40 Hz could be related to specific odor coding. Learning-related changes in these two frequency bands may inform us whether the learned signal may be more specifically coded and whether the fast and general odor signal may become less important.

Locking between spikes and LFP oscillations was observed particularly for the 45-60 Hz band, which shows a large power increase during the initial phase of the stimulus presentation. The locking was unit-specific, occurred only during subsets of unit and LFP responses, and was dominant for a certain frequency band, indicating an AL-intrinsic mechanism. It is thus possible that these oscillations reflect a transient synchronization of PN subgroups, as has been suggested for *Schistocerca americana* (Laurent, 2002). Such transient synchronization could play a role in the fine tuning of odor representations, and/or could provide a coincident input to Kenyon cells (Finelli *et al.*, 2008).

Learning-related rate changes in population responses

The goal of this study was to investigate whether learning-related plasticity exists in the AL of the bee and how it might be expressed at the levels of neuronal spiking and population activity. The AL of the honeybee is believed to be a component of a distributed network storing olfactory information (review: Menzel & Müller, 1996). Evidence comes from four observations. (1) The reward neuron VUMmx1 converges with the olfactory pathway in the AL besides its convergence sites in the mushroom body and the lateral protocerebrum (Hammer, 1993); (2) local injection of the putative transmitter of VUMmx1, octopamine, substitutes for the unconditioned stimulus (US, sucrose) when injected into the AL immediately after an odor presentation (Hammer & Menzel, 1998); (3) local uncaging of cAMP in the AL (cAMP is a second messenger known to promote the transfer from short- to long-term memory in bees), will shift short to long-term memory when it is uncaged shortly after a single learning trial (Müller, 2000); (4) optical imaging of activity-dependent Ca^{2+} signals of glomeruli in the AL of the honeybee and *Drosophila* reveal an increase in the Ca^{2+} signal for the CS+ in the case of the bee (Faber *et al.*, 1999), and in *Drosophila* a recruitment of previously inactive output synapses of the PNs within the glomeruli for the CS+ (Yu *et al.*, 2004).

Our results show significant unit rate changes, both increases and decreases, for the three groups of odors used in differential conditioning (CS+, CS-, Ctrl). This result is not surprising if one considers the complex interactions between excitatory and inhibitory neurons within the AL. Strengthening the response in a subgroup of AL neurons may well lead to both an increase of excitation and inhibition. Although we cannot strictly relate our extracellular recordings to any type of neuron (see above) and thus cannot separate the strengthening and weakening of responses to any neuroanatomical subtype, we still presented arguments above in favor of the interpretation that our electrodes predominantly picked up spikes of PNs. In this case, our results would mean that PNs - possibly even those receiving input within the same glomerulus - may change their response strengths in opposing directions. This situation might explain why optophysiological recordings from glomeruli showed in some experiments learning-induced changes and in others no such changes were observed. In the study by Peele *et al.* (2006) which gave no indication of learning related Ca^{2+} activity changes, special care was taken to fill many IACT neurons with the Ca^{2+} indicator dye. This could have well lead to the filling of several PNs projecting out of one glomerulus, and if they changed their responses in opposing directions, as concluded here, no effect will be seen at the glomerular level. Interestingly, in a parallel study (Weidert *et al.*, 2004) both enhancement and weakening of the Ca^{2+} responses to the learned odor (CS+) was found in IACT postsynaptic sites. In this latter study less intensive dye filling has been applied. These controversial results have yet to be resolved, and our findings here may provide a basis for this.

Rate changes were not only seen for the CS+ but also for CS- and to a slightly lower degree for Ctrl. Although average response rates decreased only slightly over the course of the sequential stimulations one might still argue that the observed rate changes may not reflect a differential learning effect, but rather non-associative phenomena (habituation, sensitization) or a decrease in the animals' fitness throughout the experiment. Habituation is a comparatively slow process in honeybee odor processing (Chandra *et al.*, 2000), and only more than 20 sequential odor stimuli lead to a rather weak habituation effect. Multiple sucrose stimuli habituate the proboscis extension response only after more than 15 sucrose applications (Braun & Bicker, 1992). Sucrose stimuli sensitize the animal which leads in behavior to a transient increase of response probability to different stimuli (Menzel, 1990). The time course of US induced sensitization lies in the range of a few minutes. Since the post phase started 5 minutes after the end of conditioning it is unlikely that US sensitization would still play a role, although it cannot be excluded that the time course of neural correlates of sensitization may last longer. The rather high degree of absolute rate changes for

Ctrl may indicate a non-associative sensitization effect besides a generalization effect. A slow weakening of the animals' fitness can be ruled out because the shifts in unit ensemble responses were strong during the differential conditioning phase, whereas both pre-test and post-test phases exhibited only small changes.

In differential conditioning bees learn both the CS+ and the CS- (Bitterman *et al.*, 1983; Menzel, 1990). The forward paired CS+ is responded to more strongly, and inhibitory learning of the backward paired CS- is uncovered by a resistance of acquisition to subsequent forward pairing (Hellstern *et al.*, 1998). The direction of neuronal change as indicated by spike rate responses is obviously not related to the distinction between excitatory forward and inhibitory backward learning during differential conditioning.

Besides these quantitative changes we also observed qualitative changes in the course of differential conditioning. About 30% of units switched their phasic ON responses, more frequently for the CS+ than for the CS-. Several units were recruited or lost their response properties during conditioning. Daly *et al.* (2004) observed in *Manduca sexta* recruitment of neuron activity for CS+, and a loss of neuron activity for the CS-. Recruitment of neural responses was also seen in *Drosophila* olfactory conditioning (Yu *et al.*, 2004). Interestingly, in our experiments all these changes occurred abruptly between the second and third training trial indicating a sudden reorganization of the AL network after more than one learning trail. Single and multiple learning trials lead to different forms of memory in the honeybee (Menzel, 1999), single learning trials only to short-term memory and multiple learning trials to long-term memory via several phases of consolidation. Since we saw the spike rate changes during acquisition but the changes in LFP power only in the post-test phase we might speculate that these two measures indicate different components of associative plasticity. The first might monitor associative plasticity initiated predominantly in the AL by pairing effects; the latter may result from consolidation phenomena which could include downstream neuropiles like the mushroom body. The mushroom bodies in bees are known to be involved in memory formation during 5–7 minutes after conditioning, whereas the AL appears to contribute to memory formation only during about one minute after conditioning (Menzel *et al.*, 1974). It will be interesting to ask in future experiments which neural correlates of associative changes are related to learning phenomena and which to those of memory formation.

The PCA analysis supports our conclusion that the learning effect in differential conditioning is spread across a population of neurons, changing their responses relative to each other both by a rise and a fall in response strength (Fig. 4). Such network properties may reflect a tendency to rebalance the effects caused by learning-related enhancement of excitation in a subpopulation

of neurons, and may indicate a shift in the stimulus sensitivity profile toward the learned stimulus as it was found for the frequency tuning of neurons in the rat primary auditory cortex (Froemke *et al.*, 2007). Since the spatial code of odor-induced activity in the AL is similar in different bees the recordings in 21 different animals can be viewed as simultaneous recordings at 21 different recording sites in one animal. According to this view, the clustering of stimuli and experimental phases reveals changes in the ensemble. Since each multi-dimensional point in PCA space contains the complete information from simultaneously recorded response rates, the differences between the clusters of the experimental phases reflect the absolute amount of rate change that has taken place in the ensemble. The analysis shows that the Euclidean distance between pre-test and post-test phases in the space of the three dominant principal components is highest for CS+, indicating a differential learning effect, and corroborating our conclusion - presented above - that the absolute rather than the positive or negative change in response rates reflects associative plasticity in the network.

Thus our analysis reveals that both unit and LFP responses monitor associative plasticity in the network. An obvious question is how these two signals are related, keeping in mind that in the common view LFP signals reflect the low-pass filtered neuronal activities in the surrounding volume. We observed that LFP power responses to CS+ increase significantly in the 15-40 Hz band and strongly decrease in the 45-60 Hz band in the course of conditioning. We do not find any clear relation between these power changes and spike rates, neither for individual nor for pooled responses. This suggests that the LFP cannot be interpreted as a mere reflection of the firing rate dynamics of the neuronal ensemble. However, one might also argue that it is difficult to observe net ensemble firing rate changes in a signal averaged over a large population of neurons, because we observe similar numbers of units that increase their rates with CS+ as compared to units that decrease their rates.

Assuming that a high degree of synchronization is reflected in enhanced oscillatory LFP components an increase in spike synchronization on the ensemble level should also be expressed in a more precise temporal relationship between spikes and these oscillations, which should occur at a similar frequency. Indeed, a similar tendency as for the learning-induced power changes is observed for changes in the precision of the phase relationship of spiking to LFP. In particular for CS+, locking sharpens for the 15-40 Hz band and reduces for the 45-60 Hz band during the differential conditioning phase (Fig. 7a). These changes for CS+ in spike-LFP locking and changes in LFP power were highly correlated across frequencies. Furthermore, we found that spikes became less precisely locked after learning almost as often as they increased their locking

(Fig. 7e). This coexistence of emergence and loss of locking is indicative of a reorganization of synchronization between neurons as a consequence of learning.

It is often assumed that the frequency of oscillations is related to the spatial extent of interaction (Buzsaki & Draguhn, 2004). In this context, we may hypothesize that the shift to lower frequencies after conditioning for both locking and power could reflect a tendency for a more global representation of the rewarded odor than before learning. Therefore, it may provide evidence for a contrast sharpening of the meaningful odors against others. Similar effects were seen in rats (Ravel *et al.*, 2003; Martin *et al.*, 2006), where odor learning leads to a shift of the dominant frequency in LFP in the olfactory bulb from the gamma band (60-90 Hz) to the beta band (15-40 Hz).

In analyzing the ensemble rate responses we found that the degree of change in the representation of the CS+ odor changed significantly more for behavioral learners as compared to the non-learners. Thus, learning-related changes in the neural network of the AL correlated with the scores of behavioral learning (probability to extend the proboscis to the CS+ alone and not to the CS-). This is in line with findings of mushroom body extrinsic neurons, the PE1, whose learning-related changes of odor responses (an associative reduction of their spiking rate) correlate with the level of behavioral learning (Okada *et al.*, 2007). Nonetheless, even in behavioral non-learners, CS+ consistently showed the strongest ensemble restructuring in our data. Thus, the rather loose correlation between scores of learning and the degree of ensemble change suggest that the performance during learning is related but not determined by the network reorganization of the antennal lobe that provides the signal for down stream networks more closely controlling behavior.

Taken together, these results indicate that olfactory learning leads to a reorganization of neural processing of odors both at the level of the particular PNs involved in transmitting the spatial glomerulus activity patterns and the network properties of many neurons in the circuit connecting the antennal lobe with the mushroom body. It will be necessary in the future to document more closely that a learned odor is indeed more precisely coded, and is more strongly represented by higher synchrony of particular PNs. In such a case the signature of a learned odor would lie in the changed balance between inputs from different PNs.

Acknowledgments

This work was supported by the Bernstein Center for Computational Neuroscience, Berlin (Grant No. 01GQ0413, to MD, SG, and RM), by the Volkswagen Foundation (MD, SG), and by the Stifterverband für die Deutsche Wissenschaft

(SG). RF was supported by the Deutsche Forschungsgemeinschaft in Sfb 515 (TP Menzel).

References

- Abel, R., Rybak, J. & Menzel, R. (2001) Structure and Response Patterns of Olfactory Interneurons in the Honeybee, *Apis mellifera*. *J. Comp. Neurol.*, **437**, 363-383.
- Bitterman, M. E., Menzel, R., Fietz, A. & Schäfer, S. (1983) Classical conditioning of proboscis extension in honeybees (*Apis mellifera*). *J. Comparative Psych.*, **97**, 107-119.
- Braun, G. & Bicker, G. (1992) Habituation of an appetitive reflex in the honeybee. *J. Neurophysiol.*, **67**, 588-598.
- Buzsáki, G. & Draguhn, A. (2004) Neuronal oscillations in cortical networks. *Science*, **304**, 1926-1929.
- Chandra, S. B., Hosler, J. S. & Smith, B. H. (2000) Heritable variation for latent inhibition and its correlation with reversal learning in honeybees (*Apis mellifera*). *J. Comp. Psychol.*, **114**, 86-97.
- Christensen, T. A., Pawlowski, V. M., Lei, H. & Hildebrand, J. G. (2000) Multi-unit recordings reveal context-dependent modulation of synchrony in odor-specific neural ensembles. *Nat. Neurosci.*, **3**, 927-931.
- Christensen, T. A., Lei, H. & Hildebrand, J. G. (2003) Coordination of central odor representations through transient, non-oscillatory synchronization of glomerular output neurons. *Proc. Natl. Acad. Sci. USA*, **100**, 11076-11081.
- Daly, K. C., Christensen, T. A., Lei, H., Smith, B. H. & Hildebrand, J. G. (2004) Learning modulates the ensemble representations for odors in primary olfactory networks. *Proc. Natl. Acad. Sci. USA*, **101**, 10476-10481.
- Denker, M., Roux, S., Timme, M., Riehle, A. & Grün, S. (2007) Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomp.*, **70**, 2096-2101.
- Faber, T., Joerges, J. & Menzel, R. (1999) Associative learning modifies neural representations of odors in the insect brain. *Nat. Neurosci.*, **2**, 74-78.
- Finelli, L. A., Haney, S., Bazhenov, M., Stopfer, M. & Sejnowski T. J. (2008) Synaptic learning rules and sparse coding in a model sensory system. *PLoS Comput. Biol.*, **4**, e1000062.

- Froemke, R. C., Merzenich, M. M. & Schreiner, C. E. (2007) A synaptic memory trace for cortical receptive field plasticity. *Nature*, **450**, 425-429.
- Galizia, C. G., Joerges, J., Küttner, A., Faber, T. & Menzel, R. (1997) A semi-in-vivo preparation for optical recording of the insect brain. *J. Neurosci. Meth.*, **76**, 61-69.
- Galizia, C. G. & Menzel, R. (2000) Odour perception in honeybees: coding information in glomerular patterns. *Curr. Opin. Neurobiol.*, **10**, 504-510.
- Gelperin, A. (1990) A taste for learning. *Am. Zool.*, **30**, 549-558.
- Georgopoulos, A. P., Schwartz, A. B. & Kettner, R. E. (1986) Neuronal population coding of movement direction. *Science*, **233**, 1416-1419.
- Goldman-Rakic, P. S. (1995) Cellular basis of working memory. *Neuron*, **14**, 477-485.
- Hallem, E. H. & Carlson, J. R. (2006) Coding of odors by a receptor repertoire. *Cell*, **125**, 143-160.
- Hammer, M. (1993) An identified neuron mediates the unconditioned stimulus in associative olfactory learning in honeybees. *Nature*, **366**, 59-63.
- Hammer, M. & Menzel, R. (1998) Multiple sites of associative odor learning as revealed by local brain microinjections of octopamine in honeybees. *Learning & Memory*, **5**, 146-156.
- Harris, K. D., Henze, D. A., Hirase, H., Leinekugel, X., Dragoi, G., Czurko, A. & Buzsáki, G. (2002) Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature*, **417**, 738-741.
- Hellstern, F., Malaka, R. & Hammer, M. (1998) Backward inhibitory learning in honeybees: A behavioral analysis of reinforcement processing. *Learning & Memory*, **4**, 429-444.
- Hurtado, J. M., Rubchinsky, L. L. & Sigvardt, K. A. (2004) Statistical method for detection of phase-locking episodes in neural oscillations. *J. Neurophys.*, **91**, 1883-1998.
- Joerges, J., Küttner, A., Galizia, C. G. & Menzel, R. (1997) Representation of odours and odour mixtures visualized in the honeybee brain. *Nature*, **387**, 285-288.
- Keverne, E. B. (1995) Olfactory learning. *Curr. Opin. Neurobiol.*, **5**, 482-488.
- Krofczik, S., Menzel, R. & Nawrot, M. P. (2009) Rapid odor processing in the honeybee antennal lobe network. *Front. Comput. Neurosci.*, **2**, 1-9.

- Laurent, G. J. (1996) Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends. Neurosci.*, **19**, 489-496.
- Laurent, G. J., Stopfer, M., Friedrich, R. W., Rabinovich, M. I., Volkovskii, A. & Abarbanel, H. D. I. (2001) Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu. Rev. Neurosci.*, **24**, 263-297.
- Laurent, G. J. (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nature Rev. Neurosci.*, **3**, 884-895.
- Le Van Quyen, M., Foucher, J., Lachaux, J., Rodriguez, E., Lutz, A., Martinerie, J. & Varela, F. J. (2001) Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J. Neurosci. Meth.*, **111**, 83-98.
- Mardia, K. V. & Jupp, P. E. (2000) *Directional Statistics*. Wiley, New York.
- Martin, C., Gervais, R., Messaoudi, B. & Ravel, N (2006) Learning-induced oscillatory activities correlated to odour recognition: a network activity. *Eur. J. Neurosci.*, **23**, 1801-1810.
- Mauelshagen, J. (1993) Neural correlates of olfactory learning in an identified neuron in the honey bee brain. *J. Neurophysiol.*, **69**, 609-625.
- Menzel, R., Erber, J. & Masuhr, T. (1974) Learning and memory in the honeybee. In Barton-Browne, L. (ed.) *Experimental analysis of insect behaviour*. Springer-Verlag, Berlin, pp. 195 – 217.
- Menzel, R. (1990) Learning, memory, and "cognition" in honey bees. In Kesner, R. P. & Olton, D. S. (eds.) *Neurobiology of comparative cognition*. Erlbaum Inc., Hillsdale, N.J., pp. 237-292.
- Menzel, R. & Müller, U. (1996) Learning and memory in honeybees: From behavior to neural substrates. *Ann. Rev. Neurosci.*, **19**, 379-404.
- Menzel, R. (1999) Memory dynamics in the honeybee. *J. Comp. Physiol. [A]*, **185**, 323-340.
- Müller, D., Abel, R., Brandt, R., Zöckler, M & Menzel, R (2002) Differential parallel processing of olfactory information in the honeybee, *Apis mellifera* L. *J. Comp. Physiol. [A]*, **188**, 359-370.
- Müller, U. (2000) Prolonged activation of cAMP-dependent protein kinase during conditioning induces long-term memory in honeybees. *Neuron*, **27**, 159-168.
- Okada, R., Rybak, J., Manz, G. & Menzel, R. (2007) Learning-related plasticity in PE1 and other mushroom body-extrinsic neurons in the honeybee brain. *J. Neurosci.*, **27**, 11736-11747.

- Peele, P., Ditzen, M., Menzel, R. & Galizia, C. G. (2006) Appetitive odor learning does not change olfactory coding in a subpopulation of honeybee antennal lobe neurons. *J. Comp. Physiol. [A]*, **192**, 1083-1103.
- Pelz, C., Gerber, B & Menzel, R. (1997) Odorant intensity as a determinant for olfactory conditioning in honeybees: Roles in discrimination, overshadowing and memory consolidation. *J. Exp. Biol.*, **200**, 837-847.
- Ravel, N., Chabaud, P., Martin, C., Gaveau, V., Hugues, E., Tallon-Baudry, C., Bertrand, O & Gervais, R. (2003) Olfactory learning modifies the expression of odour-induced oscillatory responses in the gamma (60-90 Hz) and beta (15-40 Hz) bands in the rat olfactory bulb. *Eur. J. Neurosci.*, **17**, 350-358.
- Rolls, E T., Critchley, H D., Mason, R & Wakeman, E. A. (1996) Orbitofrontal cortex neurons: Role in olfactory and visual association learning. *J. Neurophysiol.*, **75**, 1970-1981.
- Rosenblum, M. G., Pikovsky, A. S. & Kurths, J. (1996) Phase synchronization of chaotic oscillators. *Phys. Rev. Lett.*, **76**: 1804-1807.
- Sachse, S. & Galizia, C. G. (2002) The Role of Inhibition for Temporal and Spatial Odor Representation in Olfactory Output Neurons: A Calcium Imaging Study. *J. Neurophysiol.*, **87**, 1106-1117.
- Sandoz, J. C., Galizia, C. G. & Menzel, R (2003) Side-specific olfactory conditioning leads to more specific odor representation between sides but not within sides in the honeybee antennal lobes. *Neurosci.*, **120**, 1137-1148.
- Schultz, W. (1998) Predictive Reward Signal of Dopamine Neurons. *J. Neurophysiol.*, **80**, 1-27.
- Sharott, A., Moll, C. K. E., Engler, G., Denker, M., Grün, S & Engel, A K (2009) Different subtypes of striatal neurons are selectively modulated by cortical oscillations. *J. Neurosci.*, **29**, 4571-4585.
- Stopfer, M., Bhagavan, S., Smith, B. H. & Laurent, G. J. (1997) Impaired odour discrimination on desynchronization of odour-encoding neural assemblies. *Nature*, **390**, 70-74.
- Sutherland, G. R. & McNaughton, B. (2000) Memory trace reactivation in hippocampal and neocortical neuronal ensembles. *Curr. Opin. Neurobiol.*, **10**, 180-186.
- Sun, X-J., Fonta, C. & Masson, C. (1993) Odour quality processing by bee antennal lobe interneurons. *Chem. Senses*, **18**, 355-377.
- Tank, D. W., Gelperin, A. & Kleinfeld, D. (1994) Odors, oscillations, and waves: Does it all compute? *Science*, **265**, 1819-1820.

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

Wang, Z., Wong, A. M., Flores, J., Vosshall, L. B. & Axel, R. (2003) Two-photon calcium imaging reveals an odor-evoked map of activity in the fly brain. *Cell*, **112**, 271-282.

Weidert, M., Galizia, C. G. & Menzel, R. (2004) Learning induced changes of neuronal representation of odors in the antennal lobe of honeybee. Proceedings of the Berlin Neuroscience Forum, April 22 -24, 2004, Liebenwalde, Germany, 75.

Wilson, D. A., Sullivan, R. M. & Leon, M. (1987) Single-Unit Analysis of Post-natal Olfactory Learning: Modified Olfactory Bulb Output Response Patterns to Learned Attractive Odors. *J. Neurosci.*, **7**, 3154-3162.

Yu, D., Ponomarev, A. & Davis, R. L. (2004) Altered representation of the spatial code for odors after olfactory classical conditioning; memory trace formation by synaptic recruitment. *Neuron*, **42**, 437-449.

2.2 Different Subtypes of Striatal Neurons Are Selectively Modulated by Cortical Oscillations

The following article and the corresponding supplemental materials have been published as:

A. Sharott, C. K. E. Moll, G. Engler, M. Denker, S. Grün, and A. K. Engel. Different subtypes of striatal neurons are selectively modulated by cortical oscillations. *J Neurosci* 29, 4571-4585 (2009).

The published article is available at <http://www.jneurosci.org>
doi: [10.1523/JNEUROSCI.5097-08.2009](https://doi.org/10.1523/JNEUROSCI.5097-08.2009)

Behavioral/Systems/Cognitive

Different Subtypes of Striatal Neurons Are Selectively Modulated by Cortical Oscillations

Andrew Sharott,¹ Christian K. E. Moll,¹ Gerhard Engler,¹ Michael Denker,² Sonja Grün,² and Andreas K. Engel¹

¹Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, 20246 Hamburg, Germany, and ²Theoretical Neuroscience Group, RIKEN Brain Science Institute, Wako-Shi, 351-0198 Saitama, Japan

The striatum is the key site for cortical input to the basal ganglia. Cortical input to striatal microcircuits has been previously studied only in the context of one or two types of neurons. Here, we provide the first description of four putative types of striatal neurons (medium spiny, fast spiking, tonically active, and low-threshold spiking) in a single data set by separating extracellular recordings of sorted single spikes recorded under halothane anesthesia using waveform and burst parameters. Under halothane, the electrocorticograms and striatal local field potential displayed spontaneous oscillations at both low (2–9 Hz) and high (35–80 Hz) frequencies. Putative fast spiking interneurons were significantly more likely to phase lock to high-frequency cortical oscillations and displayed significant cross-correlations in this frequency range. These findings suggest that, as in neocortex and hippocampus, the coordinated activity of fast spiking interneurons may specifically be involved in mediating oscillatory synchronization in the striatum.

Introduction

The striatum is the major source of input to basal ganglia networks and receives afferents from nearly all cortical areas (Bolam et al., 2000). The entrainment of striatal neurons to thalamocortical rhythms suggests a high sensitivity of these networks to global forebrain dynamics (Gervasoni et al., 2004). Striatal projection and interneurons recorded from awake rats are entrained by slow-wave, high-voltage spindle and theta oscillations from the cerebral cortex and hippocampus in a topographical manner (Berke et al., 2004; Dejean et al., 2007). Phase locking of single striatal neurons to high frequencies has not been demonstrated (Berke et al., 2004).

Medium spiny neurons (MSNs), the projection neurons of the striatum, are GABAergic and comprise ~97% of the rat striatum (Rymar et al., 2004). They receive both cortical and thalamic synaptic inputs on the heads of their dendritic spines (Groenewegen and Berendse, 1994; Bolam et al., 2000). The remaining striatal neurons are made up of four different types of spiny interneurons. “Giant” aspiny interneurons that release acetylcholine are generally agreed to correspond to the tonically active neurons (TANs) recorded extensively in primates *in vivo* (Tepper and Bolam, 2004). The unique ion channel properties of TANs, as defined *in vitro*, suggest that these characteristic features of their spike trains could be intrinsically generated (Bennett et al., 2000; Wilson, 2005).

The remaining striatal neurons are GABAergic interneurons that can be differentiated neurochemically. Parvalbumin-positive neurons receive input from diverse cortical areas, suggesting a role in integration of cortical inputs in the striatum (Ramanathan et al., 2002). Electrophysiologically, they resemble the fast spiking interneurons (FSIs) found in the cerebral cortex and hippocampus (Kawaguchi and Kondo, 2002; Somogyi and Klausberger, 2005), where their coordinated activity underlies the entrainment of networks to global low-frequency oscillations and to local high-frequency oscillations (Klausberger et al., 2003; Bartos et al., 2007; Tukker et al., 2007). Although there is evidence that putative FSIs in the striatum are strongly entrained to low-frequency (<10 Hz) cortical oscillations (Berke et al., 2004), it is unclear whether this property is specific to FSIs and whether it also applies to higher frequency activities. Anatomical evidence suggests that the synchronized firing of these neurons could coordinate striatal activity (Kawaguchi, 1993; Ramanathan et al., 2002). GABAergic interneurons without parvalbumin, but expressing neuropeptide Y and somatostatin, can be differentiated electrophysiologically *in vitro* as they display low-threshold calcium spikes (LTSs) (Tepper and Bolam, 2004). A third type of GABAergic interneuron, expressing the calcium binding protein calretinin, has also not been characterized electrophysiologically *in vivo*, although there is anecdotal evidence that they may also display LTSs (Tepper and Bolam, 2004).

Here, we demonstrate that extracellular recordings of striatal neurons under halothane anesthesia can be convincingly separated into four putative subtypes. This enables us for the first time to separately analyze and compare their relationship to spontaneously occurring high- and low-frequency oscillations. The results provide strong evidence that putative GABAergic interneurons could control the propagation of oscillations in the corticostriatal pathway.

Materials and Methods

Animal preparation. The experiments were performed on nine male Brown-Norway rats weighing 200–300 g. Precautions were taken to

Received Oct. 19, 2008; revised Feb. 24, 2009; accepted March 4, 2009.

This work was supported by Bundesministerium für Bildung und Forschung Grant 01GQ01413 (Bernstein Center for Computational Neuroscience Berlin), Volkswagen Foundation Grant Az. 79342, Stifterverband für die Deutsche Wissenschaft Grant H140 5405 5024 111038, and European Union Marie Curie Training Network Programme Grant MRTN-CT-2005-019247. We thank Dr. Johannes Sarnthein for advice on burst analysis.

Correspondence should be addressed to Dr. Andrew Sharott, Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Martinistrasse 52, 20246 Hamburg, Germany. E-mail: asharott@uke.de.

DOI:10.1523/JNEUROSCI.5097-08.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/294571-15\$15.00/0

avoid pain and stress in accordance with the German law for the protection of animals and National Institutes of Health *Guidelines on the Care and Use of Laboratory Animals*. The animals were kept under controlled environmental conditions (ambient temperature, 20°C; 12 h light/dark cycle). A standard diet and water were allowed *ad libitum*. After initial anesthesia with ketamine (20 mg/kg, i.p.) and xylazine (5 mg/kg, i.p.), animals were tracheotomized, artificially respirated (70% O₂/30% N₂O/0.5–1% halothane), and fixed in a stereotactic frame. After trepanation and removal of the dura, the cortex was exposed bilaterally to allow the placement of two hydraulic microdrives (FHC). To ensure stability of anesthesia, electrocardiograms were constantly monitored and all animals were respired at a rate adjusted to maintain an endexpiratory CO₂ level of 3.5–4% (small animal ventilator, model 683; Harvard Instruments) as measured by a CO₂ analyzer (Capstar 100; CWE). Body temperature was monitored and maintained at 36.5–37.5°C using a homeothermic blanket (Fine Science Tools). Whereas halothane level was kept at 1% during all preparatory manipulations, it was kept constant at 0.5–0.7% throughout the recording sessions. At the end of each experiment, the animals were killed with a lethal overdose of pentobarbital (200 mg/kg, i.p.) and electrolytic lesions (20 μ A/20 s) were made to mark the recording sites within the respective target structure. Subsequently, rats were perfused transcardially with saline and then with buffered formalin. To verify the recording locations, fixed brains were sectioned and stained for Nissl and acetylcholinesterase (supplemental Fig. 1A, available at www.jneurosci.org as supplemental material). Pilot experiments were performed under ketamine/xylazine anesthesia (100 mg/kg ketamine/10 mg/kg xylazine) using the procedure, with the exception that the animals were breathing freely. Additional dosages of ketamine were given at regular intervals (every 60–90 min).

Recording and data acquisition. In each hemisphere, four tungsten electrodes (FHC; impedance, 200–800 k Ω) were independently targeted to the desired recording site. Distances between electrodes varied between 300 and 1000 μ m. Recording sites were sensorimotor cortex (Cx) and the caudate-putamen complex (CPU). All striatal cells were recorded from a limited region of the CPU (1.0 mm anterior to 1.5 mm posterior to bregma), on both medial and lateral sides of the striatum (2.5–4.0 mm lateral to bregma). The recording depth varied depending on the desired target structure from 0 to 2.0 mm (Cx) and 2.5 to 6.0 mm (CPU), respectively. Electrodes were left in place for a minimum of 5 min to allow stabilization before recording. Neural activity that could not unequivocally be attributed to either cortex or striatum was not included in this study. During the course of each experiment, neural activity was recorded from two or three different trajectories varying in the rostrocaudal plane.

Additionally, electrocorticographic (ECoG) recordings were performed using two skull screws placed above both ipsilateral and contralateral sensorimotor cortices (2.7–3.0 mm anterior to bregma, 2.0 mm lateral). A reference electrode was placed on the skull 4 mm anterior to bregma. Signals from each electrode were preamplified, amplified and differentially filtered (Multi Channel Processor; Alpha Omega). From the broadband recordings (1 Hz to 5 kHz), we extracted single-cell action potentials (single-unit activity), multiple-unit spike activity (multiunit activity) (both bandpass filtered 500 Hz to 5 kHz), and local field potentials (LFPs) (bandpass filter, 1–100 Hz). Whereas ECoG waves were used as correlates of a coherent synaptic activity in cortical neurons, striatal LFPs were used as such in striatal neurons. In some of the recordings, a notch filter was used to eliminate line noise at 50 Hz, and we therefore consistently excluded the bins at 50 Hz from the spectral analysis of these signals. Additionally, we verified in the ECoG and LFP that potentially induced phase jumps across the notch filter frequency were small (<10°) by estimating the average phase differences between neighboring frequency bins of their Fourier spectrum in nonoverlapping sliding windows of 10 s.

Recordings used for analysis had a minimum length of 184 s (mean, 782 \pm 33 s) each with between one and eight recorded single units. Recordings with obvious artifacts were discarded. Bipolar derivations of LFPs were used only when two electrodes were in the striatum of the same hemisphere. Because LFPs were recorded on only two electrodes per hemisphere, some unit recordings did not have a corresponding bipolar LFP. Single units from the same electrode were not used in any correlation-based analysis because of possible shadowing effects (Bar-

Gad et al., 2001). Unless stated otherwise, units from both hemispheres were included in all analyses. All data analysis was performed by using Matlab and its Signal Processing Toolbox (Mathworks).

Spike detection and characterization of unit subtypes. Spike detection was performed off-line by a level crossing algorithm. The threshold value was set sufficiently high above the noise level to avoid artificial triggers by noise peaks (mean threshold \pm SD, 7.4 \pm 2.7 σ). When possible, multiple single units were separated on the basis of several waveform parameters, including principal components, signal energy, peak time, and other distinguishing features (Offline-Sorter; Plexon). Spike durations were obtained by reading the time difference between cursors placed at the points at which the averaged spike waveform departed and returned to baseline. A mean of 3686 digitized spikes was averaged to characterize the spike waveform for a given cell.

After spike sorting, the width of each phase of each spike was calculated using a peak finding algorithm. Spikes were then clustered using three parameters related to the length of the waveform: the length of the total initial deflections, the length of the valley, and the sum of these two parameters. Similar parameters have been previously used to discriminate putative striatal projection neurons and interneurons (Berke et al., 2004). Clustering was performed using a K-means supervised clustering algorithm and was performed using both a value for each record and for each neuron. As waveform shapes remained highly stable across recordings (i.e., when other electrodes were moved), this made little difference to the result. Final classification used clustering using the individual records, which slightly improved the estimate.

Detection of LTS bursting. All spike trains were examined for evidence of LTS bursts. To characterize the types of bursts in the data, we used the methods of Jeanmonod and colleagues (Jeanmonod et al., 1996; Sarnthein and Jeanmonod, 2007), which examine three main criteria. In line with these studies, a burst here was defined simply as two or more spikes with interspike intervals (ISIs) of <10 ms. First, a standard ISI histogram for LTS neurons should be bimodal and have a peak <10 ms. Second, the duration of ISI was tested for a positive correlation with the ordinal of the ISI in each burst for bursts of comprising two to five spikes. Third, in the majority of bursts, the second spike should be smaller than the first. If 65% of 50 randomly selected bursts displayed this characteristic, this third criterion was considered to be met. Only neurons that passed all three criteria were considered LTS positive.

Cross-correlation analysis of spike trains. In recent years, analysis of the oscillatory properties and correlation between multiple neurons is increasingly performed directly in the frequency domain (Rivlin-Etzion et al., 2006; Witham et al., 2007). Here, we chose to analyze unit activity first in the time domain, to avoid the assumption that the correlations between neurons would be oscillatory. To this end, we compared the maximum correlations, number of significant correlations, and the Fourier transform of the cross-correlation to characterize the absolute correlation, the likelihood of correlation, and the oscillatory content of the correlation, respectively.

Raw cross-correlations were calculated for each pair of spike trains and normalized by removing the mean correlation across the lag range. The aim of the analysis was to detect pairs of spike trains in which the number of coincidences and/or oscillatory properties were significantly different than those that would be predicted by the first-order characteristics of the spike train (i.e., rate and interspike interval distribution). To this end, cross-correlations were calculated using surrogate spike trains constructed by globally shuffling the interspike intervals of both neurons and calculating their correlation 1000 times. This produced a null hypothesis distribution for each lag point. The real correlation was then converted to a *t* score (the number of SDs of the true correlation from the mean of the null hypothesis) that was used as a measure of the correlation strength, as it is dependent mainly on the temporal locking of the two spike trains. The maximum absolute *t* score within the center of the cross-correlation (0 \pm 0.25 s) was used for statistical comparison. The absolute score indicates how strongly the neurons influenced each other, regardless of the sign. A cross-correlation was considered significant at a given lag if it was outside 2 SDs of the null hypothesis (again, significantly positive and negative correlations are grouped together). This criterion was used to construct significance histograms (see Figs. 5–7), which were used to

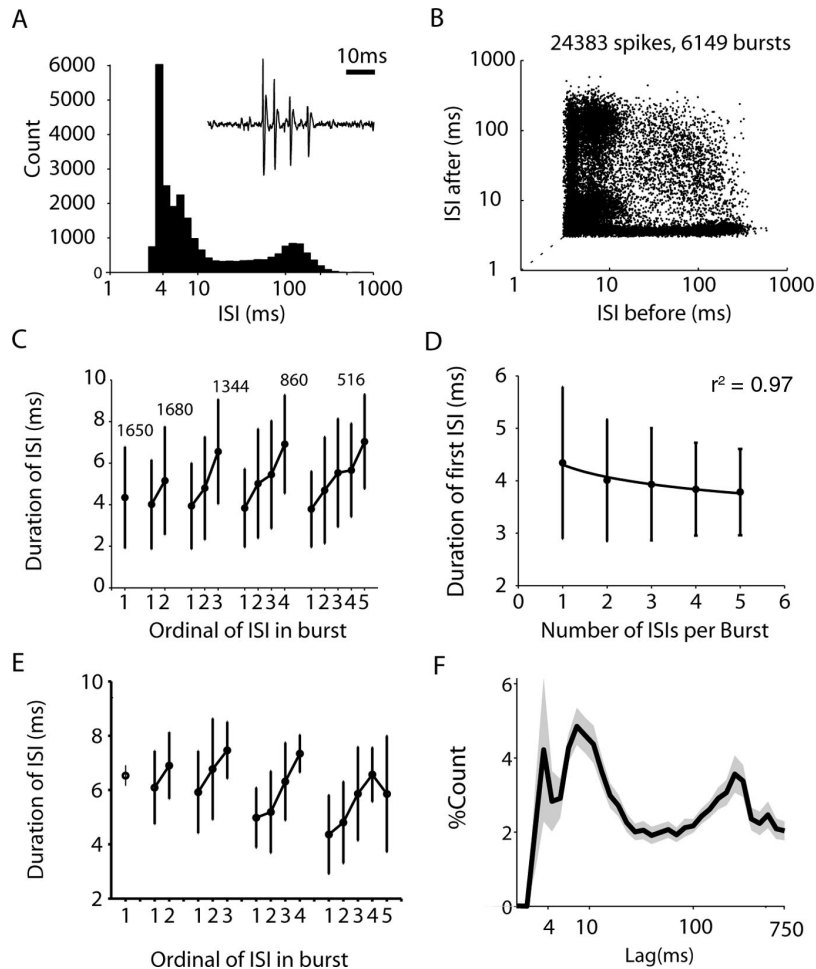


Figure 1. LTP bursting in striatal neurons under halothane anesthesia. **A–D**, Burst and interspike interval analysis of a putative parvalbumin negative interneuron displaying LTP bursts. **A**, Interspike interval histogram of a single LTP-positive neuron. Note the bimodal distribution with a peak < 10 ms. The inset shows an example of a single burst from this neuron. **B**, The ISI return plot of the same neuron shows asymmetry at the diagonal < 10 ms showing that the interburst interval is not constant. This is elaborated in **C**, which demonstrates that the interspike interval positively correlates with its ordinal in the burst shown for bursts consisting of up to five ISIs. **D**, The length of the first interspike interval decreases together with the number of spikes in the burst. **E**, **F**, Plot of the mean ordinal (**E**) and ISI histogram (**F**) showing that these features are consistent in all the neurons demonstrating LTP bursting. The shaded areas in **F** show the SEMs across neurons.

investigate the likelihood of significant correlation between a specific pair type at a given lag.

Correlation strength was calculated in this way for each pair of spike trains using two sets of parameters. To maximize the detection and visualization of low- and high-frequency activities, a 2 s and a 0.2 s window were used, respectively. In cross-correlation-based analysis, large-amplitude common elements at low frequencies tend to dominate (Challis and Kitney, 1990). To test for any correlation between spike trains at higher frequencies, cross-correlations calculated with the smaller window were high-pass filtered at 30 Hz. To smooth the low-frequency activities, a 10 Hz low-pass filter was also used for some analyses. In both cases, a two-pole Butterworth filter was applied. Additionally, the Fourier transform (as opposed to windowed power spectral estimate) is vulnerable to outlying values, a multitaper approach, with three tapers, was used to stabilize the estimate. For analysis of low- and high-frequency activity, the center 0.5 and 0.3 s of the cross-correlation were used, respectively. This process was performed on the *t* score, which ensured that

the amplitude of the oscillations in the cross-correlation were the result of temporal correlations above those expected because of the primary spiking statistics.

Unless otherwise stated, statistical comparisons between different putative neuron groups were calculated using Wilcoxon's rank sum test with Bonferroni–Holmes step down corrections for multiple comparisons. For selected analysis, results that were significant without the corrections are indicated with one asterisk and those significant with the correction are indicated with two asterisks (for details, see individual figure legends).

Spectral analysis of LFPs. For LFPs, all spectral parameters were computed using 0.25 Hz resolution [4040 points/fast Fourier transform (FFT) window] to allow separation of low-frequency activities, and with a Hanning window to prevent spectral leakage. Coherence was computed using the same parameters and the square root of the coherence was Fisher transformed to normalize the variance before any averaging or statistical analysis (Rosenberg et al., 1989). Coherence was calculated between the ECoG and the raw, monopolar LFP and the bipolar derivation of the LFPs in the same hemisphere.

The coherence between the ECoG and the bipolar derivation should provide a measure of coupling relatively free from volume conduction and global noise. However, it is possible that coherence after bipolar derivation could be contaminated by subtraction artifacts or, conversely, underestimates the real coherence. To try to counter these problems, the phase spectra was calculated for the ECoG to each monopolar and bipolar pair, together with 95% confidence limits as described by Halliday et al. (1995). A detailed investigation of the phase relations between cortex and striatum is beyond the scope of this study, as the phase delay is ambiguous in complex systems, unless used together with complementary techniques (Cassidy and Brown, 2003). Therefore, we used confidence limits of the phase spectra to establish the incidence of nonzero phase coupling across the frequency range. By definition, if at a given frequency the number of phase values with confidence intervals excluding 0 and $\pm 180^\circ$ (which could be attributable to an arbitrary polarity reversal) greatly exceeds 5% (the proportion that would be expected by chance), the data would be incompatible with a situation in which volume conduction or noise are responsible for all of the significant coherence (Magill et al., 2006). In the case of bipolar derivations, these results would also argue against coherence caused by subtraction artifacts (Magill et al., 2006). This was statistically evaluated by using Fisher's exact test to compare the sum of nonzero values across a given frequency bands to the number that would be expected by chance for all the ECoG-striatal LFP pairs within each animal.

To investigate the fundamental frequencies in the ECoG and LFP data, we used the correlation matrix method previously described by Masimore et al. (2004) that has successfully been applied to basal ganglia LFP data (Fogelson et al., 2005; Masimore et al., 2005). Briefly, power spectral parameters were calculated in short time windows (10 s/5 s overlap or 12 s/6 s overlap). The correlation coefficient was calculated for all values across time for each frequency bin against all other frequency bins from 0.5 to 100 Hz with a frequency resolution of 1 or 0.25 Hz to examine high and low frequencies, respectively. Correlations that were not significant after Bonferroni's correction were then set to zero. Significant correla-

tion that were not significant after Bonferroni's correction were then set to zero. Significant correla-

tions were separated into positive and negative correlations and Fisher transformed before averaging. Positive/negative separation was performed to clarify visualization and reveal whether there were both positive and negative correlations in different records, which would not be clear after averaging (averaged significant positive and negative correlations could result in a false zero). Correlation matrices for each record were first averaged within animals and subsequently across animals. Correlations therefore had to be highly consistent within and between animals to result in a high coefficient in the final analysis. Correlation matrices provide two types of information. First, patches on the diagonal demonstrate covariation in power over time, indicating that these frequencies form a functional band. Second, significant values outside the diagonal show correlation between the involved bands (for more details, see Results).

Analysis of spike-field phase locking. Oscillatory modes of field potentials are commonly considered as an indicator of coordinated network dynamics (Elul, 1971; Mitzdorf, 1985; Logothetis and Wandell, 2004). To explore how the spiking activity of identified neurons is related to this rhythmic activity, we evaluated the phase relationship between the ECoG/LFP signal and the spiking activity of single units. To this end, we expanded on current methods originating from phase synchronization analysis (cf. Varela et al., 2001; Friedrich et al., 2004) to study the instantaneous phase of the field potentials at spike times.

ECoG and LFP signals were filtered using a neutral-phase bandpass filter (sixth-order Butterworth), in which the width of each of the 50 defined pass bands (centered on frequencies from 2 to 126 Hz) was chosen as $\pm 15\%$ of the respective center period (exemplified in Fig. 4A for three different frequencies). In a subsequent step, the instantaneous phase of the field potentials was calculated from the analytic signal obtained via the Hilbert transformation (for details, cf. Lachaux et al., 1999; Le Van Quyen et al., 2001). In this formalism, troughs of the LFP correspond to a phase of 180° . The calculation of the analytic signal can be applied to arbitrary signals, but its interpretation as instantaneous phase is difficult if the signal contains time periods in which either the amplitude becomes too small to discriminate the oscillation from background noise, or in which the regular oscillation is disrupted (Boashash, 1992). To account for these effects, we discarded all calculated phase values that violated the monotonicity of the phase time series or exhibited instantaneous phase jumps.

We then analyzed the distributions of extracted phase values (Denker et al., 2007) using tools from circular statistics (Mardia and Jupp, 2000). In particular, we obtained the mean phase ϕ of the circular average,

$$Re^{i\phi} = \frac{1}{N} \sum_i e^{i\phi(t_i)}$$

where $\phi(t_i)$ indicates the phase of the field potential at time t_i of spike i . In addition, we used

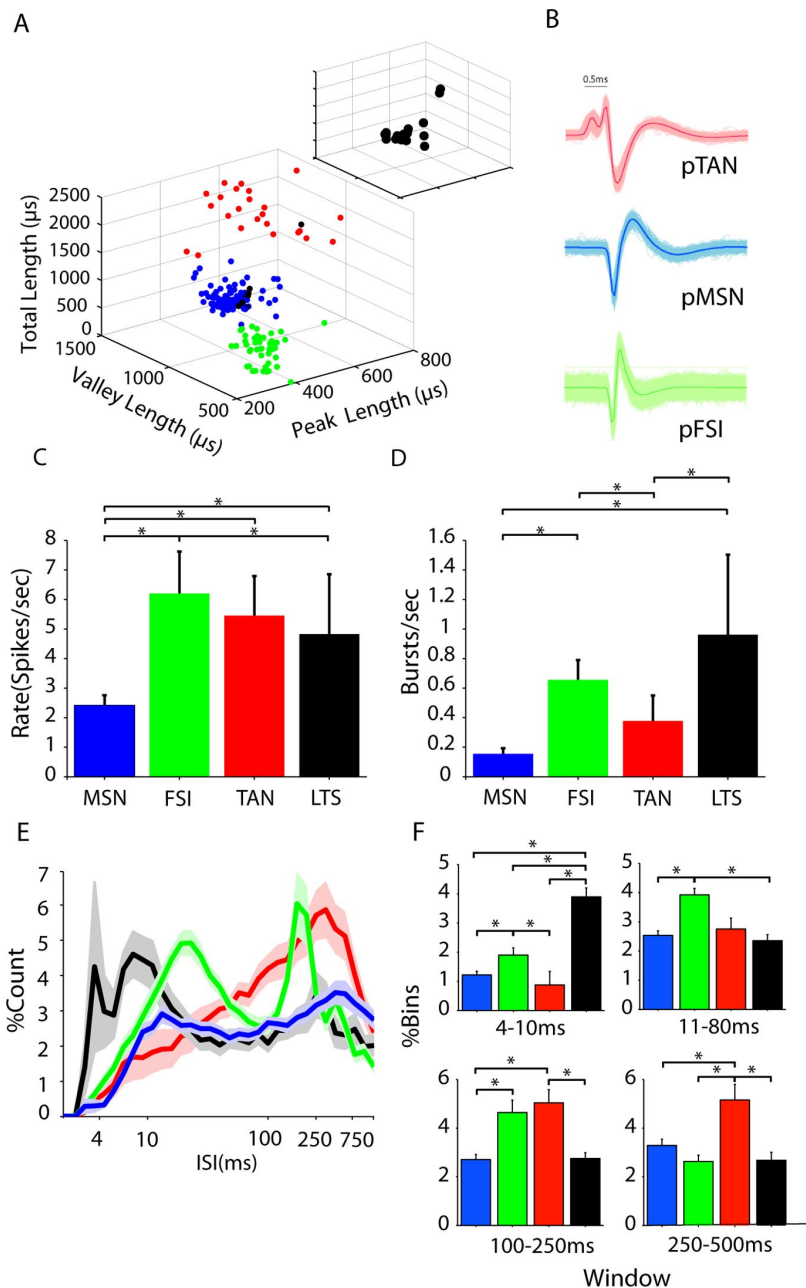


Figure 2. Separation of striatal neuron subtypes based on waveform parameters and LTS bursting. **A**, The parameters of mean waveforms for each neuron were plotted against each other, revealing three clear clusters, which were verified using a clustering algorithm. **B**, The clusters corresponded to three types of waveforms, which could be distinguished by eye after sorting. Long waveforms (red cluster/waveform) and very short waveforms (green cluster/waveform) correspond to previous descriptions of cholinergic neurons, usually referred to as tonically active neurons (pTANs) and parvalbumin-positive GABAergic interneurons, usually referred to as fast spiking interneurons (pFSIs), respectively. The majority of neurons had a waveform between these two extremes (blue cluster/waveform) and were thought to be putative medium spiny projection neurons (pMSNs). With one exception, all of the neurons identified as LTS positive (black spheres) were in this middle cluster (**A**) and are also shown separately (inset). **C**, **D**, Histograms of the rate and burstiness of putative neurons; color codes are as above. **C**, The rate of all putative interneuron types was significantly higher than that of the pMSNs. **D**, pLTS units had a significantly larger burst rate than all other putative neuron types. In addition, pFSIs were burstier than pTANs and pMSNs. **E**, The normalized interspike interval histograms averaged across all putative neuron types showed clear variations between the subtypes. **F**, LTS neurons had significantly more ISIs < 10 ms than any other type. In the other time windows analyzed, there were significant differences between all putative subtypes. Error bars indicate SEM. Asterisks indicate significance with Bonferroni–Holmes correction.

the transformation of the vector strength R to the circular SD $\sigma = \sqrt{-2 \log R}$ as a measure of the concentration of the phase distribution. For small values of σ , this measure relates to the SD of a normal distribution, whereas for flat distributions $\sigma \rightarrow \infty$.

We quantified whether the spikes recorded from each neuron showed a significant phase preference using two distinct criteria. First, we tested against the null hypothesis that the phase sample was taken from the uniform circular distribution (Rayleigh's test), which would be expected assuming a regular (e.g., filtered) field potential and an independent spike train with a Poisson spiking statistics. However, deviations from this assumption toward more regular spike trains as observed in this study could impose a certain degree of chance locking between the two signals. To examine whether phase distributions showed a stronger phase preference than expected given their underlying spiking statistics, we tested the measured value of the vector strength R against the distribution of R obtained from 1000 surrogates constructed from the original data. For each surrogate, we locally shuffled the interspike intervals of the recorded spike train in nonoverlapping windows of 10 s, while keeping the field potential trace unchanged. By doing so, we aimed to destroy the phase relation to the field potential, while retaining to first order the regularity of the spike train and respecting the overall rate profile. Unless otherwise stated, we say that a neuron has a nonuniform phase distribution if it passes the first criterion based on the Rayleigh's test, and that it is locked with respect to the surrogates if it passes the second, more restrictive criterion.

Results

Detection of low-threshold spike bursts allows separation of putative parvalbumin-negative GABAergic interneurons (pLTS)

LTS bursting is a common feature of thalamic neurons during sleep (Linás and Steriade, 2006) and can be reliably detected in extracellular recordings (Jeanmonod et al., 1996; Lacey et al., 2007). *In vitro* recordings have shown that, in the striatum, only parvalbumin-negative GABAergic interneurons display LTS bursting in response to current injection (Tepper and Bolam, 2004). Figure 1, *A–D*, shows an example of a single neuron displaying all the key features of LTS bursting. These include a bimodal ISI histogram with a peak <10 ms, a positive correlation between the ISI and ordinal of the ISI in the burst, and a negative correlation between the first ISI and the size of the burst. In total, 15 neurons (21 recordings) in the sample matched the two key markers of LTS (i.e., an increase in ISI with the ordinal of the spike in the burst and a significant number of ISIs <10 ms). Figure 1, *E* and *F*, shows that these properties were highly consistent within the LTS-positive population. In addition, these neurons had a high percentage of bursts with a smaller second spike than the first ($83 \pm 10\%$ SD). LTS-positive neurons were recorded in five of nine animals across a range of depths. These LTS-displaying neurons were considered to be putative parvalbumin-negative GABAergic interneurons (pLTS). Recordings of pLTS, and all other neuron types, were highly stable across hundreds of seconds (supplemental Fig. 1C, available at www.jneurosci.org as supplemental material), and bursts were therefore unlikely to be the result of injury discharge or other phenomenon related to recording conditions.

Clustering of unit waveforms allows separation of three unit subtypes verified by physiological criteria

Previous studies have described the different waveform properties of different striatal neurons that have been subsequently labeled (Wilson et al., 1990; Kawaguchi, 1993; Mallet et al., 2005). We plotted three waveform parameters from all recorded neurons: the peak, initial first deflections, the valley and the sum of the two. These parameters yielded three clusters, which could be clearly separated by a supervised clustering algorithm (Fig. 2*A*). The long waveforms (red) of one of the clusters correspond to previous descriptions of giant cholinergic interneurons, often called TANs, and we therefore consider these cells to be putative

Table 1. Statistical comparisons of rate and burst parameters

	Cell type	pFSI	pTAN	pLTS
Rate	pMSN	$p < 0.0001$	$p < 0.0001$	$p < 0.02$
	pFSI		NS	$p < 0.01$
	pTAN			NS
Burst (10 ms ISI)	pMSN	$p < 0.0001$	NS	$p < 0.00001$
	pFSI		$p < 0.0046$	$p < 0.005$
	pTAN			NS

Significance is noted only after Bonferroni–Holmes correction within each parameter.

Table 2. Statistical comparisons of normalized ISI histogram

Time window	Cell type	pFSI	pTAN	pLTS
4–10 ms	pMSN	0.0230	NS	<0.0001
	pFSI		0.0030	0.0001
	pTAN			0.0004
10–80 ms	pMSN	0.0000	NS	NS
	pFSI		NS	0.0001
	pTAN			NS
100–250 ms	pMSN	0.0007	0.0002	NS
	pFSI		NS	NS
	pTAN			0.0032
250–500 ms	pMSN	NS	0.0111	NS
	pFSI		0.0034	NS
	pTAN			0.0043

ISI histograms were calculated between 1 and 912 ms and normalized. Significance is noted only after Bonferroni–Holmes correction within each window.

neurons of this type (pTANs). Conversely, waveforms in the cluster with the shortest waveforms (green) match descriptions of parvalbumin-positive GABAergic interneurons, putative fast spiking interneurons (pFSIs). The majority of neurons had a waveform between these two extremes, conforming to previous reports of medium spiny projection neurons (pMSNs). The vast majority of waveforms for pLTS neurons were found in the central pMSN cluster, consistent with previous reports that their waveform is longer than FSIs (Kawaguchi, 1993) (Fig. 2*A*, inset). One pLTS waveform was in the pTAN cluster, but lacked the distinctive shape seen in the majority of pTAN waveforms (Fig. 2*A*). As there are no reports of LTS bursting in TANs, this may have been attributable to the electrode being particularly close to the soma (Gold et al., 2006), as indicated by the high signal-to-noise ratio in the recordings of that neuron. The consistency of the waveform shape across long periods of time (600–1500 s) could be clearly demonstrated in individual examples of each putative subtype in different animals (supplemental Fig. 2, available at www.jneurosci.org as supplemental material).

The second step was to verify the separation using known physiological properties of the neurons that are independent of waveform and LTS-related parameters. The mean rate of the pMSNs was significantly lower than all other putative subtypes (Fig. 2*C*, Table 1). In addition, pFSIs had a significantly higher rate than pLTS cells. Both pLTSs and pFSIs had more bursts per second than pMSNs and pTANs using the burst criteria applied for LTS analysis (Fig. 2*D*, Table 1). The mean firing rate of the pFSIs (~6 spikes/s), although on first reflection seems low, is faster than that observed in histologically identified FSIs recorded under ketamine anesthesia (Mallet et al., 2005) and within the range seen in the awake animal (Berke et al., 2004). Finally, the normalized interspike interval histogram was compared using four different windows (4–10, 11–80, 100–250, 251–500 ms). Clear differences could be seen in the profile of the ISI histogram for each putative subtype (Fig. 2*F*, Table 2). The ISI histogram for

pTANs was unimodal with a clear peak ~ 250 ms, consistent with a regular firing rate of 4–5 Hz as previously described in several preparations (Aosaki et al., 1995; Bennett et al., 2000; Apicella, 2007). In contrast, pFSIs consistently had a bimodal distribution suggestive of bursting with a longer interburst interval compared with LTS neurons with a peak at ~ 20 ms, consistent with subsequently labeled recordings under ketamine (Mallet et al., 2005) (Fig. 2D). *In vitro* studies have led to the suggestion that this is the likely firing pattern for FSIs *in vivo* (Tepper and Bolam, 2004). Significant differences were found between all combinations of putative neurons in at least one of these windows (Fig. 2D, Table 2).

Halothane anesthesia leads to multiple interrelated oscillatory activities

Spontaneous population activity in cortex and striatum, in the form of LFPs and ECoGs, has been extensively characterized under urethane and/or ketamine/xylazine anesthesia, but far less so using volatile anesthetics. We observed spectral properties of these population signals under halothane that markedly differed from those reported under ketamine-based anesthesia. Power spectra from striatal LFPs and ECoG displayed prominent peaks in the delta range (2–4 Hz) and the gamma range (40–80 Hz), which were highly consistent across animals (Fig. 3A–C). Both peaks were reduced, but clearly present in the bipolar derivations of both the striatal LFP power and the coherence between the ECoG and striatal LFP, suggesting that they were not the result of volume conduction between the two structures (Fig. 3A–C). This was supported by analysis of the phase spectra, which demonstrate that, even in the monopolar case, the percentage of coherence values with nonzero phase was well above chance across the frequency range from 1 to 100 Hz (supplemental Fig. 4, available at www.jneurosci.org as supplemental material). The bands with the highest monopolar coherence, 2–4 and 40–80 Hz (excluding 48–52 Hz, which could be affected by line noise), had the largest percentage of nonzero phase delays (~ 60 –70%), suggesting that the frequencies with the highest coherence were attributable to neuronal interaction and not volume conduction, which would be expected to have a delay incorporating 0° or $\pm 180^\circ$. For ECoG and bipolar-LFP pairs, the distribution of nonzero coupling was more evenly distributed across the frequency spectra, as previously shown for similar data (Magill et al., 2006). At these frequency ranges, the sum of

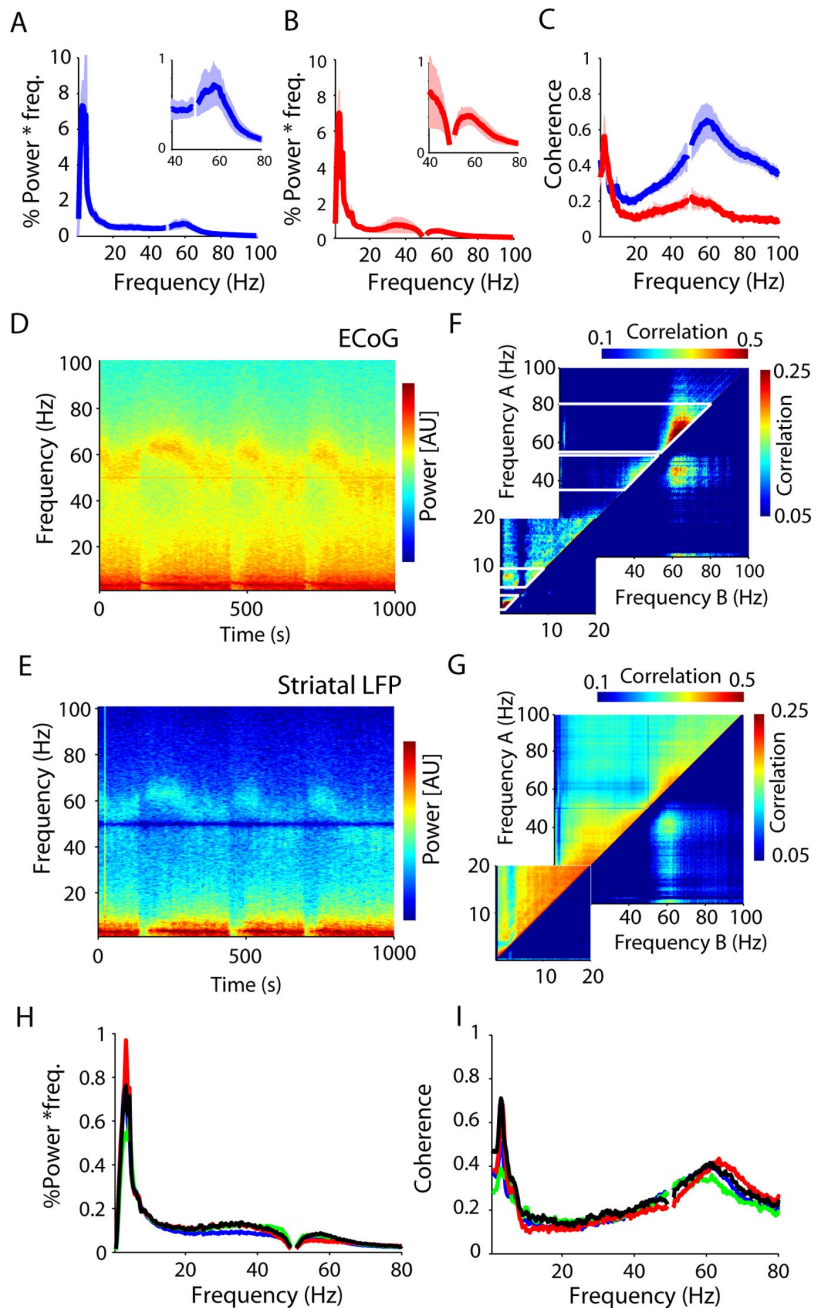


Figure 3. Spontaneous corticostriatal population activity under halothane anesthesia. **A, B**, The power spectra of the ECoG (**A**) and striatal LFP (**B**) have two prominent peaks at low (~ 4 Hz) and high frequencies (40–75 Hz). **C**, These activities are coherent between the structures whether using the monopolar (blue) or bipolar (red) striatal LFP. The shaded areas show SEMs across animals. **D, E**, Time–frequency analysis of power in the cortex (ECoG; top) and striatum (LFP; bottom) reveals that these peaks are often composed of more focal activities that wax and wane over time. **F, G**, The relationship across frequency bins can be described using correlation matrices in which the Pearson correlation coefficient is calculated for the fluctuations in power over time for each frequency with all other frequencies between 0.5 and 100 Hz. Hence the diagonal (each frequency correlated with itself) is always 1. The matrices show separately the positive correlations above the diagonal (top) and the absolute value of negative correlations below the diagonal (note that positive and negative correlations are coded by two distinct color scales). Patches above the diagonal show neighboring frequencies that covary over time and can therefore be considered as bands (marked by the white lines). Patches off the diagonal indicate positive and negative correlations between those bands. The matrix for the ECoG shows two individual bands between 35–55 and 60–80 Hz. Activity is negatively correlated between these two frequency ranges within the gamma band. **H, I**, The mean bipolar LFP power (**H**) and coherence (**I**) between the ECoG and LFP showed were similar between recordings containing each neuron type (blue, pMSN; green, pFSI; red, pTAN; black, pLTS).

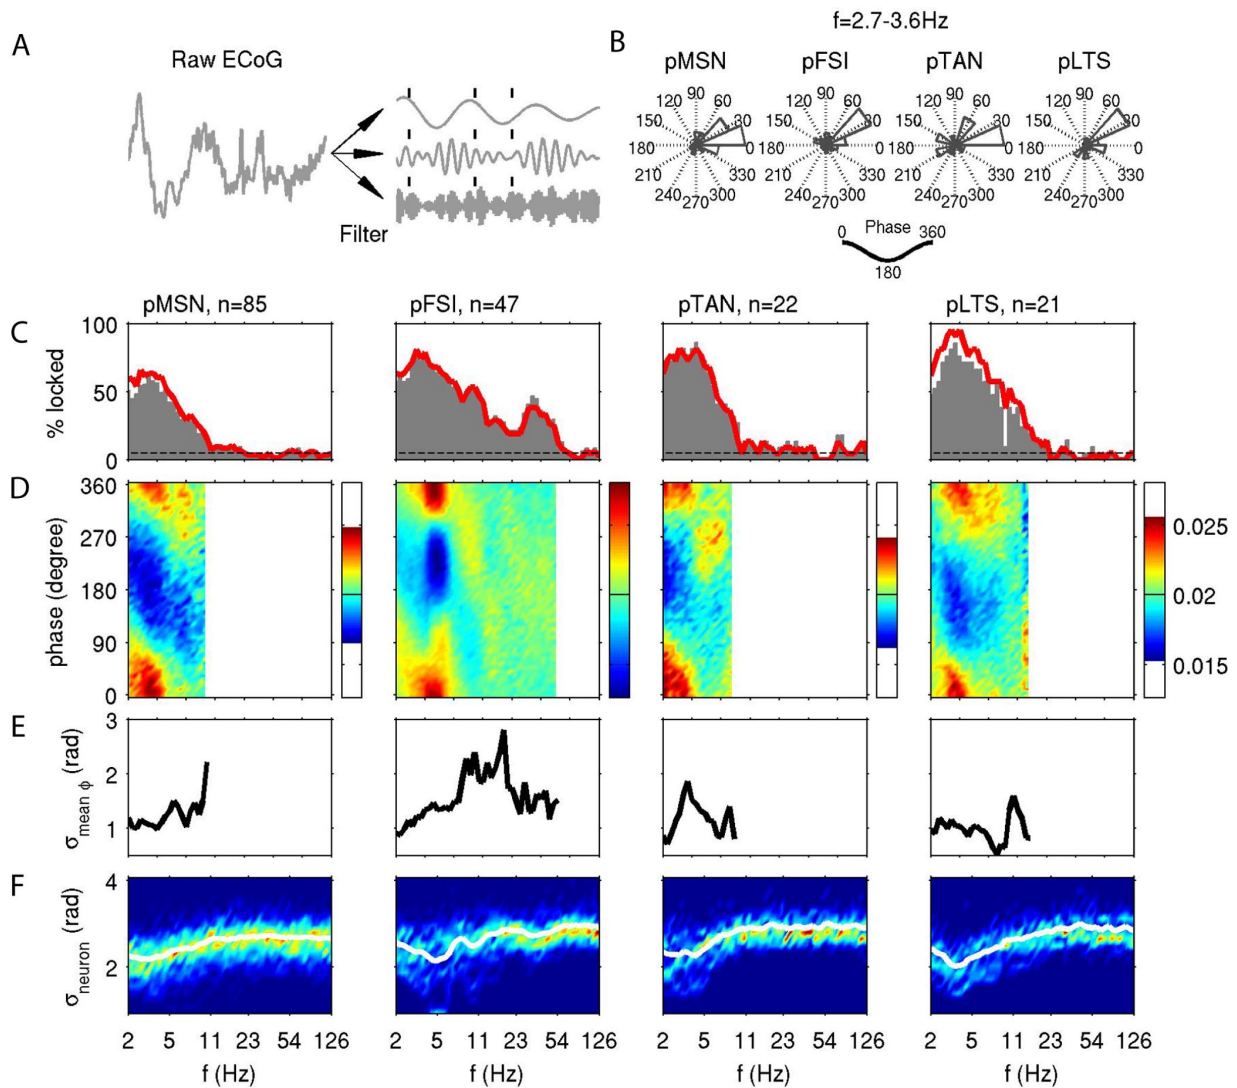


Figure 4. Phase locking of striatal neurons to ECoG oscillations. **A**, Sketch of the phase analysis. The raw ECoG signal (left) is filtered separately for different frequency bands, illustrated by three examples on the right. The instantaneous phases at spike occurrences (ticks) are assessed separately for each frequency band. **B**, Distribution of the mean phases of each neuron showing a nonuniform phase distribution ($p = 0.05$, Rayleigh's criterion) at a filter center frequency of 3.1 Hz (cutoff at 2.7 and 3.6 Hz). Here, and in the following panels, the distributions are shown separately for neurons classified as pMSN, pFSI, pTAN, and pLTS. **C**, Relative number of significantly ($p = 0.05$) locked neurons using the Rayleigh (red curve) and surrogate (gray histogram) criteria for significance evaluation, resolved by the filter center frequency. The expected percentage of false positives of 5% is indicated by the dashed lines. **D**, Frequency-resolved phase distributions computed from all spikes recorded from neurons that exhibited a nonuniform phase distribution ($p = 0.05$, Rayleigh's criterion). The color code of each graph is chosen to maximize the contrast by mapping maximum and minimum values to the same color in all graphs (red and blue, respectively). The absolute values for each graph can be read from the color bar on the right. For each frequency, the phase distribution was normalized to its sum across the 50 bins of the vertical phase axis, such that a flat distribution would show a probability of 0.02 per bin (marked by a black line in each color bar). Results are displayed for frequencies, in which $>10\%$ of neurons locked to the ECoG (compare **C**, red curve). **E**, Consistency of phase across neurons. Frequency-resolved estimate of the variability of the mean phase across neurons measured by its circular SD. Low values of this measure indicate consistency of the mean phase across neurons. As in **D**, for each frequency only neurons with a nonuniform phase distribution enter the analysis, and results are displayed for a minimum of 10% of the neurons showing such locking. **F**, Variability of locking strength across neurons. Histogram of the circular SDs of the phase distributions of each neuron resolved by frequency. Colors indicate counts per bin; their mean is indicated by a white curve. In contrast to **D** and **E**, all neurons enter the distribution regardless of whether they exhibit a nonuniform phase distribution.

nonzero values was significantly above that expected by chance for both monopolar and bipolar pairs within each individual animal (Fisher's exact test, $p < 0.0001$).

Spectral analysis of large data segments can mask fluctuations in dominant frequencies over time (Fig. 3*D,E*). To analyze such temporal relationships, we used a correlation matrix approach (Fig. 3*F,G*), which has been previously used to reveal interfrequency relationships in corticobasal ganglia networks (Masimore

et al., 2004; Fogelson et al., 2005). The correlation matrix uncovers two pieces of information. First, continuously significant bins around the diagonal show frequencies that covary over time. The matrix for cortical ECoG shows discrete bands in the lower frequencies at 2–4 and 5–9 Hz, and at high frequencies between 35–55 and 60–80 Hz (marked on Fig. 3*F*). The matrix for striatal LFP is very similar, although there is also a clear band around beta frequencies (10–30 Hz) that is not a prominent feature in the

cortex. Second, significant frequencies outside the diagonal show either positive or negative correlations between those bands identified on the diagonal. These will not be discussed in detail here, but we point out that the two prominent frequency ranges within the gamma band (35–55 and 56–80 Hz, respectively) are anticorrelated in their power, providing additional evidence that they originate in discrete activities.

A central aim of this work was to compare the oscillatory correlation between different unit subtypes. This comparison would be more meaningful if the oscillatory conditions at the level of the cortical and striatal population were not significantly different during the recording of the different neuronal types. To this end, the power and coherence of the different recordings was compared in the frequencies identified by the correlation matrices based on which units were in each recording [delta (2–4 Hz), theta (5–9 Hz), low gamma (40–55 Hz; excluding 48–52 Hz, which could be affected by line noise), and high gamma (56–80 Hz)]. After Bonferroni–Holmes correction for number of frequency band comparisons, no significant difference was found in the normalized ECoG power or monopolar striatal LFP in any frequency range ($p > 0.05$). There was also no difference in the normalized bipolar LFP power, arguably the most accurate measure of the population activity in the striatum, without any correction for multiple comparisons (Fig. 3H) (rank sum test, $p > 0.05$). Finally, no significant difference was found in the coherence between ECoG and striatal monopolar and bipolar LFPs (Fig. 3I). It therefore seems reasonable to conclude that differences between the oscillatory properties of different neuron subtypes were unlikely to reflect global state changes that can be caused by factors such as variations in the level of anesthesia across different recordings.

Phase locking of striatal neurons to cortical oscillations

We examined the phase relationship between the spiking activity of individual neurons and the oscillatory ECoG dynamics (Fig. 4A). For each neuron, we extracted the instantaneous phase at the time of spike occurrences and calculated its mean. Figure 4B shows the distributions of mean phases of the population of neurons that showed a significant nonuniform phase distribution ($p < 0.05$). Results are separated by the putative neuronal subtype, and relate to the low delta frequency band. At this frequency, all neuron types showed a strong tendency to spike shortly after the peak of the ECoG oscillation cycle, regardless of the strength of the observed phase

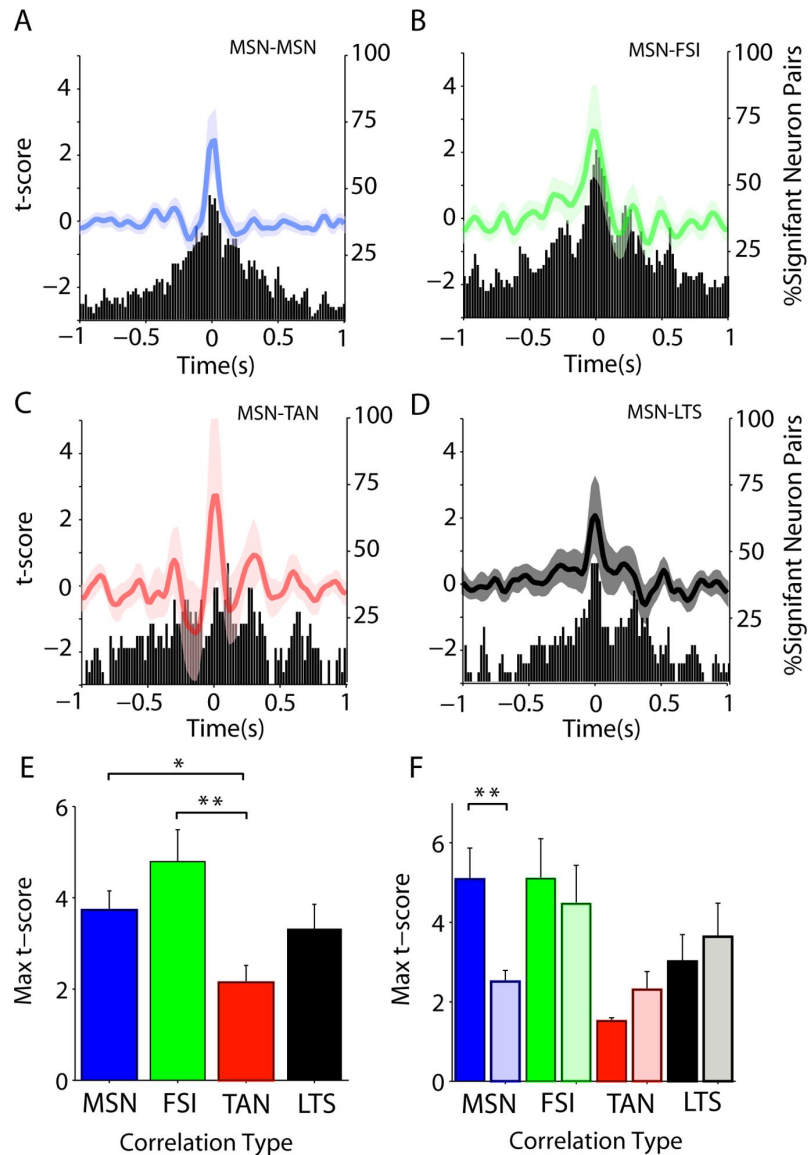


Figure 5. Slow correlation between pMSNs and putative interneurons. Correlation between different putative interneurons and pMSNs was assessed using cross-correlation. For all panels, the color code for pair types is the same: blue, pMSN–pMSN; green, pMSN–pFSI; red, pMSN–pTAN; black, pMSN–pLTS. **A–D**, The t score is the number of SDs from the surrogate correlation distribution calculated from the same neurons with shuffled interspike intervals. The mean t score for each pair type is shown by the lines (scale on left), with the shaded areas showing the SEM. The histogram shows the percentage of neurons in which 2 SDs were exceeded at a given lag regardless of the sign of the correlation. More significant correlations than would be expected by chance were observed between pMSNs and other pMSNs (**A**), pFSIs (**B**), pTANs (**C**), and pLTS cells (**D**). **E**, The maximum t score (between ± 250 ms) of the correlation was significantly higher for pMSN–pFSI correlations than for pMSN–pTAN and pMSN–pLTS correlations. **F**, When the strength of ipsilateral (dark bars) and contralateral (light bars) correlations was compared, only pMSN pairs showed a significant difference with correlations being stronger between ipsilateral pairs. Error bars indicate SEM. Asterisks indicate significance without (*) and with (**) Bonferroni–Holmes correction.

locking. For the putative GABAergic interneurons, pFSI and pLTS, a majority of neurons favored a mean spiking phase that was slightly advanced by ~ 20 – 30° from that of either pMSNs or pTANs, although significance of this finding could not be established because of the small sample sizes.

We quantified the relative number of neurons that were locked to the ECoG as a function of frequency using two distinct criteria. In a first step, we tested against the null hypothesis that

Table 3. Statistical comparisons of low-frequency power in pMSN cross-correlations

Correlation type	0–2 Hz	2–4 Hz	5–9 Hz
pMSNpMSN/pMSNpFSI	$p < 0.0005$	NS	NS
pMSNpMSN/pMSNpTAN	$p < 0.008$	$p < 0.007$	$p < 0.002$
pMSNpMSN/pMSNpLTS	NS	NS	$p < 0.01$
pMSNpFSI/pMSNpTAN	$p < 0.00001$	$p < 0.0033$	$p < 0.005$
pMSNpFSI/pMSNpLTS	NS	NS	NS
pMSNpTAN/pMSNpLTS	$p < 0.00005$	$p < 0.02$	$p < 0.0001$

the phase sample was taken from the uniform circular distribution (Rayleigh's test), which would be expected assuming a regular LFP and an independent spike train with a Poisson spiking statistics (Fig. 4C, red curve). We observed that, for each of the neuronal subtypes, 60–80% of the respective neurons showed significant phase locking to the low-frequency component with a peak at 4 Hz. In addition, pMSN, pTAN, and pLTS exhibited a second smaller peak at ~ 8 Hz (pFSI, 11 Hz). Moreover, pFSIs showed significant locking at two frequencies not observed for the other neuron types, identified by a small peak at ~ 15 Hz and strong component (46%) in the lower gamma band at ~ 35 Hz. For all other neuron types, the number of neurons with a significant phase preference at gamma frequencies was only found at chance level (5%, as indicated by the dashed line in Fig. 4C showing the expected percentage of false positives).

In a second step, we accounted for chance locking caused by potential regularity of the spike trains by testing the strength of observed phase locking against surrogates derived by local shuffling of the ISIs (for details, see Materials and Methods). We observe a similar percentage of significantly locked neurons (Fig. 4C, gray histograms) as for Rayleigh's test (red line), excluding the possibility that the observed locking is a trivial consequence of the underlying spike train statistics.

In Figure 4D, we show the corresponding phase distributions constructed from all spikes for each frequency, in which the percentage of locked neurons exceeded 10% using the test of non-uniformity (compare Fig. 4C, red curves). For low frequencies, all neuron types have comparable phase distributions related to the peaks indicated above. Although these phase distributions for the delta band frequencies significantly differed in mean and/or variance for all subtype comparisons (Mardia–Watson–Wheeler test, $p < 0.01$), they did not provide a clear criterion for the separation of the neuronal subtypes. The exceptional locking of pFSI in the gamma range is associated with a weakly modulated phase distribution. The reason for the broadening of the population-averaged phase distribution is twofold: First, the mean phases across individual pFSI neurons were less consistent than observed for other frequencies in which neurons exhibited phase locking (Fig. 4E). Second, phase distributions obtained from each neuron were less concentrated compared with those found in the delta range (Fig. 4F). In conclusion, we observed strong and similar locking of all neuronal subtypes to the delta frequency band of the ECoG, but pFSIs in addition expressed locking in the low gamma frequency band in nearly one-half of all observed neurons.

To complement these findings, we repeated the previous analysis in which the reference signal was replaced by an LFP recording (excluding the electrode used to record spiking activity). We obtained similar results (data not shown), but note two major differences: the preferred phase consistently differed from that obtained for the ECoG independent of the neuronal subtype, and was approximately shifted by 180° . Furthermore, we observed that the frequency range in which significant locking of pFSI

neurons to gamma oscillations is observed, increased to include also higher gamma frequencies than 60 Hz. The exclusive entrainment of the class of pFSI to gamma frequency oscillations in both types of field potentials provides evidence that activity of pFSI neurons is tightly related to cortical and striatal oscillations in this frequency range.

pMSNs correlate more strongly with putative GABAergic interneurons than pTANs

Cross-correlations were computed between pMSNs and all simultaneously recorded single neurons not recorded on the same electrode (Fig. 5A–D). All pair types had more significant correlations than would be expected by chance (5%), with the most correlations occurring around zero lag (Fig. 5A–D, histograms). The strength of connectivity between different neuron types was assessed by comparing the maximum t scores (the number of SDs from the surrogate null hypothesis) for each pair type (Fig. 5E). pMSNs were significantly more correlated with pFSIs ($n = 59$) than pTANs ($n = 20$; $p < 0.004$). They were also more correlated with themselves ($n = 80$; $p < 0.04$), although this was not significant after correction for multiple comparisons. pMSNs–pMSN pairs were the only type to show a significant difference between ipsilateral ($n = 38$) and contralateral ($n = 42$) pairs ($p < 0.0004$), indicating that pMSNs were more correlated if they were in the same hemisphere (Fig. 5F). In contrast, differences between pMSN and putative interneuron pairs were not significant. Low-frequency correlations between pMSNs, pFSI, and pLTS neurons could be clearly observed in raw examples (supplemental Fig. 5, available at www.jneurosci.org as supplemental material). Together, these results indicate that the locality of neuron pairs affected correlation between the putative projection neurons more than the correlation between projection and interneuron.

As all pair types showed some evidence of oscillation, we then compared the amplitude of the cross-correlations at the low frequencies in the ECoG/LFP signals (Fig. 5A–D). In all low-frequency ranges, the amplitude of the cross-correlation between pMSNs and pFSI or pLTS neurons was significantly higher than for pMSN–pTAN pairs (Table 3). pMSN–pFSI pairs were also more significantly correlated than pMSN–pMSN pairs in the lowest range. In the 5–9 Hz range, the amplitude of pMSN–pLTS was also significantly greater than pMSN–pMSN correlations. In summary, pMSNs showed the highest correlation between themselves and putative GABAergic interneurons. Correlations between putative projection and GABAergic interneurons were oscillatory, whereas those with pTANs were usually not.

Pilot recordings under ketamine anesthesia showed that striatal neurons could also be differentiated by waveform and were strongly correlated when recorded using another anesthetic preparation (supplemental Fig. 3, available at www.jneurosci.org as supplemental material).

Correlated activity is stronger between pairs of pFSIs than between pFSIs and other putative interneurons

More dramatic differences were observed in correlations between putative interneurons. pFSI pairs were often highly oscillatory across the whole lag range, and the phase of the oscillation at ~ 5 Hz was consistent enough to accumulate in the mean (Fig. 6A). Closer examination of individual pairs revealed that correlations were observed both at very low frequencies (Fig. 6B, top left) and at 5 Hz (Fig. 6B, bottom left), but that in both cases the correlation was extremely strong and centered around zero. The frequency of the correlation matched the dominant power in the ECoG (corresponding panels to the right). Correlations between

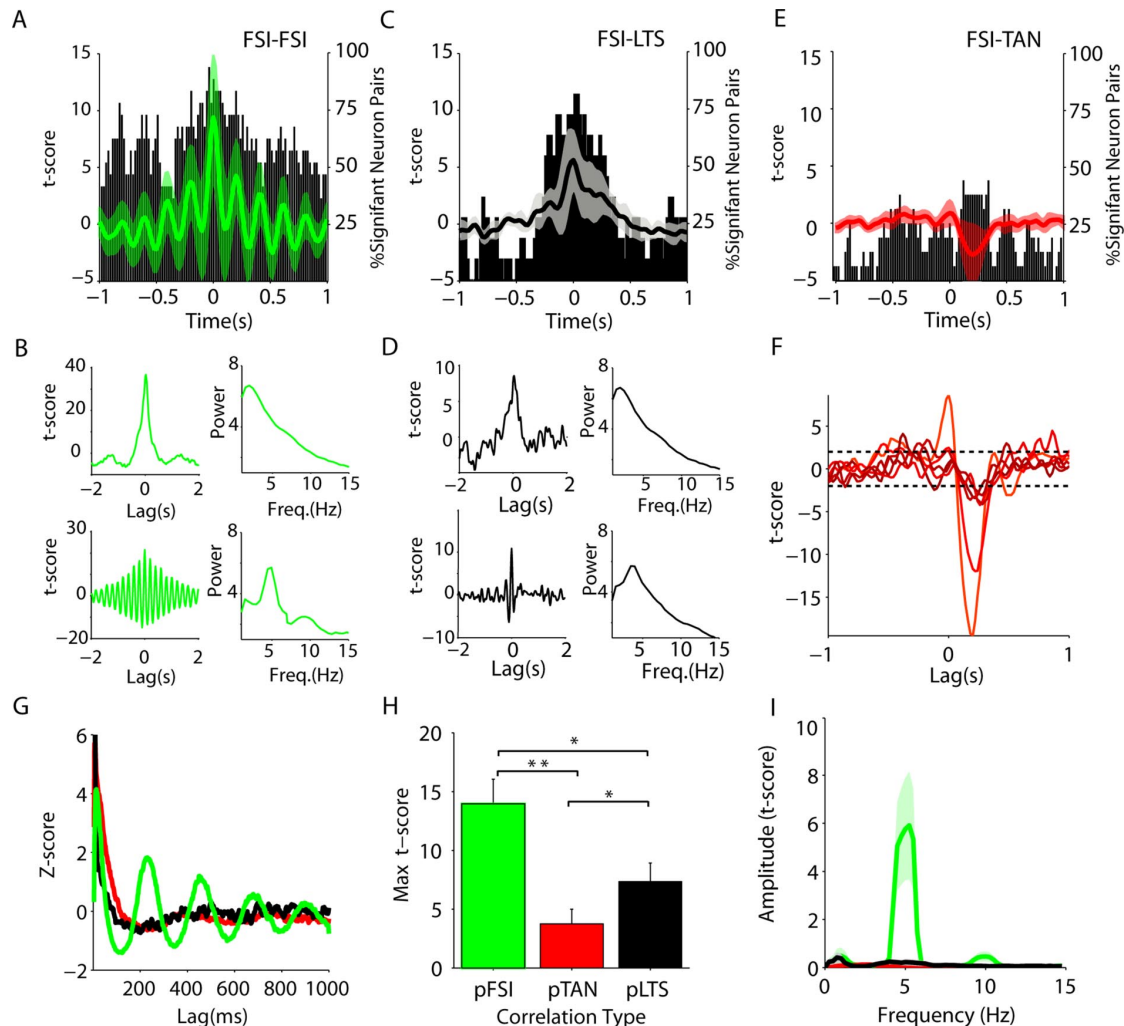


Figure 6. Low-frequency cross-correlation analysis between pFSIs and putative interneurons. Cross-correlations between different types of putative interneurons showed marked differences. For all panels, the color code for pair types is the same: green, pFSI–pFSI; red, pFSI–pTAN; black, pFSI–pLTS. As in Figure 5, the curves in **A**, **C**, and **E** denote the average t score, and the histogram shows the percentage of neurons in which 2 SDs were exceeded at a given lag regardless of the sign of the correlation. **A**, Pairs of pFSIs ($n = 17$) were almost always significantly correlated at zero phase lag and were often highly oscillatory. **B**, Two types of correlations were observed, demonstrated here by two representative example neurons. The top cross-correlation is slow with some oscillation and was accompanied by slow oscillations in the ECoG (right). The bottom correlation is highly oscillatory at ~ 5 Hz and is accompanied by power at the same frequency in the ECoG (right). ECoG in both cases is averaged across the hemispheres. Note that, in both cases, the t score is large and centered on zero. **C**, The majority of pFSI–pLTS pairs ($n = 11$) were significant near zero lag but showed little evidence of oscillation. **D**, Analysis of single pairs shows, however, that these pairs also oscillate depending on the ECoG power (same format as **B**), albeit less than the FSI pairs. Note that both correlations are centered on zero. **E**, In contrast, fewer pFSI–pTAN pairs ($n = 16$) were significantly correlated and those that were exhibited negative and nonoscillatory correlations. **F**, In six pairs in which there was a significant t score, however, the pattern of correlation was highly uniform with a significant drop in coincidences 200 ms after the firing of the pFSI (the reference). **G**, The particularly strong 5 Hz oscillations of some pFSI pairs (not observed for other putative interneuron types) was visible in the autocorrelations of simultaneously recorded putative interneurons (color code as above). **H**, The maximum t score (between ± 250 ms) of the correlation was significantly higher for pFSI–pFSI correlations than for pFSI–pTAN and pFSI–pLTS correlations. pLTS correlations were also significantly stronger than pFSI–pTAN correlations. Error bars indicate SEM. **I**, FFT of the correlation of the different neuron types confirmed the pFSI pairs were significantly more oscillatory at ~ 5 Hz. Asterisks indicate significance without (*) and with (**) Bonferroni–Holmes correction.

pFSIs and pLTS neurons were also common around zero lag (75% pairs) (Fig. 6C). Strong low-frequency correlations between pLTS and pFSI cells could be clearly observed in raw cross-correlation histograms (supplemental Fig. 5, available at www.jneurosci.org as supplemental material). As with pFSI pairs, analysis of single pairs showed similar low and 5 Hz correlation as pFSI pairs, which matched the dominant frequency in the ECoG power (Fig. 6D).

pFSIs were correlated both with themselves and pLTS neurons significantly more than with pTANs ($p < 0.0003$; $p = 0.05$) (Fig.

6H). Although strong correlations could be found between pFSIs and pTANs, they were clearly different in nature. Although the majority of putative GABAergic interneuron correlations were significantly positive around zero lag, the significant pFSI–pTAN correlations were negative with a positive latency when the pFSI was the reference (Fig. 6E, F). Despite the similar lag and oscillation of pFSI–pFSI and pFSI–pLTS correlations, pFSI–pFSI pairs were more correlated than pFSI–pLTS pairs ($p = 0.03$) (Fig. 6G), although this was not significant after correction. A large peak in the amplitude of the cross-correlation was observed between

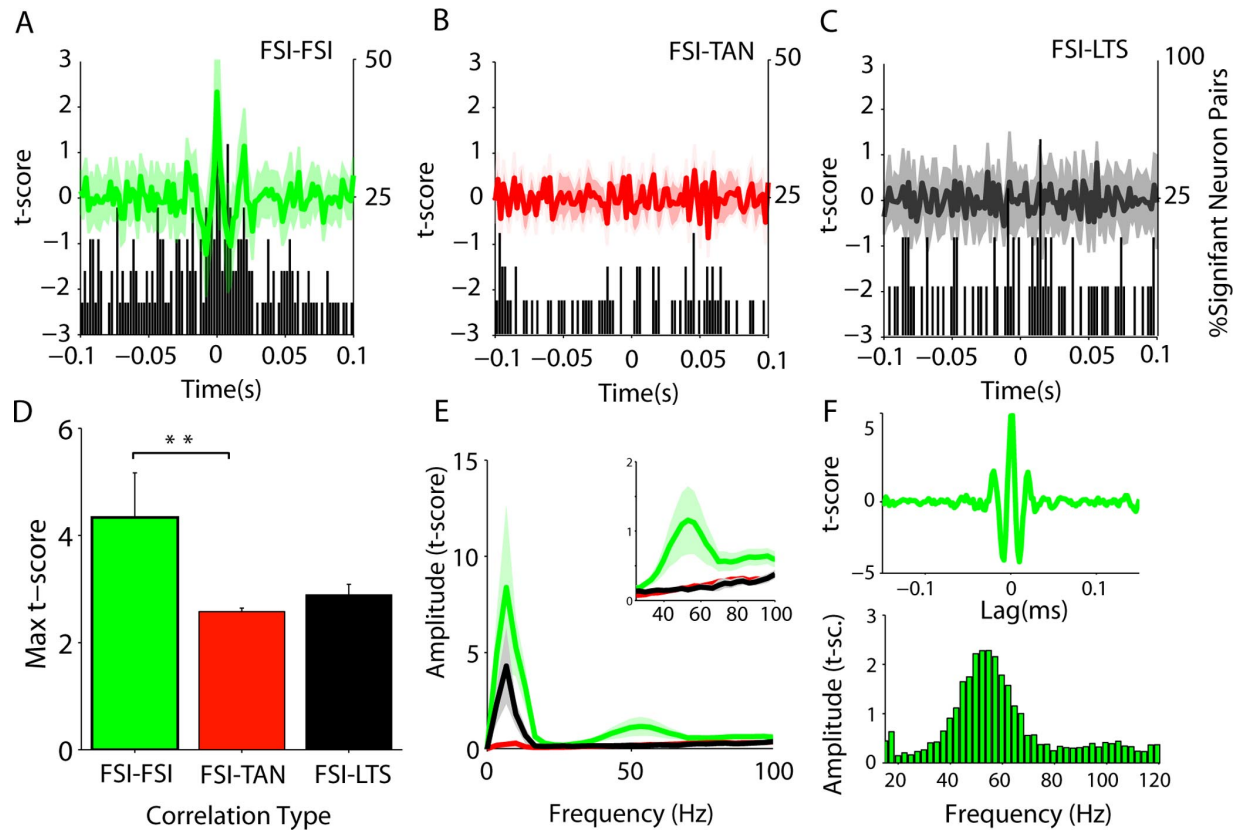


Figure 7. High-frequency cross-correlation analysis between pFSIs and putative interneurons. Correlation analyses in Figure 6 were repeated after high-pass filtering to reveal high-frequency/short timescale correlations. For all panels, the color code for pair types is the same: green, pFSI–pFSI; red, pFSI–pTAN; black, pFSI–pLTS. **A**, pFSI pairs were also correlated at high frequency (high-pass filter, 30 Hz), which was not apparent in the other correlation types, pFSI with pTAN (**B**) and pFSI with pLTS (**C**). **D**, The strength of the correlation, however, was only significantly different for pFSI–pTAN pairs. Error bars indicate SEM. **E**, As for slow correlations, however, the high-frequency power was significantly higher for pFSI pairs than for other combinations of neurons types and over a frequency identical with the gamma frequency ECoG and LFPs. **F**, Gamma oscillation in a single pair of pFSIs in the cross-correlation (top) and its FFT (bottom). Two asterisks indicate significance after Bonferroni–Holmes step down correction.

pFSI pairs at ~5 Hz and the amplitude of the cross-correlation significantly greater in the 5–9 Hz range than for pFSI–pTAN pairs (Fig. 6*H*, Table 4). Although the amplitude in this range was not significantly greater than for pFSI–pLTS pairs, single examples could be demonstrated in which a pFSI followed this frequency, whereas pLTS neurons did not (Fig. 6*G*), and pFSIs showed a uniquely strong relationship to cortical oscillations at this frequency over time (see below) (see Fig. 8). Differences between pFSI–pFSI and pFSI–pTAN pairs could also be seen in relation to oscillations of different frequencies that could be observed in single recordings (supplemental Fig. 6, available at www.jneurosci.org as supplemental material).

After filtering of low frequencies, only pFSI–pFSI correlations showed convincing evidence of high-frequency oscillations (Fig. 7*A–C*). The absolute correlation was significantly different between pairs of pFSIs than pFSI–pTAN pairs (Fig. 7*D*). The amplitude of the *t* score between 35 and 70 Hz for pFSI pairs was significantly higher than for correlations between pFSIs and pTAN ($p < 0.015$) or pLTS ($p < 0.03$) cells (Fig. 7*E*), and large gamma oscillations could be clearly observed in individual pFSI pairs (Fig. 7*F*; supplemental Fig. 6, available at www.jneurosci.org as supplemental material). The relationship between FSIs with themselves was, therefore, notably stronger than with other interneurons on both slow and fast timescales.

Table 4. Statistical comparisons of low-frequency power in putative interneuron cross-correlations

Correlation type	0–2 Hz	2–4 Hz	5–9 Hz
pFSIpFSI/pFSIpTAN	$p < 0.0004$	NS	$p < 0.002$
pFSIpFSI/pFSIpLTS	$p < 0.02$	NS	NS
pFSIpTAN/pFSIpLTS	$p < 0.03$	NS	$p < 0.001$

The oscillatory behavior of pFSIs is strongly related to cortical oscillations

As demonstrated by the phase-locking analyses, pFSIs had a particularly strong relationship to cortical oscillations. We explored this further by looking at the dynamic relationship between pFSI pairs and cortical oscillations at different frequencies. During periods of highly synchronized cortical activity in the theta range (Fig. 8*Ai*), the oscillation frequency changed by small graduations (0.25–0.5 Hz) over time. Individual pairs of pFSIs could follow these small changes remarkably closely (Fig. 8*Aii–Aiv*). In several recordings, the firing of pFSIs appeared to be highly related to periods in which the ECoG activity was dominated by this theta-like activity and the peak frequency of gamma oscillations was in transition (Fig. 8*B*). During these transition periods, the autocorrelation of pFSIs displayed gamma oscillations. In other recordings, in which 5–9 Hz activity was sustained, the cross-

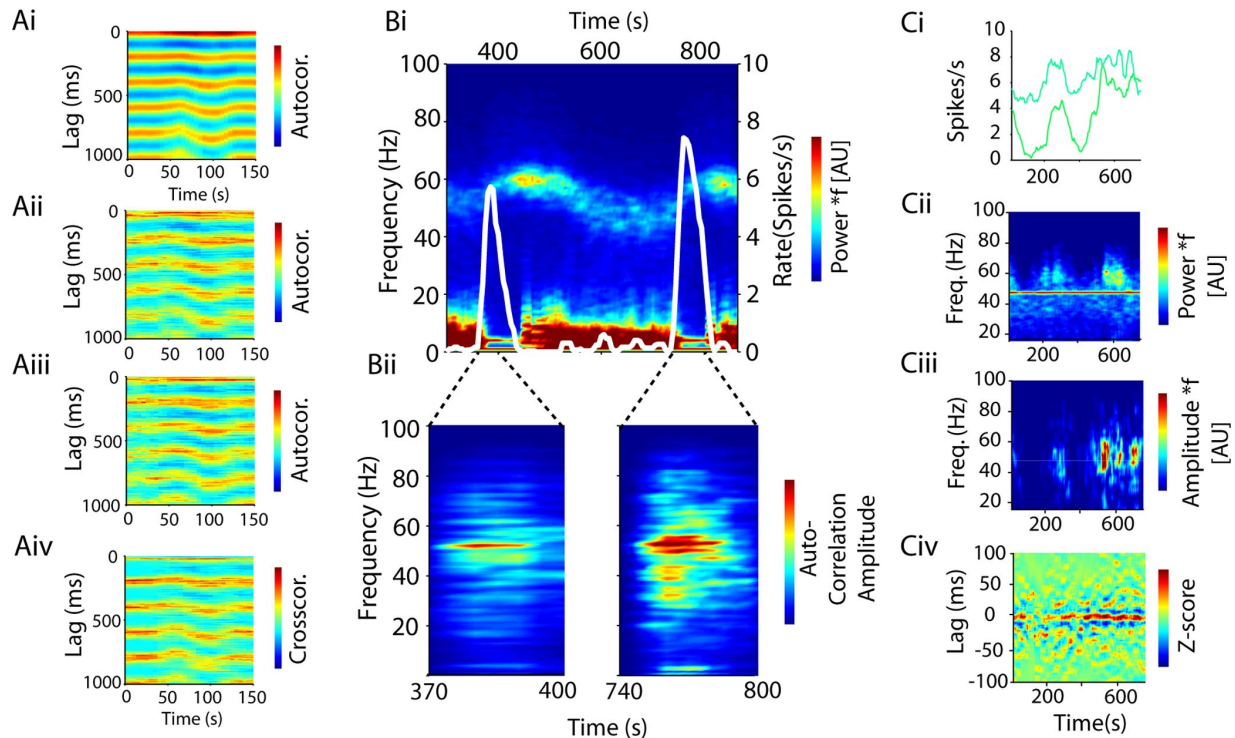


Figure 8. Oscillations in single and coupled pFSI firing are strongly related to those in the cortical ECoG over time. **Ai**, The autocorrelation of an ECoG channel plotted over 150 s. The color scale shows the movement of the lag of the peaks and troughs of the prominent 5 Hz oscillation. **Aii, Aiii**, The autocorrelations of two concurrently recorded pFSIs also have clear 5 Hz oscillations and show the same small deviations in frequency over time, which are also seen in the cross-correlation of the same units (**Aiv**). **Bi**, The time–frequency plot of an ECoG channel over 700 s shows characteristic transitions in both low and high frequencies (20 s windows). The rate of a single pFSI (superimposed white line) is elevated only during the transition periods. **Bii**, Gamma oscillations are present in the FFT of the autocorrelation of the pFSI in these periods. **Ci**, The rate of two pFSI neurons during an 800 s period with stable 5 Hz oscillations. **Cii**, The time–frequency representation of the ECoG during this period displays fluctuations in gamma oscillations. **Ciii**, The amplitude of the cross-correlation in the frequency domain of the two pFSIs display gamma oscillations during these periods, which can also be seen in the filtered cross-correlation (**Civ**).

correlation of pFSIs had power at the gamma frequency during periods of higher gamma power (Fig. 8C). Together, these results suggest that coordinated activity of these pFSIs could propagate high and low cortical rhythms to the striatum.

Discussion

These results provide the first simultaneous identification and analysis of four putative subtypes of striatal neurons in a single data set and the first evidence that the subtypes differ significantly in their phase locking to cortical population oscillations. The magnitude and rhythmicity of correlations between different neuronal subtypes varies systematically with neuron type. In particular, the majority of pFSI neurons are locked to fast cortical oscillations, and a large percentage of pFSI pairs are correlated at high frequency.

It is important to address the limitations of the methodology used here. First, there is a limit to the extrapolation of findings under halothane anesthesia to the awake animal. In practical terms, halothane enabled a high level of stability that allowed us to record relatively stationary neuronal activity over long periods of time, aiding both spike sorting and data analysis. As with other anesthetics, the exact mechanism of action of halothane is unknown; however, the major frequencies observed are consistent with those previously described in the cortex of halothane-anesthetized rats (Berridge and Foote, 1991; Berridge et al., 1996; Imas et al., 2004). Halothane is also known to block gap junctions (Johnston et al., 1980), which simplifies the interpretation of the

cross-correlation results (see below). In terms of how this compares to the state-related oscillatory activity in the awake animal, low-frequency oscillations (<4 Hz) are found only during sleep (Steriade, 2000). Despite this, the high level of synchronization during slow oscillations provides a highly useful tool for looking at neuronal interactions between the basal ganglia and cortex (Stern et al., 1998; Mahon et al., 2001; Berke et al., 2004), which can be predictive of higher frequency activities at the single-neuron level (Mallet et al., 2008). With this in mind, it is important to note that theta and gamma activities over a similar frequency range are a prominent feature in the striatum of awake rat (Berke et al., 2004; Tort et al., 2008). Therefore, although we cannot rule out the possibility that the choice of anesthetic influenced certain parameters, such as the proportion of each neuron type recorded, the majority of our findings regarding the temporal relationships of different neurons to each other and to cortex raise specific hypotheses about those relationships during those activities in the awake animal.

Second, how sure can we be that the putative neuronal subtypes described here belong to those different populations? Ideally, recorded neurons would be labeled and identified histologically using the juxtacellular method, which has been used to investigate the relationship of specific neuron types to network oscillations in hippocampus, cortex, and thalamus (Klausberger et al., 2003; Tierney et al., 2004; Lacey et al., 2007). However, to study the large numbers of simultaneous pairs of single neurons

(>200) described here, it becomes essential to use a different method. Using waveform parameters to separate neurons has proved a valuable technique for separating neurons (Csicsvari et al., 1999; Constantinidis and Goldman-Rakic, 2002; Mitchell et al., 2007). The waveform characteristics of our pFSIs and pTANs match those of labeled parvalbumin-positive GABAergic interneurons (Kawaguchi, 1993; Mallet et al., 2005) and giant cholinergic neurons (Wilson et al., 1990), respectively. This approach demands care to circumvent several potential sources of error. Most obviously, it relies on high-quality spike sorting to produce an average waveform that truly reflects a single neuron. Although it is common to use tetrodes for this purpose (Berke et al., 2004; DeCoteau et al., 2007), the electrodes used here gave excellent results and enabled us to identify more types of striatal neurons than previous studies. The shape and length of the waveform are to some extent a function of the position of the electrode in relation to the neuron (Gold et al., 2006, 2007). To address this issue, we used several physiological parameters to convincingly verify the separation of units. Regarding pLTS neurons, we provide the first description of low-threshold spike bursts in the striatum *in vivo*. There is no evidence to suggest that MSNs, FSIs, or TANs display LTS bursts. Therefore, based on *in vitro* studies (Kawaguchi, 1993), it seems most likely that pLTS neurons correspond to the subpopulation of GABAergic interneurons expressing neuropeptide Y and somatostatin. Because of the limited descriptions of calretinin expressing GABAergic interneurons, however, we cannot rule out the possibility that they neurons may also correspond to this type. Despite the considerable electrophysiological evidence presented here, this putative separation will need to be verified using histological methods.

Third, it is important to evaluate the analytical methods used here. A significant cross-correlation between two neurons does not necessarily reflect direct synaptic connectivity (Nowak and Bullier, 2000). Indeed, it is highly unlikely that many of the neurons in this sample were monosynaptically connected, given that most were recorded from electrodes ~0.5 mm apart. Rather than representing specific synaptic connections, the correlations seen here are more likely to reflect the locking to cortical population activity. It is likely that the strong correlations between striatal neurons are attributable to synchronized inputs from the cortex. The corticostriatal projection is highly convergent, with a convergence ratio of 10 between corticostriatal neurons and striatal neurons, but also sparse, with many cortical inputs being necessary to cause action potential firing in projection neurons (Ramanathan et al., 2002; Zheng and Wilson, 2002). This interpretation would explain why two GABAergic neurons were usually positively correlated around zero lag, rather than negatively correlated. Although this result seems initially counterintuitive, it is what would be expected if the correlations were driven primarily by efferents from the cortex.

The corticostriatal projection is a key processing step in the basal ganglia network. Although the results of cross-correlation analyses presented here may not reflect direct synaptic interaction, the relationships between different striatal neurons in relation to cortical activity is, therefore, of considerable importance to striatal physiology. In the hippocampus and cerebral cortex, there is reciprocal inhibitory/excitatory connectivity between interneurons and projection neurons that is likely to underlie local high-frequency oscillations (Traub et al., 2004). Oscillations in striatum, however, appear primarily driven by the cortex, and at low frequency there is ample evidence for this (Stern et al., 1998; Murer et al., 2002; Berke et al., 2004). The specificity of gamma

locking of pFSIs to cortical oscillations revealed in the present phase-locking analysis suggests that coherent gamma oscillations between cortex and striatum at the population level are entrained through this connection. Several lines of evidence from previous investigations support this finding. First, FSIs in the cortex and hippocampus are key drivers of high-frequency population oscillations (Klausberger et al., 2003; Hasenstaub et al., 2005). Second, the circuitry of striatal FSIs is suggestive of their involvement in coordinating large-scale activity (Ramanathan et al., 2002; Tepper and Bolam, 2004). FSIs form multiple synaptic contacts on MSN soma providing powerful inhibition (Koos et al., 2004). Third, the FSIs receive short latency cortical input, which is hypothesized to allow feedforward inhibition of MSNs (Mallet et al., 2005). Together, this suggests that FSIs may be involved both in the local generation of fast striatal oscillations and a key mediator of their entrainment to cortical rhythms. Gap junctions are commonly found between neighboring FSIs allowing for fast synchronization (Kita et al., 1990); however, the finding that halothane is a gap junction blocker (Johnston et al., 1980) suggests that these may not be required for the interactions among FSIs observed here.

To what extent can we rule out the contributions of other interneuron types to the propagation of oscillatory activity in the striatum? We could not test for correlations between pTAN or pLTS neurons, but, given the relatively low incidence of locking to the cortical oscillation, it seems unlikely that there would be high-frequency locking between these neuron types. In contrast, contrary to our observations, one might expect high-frequency correlations between pFSIs and pMSNs, as the latter need to be recruited to influence other basal ganglia structures. However, oscillations can manifest in the membrane potentials of neurons without initiating similar temporal patterning in the action potentials (Hasenstaub et al., 2005), including MSNs (Mahon et al., 2006). Indeed, given that FSIs make up only 0.8% of the striatum, a large summation of synchronized activity at the level of MSN membrane potential is the most likely explanation for the prominent power of oscillations in the bipolar LFP.

Many studies have demonstrated the importance of such oscillations in neuronal processing (Engel et al., 2001; Buzsáki and Draguhn, 2004) and gamma oscillations have received special attention as possible mediators of dynamic neuronal communication (Fries, 2005). A possible role of the coordinated high-frequency firing of FSIs might be to provide precise time windows for processing, integrating any MSNs that have sufficient synaptic input to fire at high frequency. At the systems level, coherent activity at gamma frequency between the striatum and cortex could facilitate the formation of temporally precise assemblies between the major processing hubs of the forebrain (Gervasoni et al., 2004) at a computationally relevant timescale (Fries, 2005; Womelsdorf et al., 2007). Our results add substantial support to the growing body of evidence that parvalbumin-positive GABAergic neurons are a crucial mediator of this process at the population level (Hasenstaub et al., 2005).

In conclusion, as in the cerebral cortex and hippocampus, the coordinated activity of fast spiking interneurons appears to underlie the emergence of high-frequency population oscillations in the striatum, fostering the entrainment of striatal rhythms by direct cortical projections. This finding suggests an important role of striatal FSIs in the propagation of oscillations within the striatum and the basal ganglia.

References

- Aosaki T, Kimura M, Graybiel AM (1995) Temporal and spatial characteristics of tonically active neurons of the primate's striatum. *J Neurophysiol* 73:1234–1252.
- Apicella P (2007) Leading tonically active neurons of the striatum from reward detection to context recognition. *Trends Neurosci* 30:299–306.
- Bar-Gad I, Ritov Y, Vaadia E, Bergman H (2001) Failure in identification of overlapping spikes from multiple neuron activity causes artificial correlations. *J Neurosci Methods* 107:1–13.
- Bartos M, Vida I, Jonas P (2007) Synaptic mechanisms of synchronized gamma oscillations in inhibitory interneuron networks. *Nat Rev Neurosci* 8:45–56.
- Bennett BD, Callaway JC, Wilson CJ (2000) Intrinsic membrane properties underlying spontaneous tonic firing in neostriatal cholinergic interneurons. *J Neurosci* 20:8493–8503.
- Berke JD, Okatan M, Skurski J, Eichenbaum HB (2004) Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron* 43:883–896.
- Berridge CW, Foote SL (1991) Effects of locus coeruleus activation on electroencephalographic activity in neocortex and hippocampus. *J Neurosci* 11:3135–3145.
- Berridge CW, Bolen SJ, Manley MS, Foote SL (1996) Modulation of forebrain electroencephalographic activity in halothane-anesthetized rat via actions of noradrenergic beta-receptors within the medial septal region. *J Neurosci* 16:7010–7020.
- Boashash B (1992) Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals. *Proc IEEE* 80:520–538.
- Bolam JP, Hanley JJ, Booth PA, Bevan MD (2000) Synaptic organization of the basal ganglia. *J Anat* 196:527–542.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Cassidy M, Brown P (2003) Spectral phase estimates in the setting of multidirectional coupling. *J Neurosci Methods* 127:95–103.
- Challis RE, Kitney RI (1990) Biomedical signal processing (in four parts). Part 1: Time-domain methods. *Med Biol Eng Comput* 28:509–524.
- Constantinidis C, Goldman-Rakic PS (2002) Correlated discharges among putative pyramidal neurons and interneurons in the primate prefrontal cortex. *J Neurophysiol* 88:3487–3497.
- Csicsvari J, Hirase H, Czurkó A, Mamiya A, Buzsáki G (1999) Fast network oscillations in the hippocampal CA1 region of the behaving rat. *J Neurosci* 19:RC20(1–4).
- DeCoteau WE, Thorn C, Gibson DJ, Courtemanche R, Mitra P, Kubota Y, Graybiel AM (2007) Learning-related coordination of striatal and hippocampal theta rhythms during acquisition of a procedural maze task. *Proc Natl Acad Sci U S A* 104:5644–5649.
- Dejean C, Gross CE, Bioulac B, Boraud T (2007) Synchronous high-voltage spindles in the cortex-basal ganglia network of awake and unrestrained rats. *Eur J Neurosci* 25:772–784.
- Denker M, Roux S, Timme M, Riehle A, Grün S (2007) Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomputing* 70:2096–2101.
- Elul R (1971) The genesis of the EEG. *Int Rev Neurobiol* 15:227–272.
- Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2:704–716.
- Fogelson N, Pogosyan A, Kühn AA, Kupsch A, van Bruggen G, Speelman H, Tijssen M, Quartarone A, Insola A, Mazzone P, Di Lazzaro V, Limousin P, Brown P (2005) Reciprocal interactions between oscillatory activities of different frequencies in the subthalamic region of patients with Parkinson's disease. *Eur J Neurosci* 22:257–266.
- Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7:862–871.
- Fries P (2005) A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends Cogn Sci* 9:474–480.
- Gervasoni D, Lin SC, Ribeiro S, Soares ES, Pantoja J, Nicoletis MA (2004) Global forebrain dynamics predict rat behavioral states and their transitions. *J Neurosci* 24:11137–11147.
- Gold C, Henze DA, Koch C, Buzsáki G (2006) On the origin of the extracellular action potential waveform: a modeling study. *J Neurophysiol* 95:3113–3128.
- Gold C, Henze DA, Koch C (2007) Using extracellular action potential recordings to constrain compartmental models. *J Comput Neurosci* 23:39–58.
- Groenewegen HJ, Berendse HW (1994) The specificity of the “nonspecific” midline and intralaminar thalamic nuclei. *Trends Neurosci* 17:52–57.
- Halliday DM, Rosenberg JR, Amjad AM, Breeze P, Conway BA, Farmer SF (1995) A framework for the analysis of mixed time series/point process data—theory and application to the study of physiological tremor, single motor unit discharges and electromyograms. *Prog Biophys Mol Biol* 64:237–278.
- Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* 47:423–435.
- Imas OA, Ropella KM, Wood JD, Hudetz AG (2004) Halothane augments event-related gamma oscillations in rat visual cortex. *Neuroscience* 123:269–278.
- Jeanmonod D, Magnin M, Morel A (1996) Low-threshold calcium spike bursts in the human thalamus. Common physiopathology for sensory, motor and limbic positive symptoms. *Brain* 119:363–375.
- Johnston MF, Simon SA, Ramón F (1980) Interaction of anesthetics with electrical synapses. *Nature* 286:498–500.
- Kawaguchi Y (1993) Physiological, morphological, and histochemical characterization of three classes of interneurons in rat neostriatum. *J Neurosci* 13:4908–4923.
- Kawaguchi Y, Kondo S (2002) Parvalbumin, somatostatin and cholecystokinin as chemical markers for specific GABAergic interneuron types in the rat frontal cortex. *J Neurocytol* 31:277–287.
- Kita H, Kosaka T, Heizmann CW (1990) Parvalbumin-immunoreactive neurons in the rat neostriatum: a light and electron microscopic study. *Brain Res* 536:1–15.
- Klausberger T, Magill PJ, Márton LF, Roberts JD, Cobden PM, Buzsáki G, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* 421:844–848.
- Koos T, Tepper JM, Wilson CJ (2004) Comparison of IPSCs evoked by spiny and fast-spiking neurons in the neostriatum. *J Neurosci* 24:7916–7922.
- Lacey CJ, Bolam JP, Magill PJ (2007) Novel and distinct operational principles of intralaminar thalamic neurons and their striatal projections. *J Neurosci* 27:4374–4384.
- Lachaux JP, Rodriguez E, Martinerie J, Varela FJ (1999) Measuring phase synchrony in brain signals. *Hum Brain Mapp* 8:194–208.
- Le Van Quyen M, Foucher J, Lachaux J, Rodriguez E, Lutz A, Martinerie J, Varela FJ (2001) Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J Neurosci Methods* 111:83–98.
- Llinás RR, Steriade M (2006) Bursting of thalamic neurons and states of vigilance. *J Neurophysiol* 95:3297–3308.
- Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. *Annu Rev Physiol* 66:735–769.
- Magill PJ, Pogosyan A, Sharott A, Csicsvari J, Bolam JP, Brown P (2006) Changes in functional connectivity within the rat striato-pallidal axis during global brain activation *in vivo*. *J Neurosci* 26:6318–6329.
- Mahon S, Deniau JM, Charpier S (2001) Relationship between EEG potentials and intracellular activity of striatal and cortico-striatal neurons: an *in vivo* study under different anesthetics. *Cereb Cortex* 11:360–373.
- Mahon S, Vautrelle N, Pezard L, Slaght SJ, Deniau JM, Chouvet G, Charpier S (2006) Distinct patterns of striatal medium spiny neuron activity during the natural sleep-wake cycle. *J Neurosci* 26:12587–12595.
- Mallet N, Le Moine C, Charpier S, Gonon F (2005) Feedforward inhibition of projection neurons by fast-spiking GABA interneurons in the rat striatum *in vivo*. *J Neurosci* 25:3857–3869.
- Mallet N, Pogosyan A, Márton LF, Bolam JP, Brown P, Magill PJ (2008) Parkinsonian beta oscillations in the external globus pallidus and their relationship with subthalamic nucleus activity. *J Neurosci* 28:14245–14258.
- Mardia KV, Jupp PE (2000) Directional statistics. New York: Wiley.
- Masimore B, Kakalios J, Redish AD (2004) Measuring fundamental frequencies in local field potentials. *J Neurosci Methods* 138:97–105.
- Masimore B, Schmitzer-Torbert NC, Kakalios J, Redish AD (2005) Transient striatal gamma local field potentials signal movement initiation in rats. *Neuroreport* 16:2021–2024.
- Mitchell JF, Sundberg KA, Reynolds JH (2007) Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* 55:131–141.
- Mitzdorf U (1985) Current source-density method and application in cat

- cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev* 65:37–100.
- Murer MG, Tseng KY, Kasanetz F, Belluscio M, Riquelme LA (2002) Brain oscillations, medium spiny neurons, and dopamine. *Cell Mol Neurobiol* 22:11–32.
- Nowak LG, Bullier J (2000) Cross-correlograms for neuronal spike trains. Different types of temporal correlation in the neocortex, their origin and significance. In: *Time in the brain*, pp 53–96. Amsterdam: Harwood Academic.
- Ramanathan S, Hanley JJ, Deniau JM, Bolam JP (2002) Synaptic convergence of motor and somatosensory cortical afferents onto GABAergic interneurons in the rat striatum. *J Neurosci* 22:8158–8169.
- Rivlin-Etzion M, Ritov Y, Heimer G, Bergman H, Bar-Gad I (2006) Local shuffling of spike trains boosts the accuracy of spike train spectral analysis. *J Neurophysiol* 95:3245–3256.
- Rosenberg JR, Amjad AM, Breeze P, Brillinger DR, Halliday DM (1989) The Fourier approach to the identification of functional coupling between neuronal spike trains. *Prog Biophys Mol Biol* 53:1–31.
- Rymar VV, Sasseville R, Luk KC, Sadikot AF (2004) Neurogenesis and stereological morphometry of calretinin-immunoreactive interneurons of the neostriatum. *J Comp Neurol* 469:325–339.
- Sarnthein J, Jeanmonod D (2007) High thalamocortical theta coherence in patients with Parkinson's disease. *J Neurosci* 27:124–131.
- Somogyi P, Klausberger T (2005) Defined types of cortical interneurone structure space and spike timing in the hippocampus. *J Physiol* 562:9–26.
- Steriade M (2000) Corticothalamic resonance, states of vigilance and mentation. *Neuroscience* 101:243–276.
- Stern EA, Jaeger D, Wilson CJ (1998) Membrane potential synchrony of simultaneously recorded striatal spiny neurons *in vivo*. *Nature* 394:475–478.
- Tepper JM, Bolam JP (2004) Functional diversity and specificity of neostriatal interneurons. *Curr Opin Neurobiol* 14:685–692.
- Tierney PL, Dégenétais E, Thierry AM, Glowinski J, Gioanni Y (2004) Influence of the hippocampus on interneurons of the rat prefrontal cortex. *Eur J Neurosci* 20:514–524.
- Tort AB, Kramer MA, Thorn C, Gibson DJ, Kubota Y, Graybiel AM, Kopell NJ (2008) Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a T-maze task. *Proc Natl Acad Sci U S A* 105:20517–20522.
- Traub RD, Bibbig A, LeBeau FE, Buhl EH, Whittington MA (2004) Cellular mechanisms of neuronal population oscillations in the hippocampus *in vitro*. *Annu Rev Neurosci* 27:247–278.
- Tukker JJ, Fuentealba P, Hartwich K, Somogyi P, Klausberger T (2007) Cell type-specific tuning of hippocampal interneuron firing during gamma oscillations *in vivo*. *J Neurosci* 27:8184–8189.
- Varela F, Lachaux JP, Rodriguez E, Martinerie J (2001) The brainweb: phase synchronization and large-scale integration. *Nat Rev Neurosci* 2:229–239.
- Wilson CJ (2005) The mechanism of intrinsic amplification of hyperpolarizations and spontaneous bursting in striatal cholinergic interneurons. *Neuron* 45:575–585.
- Wilson CJ, Chang HT, Kitai ST (1990) Firing patterns and synaptic potentials of identified giant aspiny interneurons in the rat neostriatum. *J Neurosci* 10:508–519.
- Witham CL, Wang M, Baker SN (2007) Cells in somatosensory areas show synchrony with beta oscillations in monkey motor cortex. *Eur J Neurosci* 26:2677–2686.
- Womelsdorf T, Schoffelen JM, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609–1612.
- Zheng T, Wilson CJ (2002) Corticostriatal combinatorics: the implications of corticostriatal axonal arborizations. *J Neurophysiol* 87:1007–1017.

Supplemental Materials

Supplemental Figure 1. Recording of putative striatal projection and interneurons. (A)

The position of all recordings was histologically verified. (Ai) Two parallel electrode tracks (marked with arrows) can be seen penetrating the dorsal striatum. (Aii) Single track with clearly defined lesions. (B) In records with multiple neurons, between 2-8 striatal units were recorded simultaneously. The recordings shown in each of the two subpanels show the electrode signals recorded simultaneously in one session. Single striatal units showed markedly different patterns of activity under the same conditions. (Bi) The top neuron fires tonically, while the middle two neurons burst, sometimes synchronously. The bottom neuron fires in only a single burst during this period. (Bii) The top two neurons show similar bursting to the middle units in Bi, with the middle neuron in Bii firing more irregularly. The bottom cell is also bursting but with much shorter intra-burst intervals. (C) Raw trace of a single neuron (subsequently identified as pLTS) at shown in 10 s epochs at 100 s intervals. Red lines indicate 5 standard deviations of the mean amplitude.

Supplemental Figure 2. Stability of waveform characteristics. Representative examples

of waveform consistency in recordings of single neurons lasting for between 600 and 1500 s. (Ai) Color representation of all the waveforms of a pMSN recorded for 630 seconds. Each column of the matrix is a single waveform with the peaks and troughs of the spike represented by the color scale. Note the consistency of both the amplitude and latency of the defining peaks and troughs. (Aii) The mean waveform $\pm 2SD$ (shaded areas). (Aiii) 5 s epochs of the raw trace of the same neuron taken arbitrarily at 100s intervals. (B-D) Identical plots for a pTAN (B), pLTS (C) and pFSI (D). Each neuron was recorded in a different animal.

Supplemental Figure 3. Putative MSNs and FSIs recorded under ketamine/xylazine

anesthesia. (Ai) Raw trace of ECoG and two striatal neurons with waveforms corresponding to those of pFSI (top) and pMSN (bottom) recorded under halothane. The pFSI fires a burst

Sharott et al., Cortical oscillations and striatal neurons

2

of spikes on every cycle of the slow oscillation, where as the pMSN only fires on some cycles. (Aii) Raw cross correlation between the two neurons in Ai displays low frequency locking. White curves show the low pass filtered correlogram. (Aiii) Subtracting the low pass filtered correlation reveals higher frequency oscillations, corresponding to spindle frequency. (Bi) Three waveforms from striatal neurons recorded in a different animal. The lengths of the top two correspond to those of pFSIs and the bottom to that of a pMSN. (Bii) Raw cross correlation between the two pFSI neurons. (Biii) As in Aiii, subtracting the low pass filtered correlation reveals higher frequency oscillations.

Supplemental Figure 4. Phase relations between ECoG and striatal LFPs. Histograms showing the percentage of values across frequencies from 1 to 100Hz (excluding 48-52 Hz which could be affected by line noise) where the phase value of the coherence spectrum between ECoG and striatal monopolar (A) and bipolar (B) LFPs was significantly different from 0° and $\pm 180^\circ$. Shaded areas show the SEM across animals.


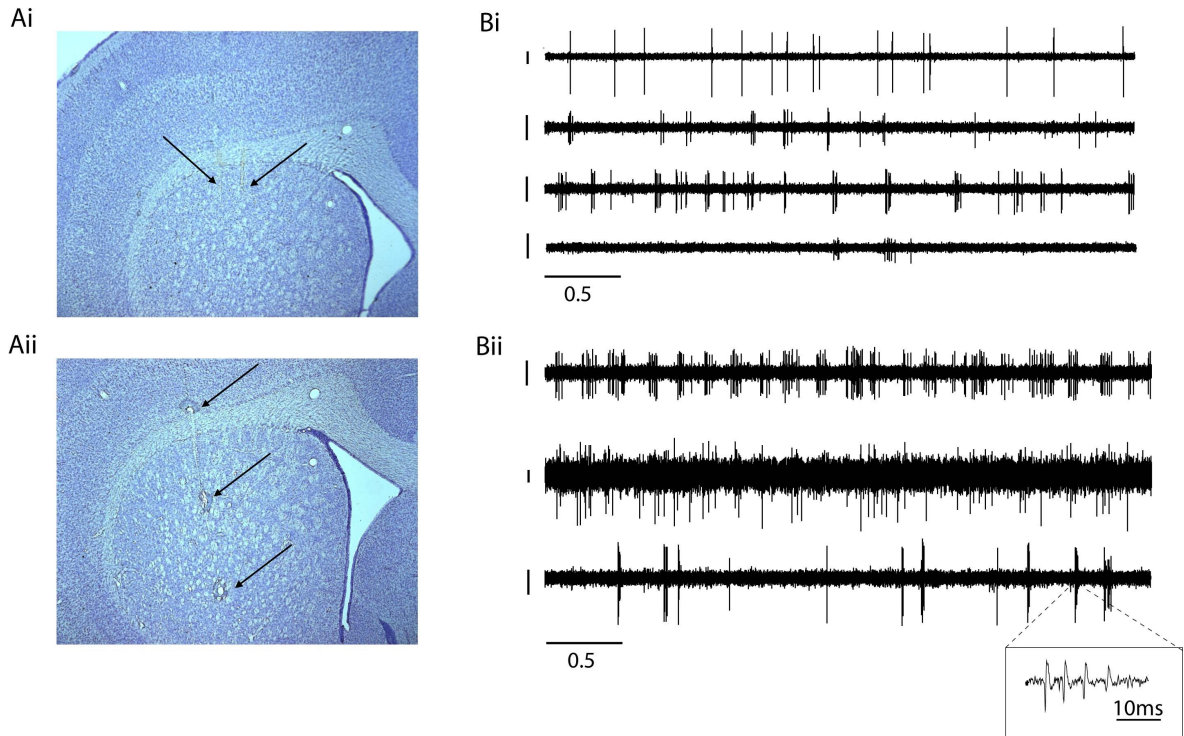
Supplemental Figure 5. Representative raw cross correlations between three pMSNs (labeled a-c), one pLTS and one pFSI recorded from five different electrodes across both hemispheres in a record of around 800s. (A) The pFSI/pLTS correlation is strongly centered around zero lag and oscillatory at around 3-4Hz. (B-D) The size and oscillatory strength of the correlations between the pFSI and pMSNs varies but in all cases is centered on zero lag and also displays the 3-4 Hz oscillation. (E-G) Correlations between the pLTS neuron and the pMSNs are also centered on zero lag but with less evidence of oscillation. (H-J) Correlations between pMSNs are the weakest, partly due to the low firing rate, but note that the pair with the lowest coincidences over the entire 2 s window is still sharply centered on zero lag (H).

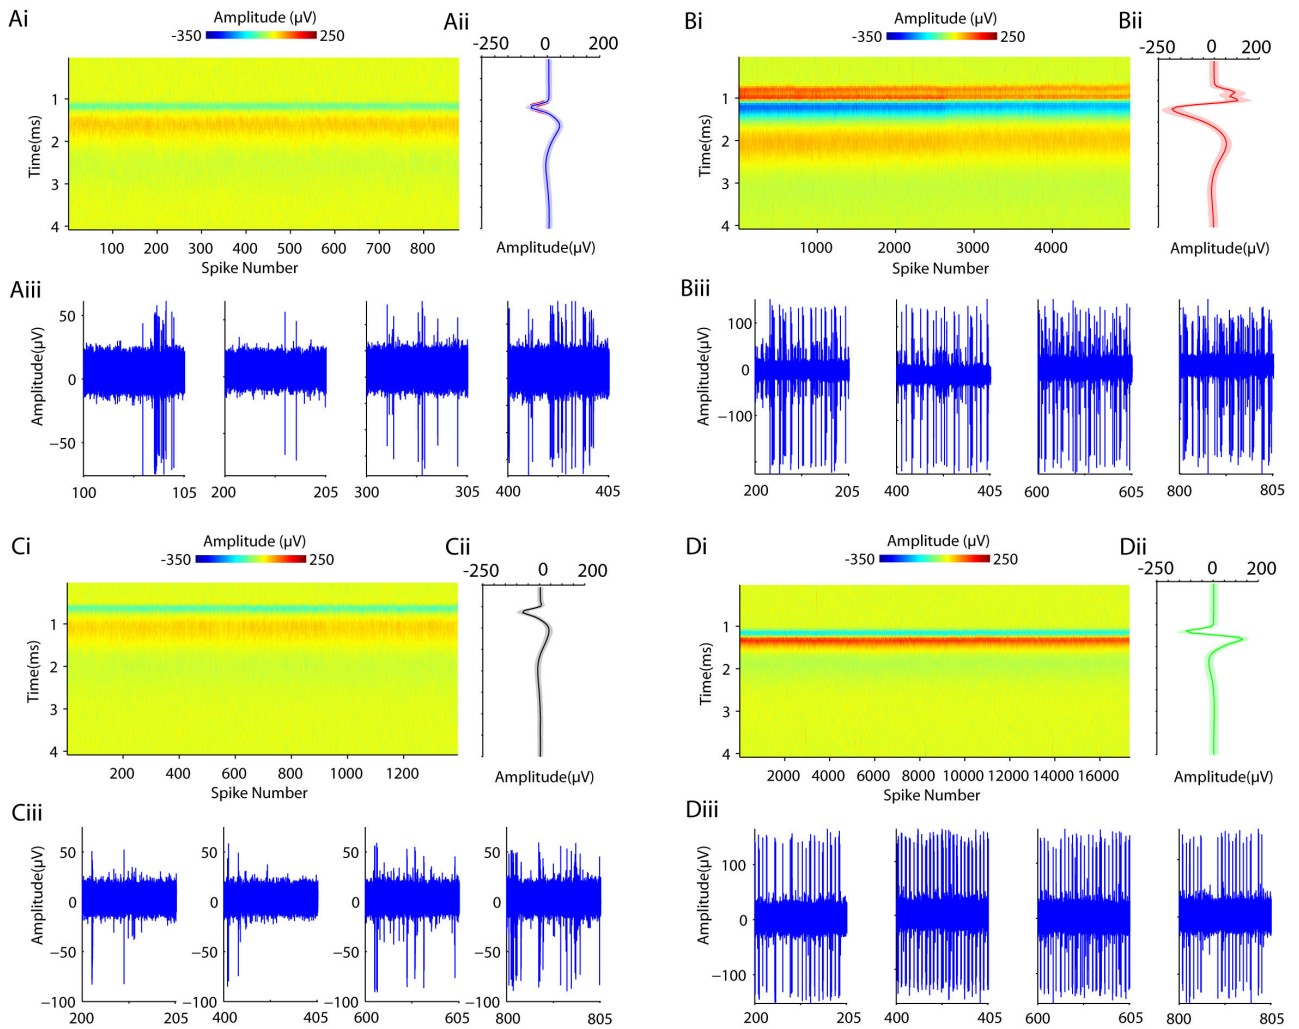
Supplemental Figure 6. Inter-pFSI correlations can be stronger and more oscillatory than pFSI-pTAN correlations under the same conditions. The figure shows raw cross

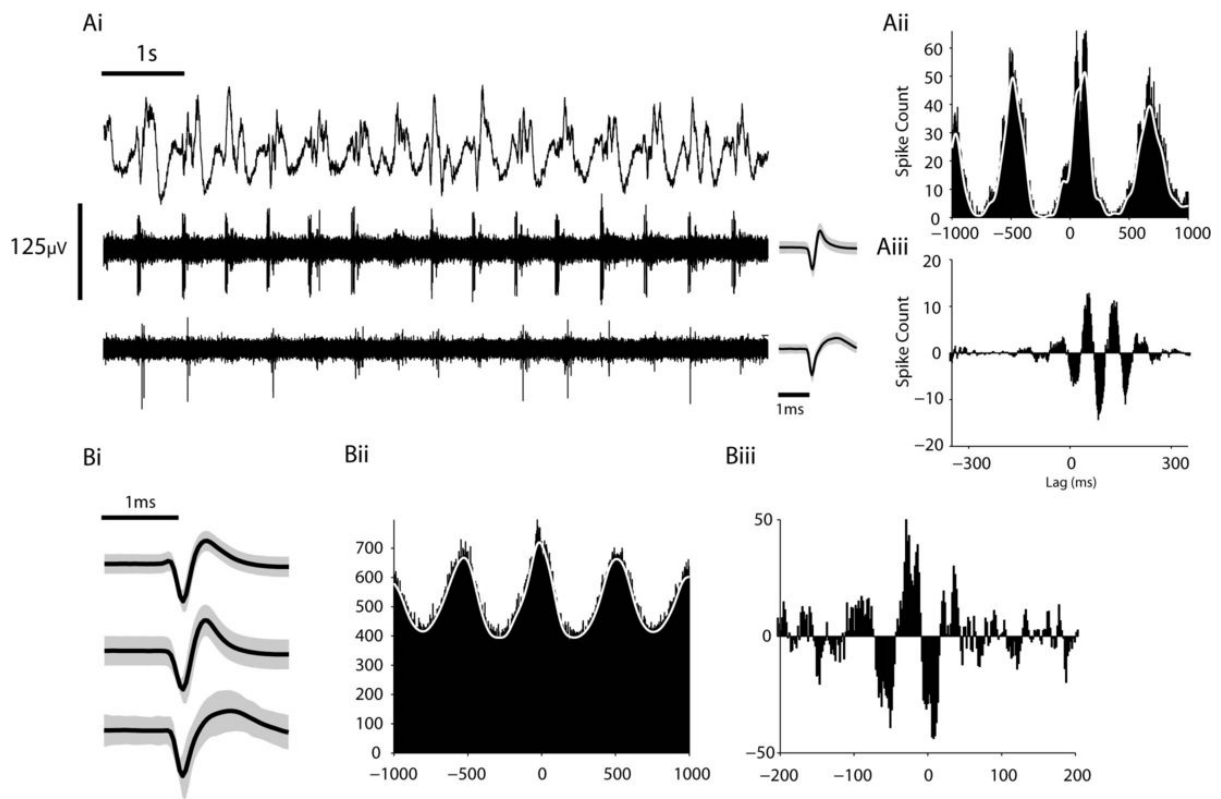
Sharott et al., Cortical oscillations and striatal neurons


3

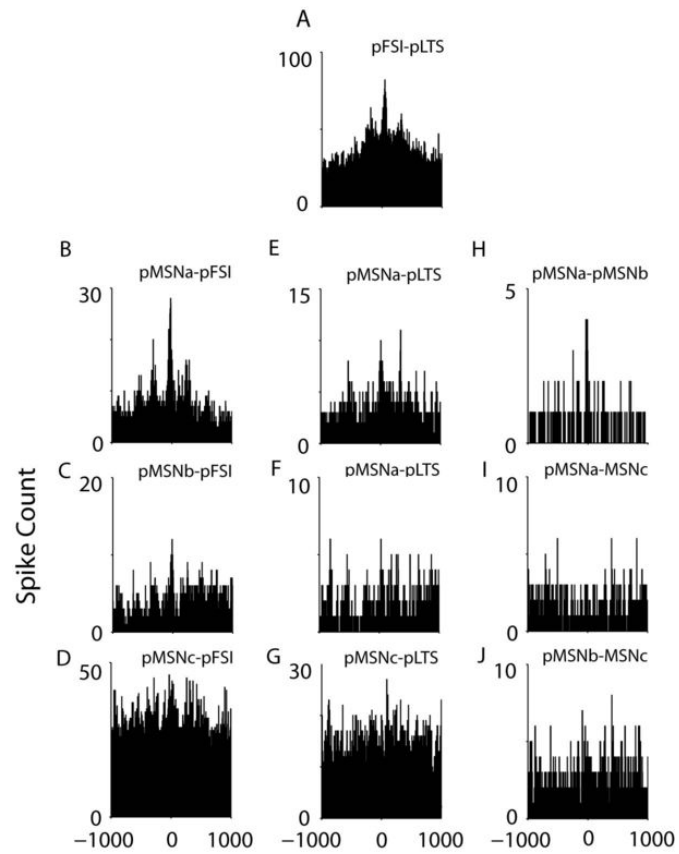
correlation histograms between three pFSIs and one pTAN recorded from four different electrodes in the same hemisphere in a record of around 600s. (A-C) Inter-pFSI correlations are extremely oscillatory at low frequency and also at higher frequencies. White curves show the low pass filtered correlogram. (D-F) Subtracting the low pass filtered correlation reveals high frequency oscillations in the bottom two pairs. (G-L) Neither type of correlation is seen between any of the pFSI-pTAN pairs, despite that there are as many or more coincident spikes in the same time range.

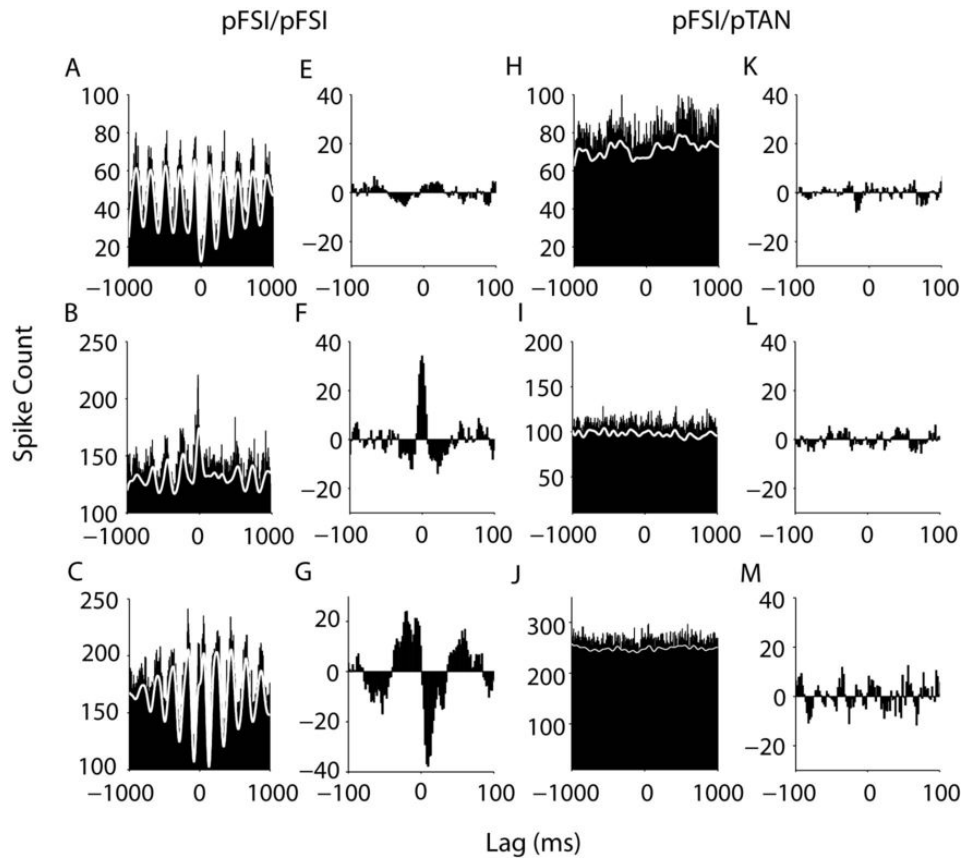












Chapter 3

Relating Local Field Potentials to Neuronal Assemblies

3.1 Phase Synchronization between LFP and Spiking Activity in Motor Cortex During Movement Preparation

The following article has been published as:

M. Denker, S. Roux, M. Timme, A. Riehle, and S. Grün. Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomputing* 70, 2096-2101 (2007).

The published article is available at <http://www.elsevier.com>
doi: [10.1016/j.neucom.2006.10.088](https://doi.org/10.1016/j.neucom.2006.10.088)

Available online at www.sciencedirect.com

Neurocomputing 70 (2007) 2096–2101

NEUROCOMPUTING

www.elsevier.com/locate/neucom

Phase synchronization between LFP and spiking activity in motor cortex during movement preparation

Michael Denker^{a,b,*}, Sébastien Roux^c, Marc Timme^{d,e}, Alexa Riehle^c, Sonja Grün^{b,f}

^aNeuroinformatics, Institute for Biology-Neurobiology, Free University Berlin, Germany

^bBernstein Center for Computational Neuroscience, Berlin, Germany

^cMediterranean Institute of Cognitive Neuroscience, CNRS and University Aix-Marseille II, Marseille, France

^dMax Planck Institute for Dynamics and Self-Organization, Göttingen, Germany

^eBernstein Center for Computational Neuroscience, Göttingen, Germany

^fComputational Neuroscience Group, RIKEN Brain Science Institute, Wako City, Japan

Available online 9 November 2006

Abstract

A common approach to measure and assess cortical dynamics focuses on the analysis of mass signals, such as the local field potential (LFP), as an indicator for the underlying network activity. To improve our understanding of how such field potentials and cortical spiking dynamics are related, we analyzed the phase and amplitude relationships between extracellular recordings from motor cortex of monkey in a delayed pointing task. We applied methods from phase synchronization analysis to extract the instantaneous phase of the LFP time series and to characterize the degree of phase coupling between the spike train and oscillation cycles in a frequency-independent manner. In particular, we investigated the dependence of observed phase preferences on the different periods of a behavioral trial. Furthermore, we present evidence to support the hypothesis that increased LFP oscillation amplitudes are related to a stronger degree of synchronization between the LFP and spike signals. However, neurons tend to keep a fixed phase relationship to the LFP independent of the amplitude or the choice of the electrode used to record the LFP.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Phase synchronization; LFP; Motor cortex; Monkey; Spike-field locking; Synchronization

1. Introduction

Cortical spiking activity is often recorded in parallel with the local field potential (LFP) obtained by low-pass filtering of the electrode signal. The LFP typically exhibits prominent oscillatory features and has been shown to display modulations that contain information about relevant behavior [11,18]. To date, the detailed mechanisms that relate LFP oscillations and network dynamics are not well understood. A current hypothesis describes the magnitude of LFP oscillations as an indicator for the presence of synchronous synaptic activity of many neurons

in a large volume around the electrode (e.g., [3,4,12]). Such cooperate network activity has been reported as a possible mechanism in the processing of information in motor cortex [2,15]. Assuming that increases in the magnitude of LFP oscillations indicate that a larger part of the neural population contributes to such coordinated activity, a neuron receiving this input would respond with increased temporal precision. This relation suggests a connection between the magnitude of LFP amplitudes and a preference for spike occurrences in a distinguished phase of the oscillation cycle. Previous studies have demonstrated the possibility to exploit such relationships between instantaneous LFP phase and spike time to extract information carried by the networks (see, e.g., [5]).

In this study we analyzed recordings from primary motor cortex of a behaving monkey in a time discrimination task

*Corresponding author. Neuroinformatics, Institute for Biology—Neurobiology, Free University Berlin, Germany.

E-mail address: mdenken@brain.riken.jp (M. Denker).

[17]. Preceding the actual execution of the movement, the experimental design involved two periods in which different amounts of prior information about the upcoming movement were available. Previous work provides evidence that the dynamics observed in motor cortex are closely related not only to the movement itself, but also to movement preparation [14,18].

Here, we employed newly developed methods (Denker et al., in prep.) that directly utilize the instantaneous phase of the LFP time series to measure the degree of phase coupling between spike trains and oscillation cycles. This approach provided an amplitude-independent measure of the current position on the oscillation cycle irrespective of the instantaneous frequency. We show that the observed locking periods of single LFP-neuron pairs exhibited a variety of different properties in terms of duration, strength and phase preference. In particular, we investigated the dependence of the precision of phase coupling on the different periods of a behavioral trial. Furthermore, we tested if periods of increased LFP amplitudes are related to a stronger degree of synchronization between the spike train and LFP time series. In addition, the observed phase preferences were compared between LFPS recorded at local or distant electrodes.

2. Methods

2.1. Behavioral task and recording

A rhesus monkey was trained to perform arm movements from a center position to one of the two possible peripheral targets left and right of the center, depending on the prior information provided by the preparatory signal (PS) and a time discrimination performed by the monkey [17]. The peripheral targets were presented simultaneously at PS, one in red, the other in green, the side of the color was chosen at random. An auditory response signal (RS) was then presented after either a short delay (600 ms) or a long delay (1200 ms). The monkey learned to associate to each color a delay and responded accordingly to the respective target. Each type of experimental configuration was presented with equal probability. As a result, during the first preparatory period (PP1) of 600 ms length the probability for the monkey to move at the end of this delay was 0.5, whereas when RS did not occur at that time (i.e., in long trials), the probability turned to 1 during the second preparatory period (PP2) after the expected signal (ES) at 600 ms. For analysis, all trials were aligned to PS occurrence.

LFPS and spikes of single neurons, detected by an online sorting algorithm (MSD, Alpha Omega, Nazareth, Israel), were obtained simultaneously from multi-electrode recordings of two to four electrodes in primary motor cortex. The inter-electrode distance was $\lesssim 400 \mu\text{m}$. LFPS were sampled at a resolution of 250–500 kHz and hardware filtered from 1 to 100 Hz. In total, we analyzed nine recording sessions,

which yielded 22 single neurons in approximately 40 trials per experimental condition.

2.2. Data analysis

We analyzed the spike trains and their relationship to the LFP by using methods that originate from phase synchronization analysis (see, e.g., [8,16,20]) and adapted them to the treatment of signals described by discrete points in time, such as spikes [6,7,19]. In a first step, we prefiltered the LFP signals to a frequency band that contained the dominant oscillatory component (phase-preserving implementation, MATLAB). In our case, we used a 10–22 Hz band, a prominent frequency band observed during an instructed delay in motor cortex (see [9] for a review). This frequency band was primarily observed during the preparatory periods between PS and RS (cf. Fig. 1a).

In a second step, we extracted the instantaneous phases of the remaining oscillatory LFP component. The phase $\phi(t)$ is a function with values in the range $[0; 2\pi]$ and serves as an indicator of the current position in the oscillation cycle, independent of the instantaneous cycle frequency (cf. Fig. 1b). For harmonic signals the phase is simply given by the argument, e.g., $\sin(\phi(t))$. Several methods have been proposed to obtain the instantaneous phase for oscillatory time series (see, e.g., [13]). Here we calculated the phase $\phi(t)$ as the angle of the complex-valued analytic signal $\tilde{L}(t) = A(t) \exp[i\phi(t)]$ corresponding to the LFP signal. The analytic signal was obtained as $\tilde{L}(t) = L(t) + i\mathcal{H}[L(t)]$, where $L(t)$ is the original real-valued LFP time series and $\mathcal{H}[\cdot]$ denotes the Hilbert transform. In addition, the amplitude $A(t)$ of the analytic signal provides a measure for the magnitude of the LFP, i.e., its envelope.

We collected the instantaneous phase $\phi(t_i)$ of the LFP at the individual spike times t_i across trials in six different time windows related to the experiment: the complete range from PS to RS (Total), and 400 ms windows around PS, around ES (long trials only), during PP1, during PP2 (long trials only) and after RS (Mvt). Spikes that occurred at extremely low LFP amplitudes were omitted when the phase signal was not monotonous and smooth. We analyzed phase distributions in these time windows using methods from circular statistics (cf. [10]). This allowed us to quantify the degree of non-uniformity in the phase distribution, and to test if a measured phase preference is statistically significant compared to a uniform distribution. The underlying assumption for a uniform phase distribution is that $2NR^2$ is χ^2 -distributed with two degrees of freedom, where N is the number of spikes and $R = N^{-1} \sum_{i=1}^N \exp[i\phi(t_i)]$ is the vector strength.

In a further step, we categorized individual spikes according to the envelope, or amplitude, $A(t_i)$ of the LFP signal at spike occurrence. Note that this is not equivalent to the amplitude $L(t)$ of the LFP itself, but rather measures the instantaneous oscillation amplitude, given by $\tilde{L}(t)$. We subdivided each spike train into two exclusive groups, where one group contained the spikes that occurred at high

2098

M. Denker et al. / Neurocomputing 70 (2007) 2096–2101




Fig. 1. Detecting instantaneous phase relationships between LFP and spikes. (a) Trial-averaged power spectrum of one LFP recording in one possible experimental condition (long delay) calculated in sliding windows of 200 ms. During the preparatory periods (PP1 and PP2) a prominent oscillatory component around 16 Hz is visible. (b) Filtered LFP signal (gray) of three example trials. The extracted phase signal $\phi(t)$ (black) from 0 to 2π is sketched above each LFP (troughs correspond to $\phi(t) = \pi$). Ticks indicate the location of spikes of a neuron recorded from the same electrode. The shaded region marks the time interval where movement started. (c) Time-resolved histogram of the phase distribution $\phi(t_i)$ for the LFP-neuron pair shown in (b), but across all trials (sliding windows $\Delta t = 100$ ms (50 ms overlap), 16 phase bins). For each time window, the distribution was normalized with respect to its maximum. The neuron shows a sustained preference for spiking during the falling flank ($0 \leq \phi \leq \pi$) of the LFP.

amplitudes (“Hi”), and the second group contained spikes that occurred at low amplitudes (“Lo”). The threshold θ for the separation of spikes into these groups is expressed as the ratio of number of spikes in the Lo group compared to the original number of spikes. Consequently, $1 - \theta$ denotes the relative number of spikes assigned to the Hi group. Subsequent phase analysis was then performed separately and independently on each of these two groups in the manner described above.

Robustness of our results was tested by a bootstrap procedure. The mean and standard deviation of the bootstrap measurements served as an estimate of the variability across the data set.

3. Results and conclusions

Single neurons showed a variety of locking behavior with respect to the LFP. In Fig. 1c the phase relationship of one neuron and the LFP recorded simultaneously from the same electrode is shown in a time-resolved manner. This particular neuron exhibited a constant tendency towards spiking on the falling flank ($0 - \pi$) of LFP cycles. However, few neurons displayed such a clearly pronounced, sustained phase preference.

To investigate the dependence between spike timing and LFP phase in relation to behavior, we calculated phase distributions during the six time windows described above related to specific periods in the experiment. In Fig. 2a we visualized the relative number of LFP-neuron pairs that showed a significant phase preference in these time

windows for any one of the experimental configurations (black bars). Note that due to the pooling of experimental configurations individual LFP-neuron pairs may enter the distribution more than once. We observe that about one fourth of the pairs shows a phase preference when considering the complete trial (Total). The fact that the other time windows showed fewer pairs is an indication of our earlier observation that sustained phase preferences across the experiment are rare. During movement, only few neurons showed a pronounced phase preference, as might be expected from the diminished LFP magnitude in the investigated frequency band during movement onset, where the LFP is dominated by a strong, slow triphasic component (movement related potential, [18]).

In particular, looking at the mean phase of LFP-neuron pairs with a significant phase preference confirmed that if a phase preference is detected, it occurred in roughly 60–70% of the cases on the falling flank of the LFP as opposed to the rising flank (not shown). Furthermore, we did not detect a dependence of the preferred phase on the experimental conditions of the task.

The envelopes of the LFP oscillation were typically modulated in time on a time scale much longer than individual oscillation cycles (cf. e.g., Fig. 1b). This modulation shows no correlation or fixed locking to trial onset or the behavioral task. The instantaneous phase typically remained well defined even for low amplitudes. To investigate the dependence of phase preferences on the amplitude, we separated individual spike trains into Hi and Lo groups of equal numbers of spikes ($\theta = 0.5$). Analysis of

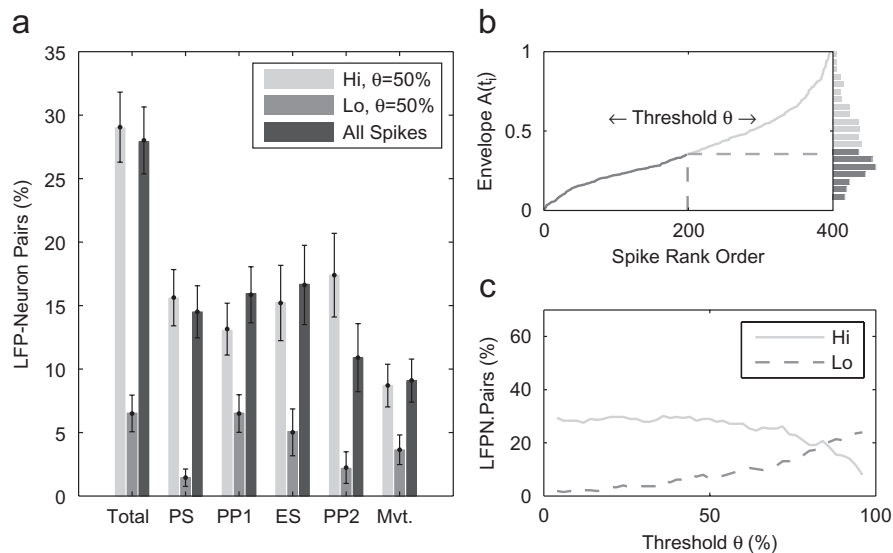


Fig. 2. Relation of phase locking to the magnitude of LFP oscillations. (a) Relative number of LFP-neuron pairs (all possible combinations) that exhibit a non-uniform phase distribution ($p = 0.01$) during the six defined time periods (Total, PS, PP1, ES, PP2 and Mvt) and across experimental configurations. Statistics was performed on the original spike train of each trial (black bars), as well as diluted spike trains which contain only those spikes that occur during the highest (Hi, light gray bars) and lowest (Lo, dark gray bars) LFP envelopes ($\theta = 0.5$, i.e., groups Hi and Lo have equal number of spikes). Error bars were obtained by 100 bootstraps that each comprised of 70% of the data. (b) Values of the envelope $A(t)$ at spike occurrences (normalized to maximum) of one typical LFP-neuron pair, sorted in ascending order. Dashed line shows the threshold of $\theta = 0.5$ separating the Hi and Lo groups containing equal number of spikes (shading as in panel (a)). Right: corresponding distribution of envelopes $A(t)$. (c) Effect of changing the threshold θ on the relative number of LFP-neuron pairs with non-uniform phase distributions in the Hi (light gray, solid) and Lo (dark gray, dashed) groups during PS–RS (Total). Low values of θ correspond to most spikes attributed to the Hi group and only few to the Lo group and vice versa. A large number of LFP-neuron pairs that displayed a phase preference for the spikes of the Hi group persisted even for high values of θ .

these groups revealed that primarily spikes occurring at high LFP envelopes (group Hi) lead to the phase preferences of a given LFP-neuron pair (Fig. 2a). In fact, uniformity in the phase distributions of spikes in the Lo group led to a smaller number of significant LFP-neuron pairs in the original spike train as compared to the Hi group (see, e.g., time period PP2). The corresponding amplitude distributions (see Fig. 2b for an example) were unimodal, but skewed towards low amplitudes.

In Fig. 2c we investigated the dependence of phase preferences in the complete preparatory period (Total) as a function of the threshold θ . Only at a threshold of $\theta = 0.8$ (i.e., the top 20% of spikes associated with the highest LFP envelopes are in the Hi group, while the remaining spikes are categorized as Lo) we find an equal number of LFP-neuron pairs that display a phase preference in both groups. This finding suggests not only that spikes at high LFP magnitudes exhibit a stronger tendency for locking to the LFP cycle, but moreover that observed phase preferences are due to spikes occurring at the highest amplitudes.

We investigated the distribution of the mean phase of all LFP-neuron pairs that showed a significant phase preference pooled across experimental configurations. The analysis was done independently for the two groups Hi and Lo (at $\theta = 0.5$). In addition, we grouped according to

whether spikes and LFP were recorded from the same or different electrodes. Fig. 3 shows the results for the Hi group. In both possible configurations (same/different electrode) the mean of the histograms was centered on the falling flank of the LFP (cf. also [1]). In particular, this peak was more pronounced for recordings on the same electrode as compared to different electrodes. Due to the low number of LFP-neuron pairs that showed a significant phase preference, the distribution for the Lo group lead to weak statistics. Nevertheless, the distribution of phase preferences for both electrode combinations showed a similar tendency as for the Hi group. We conclude that individual LFP-neuron pairs that exhibit a phase preference tend to have a fixed preferred phase independent of the LFP amplitude. This preference is more pronounced when spikes and LFP are taken from the same electrode. Although simultaneously recorded LFPS are typically highly correlated, our finding may be understood under the assumption that the LFP is a reflection of synaptic input activity preceding a spike. Therefore, when LFP and spikes are taken from the same electrode the LFP would be more closely related to this local input and therefore expected to exhibit a more precise temporal relationship to the spike times.

In summary, our analysis demonstrates that spikes of neurons in motor cortex have a tendency for a phase

2100

M. Denker et al. / Neurocomputing 70 (2007) 2096–2101

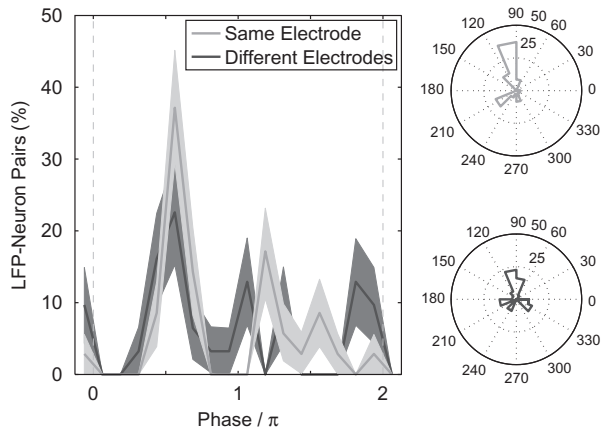


Fig. 3. Distribution of mean phases of LFP-neuron pairs that showed a non-uniform phase distribution ($p = 0.01$) for the spikes of the Hi group ($\theta = 0.5$) during PS-RS (Total) across experimental configurations. Data from pairs recorded from the same electrode are shown in light gray, whereas pairs taken from different electrodes are shown in dark gray (equal number of pairs as for same electrode). Error margins were obtained by boot-strapping as performed in Fig. 2. To aid the visualization of the distributions, polar plots on the right show the same distributions as on the left.

preference. The preferred phase relationship is specific for the individual LFP-neuron pair. Nonetheless, across pairs this phase preference showed a tendency to be located on the falling flank of the LFP oscillation. Furthermore, our analysis revealed that periods of higher LFP amplitudes tend to show an increase in the precision of locking between spikes and LFP. We hypothesize that this effect is a result of neurons becoming entrained to take part in a cooperative network activity that is manifested in an increased oscillatory activity.

Acknowledgments

We thank Dr. Ulrich Ebert for fruitful discussions. This work was carried out while Sonja Grün was based at the Free University, Berlin, Germany. This work was supported by the Volkswagen Foundation, the Stifterverband für die Deutsche Wissenschaft and the Bundesministerium für Bildung und Forschung (Grant 01GQ0413 to BCCN Berlin and Grant 01GQ0430 to BCCN Göttingen).

References

- [1] S.N. Baker, E.M. Pinches, R.N. Lemon, Synchronization in monkey motor cortex during a precision grip task. II. Effect of oscillatory activity on corticospinal output, *J. Neurophysiol.* 89 (2003) 1941.
- [2] S.N. Baker, R. Spinks, A. Jackson, R.N. Lemon, Synchronization in monkey motor cortex during a precision grip task. I. Task-dependent modulation in single-unit synchrony, *J. Neurophysiol.* 85 (2001) 869.
- [3] R. Eckhorn, A. Obermueller, Single neurons are differently involved in stimulus-specific oscillations in cat visual cortex, *Exp. Brain Res.* 95 (1993) 177.
- [4] R. Elul, The genesis of the EEG, *Int. Rev. Neurobiol.* 15 (1971) 227.
- [5] R.W. Friedrich, C.J. Habermann, G. Laurent, Multiplexing using synchrony in the zebrafish olfactory bulb, *Nat. Neurosci.* 7 (2004) 862.
- [6] K.D. Harris, D.A. Henze, H. Hirase, X. Leinekugel, G. Dragoi, A. Czurko, G. Buzsaki, Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells, *Nature* 417 (2002) 738.
- [7] J.M. Hurtado, L.L. Rubchinsky, K.A. Sigvardt, Statistical method for detection of phase-locking episodes in neural oscillations, *J. Neurophysiol.* 91 (2004) 1883.
- [8] J.P. Lachaux, E. Rodriguez, J. Martinerie, F.J. Varela, Measuring phase synchrony in brain signals, *Hum. Brain Mapp.* 8 (1999) 194.
- [9] W.A. MacKay, Wheels of motion: oscillatory potentials in motor cortex, *Motor Cortex in Voluntary Movements: a Distributed System for Distributed Functions*, CRC Press, Boca Raton, FL, 2005.
- [10] K.V. Mardia, *Statistics of Directional Data*, Academic Press, New York, 1972.
- [11] C. Mehring, J. Rickert, E. Vaadia, S.C. de Oliveira, A. Aertsen, S. Rotter, Inference of hand movements from local field potentials in monkey motor cortex, *Nat. Neurosci.* 6 (2003) 1253.
- [12] V.N. Murthy, E.E. Fetz, Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys, *J. Neurophys.* 76 (1996) 3968.
- [13] R.Q. Quiroga, A. Kraskov, T. Kreuz, P. Grassberger, Performance of different synchronization measures in real data: a case study on electroencephalographic signals, *Phys. Rev. E* 65 (2002) 041903.
- [14] A. Riehle, Preparation of movement: one of the key functions of motor cortex, *Motor Cortex in Voluntary Movements: a Distributed System for Distributed Functions*, CRC Press, Boca Raton, FL, 2005.
- [15] A. Riehle, S. Grün, M. Diesmann, A. Aertsen, Spike synchronization and rate modulation differentially involved in motor cortical function, *Science* 278 (1997) 1950.
- [16] M.G. Rosenblum, A.S. Pikovsky, J. Kurths, Phase synchronization of chaotic oscillators, *Phys. Rev. Lett.* 76 (1996) 1804.
- [17] S. Roux, M. Coulmance, A. Riehle, Context-related representation of timing processes in monkey motor cortex, *Eur. J. Neurosci.* 18 (2003) 1011.
- [18] S. Roux, W.A. Mackay, A. Riehle, The pre-movement component of motor cortical local field potentials reflects the level of expectancy, *Behav. Brain Res.* 169 (2006) 335.
- [19] C. Schäfer, M.G. Rosenblum, H.H. Abel, J. Kurths, Synchronization in the human cardiorespiratory system, *Phys. Rev. E* 60 (1999) 857.
- [20] F. Varela, J.P. Lachaux, E. Rodriguez, J. Martinerie, The brainweb: phase synchronization and large-scale integration, *Nat. Rev. Neurosci.* 2 (2001) 229.



Michael Denker studied physics at the Georg August University in Göttingen, Germany, and at the University of California, San Diego, USA. He received his diploma degree in 2002 at the Department of Nonlinear Dynamics, Max Planck Institute for Flow Research, Göttingen, with work on the analysis of collective dynamics in complex networks of pulse-coupled oscillators. From 2003 to 2006 he worked as a PhD student at the Group for Neuroinformatics at the Institute for Biology, Free University, Berlin, and has been a member of the Bernstein

Center for Computational Neuroscience, Berlin. Since October 2006 he is employed at the RIKEN Brain Science Institute in Wako, Japan. His current research interests focus on the analysis of relationships between neuronal activity on different spatial and temporal scales and the study of synchronized activity in neuronal systems.



Sébastien Roux received the MBio degree in cellular biology and physiology from the University of St. Jérôme, Marseille, France, in 2001. He is currently finishing a PhD degree in neurosciences from the Mediterranean University, Marseille, France, and at the Mediterranean Institute of Cognitive Neurosciences, CNRS, Marseille, France, from 2002 to 2006. He is electrophysiologist and his research interests focus on the neural basis of movement preparation and execution in monkey motor cortex by

using multiple single unit, LFP and EMG recordings.



Marc Timme studied physics at the University of Würzburg, Germany, at the State University of New York at Stony Brook, USA, and at the University of Göttingen, Germany. He received an MA in physics in 1998 (Stony Brook) and a doctorate in theoretical physics in 2002 (Göttingen). He worked as a postdoctoral researcher in the Department of Nonlinear Dynamics, Max Planck Institute for Flow Research, Göttingen, and as a research scholar at the Center for Applied Mathematics, Cornell University, Ithaca, USA. As of October 2006, Marc is a junior research group leader (associate professor level) at the Max Planck Institute for Dynamics and Self-Organization in Göttingen. He is also a founding member and a principal investigator at the Bernstein Center for Computational Neuroscience (BCCN) Göttingen. His research interests include the nonlinear dynamics and statistical physics of networks with a focus in theoretical neuroscience, and a particular interest in spiking neural networks.

As of October 2006, Marc is a junior research group leader (associate professor level) at the Max Planck Institute for Dynamics and Self-Organization in Göttingen. He is also a founding member and a principal investigator at the Bernstein Center for Computational Neuroscience (BCCN) Göttingen. His research interests include the nonlinear dynamics and statistical physics of networks with a focus in theoretical neuroscience, and a particular interest in spiking neural networks.



Alexa Riehle received the BSc degree in Biology (main topic: deciphering microcircuitries in the frog retina) from the Free University, Berlin, Germany, in 1976, and the PhD degree in neurophysiology (main topic: neuronal mechanisms of temporal aspects of color vision in the honey bee) from the Biology Department of the Free University, Berlin, Germany, in 1980. From 1980 to 1984, she was a postdoctoral fellow at the CNRS in Marseille, France (main topic: neuronal mechanisms of elementary motion detectors in the fly visual system). In 1984, she moved to the Cognitive Neuroscience Department at the CNRS, Marseille, France, and is since then mainly interested in the study of cortical information processing and neural coding in cortical ensembles during movement preparation and execution in non-human primates.

In 1984, she moved to the Cognitive Neuroscience Department at the CNRS, Marseille, France, and is since then mainly interested in the study of cortical information processing and neural coding in cortical ensembles during movement preparation and execution in non-human primates.



Sonja Grün studied physics at the Eberhard-Karls University in Tübingen, Germany, and graduated with modeling work on sound localization at the Max Planck Institute for Biological Cybernetics in Tübingen. She did her PhD work in the field of computational neuroscience at the Ruhr University in Bochum, Germany, and at the Weizmann Institute in Rehovot, Israel, and obtained her PhD (Physics) at the Ruhr University in Bochum, Germany. In her postdoctoral work at the Hebrew University in Jerusalem, Israel, she

did electrophysiological work in awake behaving monkeys. From 1998 to 2002 she was a senior fellow at the Max Planck Institute for Brain Research in Frankfurt/M, Germany. From 2002 to 2006 she was assistant professor for Neuroinformatics/Theoretical Neuroscience at the Free University, Berlin, Germany and was a founding member of the Bernstein Center of Computational Neuroscience in Berlin. Since 9/2006 she is the head of a research unit at the RIKEN Brain Science Institute in Wako, Japan. Her main interests are in statistical neuroscience which includes modeling of stochastic processes and the development of data analysis techniques for multiple parallel neuronal time series.

3.2 Local Field Potentials Reflect Neuronal Assemblies

The following manuscript will be submitted as:

M. Denker, S. Roux, H. Lindén, M. Diesmann, A. Riehle, and S. Grün. Local field potentials reflect neuronal assemblies.

Local Field Potentials Reflect Neuronal Assemblies

Michael Denker¹, Sébastien Roux², Henrik Lindén³,
Markus Diesmann¹, Alexa Riehle⁴, Sonja Grün^{1,5}

July 17, 2009

¹ RIKEN Brain Science Institute, Wako City, Japan

² Bernstein Center for Computational Neuroscience, Albert-Ludwigs
University, Freiburg, Germany

³ Dept. of Mathematical Sciences and Technology, Norwegian University of
Life Sciences, Ås, Norway

⁴ Mediterranean Institute of Cognitive Neuroscience (INCM), CNRS -
University Aix-Marseille 2, Marseille, France

⁵ Bernstein Center for Computational Neuroscience, Berlin, Germany

Abstract

Ongoing efforts to unravel the mechanisms governing brain processing have led to the proposal that information is conveyed by the functionally coordinated discharge patterns of specific neuronal subgroups (assemblies). The oscillatory nature of the cortical local field potential (LFP) is commonly interpreted as a reflection of synchronized network activity, but its relationship to observed transient coincident firing of neurons on the millisecond time-scale remains unclear. Here we present experimental evidence to reconcile the notions of synchrony at the level of neuronal spiking and at the mesoscopic scale in the assembly coding framework. We demonstrate that the phase coupling of spikes to the LFP is stronger for spikes that are synchronous with other spikes, in particular in periods of strong LFP oscillations. However, spike coincidences are better entrained to the LFP than expected on the basis of the locking of individual spikes only in time intervals where the large number of coincidences indicates that the neurons coordinate their spike emission. A quantitative model explains the LFP dynamics by the orchestrated spiking activity in neuronal assemblies that constitute the observed excess synchrony. From the pairwise correlation analysis, we infer that neurons participate in several

assemblies but contribute only a fraction of their spikes to temporally precise assembly activity, suggesting a dual coding scheme of rate and synchrony. Revealing that transient spike synchronization correlates not only with behavior, but with a mesoscopic brain signal corroborates its relevance in cortical processing, and opens new perspectives for decoding population signals in brain-machine interfaces and diagnostics.

Introduction

Ongoing efforts to unravel the mechanisms governing brain processing have led to the proposal that information is conveyed by the functionally coordinated discharge patterns of specific neuronal subgroups (assemblies). The signature of such an assembly coding scheme is behavior-specific synchronous spiking (1) with millisecond precision observed in parallel recordings of neuronal activity (2,3). Despite recent advances in tackling the experimental (4-6) and theoretical (7,8) difficulties in finding such signatures, the local field potential (LFP) is recorded as an alternative mesoscopic measure of neuronal population dynamics. The LFP is a spatially weighted average of the synaptic transmembrane currents (9,10) and its oscillatory structure is hypothesized to reflect predominantly synchronized synaptic input (11). More specifically it has been conjectured that LFP oscillations may represent an alternative network-averaged signature of assembly activations (12,13) that enable the timed binding of features coded by different assemblies (14).

Indeed, features of the LFP signal correlate with external stimuli (15), behavioral aspects (16), internal processes (17-19), and attentional modulation (20). To date, it has been established that single spikes can become transiently entrained to the LFP in a rhythmic or non-oscillatory fashion (21,22). The degree of phase locking between neurons and the LFP depends in general on the strength of LFP oscillations (23), and the phase relationships are functionally informative (24-26). The average postsynaptic effect in the LFP triggered on spikes initiated across a patch of cortex is predictive of the LFP at a given site (27). In LFP spindles spiking activity often becomes oscillatory (22), and firing rate profiles correlate with gamma band LFP power when the level of interneuronal correlation on slow time scales is high (28). In contrast, distinct spike patterns across neurons and their phase relationship to LFP oscillations encode a substantial amount of surplus of information about the stimulus compared to information contained in the firing rate alone (29). Despite these advances, the critical link between the dynamics of precise interneuronal spike-spike correlations and the LFP on a trial-by-trial basis is missing. Furthermore, in the absence of a network oscillation in the spiking activity (Fig. S1B-E, cf. (30)) there

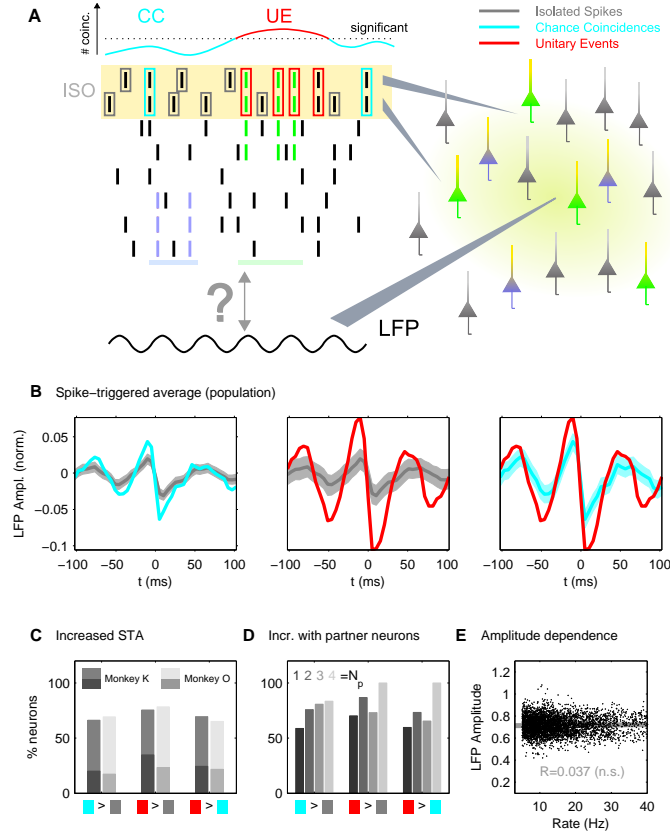


Figure 1: The magnitude of the spike-triggered average (STA) depends on the occurrence of synchronized spiking activity. (A) Sketch of the analysis. Spikes of two neurons (yellow background) and an LFP are recorded from electrodes separated by approximately $400 \mu\text{m}$ (right). Spikes are classified as isolated (ISO, gray), chance coincidence (CC, cyan), or Unitary Event (UE, red) depending on their precise synchronization with a spike of a second neuron recorded in parallel. In contrast to CCs, UEs identify coincidences in transient epochs where the high number of observed coincidences (top left) significantly exceeds the prediction based on the firing rates. In UE epochs, synchrony between both neurons in excess of the chance contribution is explained by their specific co-activation in a neuronal ensemble, termed assembly. Two assemblies are sketched in green and blue but the recorded neurons participate only in the green one. We investigate the relationship of the two types of observed spike synchrony (CC and UE) to the LFP. (B) STA of the LFP averaged over all 123 neurons ($n=297484$ spikes total) for the three disjunct sets of spikes. The left panel compares STAs of ISO (dark gray curve, $n=240455$) to CC (cyan curve, $n=44913$). To account for the difference in variability due to sample sizes, the STA of ISO is repeatedly recomputed using only 44913 random trigger spikes. The light gray band encloses at each point in time 95% of all recomputed STAs. The middle and right panel compare STAs of UE ($n=12116$) to ISO and CC, respectively. (C) Relative number of neurons per animal (vertical) with the STA of one spike set exceeding (in area) the STA of the other set (horizontal, color codes). The STA of the first set qualifies as larger if it exceeds the other STA in 50% of 1000 recomputations (superimposed darker bars: 95%, i.e. $\alpha=5\%$). (D) The four bars distinguish STAs obtained for neurons with the same number N_p of partner neurons used in coincidence detection. Same criteria (50%, both animals) as in (C). (E) The correlation of LFP amplitude and spike rate is not significant ($\alpha=0.01$, coefficient R).

is no intuitive correspondence between spike synchrony and spatially extended LFP oscillations (Fig. S1A, cf. (31)). Recent studies succeeded to directly relate synchronized subthreshold membrane potential oscillations to LFPs, but did not find such a simple relationship for synchronized action potentials of the same neurons (18). This discrepancy between subthreshold dynamics and spiking activity is in agreement with theoretical work linking subthreshold and suprathreshold dynamics in balanced cortical states (32). Therefore, synchrony of action potentials is expected to be dominated by processes that deliver strong, precisely timed, and specific inputs to the cells (18), such as required for the reliable activation of neuronal assemblies.

Here, we uncover the missing link between observed precise spike synchrony and LFPs by directly relating these two observables. For this purpose, we concentrate on timing experiments in motor cortex of monkey, where both spike synchrony (2) and LFP oscillations in the beta band (17) have been shown to be behaviorally relevant in tasks involving movement preparation. Using the Unitary Events analysis (33,34), we identify transient periods where the spiking activity of simultaneously recorded sets of neurons shows a surplus of coincidence events compared to the number expected on the basis of the firing rates (for a schematic illustration, see Figure 1A). During these periods the excess synchrony is attributed to the synchronous firing of both observed neurons as part of a network process that activates a specific subset of neurons: the assembly (Figure 1A depicts the spikes of two different assemblies in green and blue). We show that synchronous spikes originating in part from transient assembly activation exhibit pronounced phase-locking to the LFP, whereas the locking of chance coincidences is predicted on the basis of spike-LFP coupling of non-coincident spikes. This finding provides direct evidence for the hypothesized relation that precise spike synchrony constitutes a major temporally and spatially organized component of the LFP, embedding the notions of synchrony in LFP oscillations and transient spike synchrony into a single framework. Furthermore, these results enable us to estimate the fractional contribution of assemblies to the neuronal activity, exploiting signatures of assembly activity on the spike and the LFP level. Revealing that assembly activity detected as a significant surplus of spike synchrony correlates not only with behavior (2), but also with a mesoscopic brain signal corroborates its relevance in cortical processing. Our findings reinterpret the dynamical features of the LFP in terms of brain processing and open a new perspective for decoding of an accessible and reliable signal in brain-machine interfaces and diagnostics.

Results

We analyze spike data of 143 single units and simultaneously recorded LFP data from motor cortical areas in two monkeys during the instructed delay (preparatory period, PP) of two motor tasks (see SI Methods). LFPs and spikes were recorded from different electrodes spaced at 400 μm (see sketch of analysis in Figure 1A) to exclude trivial signal correlations induced by volume conduction effects (cf., e.g. (35)). We identify precise spike synchrony (36) between all neuron pairs of a given neuron and classify the spikes recorded from each neuron (**all spikes**) exclusively into one of three sets: **isolated spikes (ISO)**, **chance coincidences (CC)**, and **Unitary Events (UE)**. Spikes involved in pairwise coincidences (within 3 ms) are classified as CC if they occur during time periods where the observed coincidence rate is explained by the instantaneous trial-by-trial rates of the two involved neurons, and as UE if their number significantly exceeds the expectation (see Methods). However, no method is available able to identify the coincidences that result from the process causing the excess of synchrony during an UE period, i.e. those coincidences stemming from the activation of the assumed assembly. Therefore, a substantial fraction (see Discussion for an estimate) of coincidences in the UE group is due to chance coincident spiking (e.g., rightmost UE coincidence in Figure 1A). Spikes not classified as CC or UE with respect to any of the simultaneously recorded neurons (2-5) are classified as ISO. Consequently each spike is labeled according to the type of event it belongs to, and an individual spike train may contain spikes of different categories (compare gray, cyan, and red boxes in Figure 1A, respectively).

As a first step, Figure 1B compares the spike-triggered averages (STAs) of the LFP for the three sets, where each STA is pooled across all neuron-LFP pairs. We observe that the magnitude of the STAs of both chance coincidences (left, cyan) and Unitary Events (middle, red) significantly exceed that of the isolated spikes (gray). Moreover, the spike-triggered average of UE is larger than that of CC (right). The oscillatory structure of the STAs exhibits a strong beta frequency component, and the STAs are typically centered on the downward slope of the oscillation cycle. Non-averaged, single-neuron STAs do not exhibit these differences between the three groups as strongly (see Figure S2A for a typical example). The reason for this is two-fold: First, individual sessions have a substantially higher sampling variance, especially considering the typically low number of UE spikes. Second, STA shapes result from the combination of three effects: instantaneous LFP frequency, spike-LFP phase locking and oscillation amplitude. Nevertheless, the STA increase, in particular for UE spikes, is observed in a significant number of single neurons of both monkeys

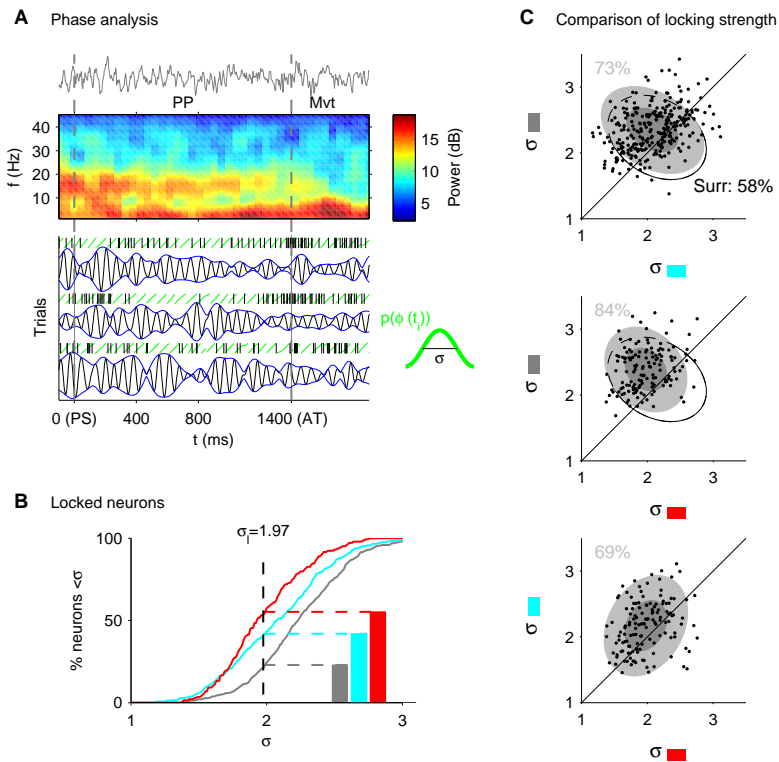


Figure 2: LFP-spike phase coupling reveals locking increase for coincidences. (A) Determination of phase and amplitude (example neuron). Top: single LFP trial; middle: trial-averaged power spectrogram. The beta activity during the preparatory period (PP, between PS and AT) disappears with movement (Mvt). Bottom: Phase (green) and amplitude (blue) of the beta-filtered LFP (upper trial shown in the top graph) extracted at the spike times (ticks). Resulting spike-triggered phase distributions (green) are characterized by their circular standard deviation σ . (B) Percentage of neurons in ISO (gray curve), CC (cyan), and UE (red) with a circular standard deviation of the phase distribution below σ (horizontal axis). For the average $\sigma_l = 1.97$ of the set of significantly locked neurons (all spikes, $\alpha = 0.05$) the percentages are also shown as bars. (C) Comparisons of the circular standard deviations σ of the three sets in the individual neurons: ISO vs. CC (top, $n = 291$), ISO vs. UE (middle, $n = 143$), and CC vs. UE (bottom, $n = 137$). Each dot represents one neuron in one experimental configuration. The percentages show the relative number of data points above the diagonal. The light (dark) gray ellipse covers 2 (1) standard deviations of the sample variance (outlined ellipse: surrogate data ISO vs. CC with shuffled ISIs).

(Figure 1C) and is consistently more pronounced for experiments where we were able to evaluate a larger number of partner neurons N_p for potential coincidences (Figure 1D), thus better identifying CC and UE groups.

Two mechanisms could underlie the differences in the STAs: changes in LFP amplitude or changes in the locking between LFP and spikes. However, the LFP amplitude does not co-vary with spike rate (Figure 1E). Therefore increased amplitudes and the disproportionate increase of the chance coincidence count during periods of elevated rates is an improbable cause of the STA increase for CC. In addition, spike histograms triggered on the peaks of the LFP oscillations (Figure S2B) reveal that spikes do not only tend to prefer the falling phase, but also avoid the rising phase of the LFP. This suggests that the three sets of spikes differ in the degree of phase coupling to the LFP rather than in the accompanying amplitude of the LFP.

Nevertheless, in order to clearly differentiate between these mechanisms, it is necessary to formally disentangle the dependence of spike timing on the amplitude of the LFP from its dependence on the phase. Figure 2A explains the procedure (for details see SI Methods). For both monkeys we consistently observe a prominent beta oscillation (in both monkeys around 15 Hz) of the LFP during the preparatory period that stops with movement onset (Mvt). Therefore we focus on the beta frequency band and extract the instantaneous phase and amplitude (envelope) of the field potential for each spike time. Compared to the STA analysis, even individual neurons exhibit clear and specific differences between ISO, CC, and UE in both measures (Figure S3, same example neuron as in Figure S2). We are now prepared to study the two contributions in detail across the population.

Figure 2B shows that across the population of neurons CC are systematically better locked (decreased circular standard deviation σ of the phase distribution) than ISO, and UE better than CC. As a suitable reference value to compare the fraction of locked neurons in the 3 sets we extracted the average locking strength $\sigma_l=1.97$ obtained for those neurons that are significantly locked if all spikes are considered (surrogate test). In the following we investigate how the systematic differences in locking strength between the three sets of spikes are affected by the intrinsic spike-LFP relationship of the neurons, i.e. if a neuron in general tends to lock well to the LFP or not. Differentiating groups of strongly (38%) and weakly (62%) locked neurons (i.e., significantly locked and unlocked neurons considering all their spikes) does not introduce a bias by affecting the percentage of neurons that exhibit CC and UE (Figure S4A). Both groups exhibit the same general pattern of locking in the three groups (Figure S4B) shown in Figure 2B. As expected, the percentage of neurons better locked than σ_l in the ISO group differs considerably (51% vs. 5%) between strongly and weakly locked

neurons. However, this difference between strongly and weakly locked neurons is less pronounced for CCs (62% vs. 28%) and further decreases for UEs (67% vs. 46%). The conservation of the locking of UE spikes in strongly and weakly locked neurons compared to the declines for ISO and CC hints at different dynamical origins of the spikes in CC and UE.

Figure 2C confirms that individual neurons are consistent with the findings for population ratios (Figure 2B). The scatter plots of the circular standard deviation reveal that in 73% of the recorded neurons CC spikes are better locked than ISO spikes, and in 84% of the neurons UE spikes are better locked than ISO spikes. Finally, in 69% of all neurons UE spikes are better locked to the LFP than CC spikes. In contrast to the experimental data, only 58% of surrogate spike trains that retain the original inter-spike interval statistics show an increase in phase locking for coincident spikes (outlined ellipse).

Because of the consistency in the population, in the following we focus on the phase locking of strongly locked neurons. The selection is conservative because it reduces the differences in locking between the three sets of spikes (Figure S4B). Comparable results are obtained for the complete set of recorded neurons. The phase distributions in Figure 3A (top) show that locking of spikes to the LFP is strongest for Unitary Events, and weakest for isolated spikes.

The phase distribution exhibited already by isolated spikes modulates the spiking probability in time. Given the high level of synchrony between LFPS (Fig. S1A), one may therefore argue that the increased modulation of the phase distribution of CC trivially results from the individual phase locking distribution of the two neurons forming the coincidence (predictor assuming independence of neurons, see SI Methods). Interestingly, the phase distribution of CC is indeed largely in agreement with this predictor (black curve), while that of UE is not. Hence, despite the impossibility to remove the substantial fraction of chance coincidences from the UE group, the locking of UE cannot be explained on the basis of the intrinsic phase locking of the neurons forming the coincidences.

Earlier studies (22,23) demonstrate that spikes occurring during periods of high LFP amplitudes exhibit a stronger locking to the LFP (see Supporting Text and Figure S5). This raises the question of whether coincidences, and in particular Unitary Events, predominantly occur at high LFP amplitudes. Figure 3A (density plots) shows the number of spikes as a function of both LFP phase and amplitude for each of the three sets ISO, CC, and UE. Here, CC and UE occur at similar amplitudes as ISO, even though the amplitude distributions (left) reveal a small shift towards high amplitudes for CC and UE. The phase distributions (top panels), however, clearly show a progressive increase in the degree of phase locking from ISO to CC to UE. Finally, observing that UEs exhibit similar amplitudes as CC, we can ask the reverse question of whether at

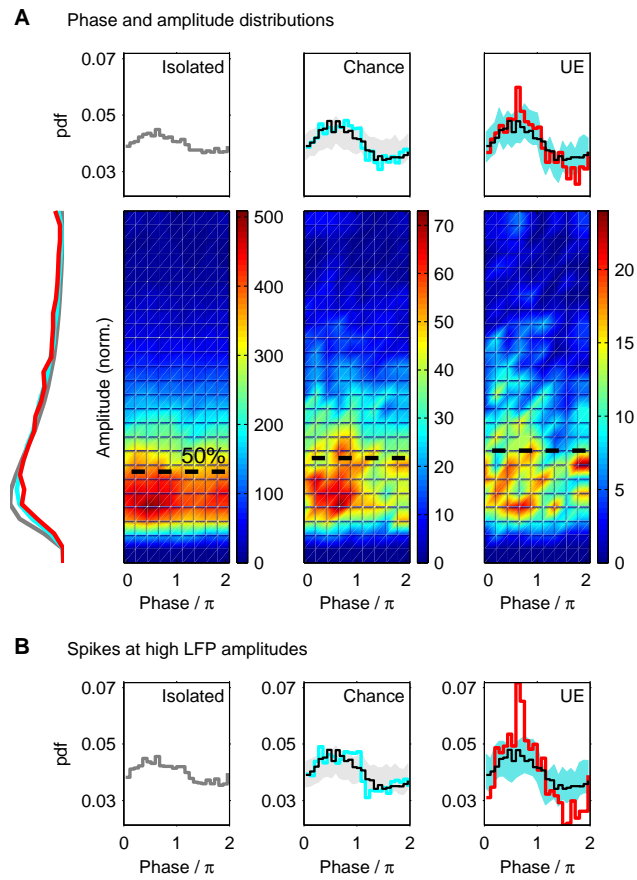


Figure 3: Relation of spike synchrony to the interplay of phase and amplitude. (A) Joint histograms of the phase and amplitude for ISO (left), CC (middle), and UE (right) pooled across the population (color bars indicate counts; phase π indicates LFP troughs). The top and left projections display the phase and amplitude distributions, respectively. The top middle and top right graph compare the phase distribution to the distribution shown in the graph to the left: The shaded areas enclose at each phase 95% of 1000 phase distributions randomly chosen from the set to the left with the same number of spikes as in the current set. Black curves are the predictions based on the phase distributions of the individual neurons. The histograms include the neurons that have a minimal spike count (total of 25 spikes and a mean rate of 5 Hz per trial) and for which the phase distribution of all spikes is significantly locked ($\alpha=0.05$). (B) Phase distributions of the three sets, considering only 50% of spikes at the highest LFP amplitudes (above dashed black line in (A)).

high amplitudes ISO, CC and UE still exhibit the systematic increase in locking. Figure 3B shows that for the 50% of the spikes occurring at the largest LFP amplitudes (above black dashed line in Fig. 3A) the effect of improved phase locking for the UE group is strongly amplified. In contrast, the ISO and CC phase distributions do not change. This finding reveals that those coincidences in UE periods that are responsible for the increased locking of UE are those that occur during strong LFP oscillations.

Discussion

In this report we explicitly reveal how the neuronal population dynamics reflected by LFP oscillations relate to the synchronous discharge of individual neurons in motor cortex. Spikes which are emitted at the same time as spikes of other neurons exhibit a better phase locking to the dominant beta-range LFP oscillation than those which occur in isolation. However, in time periods where the number of spike coincidences is at chance level, the quality of the locking is to a large extent explained by a predictor assuming independence of the spikes constituting a coincidence. In contrast, the pronounced locking to the LFP in time periods with a significant excess of coincident spikes (Unitary Events) cannot be explained in this way. The probability of the occurrence of coincident spikes is only weakly coupled to changes in the magnitude of the LFP signal. Nonetheless, spikes that coincide with episodes of high LFP amplitudes are on average better locked to the LFP than those at low amplitudes. A separate analysis of these two factors, spike synchrony and LFP magnitude, demonstrates that both affect the strength of the spike-LFP coupling largely independent of each other. What conclusions about network dynamics and possible coding mechanisms do these results imply, in particular in the light of the distinctive role of Unitary Events?

Several authors have elucidated the functional role of LFP oscillations in motor cortex in the beta and lower gamma range. These oscillations are only loosely correlated across trials, i.e. their phase is not time-locked to any external (e.g. stimulus) or internal (e.g. movement onset) event. Oscillatory beta range LFP activity in motor cortex is a unique feature of experimental protocols including a waiting period before movement execution and has been described in relation to attentional processes, movement preparation and motor maintenance (12,15,17,37-39). The oscillations terminate at movement onset and may well represent a top-down modulatory input from higher sensory areas (e.g., (40)). Furthermore, there is a large body of knowledge about delay-related spiking activity in motor cortical areas and its functional implication in sensorimotor integration and movement preparation (for a review, see (41)). Finally, tran-

sient spike synchrony observed among individual neurons is remarkably well related to timing-related aspects of the behavioral task (2) but does not depend on the mean firing rate of the participating neurons (42). However, only a few studies relate LFP oscillations to correlations of the spiking activity (22,28). Reports in various brain areas demonstrate single neurons which selectively participate in oscillatory periods of the LFP by phase locking (20,21,37,43), where occasionally the autocorrelations of the spike trains become oscillatory (22,40). In conclusion, the apparent complexity of the simultaneous coding of neuronal activity for different aspects of motor cortical processing challenges the idea that LFP oscillations and the emergence of transient UEs are two reflections of only one single functional process performing the planning and preparation of movements.

It is reasonable to assume that synchrony on a spike-by-spike level, and population oscillations expressed by the LFP both originate from network processes that involve the pulsed, synchronous co-activation of specific subsets of neurons. One may argue that in this case we should observe an even more distinct relationship between the two measures. However, our techniques to detect synchrony related to the activation of neuronal assemblies are limited. The Unitary Event analysis assesses indirectly which coincidences are more likely to originate from such activations based on the comparison of the time-resolved rate of observed and expected coincidences. Nevertheless, the set of UEs may be composed of coincidences resulting from assembly activation and a considerable fraction of chance coincidences (see estimate below). Therefore, although the difference in locking precision between significant (UE) and non-significant (CC) time segments seems small at first glance, in this light it is even more surprising that we are able to observe an enhanced phase locking for the UEs. The argument implies that the subset of coincidences caused by assembly activation has a tight locking to the LFP. This conclusion is supported by previous work demonstrating that coherent membrane potential oscillations do not generate synchronized output spikes, and that brief, simultaneous synaptic inputs to a cell are the likely drive for action potential generation (18).

Unitary Events prefer a particular phase of the LFP oscillation, a signal which is rather homogeneous across the motor cortex (17,31). This finding renders unlikely a model of processing where assemblies can be simultaneously active and still distinguished (multiplexed) by locking to different phases of the oscillatory cycle (e.g., (26)). Moreover, in such a model the waxing and waning of the LFP oscillation would likely show phase shifts as different assemblies become active. Our results insinuate that neurons participate in different assemblies at different times (see also (2)), but predominantly at the same phase of the LFP (cf., (13)). We observe the phenomenon in 20-30% of the neurons

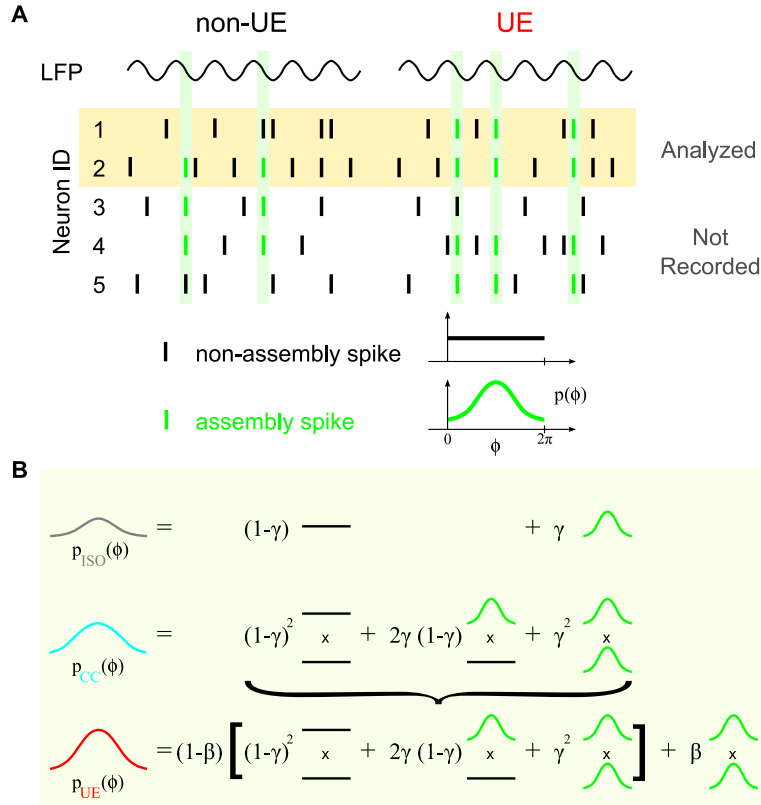


Figure 4: A conceptual model relating increased LFP locking and assemblies. (A) Sketch of the LFP (top) and the simultaneous spiking activity of five neurons (middle), of which only two are recorded (yellow background). Based on the latter, time periods where coincidences occur at chance level (non-UE, left) are distinguished from those with excess synchrony (UE, right). Each spike is either part of an assembly of co-active neurons (green) or not (black). In this simplified scenario, one assembly is active on the left, and a different one on the right; both observed neurons contribute to the latter. Only assembly spikes exhibit locking to the LFP, expressed by a non-uniform phase distribution $p(\phi)$ (green). (B) Two ratios β and γ determine the composition of the phase distributions for ISO, CC, and UE (left) of assembly and non-assembly spikes. γ determines the overall probability that a spike is part of an assembly activation (top, ISO). $p_{\text{CC}}(\phi)$ (middle) results from the combinatorics of two independent spike trains (ISO). $p_{\text{UE}}(\phi)$ (bottom) differs from $p_{\text{CC}}(\phi)$ by the relative excess β of assembly spikes in UE periods. A conservative (minimal) estimate of β , i.e., maximally locked $p^2(\phi)$, is obtained by substituting $p_{\text{UE}}(\phi)$ and $p_{\text{CC}}(\phi)$ in the bottom equation by the experimental distributions. γ is determined from either of the top two equations by using $p(\phi)$.

in agreement with estimates from other studies (e.g., (22)). However, even in this category of neurons we can attribute only a fraction of spikes to assembly activation. One hypothesis is that the motor cortex is involved in parallel coding schemes, where synchronous assembly activity can be dissociated from the rate-based continuous-time coding.

To better understand the implications for the organization of cortical processing we consider a conceptual model where spikes of a neuronal assembly are locked to the LFP (Figure 4) based on (i) the assumption that UEs reflect assembly activity (2) and (ii) our observation that UEs have the strongest locking to the LFP. A potential mechanism is that assembly spikes originate from synchronous synaptic input to local groups of neurons. The simplest explanation for the finding that ISO and CC also exhibit locking, albeit weaker than UE, is that the spikes of a neuron are composed of a mixture of non-assembly (unlocked) and assembly spikes (locked). The latter are not identified as UE due to the lack of corresponding partner neurons in the recording (Figure 4A). Consequently, the phase histogram of the ISO spikes is a superposition of the histograms of non-assembly and assembly spikes, with a factor γ determining their ratio (Figure 4B, top row). Chance coincidences are composed of spikes from independent sources (Figure 4B, middle row) but the combinatorics of non-assembly and assembly spikes enhances the locking. Finally, periods identified as UE contain excess coincidences (Figure 4B, bottom row) resulting from the activation of an assembly in which both neurons participate. Their relative contribution β leads to an enhanced locking of UE compared to CC. The structure of the model allows us to derive conservative estimates of the parameters γ and β from the experimental phase histograms. We find that outside of UE periods $\gamma=11\%$ of the spikes of a neuron participate in an assembly, and $\beta=23\%$ of the coincidences in UE periods result from the joint participation in an assembly. Even though this is clearly a highly simplified model, it provides a first quantitative bridge between functionally relevant spike synchrony (2,13,44) and the LFP as a robust mesoscopic measure of brain activity (45).

Our results show that neuronal mass signals like the LFP convey specific information about network processes. We directly demonstrate in the brain of a behaving animal that the LFP is related to spike synchronization. Nevertheless, there is a substantial fraction of spikes without an apparent relationship to the LFP. Thus the two measures are observables of the same neuronal network but do not necessarily carry the same information. Taken together, we interpret our results as evidence that LFP (beta) oscillations, especially at high amplitudes, are reflections of the activation of neuronal assemblies which propagate a synchronous volley through the network. With massively parallel recordings becoming available we may be able to disambiguate the superposition of multiple

neuronal assemblies. This gives us confidence that by improving our understanding of the various components of the LFP signal we will eventually be able to use the LFP as an antenna delivering news from several communicating network stations.

Methods

Coincidence detection and Unitary Event Analysis

From simultaneously recorded spike data of individual sessions we extract all unique pair combinations of spike trains that are recorded from distinct electrodes. Coincident spike events are derived by counting coincident spike occurrences of both neurons (compare Figure S6). To allow coincidences with a temporal jitter up to a maximal coincidence width of $b=3$ ms, we apply the 'multiple-shift' approach (36,42). In this method exact coincidences (within the time resolution h of the data) are detected for various shifts of the second spike train against the first (reference) spike train. The shift is stepwise increased from 0 in steps of $h=1$ ms up to $b=3$ ms (for positive and negative shifts). Detected coincidences are marked for further analysis.

To account for the non-stationarity of the neurons' firing rates, and to capture the dynamics of correlation, we perform the Unitary Event (UE) analysis in a sliding window fashion (34). This is done by moving a window of fixed duration (here: $T_w=100$ ms) along the data to cover the duration of a trial, i.e. the duration of the PP. The length of the time window is chosen large enough to include at least one complete cycle of the beta oscillation. The window is advanced in steps corresponding to the time resolution h of the data. The first window position is centered at trial onset, and the last window at the end of the delay period.

Within each window position the total number of empirical coincidence counts n_{emp} is derived by summing the exact coincident spike events from each shift l and from all M trials j : $n_{emp} = \sum_{j=1}^M \sum_{l=1}^L n_{emp}^{j,l}$, with $L=2(b/h)+1$. To derive UEs this count is compared to the number of coincidences that would occur by chance given the firing rates of the neurons. This involves the following calculations. To account for non-stationary rates across trials (46), the relevant measures are obtained from the single trial and only subsequently summed across trials. Thus, within the analysis window the expected number of coincidences is calculated on the basis of the trial by trial firing probabilities $p_{i,j}$ which are estimated by the spike count $c_{i,j}$ of neuron i in trial j divided by the number of bins N within a window: $p_{i,j}=c_{i,j}/N$ with $N=T_w/h$. The joint probability for finding a coincidence by chance per trial is calculated by the product of the

single neuron firing probabilities $p_{12,j} = p_{1,j} p_{2,j}$. The expected number of coincidences per trial j results from multiplying this probability with the number of bins N that are included in the analysis window and the number of shifts L : $n_{\text{exp}}^j = NL p_{12,j}$. The total number of expected coincidences within the window is derived from the sum of the expected numbers per trial: $n_{\text{exp}} = \sum_{j=1}^M n_{\text{exp}}^j$.

Finally we compare the empirical n_{emp} to the expected number n_{exp} of coincidences to detect significant deviations. To this end, we calculate the joint- p -value jp , i.e. the probability of measuring the given number of empirical coincidences (or an even larger number) under the null-hypothesis of independent firing. The distribution under this null-hypothesis representing the probability to find a given number of coincidences is given analytically assuming Poisson processes (33). The latter assumption is shown to yield a conservative estimate for cortical spike trains considering their non-Poisson and non-renewal properties (8). Then the significance of n_{emp} yields (33): $jp(n_{\text{emp}}|n_{\text{exp}}) = \sum_{r=n_{\text{emp}}}^{\infty} \frac{n_{\text{exp}}^r}{r!} e^{-n_{\text{exp}}}$. If its value is below an *a priori* threshold (here chosen as 5%) coincident firing is classified as significant and identified as Unitary Events. Spikes are labeled as UE if they are part of at least one sliding window identified to contain significant excess synchrony (for an illustrated summary of this analysis approach, see (44)). In addition, we require such time windows to exhibit a minimum firing rate of 5 Hz for each neuron. Spikes that are part of coincident events but not identified as UE with respect to any of the neurons recorded in parallel are labeled as chance coincidences (CC), all remaining spikes as isolated spikes (ISO).

A thorough description of the experimental protocol and the methods used to relate spikes and LFP, can be found in SI Methods.

Acknowledgments

We thank Moshe Abeles, Walter Freeman, and George Gerstein for valuable comments on an earlier version of the manuscript. Partially supported by the Volkswagen Foundation, the Stifterverband für die Deutsche Wissenschaft, the BMBF (grant 01GQ0413 to BCCN Berlin), the Helmholtz Alliance on Systems Biology, the French National Research Agency (ANR-05-NEUR-045-01), DAAD, The Research Council of Norway (eScience Programme) and EU grant 15879 (FACETS).

References

1. Gerstein GL, Bedenbaugh P, Aertsen, MH (1989) Neuronal assemblies. *IEEE Trans Biomed Eng* 36: 4-14.
2. Riehle A, Grün S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278: 1950-1953.
3. Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y et al. (1995) Dynamics of neuronal interaction in monkey cortex in relation to behavioral events. *Nature* 373: 515-518.
4. Euston DR, Tatsuno M, McNaughton BL (2007) Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* 318: 1147-1150.
5. Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11: 823-833.
6. Nicolelis MA, Ghazanfar, AA, Faggin BM, Votaw S, Oliveira LM (1997) Reconstructing the engram: simultaneous, multisite, many single neuron recordings. *Neuron* 18: 529-537.
7. Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci* 7: 456-461.
8. Grün S (2009) Data driven significance estimation of precise spike correlation. *J Neurophysiol* 101: 1126-1140.
9. Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and EEG phenomena. *Physiol Rev* 65: 37-100.
10. Viswanathan A, Freeman RD (2007) Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nat Neurosci* 10: 1308-1312.
11. Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. *Annu Rev Physiol* 66: 735-769.
12. Donoghue JP, Sanes JN, Hatsopoulos NG, Gaál G (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol* 79: 159-173.
13. Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24: 49-65.
14. Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W et al. (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol Cybern* 60: 121-130.
15. O’Leary JG, Hatsopoulos NG (2006) Early visuomotor representations revealed from evoked local field potentials in motor and premotor cortical areas.

J Neurophysiol 96: 1492-1506.

16. Scherberger H, Jarvis MR, Andersen RA (2005) Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron* 46: 347-354.

17. Murthy VN, Fetz EE (1996) Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol* 76: 3949-3967.

18. Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454: 881-885.

19. Roux S, Mackay WA, Riehle A (2006) The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behav Brain Res* 169: 335-351.

20. Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291: 1560-1563.

21. Eckhorn R, Obermueller A (1993) Single neurons are differently involved in stimulus-specific oscillations in cat visual cortex. *Exp Brain Res* 95: 177-182.

22. Murthy VN, Fetz EE (1996) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol* 76: 3968-3982.

23. Denker M, Roux S, Timme M, Riehle A, Grün S (2007) Phase synchronization between LFP and spiking activity in motor cortex during movement preparation. *Neurocomp* 70: 2096-2101.

24. Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7: 862-871.

25. Harris KD, Henze DA, Hirase H, Leinekugel X, Dragoi G et al. (2002) Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* 417: 738-741.

26. Womelsdorf T, Schoffelen J, Oostenveld R, Singer W, Desimone R et al. (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316: 1609-1612.

27. Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nat Neurosci* 12: 70-76.

28. Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A et al. (2007) Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Curr Biol* 17: 1275-1285.

29. Kayser C, Montemurro MA, Logothetis NK, Panzeri S (2009) Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61: 597-608.

30. Nawrot MP, Boucsein C, Rodriguez-Molina V, Riehle A, Aertsen A et al.

- (2008) Measurement of variability dynamics in cortical spike trains. *J Neurosci Meth* 169: 374-390.
31. Rubino D, Robbins KA, Hatsopoulos NG (2006) Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci* 9: 1549-1557.
32. Tetzlaff T, Rotter S, Stark E, Abeles M, Aertsen A et al. (2008) Dependence of neuronal correlations on filter characteristics and marginal spike train statistics. *Neural Comput* 20: 2133-2184.
33. Grün S, Diesmann M, Aertsen A (2002) Unitary events in multiple single-neuron spiking activity: I. Detection and Significance. *Neural Comput* 14: 43-80.
34. Grün S, Diesmann M, Aertsen A (2002) Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput* 14: 81-119.
35. Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL et al. (2009) Local origin of field potentials in visual cortex. *Neuron* 61: 35-41.
36. Grün S, Diesmann M, Grammont F, Riehle A, Aertsen A (1999) Detecting unitary events without discretization of time. *J Neurosci Meth* 94: 67-79.
37. Baker SN, Olivier E, Lemon RN (1997) Coherent oscillations in monkey motor cortex and hand muscle EMG show task-dependent modulation. *J Physiol* 501: 225-241.
38. Murthy VN, Fetz EE (1992) Coherent 25- to 35-Hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc Natl Acad Sci U S A* 89: 5670-5674.
39. Sanes JN, Donoghue JP (1993) Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proc Natl Acad Sci U S A* 90: 4470-4474.
40. Lebedev MA, Wise SP (2000) Oscillations in the premotor cortex: single-unit activity from awake, behaving monkeys. *Exp Brain Res* 130: 195-215.
41. Riehle A (2005) Preparation for action: one of the key functions of the motor cortex. In: *Motor cortex in voluntary movements: a distributed system for distributed functions*, eds Riehle A, Vaadia E (CRC-Press, Boca Raton, FL), pp. 213-240.
42. Grammont F, Riehle A (2003) Spike synchronization and firing rate in a population of motor cortical neurons in relation to movement direction and reaction time. *Biol Cybern* 88: 360-373.
43. Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep phases. *J Neurosci* 19: 4595-4608.
44. Maldonado P, Babul C, Singer W, Rodriguez E, Berger D et al. (2008) Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *J Neurophysiol* 100: 1523-1532.

45. Mehring C, Rickert J, Vaadia E, de Oliveira SC, Aertsen A et al. (2003) Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci* 6: 1253-1254.
46. Grün S, Riehle A, Diesmann M (2003) Effect of cross-trial nonstationarity on joint-spike events. *Biol Cybern* 88: 335-351.

Supporting Information

Methods

Ethics Statement

Care and treatment of the animals during all stages of the experiments conformed to the European and French government regulations, according to the Weatherall report ('The use of non-human primates in research', December 2006).

Experimental design and electrophysiological recordings

All data were taken from recordings partially presented elsewhere (1). Two rhesus monkeys (monkey K and monkey O) were trained to perform arm movements from a center position to one of two possible peripheral targets left and right of the center in two different tasks involving an instructed delay. In the first, a choice reaction time task (chRT), both peripheral targets were presented simultaneously as a preparatory signal (PS), one in red and the other in green. The animal learned to attribute to each color one of two possible delay durations. If the (directionally non-informative) auditory response signal (RS) occurred after a short delay, the monkey had to select the red target, after a long delay the green one. Both the laterality of the colored targets and the presentation of the two durations were varied at random with equal probability. In contrast, in the second self-paced movement task (SELF), the presentation of only one peripheral target, either in red or green, either at the left or the right, required a self-initiated response after estimating one of the two delays as coded by PS. In both tasks (1), four different timing patterns were used to identify the short and long delay, respectively: (i) 500 ms and 1000 ms (monkey K); (ii) 500 ms and 1200 ms (monkey K); (iii) 600 ms and 1200 ms (monkey O); (iv) 1000 ms and 1400 ms (monkey O).

In this study we exclusively analyzed the delay activity, i.e. activity recorded during the preparatory period (PP) starting at PS and ending with either RS in the chRT task or the earliest allowed response time (AT) in the SELF task. Therefore, the trials were aligned to PS occurrence for the analysis. The neural activity related to movement execution, i.e. after RS or AT, respectively, is not analyzed. For both tasks, only correct trials were considered, in which the monkey responded within a time window (after the end of PP) of maximally 300 ms (monkey O) and 500 ms (monkey K) and in which movements were performed in the required movement direction.

In order to exclude effects due to pooling of neuronal activities of different behavioral contexts and different tasks, their activity was analyzed separately

for the four possible behavioral conditions (combinations of short or long delay duration and left or right upcoming movement direction) and each experimental session. For the sake of simplicity, we refer in this manuscript to a recorded neuron by the combination of its identity and the behavioral context during which it was recorded. In this sense, data recorded from the same neuron may enter a population average up to eight times (maximum of four different conditions in two tasks).

Data acquisition and data analysis

LFPs and spikes were recorded simultaneously in primary motor cortex using a multielectrode device of 2-4 electrodes (MT-EPS, Alpha Omega). Spikes of single neurons were detected by an online sorting algorithm (MSD, Alpha Omega, Nazareth, Israel). The inter-electrode distance was on the order of 400 μm . LFPs were sampled at a resolution of 250-500 Hz and hardware filtered (band pass, 1-100 Hz). In total, we analyzed 53 recording sessions (monkey K: 25; O: 28), which yielded 143 single neurons or 570 combinations of neurons and behavioral conditions. On average 33 ± 11 trials were recorded per experimental condition. In analyses that combine spikes and LFP, each neuron enters only once, and we never combined LFP and spikes that were recorded on the same electrode to exclude the possibility of spike artifacts in the signal. We confirmed that simultaneously recorded LFPs are highly synchronous in the frequency regimes of interest. Likewise, coincident activity between neurons was analyzed only from neurons recorded from different electrodes, totaling 123 analyzed pairs of neurons. All data analysis was performed using the Matlab software environment (The Mathworks Inc., Nattick MA).

Spectral analysis

Power spectra are used to assess the dominant frequencies in the LFP during the task. All power spectra are calculated using a Hamming window as taper. To illustrate the temporal modulation of power in different frequency bands, we use a time-resolved spectral analysis using 200 ms windows with a 50 ms overlap.

Spike-triggered averages

Spike-triggered averages (STAs) are computed by averaging LFP segments from time windows of 200 ms centered at each spike time. For the STA analysis, LFPs are filtered between 2-80 Hz to remove DC components. To compare STAs across recordings, in which electrode signals often differ in their absolute amplitude values, we z-transform each LFP before further analysis by subtracting its mean

(calculated across trials) and dividing by its standard deviation. In order to quantify the magnitude (or size) of an STA, we calculate the total area the STA encloses with the time axis. Similar results to those presented here (not shown) are obtained using alternative measures of the STA magnitude, such as the area under its envelope, or the maximum of its absolute value. The magnitude of the STA is in general dependent on the number of trigger spikes. In order to compare STAs obtained from two sets of trigger spikes of different number of spikes n_1 and n_2 ($n_1 > n_2$) we construct 1000 STAs of set 1, each computed from n_2 randomly selected spikes. We define the STA of set 2 to be larger than that of set 1 if the magnitude of set 2 exceeds 50% of the re-computations of set 1, and significantly larger (at a level of 5%) if it exceeds 95% of the re-computations.

Peak-triggered spike histograms

We evaluate the population-averaged spiking discharge triggered on the peaks of the LFP oscillation (43). To this end we detect maxima of the LFP separated by a minimum time interval of 33 ms, which corresponds to a maximal oscillation frequency of 30 Hz. The spike histogram is calculated from data within a window of 200 ms around each peak, and averaged across all individual peaks in all neurons (see (2) for a different technique to relate spike times to EEG time course based on amplitude). Simultaneously, we also compute the peak-triggered LFP by averaging the z-transformed LFP aligned on its peaks.

Rate-amplitude correlation

To assess the degree of correlation between LFP oscillation strength and spike rates, we calculate the mean value of the rectified, z-transformed LFP along each trial with sliding windows of 200 ms length and 100 ms overlap. These values are then correlated with the rate profile of the neuron estimated as the spike count across trials in the same windows. Similar results as those shown here are obtained using alternative measures of LFP strength, including the mean value of the envelope of the beta-filtered signal (compare phase-locking analysis), or by using the total signal power in the beta range (10-22 Hz).

Phase analysis

After examination of the dominant beta frequencies on a session-by-session basis, LFPS of both monkeys are filtered with a zero-phase 10-22 Hz band pass filter (Butterworth, 8-pole). Short filter transients in the time domain allow for good estimates of the instantaneous LFP amplitude. In a subsequent step, we calculate the instantaneous phase of the LFP from the analytic signal $\xi(t) = x(t) + i \tilde{x}(t)$ obtained via the Hilbert transformation

$$\tilde{x}(t) = \frac{1}{\pi} \text{P.V.} \int \frac{x(\tau)}{t - \tau} d\tau$$

of the original signal $x(t)$, where P.V. denotes that the integral is to be taken as Cauchy principal value (3). In this formalism, troughs of the LFP are identified by a phase of π . The calculation of the analytic signal can be applied to arbitrary signals, but its interpretation as instantaneous phase is difficult where either the signal amplitude becomes too small to discriminate the oscillation from background noise, or where the regular oscillation is disrupted (4). To account for these effects, we discard phase values which violate the monotonicity of the phase time series or exhibit instantaneous phase jumps. To further corroborate our results, we exclude from our analysis those 10% of spikes per neuron that occur at the lowest LFP amplitudes.

We analyze the distributions of extracted phase values at the times of spike occurrences (23) using tools from circular statistics (5). The mean phase ϕ is obtained via the circular average $R e^{i\phi} = N^{-1} \sum e^{i\phi(t_i)}$, where $\phi(t_i)$ indicates the phase of the field potential at time t_i of spike i . Furthermore, we utilize the transformation of the vector strength R to the circular standard deviation $\sigma = \sqrt{-2 \log R}$ as a measure of the concentration of the phase distribution. For small values, σ relates to the standard deviation of a normal distribution, whereas for flat distributions it behaves as $\sigma \rightarrow \infty$. In all phase analysis, we discard neurons that fire in total (across trials) less than 25 spikes.

Additionally, we employ two measures to quantify whether spikes recorded from individual neurons show a significant phase preference to the LFP. For the first, we test against the null hypothesis that the phase sample is taken from the uniform circular distribution (Rayleigh test (5)), which is expected by assuming a regular (e.g., filtered) field potential and independent random spiking. However, spike trains that have a certain regular structure in time may display intrinsic locking to the LFP. To measure the degree of genuine locking that is not explained by the regularities of the two signals, we calculate as the second measure the degree of locking R in 1000 surrogates, each created by shuffling the inter-spike intervals of the spikes on a trial-by-trial basis (random placement of the first spike). This procedure preserves to first order the regularity manifested in the inter-spike interval distribution. A comparison with the measured value R yields the p-value for this surrogate test. Since the construction of such surrogates can only be performed on the complete spike train, this measure could not be sensibly applied to the subsets of spikes in our analysis (i.e., ISO, CC, UE, as well as Lo and Hi in the amplitude analysis presented in the Supporting Information).

The phase distribution of coincident spikes may be trivially sharpened due

to a preferred phase occurrence of individual spikes. To correct for that we calculated the expected phase distribution of coincident spikes (compare black curve in Figure 3 and S3A). To this end, we calculate the joint phase probability distribution of a neuron pair by the phase-by-phase multiplication of the occurrence probabilities of spikes at these phases. The predictor for the whole population is the average of the pair-wise phase distributions weighted by the relative number of coincidences between the two neurons.

In contrast to this predictor which considers the phase of spikes irrespective of the spike interval distribution, we also construct a predictor based on the reverse scenario. For each pair of simultaneously recorded neurons the interspike intervals of the spike trains of each neuron are shuffled on a trial-by-trial basis to create a set of 1000 surrogate pairs. For each surrogate, the variance σ is evaluated separately for the resulting sets of non-coincident and coincident spikes. Thus, we obtain for each neuron the variances σ of phase locking of coincident and non-coincident spikes for the original data and for the 1000 surrogates, allowing us to compare their distributions (Figure 2C).

Supporting Text

Magnitude of global oscillations influences spike locking. At a given time the amplitude of the LFP oscillation is defined by its envelope (blue curves in Figure 2A). To examine the dependence of spike locking on the amplitude of the LFP (23), we form two exclusive sets of spikes, termed 'Hi' and 'Lo', based on whether a spike occurs at an amplitude above or below a certain value, respectively (Figure S5A). We account for the session-by-session variability of the LFP amplitude by defining the threshold θ in terms of the fraction of spikes an individual neuron contributes to the Lo category (Figure S5B).

For threshold ranges between 0.2 and 0.8 we observe that the percentage of significantly locked neurons (Rayleigh test, $\alpha=0.05$) of the Hi set is only decaying slightly from 41% to 34% (Figure S5C). This percentage is in the same range as the percentage of locked neurons considering all spikes (Figure 2B). We emphasize that even for high thresholds, where only few spikes are included, the locking of neurons can be explained using Hi spikes only. In contrast, when considering spikes of the Lo set, the percentage of locked neurons starts at 5% and increases approximately linearly with θ at a much steeper slope, meaning that at increasingly higher amplitudes more and more spikes are included in the Lo set. This shows that locking of spikes to the local field potential is largely due to spikes that occur at high LFP amplitudes.

References

1. Roux S, Mackay WA, Riehle A (2006) The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behav Brain Res* 169: 335-351.
2. Eeckman, FH, Freeman, WJ (1990) Correlations between unit firing and EEG in the rat olfactory system. *Brain Res* 528: 238-244.
3. Le Van Quyen V, Foucher J, Lachaux J, Rodriguez E, Lutz A et al. (2001) Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J Neurosci Meth* 111: 83-98.
4. Boashash B (1992) Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals. *Proc IEEE* 80: 520-538.
5. Mardia, KV, Jupp PE (2000). *Directional Statistics*, (John Wiley & Sons Ltd., Chichester).

Supporting Figures

The following pages contain supporting Figures S1-S6 and their captions.

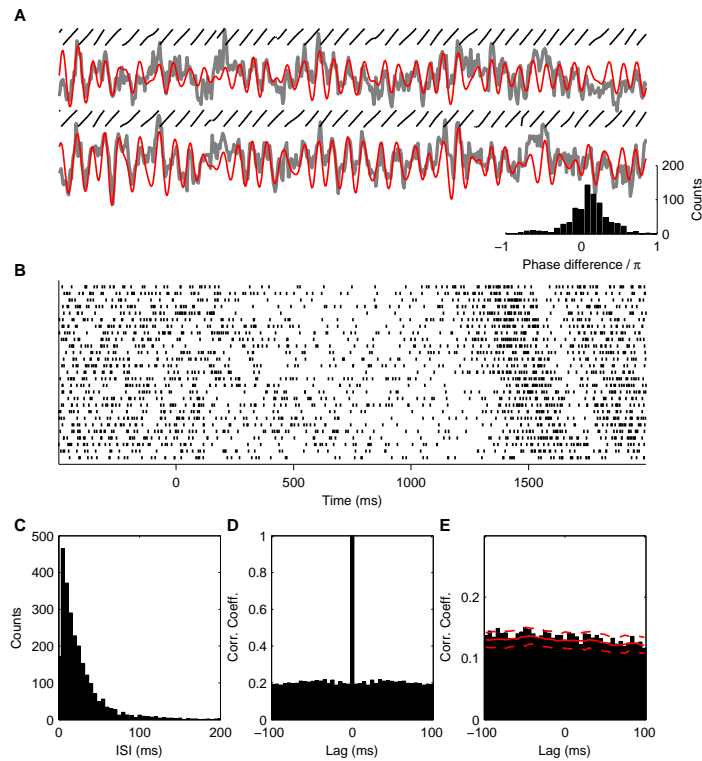


Figure S1: Characteristics of LFP and spiking dynamics. (A) Two single-trial LFPS recorded simultaneously (gray) at different electrodes (during long trials with movement to the right in the SELF task). Superimposed are the beta-filtered (10-22 Hz) signals (red) and their instantaneous oscillation phase (black lines). The histogram visualizes the phase differences between the two signals across all time bins. (B) Spike raster of one example neuron recorded in parallel to the LFP shown above (same neuron as in Figure 2). (C,D) Neither the trial-averaged inter-spike interval distribution (C) nor the normalized auto-correlograms (D) indicate an oscillatory nature of the neuron. (E) The cross-correlogram with a different neuron recorded in parallel (neuron 1 in Figure S6) remains flat. Red lines indicate mean (solid) and 5% confidence intervals (dashed) of cross-correlograms obtained from surrogate spike trains where each spike was jittered uniformly in window of ± 20 ms around its original position.

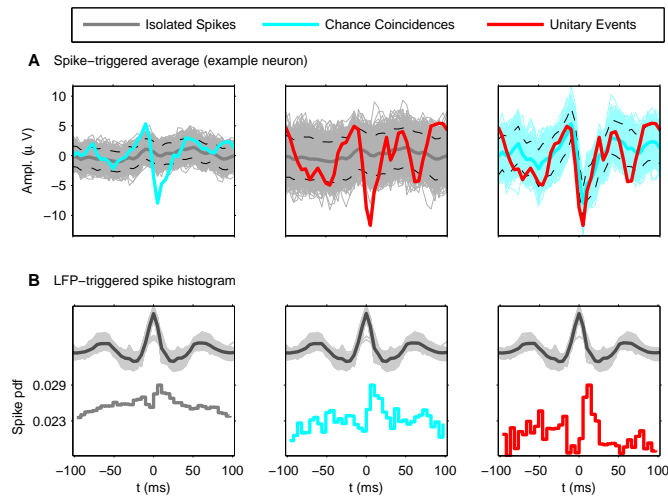


Figure S2: Relationship of LFP and synchronized spiking behavior in a single neuron and LFP-triggered PSTHs of synchronized activity. (A) STA of the LFP (filtered between 2-80 Hz to remove DC components) of one neuron (same neuron as in Figures 2 and S1) for three disjunct sets of trigger spikes: not coincident with spikes from simultaneously recorded other neurons (isolated spikes, ISO, gray), in coincidences predicted by rate (chance coincidences (within 3 ms), CC, cyan), and involved in significant coincidences (Unitary Events, UE, red). The left panel compares the STA of ISO (dark gray curve, $n=4098$) to the STA of CC (cyan curve, $n=533$). To account for the difference in variability due to sample sizes, the STA of ISO is recomputed using only 533 random trigger spikes. The light gray band results from the superposition of 1000 re-computations of which 95% are enclosed by the dashed curves at each point in time. Similarly, the middle and right panel compare the STA of UE (red curve, $n=150$) to the STA of ISO and CC, respectively. (B) Bottom: Population-averaged LFP-triggered histogram of ISO (left), CC (middle), and UE (right). The trigger times are the largest local maxima of the LFP that are separated by a minimum distance of 33 ms. The spikes of a neuron are triggered on exactly one LFP channel. Top: LFP averages for each neuron contributing to the histogram (light gray curves) based on the same trigger. The dark gray curve is the average of the single neuron LFP averages.

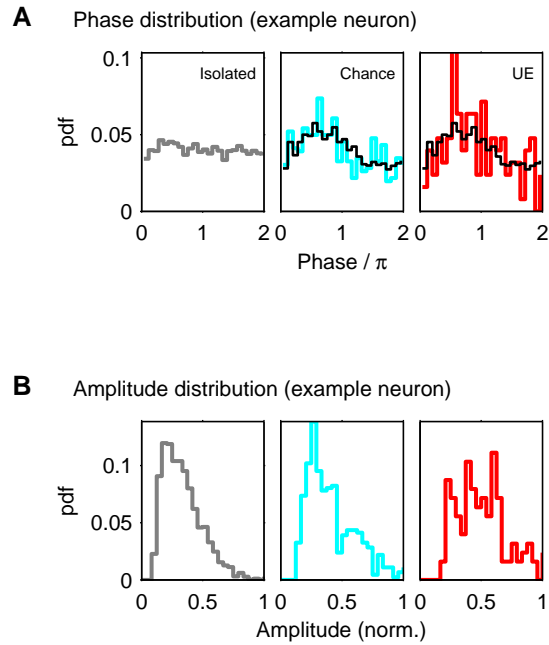


Figure S3: Phase and amplitude distributions (normalized to unity area) of the same example neuron as in Figures 2, S1, and S2 shown separately for ISO (left), CC (middle), and UE (right). (A) The modulation of the phase distribution increases from left to right. Phase π is the location of the trough of the LFP oscillation. The black curve in the middle and the right panel is the expected phase distribution of coincidences predicted from the phase distributions of the contributing neurons (see Materials and Methods). (B) Simultaneously to the increased locking, the amplitude distribution shifts to higher values.




Figure S4: The increased locking of UEs is independent of the overall degree of locking of the neuron. (A) Fraction of neurons exhibiting (threshold of 25 spikes) ISO, CC and UE separately for the sets of strongly (left) locked and weakly (right) locked neurons (criterion: surrogate test ($\alpha=0.05$) on original spike train containing all spikes). (B) Percentage of neurons with a locking stronger than σ_l in each of the two groups (strongly and weakly locked). For the selected value of $\sigma_l=1.97$ (average locking strength of strongly locked neurons) the percentages are shown as bars.

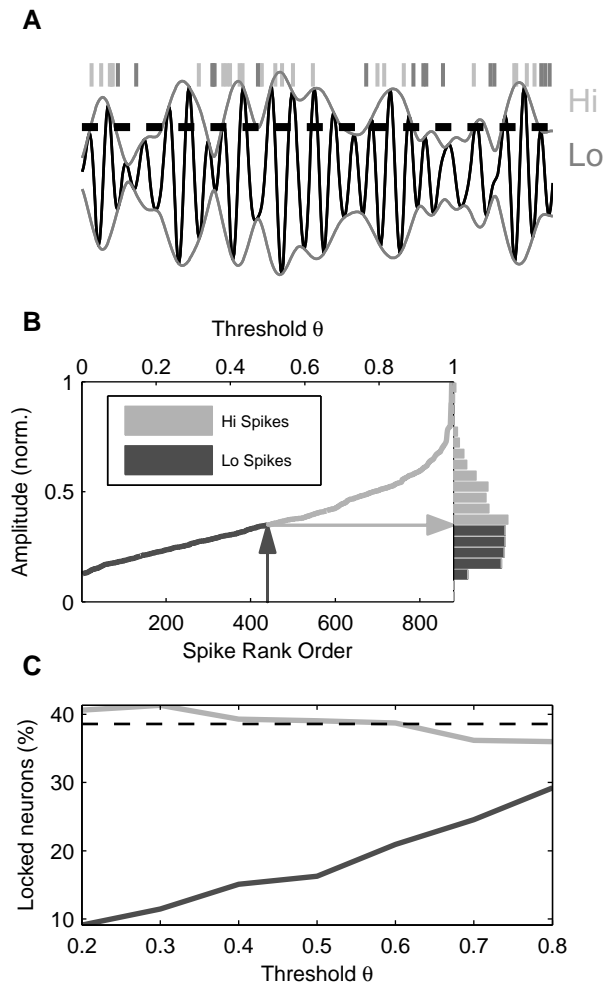


Figure S5: Influence of oscillation magnitude on locking of spikes to LFP. (A) Spikes in periods with an LFP magnitude (i.e. envelope of LFP, light gray curve) above a certain threshold (dashed line) are termed the 'Hi' set (light gray ticks) and the remainder the 'Lo' set (dark gray ticks). (B) Separation of spikes into Hi and Lo for the same example neuron as in Figures S1 and S2. Spikes are rank ordered according to LFP magnitude, the histogram on the right shows the distribution of the respective magnitudes. The threshold θ is defined as the relative number of spikes labeled as Lo. The dark gray arrow illustrates a threshold choice of $\theta=0.5$, and corresponds to a data dependent relative amplitude (light gray arrow). Spikes at extremely low LFP amplitudes (lowest 10%) do not enter the analysis. (C) Percentage of neurons with significant (Rayleigh test, $\alpha=0.05$) phase-locking of the Hi spikes (light gray curve) and of the Lo spikes (dark gray curve) as a function of magnitude threshold. Even for large θ , the set of Hi spikes shows significant locking in 34% of the neurons, although it consists of only few spikes. The dashed line shows as a reference the percentage of locked neurons computed if spikes are not separated into Hi and Lo (i.e. all spikes). Thus the locking of neurons is mainly explained by the locked Hi spikes, and their locking is approximately independent of θ .

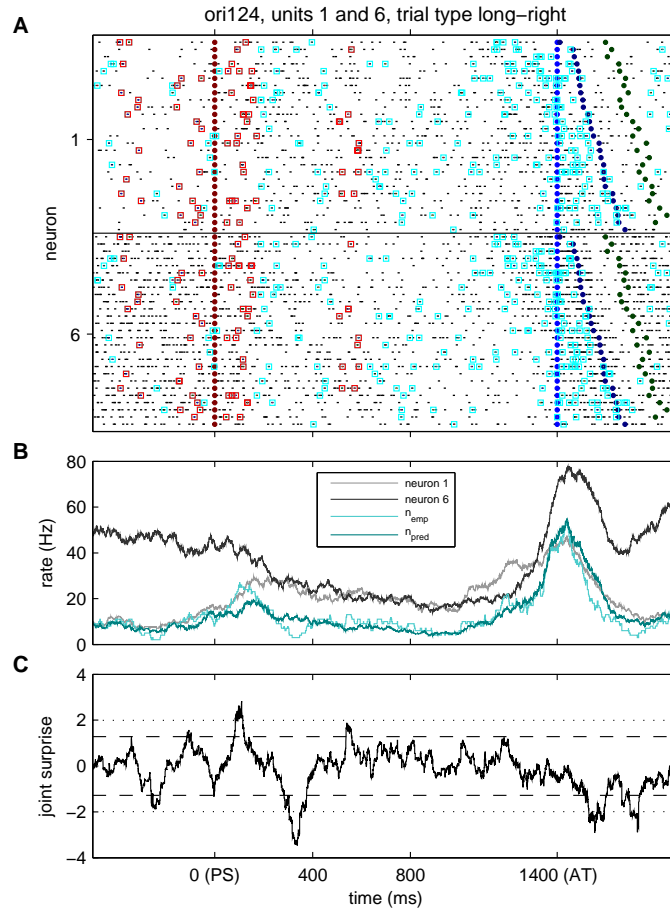


Figure S6: Detection of Unitary Events. (A) Spike rasters for the same neuron (neuron 6) shown in Figures S1-S3 and one simultaneously recorded neuron (neuron 1). Each line in the rasters corresponds to one trial. Simultaneously recorded activities of the two neurons are shown on lines of the same height in the respective raster. Spikes are indicated by black dots, coincident spikes and Unitary Events are surrounded by a cyan or red square, respectively. Data shown are recorded during the self-paced task with long time delay (see Methods and Materials for experimental details). The corresponding behavioral events are marked in the rasters with differently colored filled circles: occurrence of the preparatory stimulus PS (dark red), allowed movement time AT (light blue), movement initiation (dark blue) and end of movement (dark green). (B) Firing and coincidence rates. The firing rates of the two neurons are shown in dark gray (neuron 6) and light gray (neuron 1), together with the rate of the empirical coincidences (light cyan) and the coincidence rate expected from the neurons' firing rates (dark cyan), calculated as the sum of the trial-by-trial rates. All rates are estimated in sliding windows of 100 ms width shifted by 1 ms. (C) Significance of empirical coincidences. The joint surprise (dark gray curve) results from the comparison of the empirical and the expected coincidence counts. Significant excess coincidences (i.e. UEs) are detected if the joint surprise is larger than the 5% level (dashed line). For comparison, the 1% level is also indicated (dotted line). UEs are found during a short period before PS occurrence, shortly after PS, and at 600 ms after PS. The latter is one of the short delay times that monkey was exposed to in parallel to the shown delay scheme. Note that although there is a considerable increase of coincident events in relation to the arm movement, they occur at chance level.

3.3 Estimating the Contribution of Assembly Activity to Cortical Dynamics from Spike and Population Measures

The following manuscript is submitted to the Journal of Computational Neuroscience as:

M. Denker, A. Riehle, M. Diesmann, and S. Grün. Estimating the contribution of assembly activity to cortical dynamics from spike and population measures.

Publisher homepage: <http://www.springer.com>

Estimating the Contribution of Assembly Activity to Cortical Dynamics from Spike and Population Measures

Michael Denker¹, Alexa Riehle², Markus Diesmann^{1,3}, Sonja Grün¹

August 1, 2009

¹RIKEN Brain Science Institute, Wako City, Japan

²Mediterranean Institute of Cognitive Neuroscience (INCM), CNRS - Univ. Aix-Marseille 2, Marseille, France

³Brain and Neural Systems Team, RIKEN Computational Science Research Program, Wako City, Saitama, Japan

Abstract

The hypothesis that cortical networks employ the coordinated activity of groups of neurons, termed assemblies, to process information is debated. Results from multiple single-unit recordings are not conclusive because of the dramatic undersampling of the system. However, the local field potential (LFP) is a mesoscopic signal reflecting synchronized network activity. This raises the question whether the LFP can be employed to overcome the problem of undersampling. In a recent study in the motor cortex of the awake behaving monkey based on the locking of coincidences to the LFP we determined a lower bound for the fraction of spike coincidences originating from assembly activation. This quantity together with the locking of single spikes leads to a lower bound for the fraction of spikes originating from any assembly activity. Here we derive a statistical method to estimate the fraction of spike synchrony caused by assemblies – not its lower bound – from the spike data alone. A joint spike and LFP surrogate data model demonstrates consistency of results and sensitivity. Combining spike and LFP signals, we obtain an estimate of the fraction of spikes resulting from assemblies in the experimental data.

1 Introduction

A common hypothesis concerning the processing of information by cortical networks involves the propagation of activity through synchronously firing groups of neurons, termed assemblies. Despite the inherent undersampling of state of the art multiple single-unit recordings, experimental studies indirectly substantiate the assembly idea with findings of behavior related significant synchronous spiking (e.g., Riehle et al., 1997). Independently thereof, a signal on the population level, like the mesoscopic local field potential (LFP), typically exhibits temporally structured oscillations commonly interpreted as correlated network activity. Synaptic transmembrane currents have been identified as the primary contributor to LFP generation (Mitzdorf, 1985; Logothetis and Wandell, 2004). Surprisingly, although synchronized membrane potential oscillations of neurons in the vicinity (Katzner et al., 2009) of the recording electrode show strong correlations with the LFP (Poulet and Petersen, 2008), these do not induce the same degree of coincident spiking between the same neurons. Their correlated spiking is on average independent of the former two signals (Tetzlaff et al., 2008; Poulet and Petersen, 2008).

In a recent study (Denker et al., submitted) we were able to demonstrate in data of the motor cortex of the awake behaving monkey the missing link between significant spike synchrony and the LFP. A conceptual model enabled us to derive lower bounds for the fraction of spike coincidences β originating from observed assembly activity and the fraction of spikes of a neuron γ caused by assembly activity whether observed or not. The results were obtained by comparing the locking of spike coincidences to the LFP in time periods with significant synchrony to the locking outside of these periods.

In the remainder of this section we first introduce the Unitary Events analysis method (Sec. 1.1) used to quantify spike synchrony. We then review our earlier findings on the phase locking (Sec. 1.2) and the model-based interpretation of the data (Sec. 1.3). In the present work we remove the limitation to lower bounds of β and γ by evaluating a parameter-free model of the composition of spike coincidence counts. Sec. 2 introduces the model and derives an estimator for β only based on the configuration of spikes in the respective time interval. In Sec. 3 we explain that γ – in contrast to β – cannot be extracted from the spike trains alone. Nevertheless equipped with the estimate of β we are in the position to compute γ using the previously found relationships between the phase distributions of the spike-LFP coupling (Sec. 1.3). This highlights the importance of the LFP in overcoming the undersampling problem. We demonstrate the consistency of the concept with the help of a joint spike-LFP toy model. In the physiological range, the parameters can reliably be determined

and the lower bounds obtained from the phase locking are compatible with the parameter estimates yielded if the spike statistics is considered in addition. Finally, we utilize our new tool to reanalyze the experimental data (Sec. 4) for β and γ . Although highly simplified the toy model enables us to discuss (Sec. 5) various aspects of the distribution of parameters obtained from the experimental data set.

1.1 Analysis of spike synchrony

In the following we will briefly summarize the results obtained in Denker et al. (submitted) which serve as a starting point for our analysis. In a first step we analyzed simultaneously recorded single units from monkey motor cortex (cf., e.g., Roux et al., 2006) for excess spike synchrony by employing the Unitary Events analysis method (Grün et al., 2002b; Grün, 2009). The method compares the empirically measured spike coincidences to the number expected by chance given by the product of the neurons' firing rates. The expected number defines the mean of the a Poisson distribution realizing the null-hypothesis of statistical independence. If the p-value of the empirically found number of coincidences evaluated by comparison to this distribution exceeds the significance level, the coincidences are considered significant and are termed Unitary Events (UE). This analysis is performed in a sliding window fashion (windows of 100 ms) to account for the dynamics of correlations and the non-stationarities in the firing rates in time. Non-stationarity across trials is accounted for by calculating the expected number of coincidences as the sum of trial-by-trial expectancies (Grün et al., 2003). Coincident spike events with a temporal precision of up to 3 ms were collected by the multiple-shift method (Grün et al., 1999), which adapts the expected number accordingly. As a result we are able to detect the dynamics of spike synchrony and can identify time intervals that exhibit UE, i.e. excess spike synchrony. Non-UE time periods contain coincidence spike events at chance level only.

1.2 Analysis of phase-locking

In a next step we investigated the relation of spikes to the LFP. In particular, we were interested if spikes in synchronous events, such as chance coincidences (CC) or UE, have a different relation to the LFP than spikes that are not involved in a coincidence, i.e. isolated spikes (ISO). Therefore we classified each spike of a data set into one of these three classes and analyzed the classes separately (see Fig. 1(a) for an illustration). Spike triggered averages (STA) of the LFP revealed for all classes an oscillatory structure at about 17 Hz with spikes occurring preferentially at the decaying amplitude. However, the amplitude of

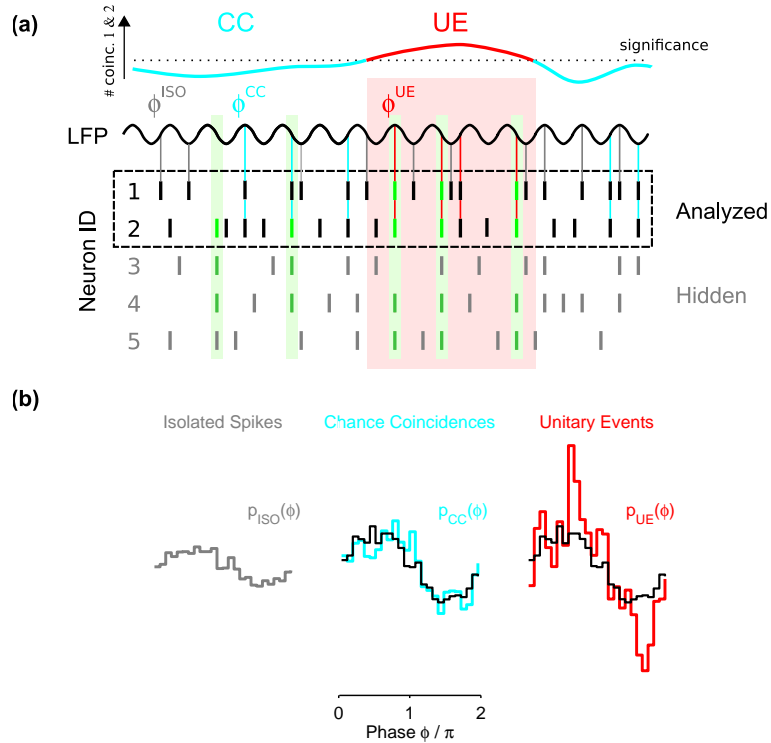


Figure 1: Strongest phase locking of spikes to the LFP during time periods of excess spike synchrony. (a) Sketch of the analysis: the number of pairwise synchronous spikes between two neurons (here, neuron 1 and 2) is evaluated in sliding windows of $T = 100$ ms (top graph). Synchronous spikes between the two neurons are classified as Unitary Events (UE; red) in windows where the observed number of coincidences exceeds a minimal number required to reach significance (dotted line). Outside of these windows coincidences are classified as chance coincidences (CC; cyan). Spikes that are not part of an observed coincident event are classified as isolated spikes (ISO; gray). UE periods are interpreted to originate from activation of a neuronal assembly (green) in which both observed neurons participate. Here, the assembly consists of the observed neurons 1 and 2 and the unobserved neurons 4 and 5. In contrast, during CC periods each of the observed neurons may individually participate in other assembly activations within the network, but not the same one (here, neuron 2 participates in an assembly that does not include neuron 1). The instantaneous phase ϕ of the LFP is extracted at spike times and pooled according to the spike classification (ϕ^{ISO} , ϕ^{CC} , and ϕ^{UE}). (b) Distributions of the LFP phase at spike occurrences for the three sets ISO (left), CC (middle), and UE (right) in experimental data from monkey motor cortex (see Sec. 4.1). Phase π corresponds to the trough of the LFP oscillations. The black curve in the middle and right graph represents the expected phase distribution of coincident spikes assuming independent neurons.

the STAs derived for the different spike classes were strikingly different: consistent across the whole data set the STA triggered on UE spikes exhibited the largest amplitude, for CC spikes a smaller amplitude and for ISO the smallest.

However, the STA analysis cannot uncover the reason for the differences in STA amplitude, since it may be due to differences in the LFP amplitudes or due to different degrees of phase locking between spikes and the LFP. In order to disentangle these aspects we performed a phase-amplitude analysis of the LFP. Using a Hilbert transform of the LFP we gained the instantaneous phase and amplitude as time dependent functions, and extracted the respective measures at spike times. The phase distribution exhibits the phase preferences by non-uniformity of the phase histograms (see Fig. 1(b)). We observed that UE spikes expressed the strongest degree of phase locking, CC less (at the predicted chance level) and even less for ISO. This result is consistent across the whole population of recorded neurons, and 69% of the neurons exhibit a stronger phase locking of UE than CC. In contrast we found much smaller or even negligible differences of the LFP amplitude measured as its envelope at spike times.

1.3 Interpretation and conceptual model

Our analysis of the relation of spikes to the LFP revealed several unexpected results. Firstly, we found a difference in the locking degree for spikes involved in different categories of coincidences, i.e. UE vs CC. Following the hypothesis that active assemblies are expressed by coordinated spiking activity, UE coincidences are interpreted as a signature of such active assemblies. Our new finding on the phase locking of UE spikes to the LFP leads to the interpretation that assembly activity occurs in a pulsed fashion, locked to the LFP oscillation. The fact that spikes of chance coincidences also occur phase locked seems contradictory. However, their degree of locking is fully explained by the (weak) locking of non-UE spikes, which trivially leads to an enhanced locking of such spikes if occurring coincidentally. Thus the question is rather, why isolated spikes outside UE periods exhibit phase locking at all.

To get a better understanding of these puzzling observations, we developed a conceptual model which consistently explains all our findings (Fig. 2(a)). One basic assumption of the model is that spikes involved in assemblies exhibit locking to the LFP while non-assembly spikes do not. According to our findings that individual spike trains typically contain a mixture of all spike categories, it is reasonable to assume that an individual spike train is composed of assembly and non-assembly spikes. Due to the severe undersampling of the system caused by the limited number of recording electrodes, it is highly likely that also ISO spikes contain assembly spikes, but the corresponding partner neurons could

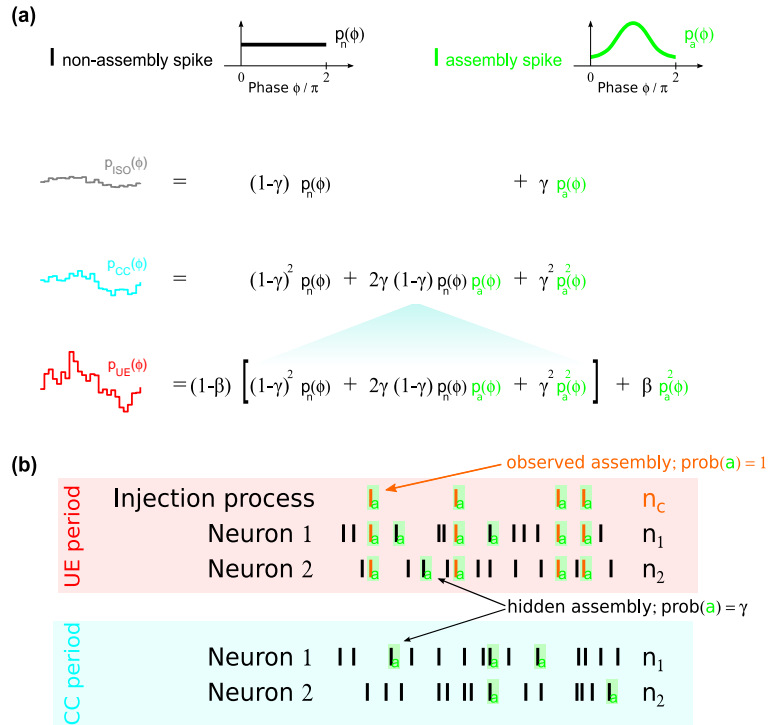


Figure 2: Conceptual model to explain the differences in spike-LFP locking and definition of a suitable stochastic spike-LFP model. (a) A individual spike train is assumed to be composed of assembly spikes that exhibit locking to the LFP (non-uniform phase distribution $p_a(\phi)$, top right) and of non-assembly spikes that do not lock to the LFP (uniform phase distribution $p_n(\phi)$, top left). The observed phase distributions for isolated spikes $p_{\text{ISO}}(\phi)$, for chance coincidence spikes $p_{\text{CC}}(\phi)$, and Unitary Event spikes $p_{\text{UE}}(\phi)$ are expressed as combinations of $p_n(\phi)$ and $p_a(\phi)$ by the three respective equations. (b) A stochastic spike-LFP model is used to probe the consistency and sensitivity in estimating the parameters β and γ from data. UE periods (red) are modeled as independent spike trains of a given background rate, and excess coincidences due to an additional process from which identical spikes are injected into both parallel spike trains (labeled by green letter 'a'). The latter is modeled as Poisson process with a number of n_c spikes. The background processes of each neuron contain $n_1 - n_c$ and $n_2 - n_c$ spikes, respectively. A fraction γ of background spikes may also be part of an unobserved assembly, also marked with the green letter 'a'. Labeled spikes exhibit locking to the LFP according to $p_a(\phi)$, non-labeled spikes do not lock (uniform phase distribution $p_n(\phi)$). A CC period (cyan) is realized as two independent Poisson processes with the same spike counts n_1 and n_2 for neuron 1 and 2, assembly probability γ as for the UE period.

not be identified. Consequently, the phase histogram of ISO must contain a mixture of unlocked and locked spikes. Assuming a uniform phase distribution $p_n(\phi)$ for non-assembly spikes and an unknown non-uniform phase distribution $p_a(\phi)$ for assembly spikes, the phase distribution of the ISO spikes $p_{\text{ISO}}(\phi)$ can be expressed by a weighted combination of the two components:

$$p_{\text{ISO}}(\phi) = (1 - \gamma)p_n(\phi) + \gamma p_a(\phi) \quad (1)$$

Thus also ISO spikes show locking, and its strength is expressed by the factor γ .

CC result from pairs of independent spike trains that express coincidences by chance. Under the above assumption that spike trains are composed of assembly and non-assembly spikes, chance coincidences are the result of combinations of different spike types: unlocked-unlocked spikes, locked-unlocked spikes (two possible combinations), and locked-locked spikes. Thus, $p_{\text{CC}}(\phi)$ is also expressed as a composition of $p_n(\phi)$ and $p_a(\phi)$:

$$p_{\text{CC}}(\phi) = (1 - \gamma)^2 p_n(\phi) + \gamma(1 - \gamma)p_n(\phi)p_a(\phi) + \gamma^2 p_a^2(\phi) \quad (2)$$

In case of time periods that contain UEs, coincidences resulting from an active assembly are present, but likely intermixed with chance coincidences that are not related to this assembly. Thus, the phase distribution of UE is assumed to be composed of a contribution of chance locking of a percentage of $(1 - \beta)$, and to a percentage of β of assembly coincidences (see sketch in Fig. 2(a)):

$$p_{\text{UE}}(\phi) = (1 - \beta) \cdot [(1 - \gamma)^2 p_n(\phi) + \gamma(1 - \gamma)p_n(\phi)p_a(\phi) + \gamma^2 p_a^2(\phi)] + \beta p_a^2(\phi) \quad (3)$$

As a result, the enhanced locking of UE as compared to CC is a consequence of the presence of assembly coincidences i.e. $\beta > 0$.

2 Estimation of assembly activations from spike statistics

In terms of spike coincidences, β is the ratio of the number of coincidences resulting from assembly activation and the total number of coincidences in a

given time window

$$\beta = \frac{n_c}{n_{\text{emp}}}. \quad (4)$$

The task therefore is to construct an estimate of n_c on the basis of the known properties of the observed spike trains. For two independently spiking neurons with n_1 spikes in one spike train and n_2 spikes in the other, the probability to observe k coincidences is given by the hypergeometric distribution (Grün et al., 2003)

$$\mathcal{H}_{T_h}^{n_1, n_2}(k) \quad (5)$$

where we introduced the shorthand

$$T_h = T/h. \quad (6)$$

and h is the resolution of the discretized time axis. We define $\mathcal{H} = 0$ outside of $[0, T_h]$. The expected number of coincidences is

$$n_{\text{exp}} = \frac{1}{T_h} \cdot n_1 n_2. \quad (7)$$

The expectation value (7) is just the probability to observe a coincidence in a given bin $(n_1/T_h)(n_2/T_h)$ multiplied by the number of available time bins T_h . In the presence of n_c deterministic coincidences, however, only a reduced number of spikes of the two neurons is available to form chance coincidences

$$n_{\text{exp},c} = n_c + \frac{1}{T_h - n_c} \cdot (n_1 - n_c)(n_2 - n_c).$$

The construction of the respective spike trains is illustrated in Fig. 2(b). Let us now assume that we measure the number of coincidences $\langle n_{\text{emp}} \rangle$ averaged over many repetitions of the experiment with exactly the same values for n_c , n_1 , n_2 . In this case we can equate the empirical average with the expectation value $\langle n_{\text{emp}} \rangle = n_{\text{exp},c}$. In this expression n_c is the only unknown variable. Hence, we can express n_c in terms of n_1 , n_2 , and $\langle n_{\text{emp}} \rangle$. Given only the triplet of a single realization $(n_1, n_2, n_{\text{emp}})$ we can still hope that the measured n_{emp} is typical and write

$$n_{\text{emp}} = n_c + \frac{1}{T_h - n_c} \cdot (n_1 - n_c)(n_2 - n_c). \quad (8)$$

Multiplication with $(T_h - n_c)$

$$(T_h - n_c) n_{\text{emp}} = (T_h - n_c) n_c + (n_1 n_2 - n_c n_2 - n_c n_1 + n_c^2)$$

shows that the quadratic terms in n_c cancel

$$-n_c n_{\text{emp}} - T_h n_c = -T_h n_{\text{emp}} - n_c (n_1 + n_2) + n_1 n_2$$

and collecting terms with n_c gives

$$n_c = \frac{T_h n_{\text{emp}} - n_1 n_2}{T_h + n_{\text{emp}} - (n_1 + n_2)}. \quad (9)$$

Using (9) we can already compute β for all experimental UE periods with their individual triplets $(n_1, n_2, n_{\text{emp}})$ and obtain the distribution of β as well as its mean β^{UE} (see Sec. 4). Clearly (9) is an approximation, the expression is negative for small n_{emp} , in particular when no coincidence has been observed ($n_{\text{emp}} = 0$).

Next we consider a slightly more realistic model with a variable number of assembly activations n_c . The purpose of the model is to help us to understand the conditions under which we can reliably extract β . In the absence of any knowledge about the process generating the additional coincidences n_c we assume that each value consistent with the observed number of coincidences is equally likely. Thus, each possible value of n_c occurs with probability $1/(n_{\text{emp}} + 1)$. For a particular n_c , however, the probability to be consistent with the n_{emp} observed coincidences is again given by the hypergeometric distribution (5)

$$\mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(n_{\text{emp}} - n_c).$$

The probability that a particular n_c underlies the observation therefore is

$$\frac{1}{n_{\text{emp}} + 1} \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(n_{\text{emp}} - n_c)$$

and the total probability to observe n_{emp} coincidences is

$$\sum_{i=0}^{n_{\text{emp}}} \frac{1}{n_{\text{emp}} + 1} \mathcal{H}_{T_h - i}^{n_1 - i, n_2 - i}(n_{\text{emp}} - i).$$

Consequently, the expected n_c given the observed triplet $(n_1, n_2, n_{\text{emp}})$ is

$$n_c = \sum_{i=0}^{n_{\text{emp}}} \frac{i \cdot \frac{1}{n_{\text{emp}} + 1} \mathcal{H}_{T_h - i}^{n_1 - i, n_2 - i}(n_{\text{emp}} - i)}{\sum_{j=0}^{n_{\text{emp}}} \frac{1}{n_{\text{emp}} + 1} \mathcal{H}_{T_h - j}^{n_1 - j, n_2 - j}(n_{\text{emp}} - j)}$$

which reduces to

$$n_c = \frac{\sum_{i=0}^{n_{\text{emp}}} i \cdot \mathcal{H}_{T_h - i}^{n_1 - i, n_2 - i}(n_{\text{emp}} - i)}{\sum_{i=0}^{n_{\text{emp}}} \mathcal{H}_{T_h - i}^{n_1 - i, n_2 - i}(n_{\text{emp}} - i)}. \quad (10)$$

Note that the denominator is not unity as the sum extends over probabilities from different hypergeometric distributions.

2.1 Correspondence of models of assembly activations

The intuitive approximation (9) is compatible with the more detailed model result (10). To see this, we first approximate the individual hypergeometric distributions in (10) by the one for the average n_c and extend the range of the sums to the maximum value $T_h - n_c$:

$$n_c = \frac{\sum_{i=0}^{T_h - n_c} i \cdot \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(n_{\text{emp}} - i)}{\sum_{i=0}^{T_h - n_c} \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(n_{\text{emp}} - i)}.$$

This reduces the denominator to unity. With the substitution $k = n_{\text{emp}} - i$ this reads

$$n_c = \sum_{k=n_{\text{emp}} - (T_h - n_c)}^{n_{\text{emp}}} (n_{\text{emp}} - k) \cdot \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(k).$$

The sum only has contributions for $k \geq 0$ and we again extend the range to the maximum value

$$n_c = \sum_{k=0}^{T_h - n_c} (n_{\text{emp}} - k) \cdot \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(k).$$

This suggests a decomposition of the sum to

$$n_c = n_{\text{emp}} \sum_{k=0}^{T_h - n_c} \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(k) - \sum_{k=0}^{T_h - n_c} k \cdot \mathcal{H}_{T_h - n_c}^{n_1 - n_c, n_2 - n_c}(k)$$

where the first sum is unity and the latter term is the expectation value of k . Hence,

$$n_c = n_{\text{emp}} \cdot 1 - \frac{1}{T_h - n_c} \cdot (n_1 - n_c)(n_2 - n_c)$$

which is (8) implying the intuitive result (9).

2.2 Accuracy of the count of assembly activations

Before we turn to the experimental data in Sec. 4 we need to assess the accuracy of our estimator of n_c . To this end we construct a surrogate data set with parameters adapted to the experimental data. The data are organized into 32 blocks containing an increasing number of assembly activations from $n_c = 0$ to 31. A block is composed of $M = 2700$ time windows, consistent with the number of UE windows found in the experimental data. A time window has

a duration (6) of $T_h = 5000$ at a resolution of $h = 1$ ms, corresponding to the 100 ms segments of UE analysis covering 50 trials. Each time window contains surrogate spike trains of two neurons with totals of exactly $n_1 = n_2 = 100$ spikes (corresponding to a rate of 20 Hz). In each spike train the $n_i - n_c$ non-assembly spikes are uniformly distributed over the time window.

Fig. 3(a) shows the distribution of injected coincidences n_c in the data set organized by the total number of coincidences in each window n_{emp} . Here and in the following only windows containing a significant number of coincidences ($n_{\text{emp}} > n_\alpha$) are analyzed (cf. Sec. 1.1). The estimator (9) of n_c assigns a unique value to each value of n_{emp} because the spike counts are identical for all windows. Therefore at a given n_{emp} the estimate of n_c is identical for each window and thereby identical to the average over all windows. We observe that the estimate well describes the mean of the actual distribution of n_c at a particular n_{emp} . The average value of n_c is replotted in Fig. 3(b). The panel demonstrates that the approximative (9) and the exact (10) estimator only start to deviate for n_{emp} below n_α . Fig. 3(c) uses the same representation as Fig. 3(a) to illustrate the excellent correspondence between the distribution of the actual β and the estimated values. The total distribution of β (Fig. 3(d)) in the model is asymmetrical because of the lack of a typical n_c and the constant spike count. In conclusion, our measure is a faithful estimator of the average β in the data.

3 Estimating the assembly participation probability in a joint spike-LFP model

In the following section, we extend the simple spike model defined in Sec. 2 to include a representation of the LFP locking and verify numerically that the results obtained in the previous section yield a reliable estimate of our model parameters β and γ .

3.1 Combined spike-LFP model

First, we choose a parameter γ^{set} as the assumed probability that any spike in the network is part of an assembly activation, in agreement with our conceptual model. In the following we simulate for each choice of n_c the set of $M = 2700$ UE time windows as combinations of an injection process and a background process for two neurons as described in Sec. 2.2. In addition, CC time windows are generated by a background process alone, again with equal spike counts $n_1 = n_2 = 100$ (see Fig. 2(b) for an illustration). Again, only significant UE windows ($n_{\text{emp}} > n_\alpha$) and non-significant CC windows ($n_{\text{emp}} \leq n_\alpha$) are retained. To model the experimental results, we assign a label 'a' to all spikes

that originate from an assembly activation (Fig. 2(b)). By definition of our assembly process, every spike that originates from the injection process receives a label. In addition, a random proportion γ^{set} of spikes from the background process is labeled (in both, UE and CC time windows). In our simulations, the overall probability for a spike to belong to an assembly is set to $\gamma^{\text{set}} = 0.1$. Next, we define two distributions $p_n(\phi)$ and $p_a(\phi)$, where the latter has a larger modulation depth, which describe the locking of non-assembly and assembly spikes, respectively (Fig. 2(a)). Here, $p_n(\phi)$ is modeled as a uniform distribution, whereas for $p_a(\phi)$ the uniform distribution is weakly modulated by a Gaussian. Classification of spikes into the groups ISO (taken here as the single spikes in CC windows), CC and UE allows us to calculate the simulated phase distributions $p_{\text{ISO}}(\phi)$, $p_{\text{CC}}(\phi)$, and $p_{\text{UE}}(\phi)$ as mixtures of $p_n(\phi)$ and $p_a(\phi)$.

3.2 Estimating the minimal β from phase distributions

The setup allows us to follow the same analysis steps that we will perform on the experimental data in the following section. From the conceptual model introduced in Sec. 1.3, which is formally expressed in (1)-(3), we infer the lower bound β_{min}^ϕ of coincidences originating from assemblies during an observed UE period. Substituting the measured population phase distribution of chance coincidences $p_{\text{CC}}(\phi)$ of (2) into (3) yields an expression relating the known phase distribution $p_{\text{UE}}(\phi)$ of UE coincidences to the parameter β and the squared phase distribution of assembly spikes $p_a^2(\phi)$:

$$p_{\text{UE}}(\phi) = (1 - \beta) \cdot p_{\text{CC}}(\phi) + \beta \cdot p_a^2(\phi). \quad (11)$$

By systematic variation of the parameter β we thus obtain a corresponding phase distribution $p_a^2(\phi)$ by solving the equation separately for each bin of the respective distributions. However, for small values of β the assembly distribution $p_a^2(\phi)$ must exhibit a strong modulation to compensate for the large difference between $p_{\text{UE}}(\phi)$ and $(1 - \beta)p_{\text{CC}}(\phi)$. In the extreme case, it can become necessary for a bin of $p_a^2(\phi)$ to contain negative values, ruling out that particular choice of the parameter β . From this consideration, we define β_{min}^ϕ as the lowest value of β that leads to a phase distribution $p_a^2(\phi)$ with non-negative entries. This lower bound depends on the choice of the assembly phase distribution $p_a(\phi)$ we initially introduced into our model: the smaller the difference in locking between assembly spikes $p_a(\phi)$ and non-assembly spikes $p_n(\phi)$, the lower are the values that β_{min}^ϕ may assume. In Fig. 3(d) we show that in our spike-LFP model we obtain a value for β_{min}^ϕ well below the mean β^{UE} extracted from the spike analysis (Sec. 2). We note that when generating the model data with a more strongly modulated Gaussian to model $p_a(\phi)$, the minimum β_{min}^ϕ approaches

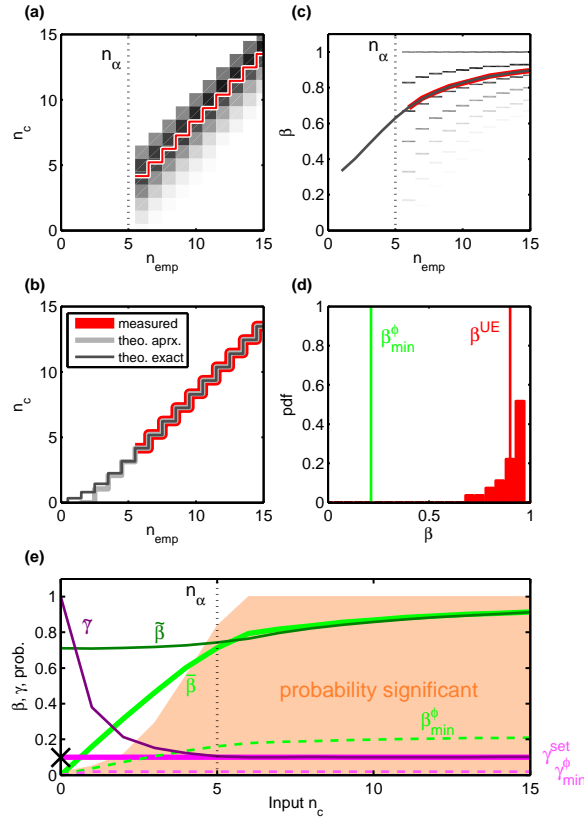


Figure 3: Consistency and sensitivity of the estimation procedure in surrogate data. (a) Estimated (red curve) and actual (gray shading indicates probability distribution) number of injected coincidences n_c as a function of the empirical number n_{emp} of coincidences per time window. Values for all significant (UE analysis, $\alpha = 0.05$, all panels unless otherwise stated) windows of 2700 windows per value for $n_c = 0, 1, \dots, 31$ were pooled and reordered according to n_{emp} . The dotted line represents the minimum number n_α of coincidences needed for a window to become significant. Since the total number of spikes n was fixed, the estimated n_c for given n_{emp} is constant. (b) Approximate (light gray) and exact (dark gray) theoretical prediction of n_c . For comparison, the red curve shows the measured n_c (same curve as in panel (a)). (c) Estimated (red line) and actual (gray distribution) values of $\beta = n_c n_{emp}^{-1}$ as a function of n_{emp} in analogy to panel (a). The dark gray curve shows the result from the exact theoretical prediction in panel (b). (d) Distribution of β (red bars) across all observed values of n_{emp} and the corresponding mean β^{UE} (red line). The green line shows that the lower bound β_{min}^ϕ calculated from the simulated phase distributions is well below the mean for our particular choice of $p_a(\phi)$. (e) Comparison of theoretical (light) and measured (dark) values of β (green) from the UE analysis and corresponding γ (magenta) from the phase analysis as a function of n_c . Dashed curves show the minimum β_{min}^ϕ and corresponding γ_{min}^ϕ . The shaded graph indicates the probability for a time window at given n_c to become significant. The deviation of estimates from their theoretical value for $n_c < n_\alpha$ (dotted line) is an expected artifact due to the pre-selection of significant UE periods. The average β^{UE} from (d) yields a good estimate $\gamma = 0.098 \approx \gamma^{set}$ (black cross on vertical axis).

β^{UE} from below (not shown).

3.3 Estimation of γ from phase distributions

We are now prepared to extract the parameter γ from the simulated phase distributions using the estimate of β^{UE} derived in Sec. 2. First, using $\beta = \beta^{\text{UE}}$ we extract $p_a^2(\phi)$ from (3) by solving the equation separately for each bin of $p_a^2(\phi)$. Taking the square root $p_a(\phi)$, we renormalize the distribution to unit area, and insert it in (1). By variation of the parameter γ , we find the value $\gamma^{\phi, \text{UE}}$ that minimizes the sum of the absolute bin-by-bin differences between the measures distribution $p_{\text{ISO}}(\phi)$ and the right side $(1 - \gamma) \cdot p_n(\phi) + \gamma \cdot p_a(\phi)$ of (1).

Using this method, we first analyze separately each of the 32 blocks of data with a fixed number of injected coincidences n_c . Let $\tilde{\beta}$ and $\tilde{\gamma}$ be the estimated values of β^{UE} and $\gamma^{\phi, \text{UE}}$ within a single block. In Fig. 3(e) $\tilde{\beta}$ and $\tilde{\gamma}$ are compared to their theoretical means $\bar{\beta}$ (calculated as the mean of n_c/n_{emp} over all significant time windows from one set of simulations with fixed n_c) and γ^{set} . The estimates are in good agreement for datasets where n_c was set above the significance threshold n_α . Below this threshold, the number of injected coincidences, and hence $\tilde{\beta}$ and $\tilde{\gamma}$, is overestimated. Clearly, in this regime contributing UE periods become significant due to an unusually high number of coincidences in the background process. Therefore, the mean approximation (9) does no longer hold. In principle it is possible to correct for the bias due to the selection of significant periods, although it is not possible to arrive at an expression in closed form for n_c . Nevertheless, in fact only a small fraction of windows actually become significant for low values of n_c (shaded graph in Fig. 3(e)). Therefore, in practice, where the true n_c is unknown, only few windows enter the analysis when n_c is below the significance threshold. For this reason, the bias does not affect the estimate when re-sorting the data according to the observed n_{emp} (Fig. 3(a)-(c)), and the original value of γ^{set} is well estimated as $\gamma^{\phi, \text{UE}} = 0.098$ from the complete dataset spanning all n_c (black cross in Fig. 3(e)). Consistently, the lower bounds β_{min}^ϕ and γ_{min}^ϕ are a lower bound on the respective estimates $\tilde{\beta}$ and $\tilde{\gamma}$ for all values of n_c . In summary, these calibrations using a simple, combined spike-LFP model demonstrate that our method and the approximation (9) are well suited to estimate the parameter γ from the phase distributions in a dataset with a realistic choice of parameters.

4 Analysis of experimental data

In this section we estimate the percentage of coincident spike events reflecting assembly activity (model parameter β) and the percentage of spikes that are part of an assembly activation (model parameter γ) from neuronal data of primary motor cortex of monkeys.

4.1 Experimental procedures

Two rhesus monkeys were trained to perform arm movements in two different tasks involving an instructed delay. In this study we exclusively analyze the delay activity during the preparatory period for the upcoming arm movement. Only correct trials were considered, in which the monkey responded within a predefined time window and in which movements were performed in the required movement direction. LFPs and spikes were recorded simultaneously in primary motor cortex using a multielectrode device of 2-4 electrodes. Spikes of single neurons were detected by an online sorting algorithm. The inter-electrode distance was on the order of $400\ \mu\text{m}$. LFPs were sampled at a resolution of 250 – 500 Hz and hardware filtered (band pass, 1-100 Hz). In total, we analyzed 53 recording sessions, which yielded 143 single neurons or 570 combinations of neurons and behavioral conditions. This selection included only those neurons which exhibited an average firing rate of 5 Hz or more. On average 33 ± 11 trials were recorded per experimental condition. In analyses that combine spikes and LFP, each neuron enters only once, and we never combined LFP and spikes that were recorded on the same electrode to exclude the possibility of spike artifacts in the signal. We confirmed that simultaneously recorded LFPs are highly synchronous in the frequency regimes of interest. For experimental details see Roux et al. (2006).

The Unitary Events analysis was applied to all simultaneously recorded pairs of neurons recorded on different electrodes thus totaling 123 pairs of neurons. Data were analyzed separately for the possible behavioral conditions and each experimental session. Defining a neuron by the combination of its identity and the behavioral context during which it was recorded, data from the same neuron may enter a population average up to eight times. Based on the results of the UE analyses each individual spike was marked as either ISO, CC or UE (see Sec. 1.1 and Denker et al. (submitted) for details). We observed a dominant component of the β -band (here around 17 Hz) in the LFP during the preparatory period and therefore filtered the LFP accordingly before applying the Hilbert transform to extract instantaneous phase and amplitude (i.e. the envelope). Finally, we derived the phase histogram pooled over the whole population for each class of spikes (see Fig. 1(b)).

4.2 Estimating the minimal β from phase distributions

As described in Sec. 3.2, we derive the lower bound of the percentage of coincidences originating from assembly activity β_{\min}^{ϕ} . Again, we make use of the equation (11) that relates the phase histogram of the UE to the model parameter β . This equation contains two components that we extract from the data, i.e. the phase distribution of the UE spikes $p_{\text{UE}}(\phi)$ and the phase distribution of the CC spikes $p_{\text{CC}}(\phi)$. We vary β in the interval $[0, 1]$ and retrieve the corresponding $p_a^2(\phi)$. Four examples are shown in Fig. 4 for different values of β . The first two examples demonstrate how choosing a small value of β may lead to distributions $p_a^2(\phi)$ that have negative entries and indicate a non-valid solution. From the systematic variation of β we derive the value $\beta_{\min}^{\phi} = 0.23$, that just leads to non-negative entries in all bins of $p_a^2(\phi)$ (see Fig. 4, bottom). This result indicates that at least about a fourth of the coincidences during a UE period are reflections of the observed assembly.

4.3 Estimate of β^{UE} from coincidence counts

As introduced in Sec. 2 the parameter β can be estimated by the comparison of the expected number of coincidences n_{exp} and the number of empirical coincidences n_{emp} . Fig. 5(a) shows n_{emp} as a function of its respective n_{exp} value for all analysis windows that were detected to contain UE for all pairs and sessions. Due to the selection of windows that contain UE, the empirical coincidence counts are larger than the minimal number of coincidences n_{α} required to be significant given the significance level α . n_{α} does not have a constant difference to n_{exp} but increases non-linearly as a function of n_{exp} since the Poisson distribution used for the evaluation of the significance becomes broader for larger expected number of coincidences (see for details in Grün et al. (2002a)).

Based on (9) we estimate the number of coincidences resulting from the active assembly n_c for each $n_{\text{exp}}, n_{\text{emp}}$ combination. We use a version of the UE analysis that also allows to detect temporally imprecise coincidences by employing the multiple-shift method. This method detects coincidences of systematically shifted spike trains up to a predefined shift, and sums the coincidence counts from all shifts. For the UE evaluation, the expected number of coincidences then has to be adjusted by a factor of $2s + 1$ with s being the maximal shift (see for details Grün et al., 1999). Thus to estimate the number of coincidences n_c that generically result from active assemblies, (9) was adjusted correspondingly. Fig. 5(b) shows n_c versus its corresponding n_{emp} and we find an average number of assembly coincidences of $n_c^{\text{UE}} = 7.72$.

Next, we can derive for each pair of (n_c, n_{emp}) an estimate of the percentage of coincidences reflecting assembly activity β (Fig. 5(c)). We find these values




Figure 4: Determination of β_{\min}^{ϕ} from experimental data. The four upper rows show the measured distributions $p_{UE}(\phi)$ (left, red), $\beta p_{CC}(\phi)$ (middle, cyan) and the resulting squared phase distribution of assembly spikes $p_a^2(\phi)$ entering the third equation in Fig. 2(b) for four different choices of β . The bottom graph shows the relationship between the choice of β and the resulting modulation depth of $p_a(\phi)$ (difference between maximum and minimum). The gray area indicates invalid choices of β where the corresponding distribution $p_a^2(\phi)$ has negative values (compare green filled areas of $p_a^2(\phi)$ in the top two rows). Hence, the minimum value for β is determined as $\beta_{\min}^{\phi} = 0.23$.

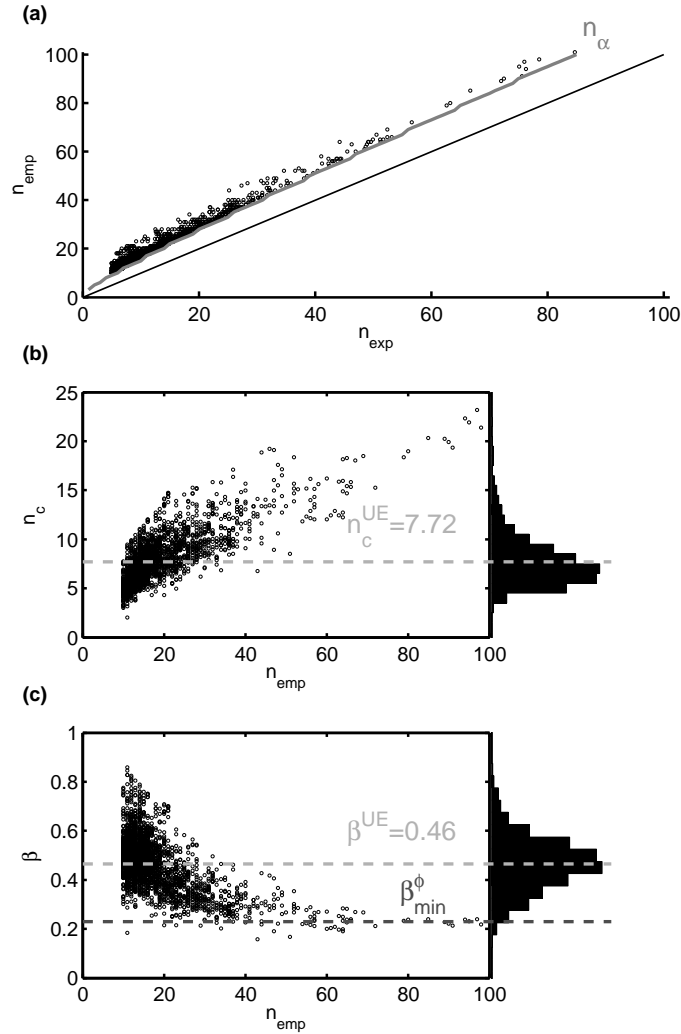


Figure 5: Distribution of estimated β and the resulting β^{UE} obtained from the data. (a) Scatter plot of the expected number of coincidences n_{exp} and the empirical number of coincidences n_{emp} . Here and in the following panels, each dot represents the values of one significant 100 ms time window ($\alpha = 0.05$, UE analysis), where each neuron exhibits a minimum average rate of 5 Hz. The gray curve indicates the critical number n_{α} of coincidences to reach significance at given n_{exp} . The black line shows the diagonal. Due to the preselection of datasets with a minimum rate, no data points for small n_{exp} are observed. (b,c) Estimated values of n_c (b) and β (c) as a function of the empirical coincidence count n_{emp} . The total distributions of n_c and β are shown on the right (black). The averages $n_c^{\text{UE}} = 7.72$ and $\beta^{\text{UE}} = 0.46$ are indicated by the light gray line. The dark gray line shows the lower bound $\beta_{\text{min}}^{\phi}$ for comparison.

peaked around mean value of $\beta^{\text{UE}} = 0.46$ with a standard deviation of 0.108. The obtained distribution of β is in good agreement with the minimum $\beta_{\text{min}}^{\phi}$ obtained in Sec. 4.2, as $\beta_{\text{min}}^{\phi}$ lies well below the mean β^{UE} of the distribution of β as expected for an estimate of a lower bound. Although $\beta_{\text{min}}^{\phi}$ in fact seems to be a bound for the complete distribution of β , this must not necessarily be the case: As a measure derived from the population estimate of phase distributions, it can only have predictive power on the population mean β^{UE} , but not the estimate β of single windows (compare Fig. 3(d)).

4.4 Estimation of γ from phase distributions

Finally, we estimate the fraction γ of spikes in the network that are part of an assembly activation (compare Sec. 3.3). Again, by systematic variation of the parameter γ in (1) we find the best fit between the measured $p_{\text{ISO}}(\phi)$ and the right side of the equation (see Fig. 6). Here, the assembly distribution $p_{\text{a}}(\phi)$ is known from Fig. 5 by setting $\beta = \beta^{\text{UE}}$, while the phase distribution of non-assembly spikes $p_{\text{n}}(\phi)$ is taken as the uniform distribution (see Sec. 5). Taking all analysis steps together, the best fit is derived for $\gamma^{\phi, \text{UE}} = 0.22$ indicating that on average 22% of the spikes of any neuron in the network originate from assembly activity (Fig. 6, bottom).

5 Discussion

Despite the complex mechanisms that contribute to the formation of the local field potential, it is well established that a primary contribution to the oscillatory LFP dynamics arises from the superposition of synchronized, slow transmembrane currents of cells close to the recording site (Mitzdorf, 1985; Logothetis and Wandell, 2004). Nevertheless, how rhythmicity in the LFP should be linked to synchrony on the spiking level has remained an open question (Poulet and Petersen, 2008) due to the unspecificity of the LFP signal and lack of a global oscillatory spiking activity. Several authors have interpreted the LFP as reflections of the specific synchronous synaptic activity responsible for the co-activation of neurons in the context of Hebbian cell assemblies (Eckhorn et al., 1988; Murthy and Fetz, 1996) which, in the simplest case, time their activations to the rhythm revealed by the LFP (Singer, 1999). Our recent experimental findings (Denker et al., submitted) demonstrate that indeed only identified assembly activity, which is identified as transient periods of significant excess spike synchrony between two neurons (UE analysis, Grün et al., 2002a,b), shows an exceptional phase relationship to LFP oscillation that exceeds expectation (Fig. 1(b)), thus confirming the hypothesis.

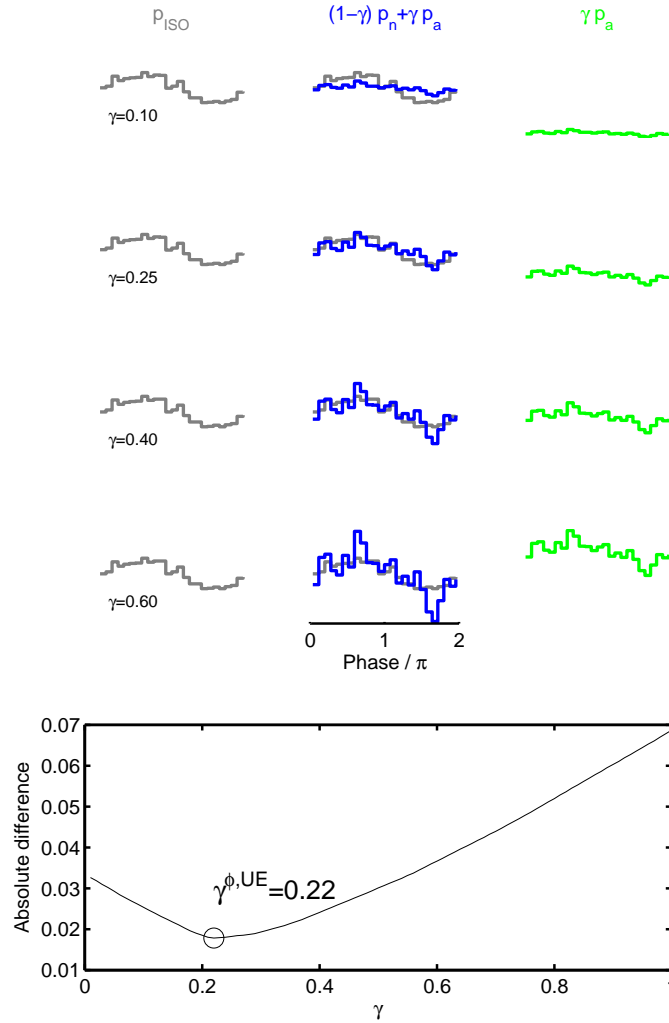


Figure 6: Calculation of the average assembly participation γ from experimental data. The top four rows show the different distributions entering the first equation in Fig. 2(b) for 4 different choices of γ . Left: measured phase distribution $p_{\text{ISO}}(\phi)$ (gray) of isolated spikes, repeated in each of the four rows. Middle: calculated distribution $(1-\gamma)p_n(\phi) + \gamma p_a(\phi)$ (blue) compared to the measured distribution $p_{\text{ISO}}(\phi)$ (gray). Right: Scaled assembly distribution $\gamma p_a(\phi)$ (green). The bottom graph shows the absolute bin-by-bin difference between $p_{\text{ISO}}(\phi)$ and the calculated distribution as γ is varied (compare distributions in the middle column). Both distributions match best for $\gamma^{\phi, \text{UE}} = 0.22$.

This link provides a handle on characterizing the spike synchronization dynamics in the context of the network oscillations. A simple model that captures the main experimental findings (Denker et al., submitted) on the spike-LFP relationship is introduced in such a way that it includes a parameter γ that measures the overall participation of individual spikes in an assembly – independent of whether it is observed or not. However, the estimation of this network parameter from the data requires knowledge of a second model parameter β , the expected relative number of coincidences stemming from an active assembly (as opposed to chance coincidences) during UE periods. Estimating this number from the spike data alone and integrating it into the conceptual model yields an average participation of single spikes to assembly activity of $\gamma = 22\%$.

The combination of measurements of synchrony on the local and mesoscopic scales enables us to access parameters of the network dynamics that remain hidden on the two individual levels of observation. The nature of the estimation process via the population phase distributions obtained from the LFP-spike locking requires a vast pooling of the experimental data in different sessions and two different monkeys in order to obtain a large sample of neurons as an appropriate representation of the network activity in motor cortex. The resulting value for γ must therefore be seen as a coarse estimate of the degree of assembly activity. Despite the finding that the parameter β shows a rather narrow distribution (Fig. 5(c)), the extent to which an individual neuron contributes to the assembly dynamics will fluctuate around the network mean γ . By quantifying the quantiles for 2 standard deviations of the experimental distribution of β we find corresponding γ values between $\gamma = 0.13$ and $\gamma = 0.30$ as an estimate of the range of γ across neurons.

In understanding the underlying dynamical structure, more interesting than the precise value of γ itself are two simple observations: first, $\gamma < 1$. Therefore, not all spikes are part of an assembly activation, and therefore some spikes must be attributed to a complementary mechanism. Intuitively, this is clear from the observation that in individual UE periods, it is always the falling phase of the LFP oscillation where increased locking is observed. Second, $\gamma > 0$, specifically at least about one fourth of all spikes must originate from assembly activity (with a non-zero lower bound γ_{\min}^{ϕ}). Thus the activation of assemblies is an ubiquitous phenomenon in the network, providing compelling evidence for the presence of an assembly coding scheme in addition to the correlation of synchrony with behavior (Riehle et al., 1997). Indeed, we typically observe a fraction of about 26% of neurons that show UEs during a given task. Therefore, combined with the large value of γ this suggests that typically any neuron is part of one or several assemblies. Nevertheless, information about the existence of higher-order synchrony between neurons is required for a further characterization of

the assembly, such as an estimate of the number of neurons that form a typical assembly.

A crucial part of our analysis is the estimation of the average relative amount of excess synchrony β during UE periods directly from the synchrony analysis. By design, all currently available methods to detect the presence of an active assembly activation, such as the UE analysis, rely on a significance test (Grün, 2009). Therefore, we must assume that the amount of excess synchrony β is influenced by the explicit choice of the significance level α used for the significance test: For a very restrictive significance criteria, only UE periods with very high excess synchrony n_c will be detected, resulting in higher values of β . Nevertheless, by choosing the same α level for the phase locking analysis, we will in turn also obtain a correspondingly more modulated distribution for $p_{\text{UE}}(\phi)$, reflecting that it contains a higher proportion of assembly spikes. Therefore, β yields a consistent phase distribution $p_a(\phi)$ of the assembly spikes independent of α .

The approximate formula (9) does not make any assumptions on preselecting significant periods of synchrony in the first place. However, we apply this estimate specifically in periods that display a significant surplus of synchrony, i.e. UE periods. If the number of coincidences n_c coming from the injection process is small, significance is only reached with an unexceptionally high amount of coincidences originating from the background. Therefore, in our stochastic spike-LFP model, in those very few trials that are significant for a given n_c , the average background rate will be underestimated and the average injection rate overestimated (Fig. 3(c)). However, restructuring the model data to the realistic case where only n_{emp} is measured (pooling across the data sets with different n_c) automatically incorporates the low frequency of significant windows with small n_c . Therefore the overall estimate of β is not significantly affected.

Our spike-LFP model (Sec. 3) is created in the spirit of providing a simple abstraction of the experimental findings in order to test the analysis under controlled conditions (equal counts per spike train with fixed injections). Due to this intentional lack of experimental detail in the model, the distribution of β that is obtained in the spike-LFP model naturally deviates from the one found in the data (compare Fig. 3(d) and Fig. 5). The model assumes an equal probability for a large range of n_c (number of injected coincidences). In real data, n_c likely follows a much more narrow distribution that does not exploit this range, resulting in lower values of β (compare Fig. 3(e)). Moreover, we assumed fixed counts n_1, n_2 for all neurons in the model. In reality, rates vary considerably across neurons. Especially for higher rates, where a higher proportion of coincidences can be assumed to originate from the background, we would expect a tendency towards lower β values.

In generating the spike-LFP model in Sec. 3 we did not explicitly model an LFP oscillation to place spikes at specific points of the field potential. In extreme situations this might be an oversimplification, where the constraints placed on the spikes due to the LFP locking influence the probability to detect coincidences. In our model, however, non-assembly spikes are associated with a uniform phase distribution, such that $p_n(\phi)$ does not impose any constraints on the spiking probability in time. The Poissonian spike interval statistics of non-assembly spikes thus remain unaffected. In contrast, the assembly spikes from the background process must in principle be adjusted with respect to a hypothetical LFP so that their phase distribution matches the assumed the distribution $p_a(\phi)$ of assembly spikes. However, as $p_n(\phi)$ is uniform, doing so will only influence the probability of finding a coincidence between two spikes that are both assembly spikes, i.e. that participate by chance in two different assemblies at the same time. Nevertheless, the probability $(\gamma^{\text{set}})^2 = 0.01$ of this to happen is negligibly small. Finally, assembly spikes from the injection process are synchronous by definition, and therefore remain unaffected by the choice of $p_a(\phi)$. Due to their small number they can always be freely placed in the time window T in accordance with $p_a(\phi)$. Taken together, the simple model introduced in Fig. 2 to represent the experimental findings in the context of the spike model includes sufficient detail without the need to explicitly model the actual positions of individual spikes with respect to an artificial LFP.

The analysis we provide in this manuscript focuses on estimating a global parameter that characterize the network dynamics. Given the overall consistency of our conceptual assembly model, the results suggest that oscillations and assemblies are tightly coupled. However, one may still speculate whether LFPS are in fact the direct cause of the assembly process, or whether they act as a supporting mechanism that coordinates synchronized firing. Whichever the initial cause of the LFP oscillation, a promising next step is to return to the level of the single neuron and integrate our knowledge on the dependence of spike synchrony on the LFP in order to discriminate which of the spikes are likely candidates to originate from the assembly process, and which are not. Two pieces of information aid us in this task: first, the knowledge of the observed UE phase distributions in relation to the ISO and CC distributions to determine the likelihood of a single spike to be part of an assembly activation. Second, the parameter γ that is informative of the relative number of spikes to assign to the assembly dynamics. Therefore knowledge about the LFP in relation to the spiking activity may well provide an independent indicator to identify periods where recorded neurons become transiently synchronized as part of an assembly activation. Moreover, the pulsed activation of assemblies triggered on the LFP oscillations cycle suggested by the experimental data stimulates the

idea that maybe even assemblies, where only one neuron is recorded, may be inferred from the phase locking statistics. Thus, incorporating a mesoscopic signals may serve to overcome problems related to the undersampling in multiple single-neuron recordings. The success of the method presented here not only corroborates the assembly hypothesis of neuronal processing, but offers a promising vista on reinterpreting the dynamical implications of observed LFP signals.

Acknowledgements

Partially funded by the Helmholtz Alliance on Systems Biology, the French National Research Agency (ANR-05-NEUR-045-01), EU grant 15879 (FACETS), and the Next-Generation Supercomputer Project of MEXT, Japan. Part of the work was carried out while SG and MDi enjoyed a scientific stay at the Norwegian University of Life Sciences, Aas in January, 2009.

List of symbols

T	Duration of each simulated time window
h	Temporal resolution of spike trains
$p_n(\phi)$	Phase distribution of spikes not participating in an assembly
$p_a(\phi)$	Phase distribution of spikes participating in an assembly
$p_{\text{ISO}}(\phi)$	Measured phase distribution of ISO spikes
$p_{\text{CC}}(\phi)$	Measured phase distribution of CCs
$p_{\text{UE}}(\phi)$	Measured phase distribution of UEs
β	Fraction of spikes originating from assembly activity during a UE period
$\beta_{\text{min}}^{\phi}$	Minimal value of β extracted from the phase distributions
γ	Fraction of spikes in the network activated in assemblies
$\gamma_{\text{min}}^{\phi}$	Value of γ that corresponds to the lower bound $\beta_{\text{min}}^{\phi}$
β^{UE}	Average value of β obtained from the UE analysis
$\gamma^{\phi, \text{UE}}$	Final estimate of γ obtained from phase distributions using β^{UE}
α	Significance level of the Unitary Event analysis
n_{α}	Number of coincidences required to reach significance at the α -level
n_i	Spike count of neuron i in a given window across trials
n	Total number of spikes per neuron and window in the stochastic spike-LFP model
n_{emp}	Empirical number of coincidences in a given time window
n_{exp}	Expected number of coincidences in a given time window based on the firing rates
n_c	Number of injected coincidences in our stochastic spike-LFP model
M	Number of windows simulated in the spike-LFP model per choice of n_c
γ^{set}	Predefined value of γ in the spike-LFP model

References

- Denker M, Roux S, Linden H, Riehle A, Grün S (submitted) Local field potentials reflect neuronal assemblies .
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? multiple electrode and correlation analyses in the cat. *Biol Cybern* 60(2): 121–130.

- Grün S (2009) Data-driven significance estimation for precise spike correlation. *J Neurophysiol* 101(3): 1126–1140.
- Grün S, Diesmann M, Aertsen A (2002a) Unitary events in multiple single-neuron spiking activity: I. detection and significance. *Neural Comput* 14(1): 43–80.
- Grün S, Diesmann M, Aertsen A (2002b) Unitary events in multiple single-neuron spiking activity: II. nonstationary data. *Neural Comput* 14(1): 81–119.
- Grün S, Diesmann M, Grammont F, Riehle A, Aertsen A (1999) Detecting unitary events without discretization of time. *J Neurosci Methods* 94(1): 67–79.
- Grün S, Riehle A, Diesmann M (2003) Effect of cross-trial nonstationarity on joint-spike events. *Biol Cybern* 88(5): 335–351.
- Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL, Carandini M (2009) Local origin of field potentials in visual cortex. *Neuron* 61(1): 35–41.
- Logothetis NK, Wandell BA (2004) Interpreting the bold signal. *Annu Rev Physiol* 66: 735–769.
- Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and eeg phenomena. *Physiol Rev* 65(1): 37–100.
- Murthy VN, Fetz EE (1996) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol* 76(6): 3968–3982.
- Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454(7206): 881–885.
- Riehle A, Grün S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278(5345): 1950–1953.
- Roux S, Mackay WA, Riehle A (2006) The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behav Brain Res* 169(2): 335–351.
- Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24(1): 49–65.

Tetzlaff T, Rotter S, Stark E, Abeles M, Aertsen A, Diesmann M (2008) Dependence of neuronal correlations on filter characteristics and marginal spike train statistics. *Neural Comput* 20(9): 2133–2184.

Chapter 4

Conclusion

In this chapter we first briefly summarize the key findings presented in [Chap. 2](#) and [Chap. 3](#) and discuss them in the context of this thesis. We then highlight several points common to all five studies, before we give an outlook on how the ideas presented in [Chap. 3](#) could be continued in the future.

4.1 Summary and Interpretation

The work presented in [Sec. 2.1](#) demonstrates learning-induced response changes during olfactory conditioning in the output of the antennal lobe (AL) of the honeybee, both on the level of the network spiking activity and the LFP. Simultaneous extracellular recordings were performed in fibers that contain mostly projection neurons which carry information from single glomeroli of the AL, the first stage of processing in the olfactory system, to second-order neuropiles. Therefore, the odor-specific spatial activity patterns of the glomeroli (e.g., [Joerges et al., 1997](#); [Hallem and Carlson, 2006](#)) and the temporal response profiles represented in the timing of individual spikes ([Laurent, 2002](#)) are likely to be relayed to the level of projection neurons. In addition, the resulting patterns of activation ([Sachse and Galizia, 2002](#); [Krofczik et al., 2009](#)) and temporal responses ([Laurent et al., 2001](#)) of the projection neurons are modulated by the interaction with a network of interneurons. A number of studies have demonstrated that learning induces plasticity in the representation of odors at the level of the AL of insects ([Faber et al., 1999](#); [Sandoz et al., 2003](#); [Yu et al., 2004](#)), suggesting that also responses at

the output level show learning-induced changes. To test this hypothesis, we subjected bees to a differential conditioning task, where administration of a rewarded odor (sucrose reward) was interwoven with an unrewarded and a control odor, and responses were compared from before the conditioning phase (pre-test) to after (post-test).

We demonstrate that learning leads to a restructuring of the rate-based odor representations in the neuronal network. Recorded units (spikes from single neurons or mixtures of spikes from several, but few neurons) often change their response characteristics in the course of conditioning. However, both rate increases and rate decreases are frequent, such that the average level of activity does not show a difference between the rewarded and unrewarded odors. However, by representing an odor by the ensemble response of neurons in the network (cf. also [Laurent, 1996](#); [Daly et al., 2004](#)), we demonstrate that the composition of rate responses across neurons that codes for an odor changes consistently in the course of conditioning. In particular, the pattern representing the rewarded odor changes most drastically. The large pool of neurons required for this analysis is obtained by pooling responses across bees and interpreting them as signals from different electrodes in the same animal. The finding of changes in the ensemble response is in agreement with previous work demonstrating recruitment and loss of responding neurons to odor conditioning in the AL ([Daly et al., 2004](#); [Yu et al., 2004](#)).

The LFPs typically display a complex response to odor stimulation, visualized in the temporal pattern of their frequency composition. Nevertheless, we frequently observe oscillations around 50 Hz shortly after odor onset. We find that for the rewarded odor only, learning induces a reduction of the power content at this frequency and an increase in power in a broad band around 15 – 40 Hz. Similar observations have been reported in the olfactory system of the rat ([Ravel et al., 2003](#); [Martin et al., 2006](#)). Therefore, both population measures, ensemble rate response and LFP, show distinct changes in response to the rewarded odor. However, we do not find a relation between these power changes and the spike rates, suggesting that the LFP reflects a change in the synchronization of neuronal ensembles in response to the rewarded odor. Using the techniques introduced in [Sec. 1.2](#), we therefore investigate the phase relationship between spiking responses and the LFP. In the unconditioned bee, most neurons exhibit significant locking to the LFP at the 50 Hz band. From pre- to post-test, however, for the rewarded odor

the change in the number of locked neurons per frequency band correlates with the respective power change in that band. In particular, the change in the frequency composition towards the lower 15 – 40 Hz band is reflected by an increase of the size of the respective phase-locked ensemble. Thus, the described changes in response to differential conditioning are indicative that plastic changes in the network are expressed on two levels, a coarse reorganization of rate responses across projection neurons, and a specific entrainment of neurons at a decreased frequency in response to the learned odor. Summarizing within the topic of this thesis, this report demonstrates how the spike-LFP relationship reveals that features of the LFP reflect a group of neurons that dynamically changes its synchronization characteristics in the response to an external stimulus.

The most pressing open question is surely, how the emergence of these synchronized responses is related to the observed restructuring of the rate responses. However, these questions cannot be answered conclusively in the present data set due to the limited number of recordings and odor presentations that provide a too sparse sampling of the network response. Partial answers may be obtained through a more simple experiment that focuses in more detail on the overlap of phase-locked patterns (David et al., 2007) to the 50 Hz oscillation and ensemble rate responses to singular odor presentations in the absence of learning.

The study introduced in Sec. 2.2 investigates the role of different neuronal subtypes in the striatum of the anesthetized rat in relaying cortical oscillations to the basal ganglia network. To present, at least four putative types of striatal neurons have been identified: the medium spiny projection neurons, in addition to the fast spiking, the tonically active, and the low-threshold spiking interneurons. In the present study we provide the first classification of recorded neurons into these four categories in a single data set. This analysis is based on waveform parameters as well as the firing statistics (rates and interspike interval distributions) of individual neurons. The resulting classification allows us to analyze to what extent individual neuron classes are involved in the generation of oscillations in different frequency bands exhibited within the striatal network.

Cortical input enters the basal ganglia mainly through the striatum. Recordings of electrocorticograms (revealing primarily cortical input) and

striatal LFP display spontaneous oscillations at both low (2 – 9 Hz) and high (35 – 80 Hz) frequencies, and oscillations are coherent between the two types of measures. Although the recordings were performed in the anesthetized animal, the use of halothane as anesthetic provides a setting where oscillations at multiple frequency bands are exhibited in a stable fashion, thus aiding our analysis. In particular, the θ and γ frequencies observed under halothane (Imas et al., 2004) are readily observed in the awake animal (Tort et al., 2008). While the involvement of striatal neurons in slow oscillations has been well described, no phase-locking of neurons to the fast γ rhythms has been reported previously (Berke et al., 2004).

In the present study, we show that only putative fast spiking interneurons exhibit significant phase locking to oscillations in the γ range. The finding is backed by a cross-correlation analysis between the different cell types that reveals significant correlations on a comparably fast time scale only for this neuron type. A further analysis that corroborates this the special role of fast spiking interneurons, previously presented by Denker et al. (2008), reveals how a locked fast spiking interneuron modulates its firing rate and strength of phase-locking in accordance with spontaneously occurring state changes in the frequency composition of the LFP oscillations. In contrast, all neuron types are entrained to the slow oscillations. The finding identifies fast spiking interneurons as the key contributors to the propagation of cortical oscillations into the striatum, a view that matches similar observations in cortex (Hasenstaub et al., 2005; Cardin et al., 2009) and hippocampus (Klausberger et al., 2003). Given the high proportion of $\approx 97\%$ of projection neurons in the striatum (Rymar et al., 2004), it is not surprising why typically the relationship between striatal γ oscillations and neuronal spiking activity remains obscured.

The finding that in the absence of behavior oscillations are driven by a single cell type further highlights the special role of interneurons (Sec. 1.1.2) in generating coherent rhythms in the high frequency range without the need for feedback from an entrained excitatory population. Thus, the results add further support for the current hypothesis that interneurons may act as a mechanism that structures the neuronal dynamics in the temporal domain (Fries et al., 2007; Berens et al., 2008), even in subcortical structures such as the basal ganglia. As a practical application, spike-LFP phase relationships may serve as an additional parameter, next to morphology, wave form

characteristics and spiking statistics, to separate neurons into their respective physiological categories. It remains to be shown to what extent these findings are transferable to the behaving animal, and how a modulation of cortical oscillatory drive is represented in the response of the phase-locked population of interneurons.

In comparison with the work presented in [Sec. 2.1](#), this study relates oscillations of the LFP to synchronization within a specific class of neurons based on their physiology, as opposed to their role in the coding of a stimulus. In both studies, however, knowledge of the spike-LFP relationship provides an handle to characterize and classify a specific ensemble of cells, whose synchronized action is reflected in the LFP.

In the subsequent three sections of [Chap. 3](#), we take the approaches presented so far one step further and show that in the behaving animal the cortical LFP reflects the transient coactivation of neuronal groups in the context of neuronal assemblies (cf. [Sec. 1.1.3](#)). In doing so, we establish the link between the precise synchronization of spikes and the oscillatory population activity of the LFP. The experimental basis for all three studies are recordings of LFP and spikes from motor cortex of monkeys in a delayed pointing task ([Roux et al., 2003, 2006](#)). The experimental design provides a time window between trial start and movement onset that has been shown to exhibit a reliable oscillation in the β range (e.g., 15 – 40 Hz) that is commonly associated with preparatory and attentional modulation ([Sanes and Donoghue, 1993](#); [Murthy and Fetz, 1996a](#)) and may also represent modulatory input from higher cortical areas ([Lebedev and Wise, 2000](#)). In addition, we observe transient spike synchronization that is correlated to behaviorally relevant time points of the trial ([Riehle et al., 1997](#); [Grammont and Riehle, 2003](#); [Kilavik et al., 2009](#)). Therefore, these experiments deliver rich data that provides signatures of synchrony on the level of single spikes and the population.

Neurons typically display a flat auto-correlation. Previous studies suggest that only during periods of enhanced LFP activity (typically occurring in spindles that encompass on the order of 5-10 oscillatory cycles), neurons may become oscillatory and entrained to the LFP oscillation ([Murthy and Fetz, 1996b](#); [Donoghue et al., 1998](#)). In addition, cross-correlations between these neurons also become oscillatory in about half of the recorded popula-

tion. In a first study (Sec. 3.1), we quantify the dependence of the locking of single spikes to the LFP as a function of the instantaneous amplitude (or envelope) of the LFP oscillations (i.e., their magnitude) using the phase analysis technique (Sec. 1.2). Consistent with previous work, we find that only a fraction of neurons (30–40%) displays significant locking to the LFP. By separating spikes into two groups based on the magnitude of LFP oscillations at which they occur, we could quantitatively verify that indeed mainly spikes associated with high LFP amplitudes are responsible for the observed locking of a given neuron. This observation is largely independent of the threshold on LFP amplitudes used to separate the two groups. Therefore, we conclude that the small percentage of the spikes that occur during episodes of the largest LFP oscillations are able to explain the phase preference of the cell's spikes. In contrast, spikes are largely independent of the LFP rhythm during low amplitudes of the population oscillations.

We investigated the dependence of observed phase locking on the different periods of a behavioral trial. This leads to a quantitative, but not qualitative difference in the results: in each period, spikes that occur at large LFP amplitudes are still best locked to the field potential oscillation. However, the percentage of phase-locked neurons decreases in total during each of the short periods of the trial compared to the complete trial, thus indicating a transient nature of the spike-LFP locking. Given that the LFP modulates its power on a roughly comparable time scale as the length of these periods, our interpretation of transient synchronization is in full agreement with the dependence of spike-LFP locking on LFP amplitudes. For all observed data sets, neurons tend to keep a fixed phase relationship to the LFP independent of the amplitude or the choice of the electrode used to record the LFP. In conclusion, we present evidence to support the hypothesis that increased LFP oscillation amplitudes are related to a stronger degree of synchronization between the LFP and spike signals. Thus, interpreting the percentages of locked neurons as probabilities to lock to the LFP, and considering the constant mean phase exhibited by locked spikes, our study links LFP amplitudes to the size of the neuronal population synchronized to the population signal.

It is commonly assumed that the oscillatory structure of mesoscopic brain signals, such as the local field potential (LFP), reflects synchronous spiking

activity (Sec. 1.1.2). However, this hypothesis was never verified directly, and what types of synchronized network activity on the level of single spikes are expressed in the LFP remains speculative. A recent study by Poulet and Petersen (2008) exposes the nature of this question: despite strong correlations of simultaneously intracellularly recorded membrane potentials and LFP signals, this correlation was not simply reflected in the highly specific occurrence of synchrony on the spike-spike level. In the central study of this thesis presented in Sec. 3.2, we are able to prove for the first time that the long-standing belief of LFPs as a reflection of synchronized spiking activity indeed has objective grounds, and provide a framework to reconcile the apparent discrepancy to Poulet and Petersen (2008).

To this end, we classified individual spikes according to their participation in synchronous events with spikes from a second neuron. We show experimentally that correlated spiking activity on a fine temporal scale (millisecond range) is indeed better locked to the phase of the LFP than single spikes. However, spike coincidences are better entrained to the LFP than expected by chance (on the basis of the locking of single spikes) only in those time intervals where the large number of coincidences indicates that the neurons coordinate their spike emission (Unitary Event analysis, see Grün et al., 2002a,b; Grün, 2009). The excess amount of synchrony observed in these periods is commonly attributed to a process that repeatedly coactivates a specific subgroup of cells where the recorded neurons are both part of, i.e., an assembly (Sec. 1.1.3). Therefore, our findings provide evidence of the hitherto unproven hypothesis (Eckhorn and Obermueller, 1993; Murthy and Fetz, 1996b; Donoghue et al., 1998; Singer, 1999) that assembly activations are the synchronized activity that is reflected in the LFP oscillation.

Our results are robust under variation of the various parameters inherent in the analysis, such as the allowed temporal jitter of coincidences or the particulars of the filtering process of the LFP. Similar, albeit less pronounced, findings already emerge using the simpler spike-triggered average (Sec. 1.2.1). However, only by means of the phase synchronization analysis it is possible to quantitatively disentangle the subtle effects of increases in locking strength due to the observation of excess synchrony and due to the amplitude effects as described in Sec. 3.1.

The experimental evidence brings up two main questions: first, how can we understand the (weakly) non-uniform distribution of spike phases that

are not involved in a synchronous event? Second, why do assemblies always prefer to lock to a specific phase of the LFP, independent of the identity of the participating neuron? To gain an understanding of the first question, we linked the concept of neuronal cell assemblies to the LFP in a unifying conceptual model based on these experimental results. Central to the model is the concept that phase-locked spikes of neurons that are not activated in an observed assembly should therefore be involved in a non-observed assembly. This simple insight highlights the LFP population signal as a measure of synchronous activations of cells within the network that are hidden from direct observation. Using simple calculations, the conceptual model therefore allows us to give an estimate of the minimal percentage of spiking activity that results from assembly activations (11%). This functional framework bridges the apparent discrepancy between synchrony on the spike and on the population levels in a quantitative way.

The second question as to why synchronized activations of neurons in assemblies are themselves synchronized to a specific phase of the LFP remains unresolved. It has been suggested that the time-locked activation of assemblies could reflect a temporal queue to multiplex coactivations of different compositions of neurons on a time grid (Singer, 1999). Thus, the LFP would provide a temporal structure on which assemblies are activated. The recent study by Womelsdorf et al. (2007) shows evidence for such a picture: Interactions between groups of neurons (given by the power correlations indicating their rate covariations) are strong only when their spiking activities exhibit a specific phase lag with respect to one another - and a specific lag to the the LFP. Our study goes beyond these findings by identifying only one specific phase lag at which neurons exhibit not only slow correlations but precise neuronal synchrony. Combined with current theories on the generation of fast oscillations by the modulatory inputs onto pyramidal cells from inhibitory circuits of interneurons (Fries et al., 2007; compare also discussion of Sec. 2.2), we could hypothesize that these oscillations act as a gating mechanism for the effective integration of synchronized input patterns (Cardin et al., 2009), that fascilitate, inhibit, or structure the functional activation of assemblies in time.

Combined with the findings detailed in Sec. 3.1, we extend our interpretation that LFP amplitudes predict the size of the population of spikes involved in assembly activity. However, a more quantitative analysis of this

hypothesis in the framework of our conceptual model is the subject of future work (see also [Sec. 4.3](#)).

In the final study presented in this thesis ([Sec. 3.3](#)), we show how to extract an estimate of the proportion of spikes γ that contributes to the assembly dynamics from the conceptual model presented above. In order to arrive at this estimate given the experimental findings in [Sec. 3.2](#), we first demonstrate formally how to estimate the relative number of coincidences β stemming from an assembly –as opposed to chance coincidences– during periods where excess synchrony is observed (i.e., a period detected by the Unitary Event analysis). Integrating this result with the measured phase distribution of coincidences during the corresponding Unitary Event periods allows to estimate the precision with which assemblies lock to the LFP. In a last step, this measure of the temporal structure of assembly activations allows to backtrack the fraction γ from the phase distributions of single spikes. The crucial step of this method is to accurately obtain the parameter β independently from our conceptual model, based on the spike data alone. A simple toy model is therefore used to calibrate our method, before it is applied to the experimental data. We find that about $\gamma = 22\%$ of spikes in the motor cortical network are involved in assemblies.

Our conceptual model represents a highly simplified description, yet by combining measurements of synchrony on the local and mesoscopic scales it reveals parameters of the network dynamics that remain hidden on the two individual levels of observation. The nature of the estimation process requires that the population activity must be represented by a large sample of spike-LFP and –in particular coincidence-LFP– combinations. To this end, experimental data in different sessions and two different monkeys are pooled, and γ can therefore only represent a rough estimate of the fraction of spikes originating from of assembly activity. Nevertheless, the fact that the parameter β (central to the estimate of γ) follows a surprisingly narrow distribution even when sampled across this pooled data set corroborates our result.

By increasing the amount of available data, a similar but more detailed analysis could be performed, in particular, by pre-grouping neurons according to rate, response behavior, or similar features and then calculating the individual probabilities to participate in an assembly. Thus, we could in

principle identify whether certain groups of neurons are more involved in the assembly dynamics than others. This type of measurement is hardly possible by evaluating assembly participation from the spiking dynamics alone due to the low probability of observing an assembly in the first place. As we have seen, in contrast the phase distributions provide information on the degree of assembly activity independent of the observation of assembly.

In summary, incorporating population signals in the analysis of spike data may serve to overcome problems related to the undersampling of the activities in multiple single-neuron recordings. Taken together, the wealth of experimental and theoretical findings presented in [Chap. 3](#) substantially strengthen the argument in favor of the assembly hypothesis of neuronal processing and provide a promising vista to reinterpret and efficiently exploit neuronal population signals in a wide range of applied areas (e.g., brain-machine interfaces, cf. [Mehring et al., 2003](#); [Hochberg et al., 2006](#)). The success of the conceptual model in predicting parameters of the underlying network dynamics shows how brain processes on the finest spatio-temporal scale are reflected in a robust population signal.

4.2 General comments

Let us now briefly discuss two concerns that are of relevance for all studies presented in this thesis. First, we point to a technical issue related to the generation of surrogate data for the phase locking analysis that was introduced in [Sec. 1.2.1](#) of the Introduction. As a reminder, the null hypothesis for the parameter R describing the degree of non-uniformity of phase distributions was generated by locally shuffling the interspike intervals of the spike train data in order to destroy the precise spike-LFP phase relations while retaining to first order the regularity of the spike data. The definition of *locally* must be adapted to data set: in the study presented in [Sec. 2.1](#), where only three trials are available, data was too scarce as to allow for shuffling windows that follow non-stationarities expressed by sudden rate changes. However, we observe that those units that showed strong rate responses had a tendency for regularity in their spiking activity during the response. Therefore, the interval shuffles nevertheless produce a far more conservative null hypothesis compared to the Poisson assumption underlying the Rayleigh test. In contrast, the extremely low firing rates encountered in [Sec. 2.2](#) and the large

number of trials in the studies shown in [Chap. 3](#) posed no problem to produce shuffles that operate on time scales that approximate the actual rate non-stationarities (see, e.g., [Fig. 1.1](#)). A completely different approach based on surrogates of the continuous signal (e.g., using techniques described in [Thiel et al., 2002](#)) also provides a promising alternative, as the regular oscillations exhibited by the LFP might provide a more stable signal that is less contaminated by non-stationary effects. We suggest that a new type of surrogate data that operates directly on the statistics of the phase description of the signal ([Hurtado et al., 2004](#)) may remedy some of the problems related to non-stationarity and higher-order serial correlations ([Farkhooi et al., 2009](#)). However, it is in the nature of such surrogates that their properties are in general more difficult to control.

Second, we return to the question raised in [Sec. 1.1.1](#) concerning possible ephaptic effects, i.e. the possibility that LFPs directly influence the sub-threshold membrane potentials, or even spike generation. As pointed out, not enough information is available to rule out the possibility of ephaptic effects. Fortunately, in the context of the current studies, we believe that this question is of little relevance. In [Sec. 2.1](#), we observe a large variation in the locking of individual neurons to the LFP, and therefore we must conclude that ephaptic effects are small enough not to induce locking by themselves. Moreover, in [Sec. 2.2](#) it is implausible that one particular neuronal subtype should be affected by extracellular fields to a much greater extent than others. Thus, while ephaptic coupling might quantitatively influence the precise strength of spike-LFP locking, they are unlikely to explain the different locking characteristics of different neuronal subtypes. Finally, although external fields in the experiments outlined in [Chap. 3](#) may play a role in determining the non-uniform phase distribution of single spikes, they are completely unable to explain the selective enhancement of the phase coupling of those coincidences that are part of a Unitary Event compared to those that occur by chance. However, ephaptic effects could lead to small corrections in the value estimated for the probability γ of spikes to belong to an assembly. In fact, we propose that it could be possible to roughly estimate the ephaptic effect by adjusting the intrinsic phase distribution of spikes not involved in assemblies (assumed as a uniform distribution $p_n(\phi)$ in the conceptual model of [Sec. 3.3](#)) such that the same value of γ simultaneously and self-consistently explains the measured phase distributions of both, single spikes and chance

coincidences (first two equations of the conceptual model).

4.3 Outlook

In this final section we give a short outlook on further possibilities to extend the work presented in [Chap. 3](#). First of all, it becomes necessary to formally reconcile the findings on the dependence of spike-LFP locking on LFP amplitude ([Sec. 3.1](#)) with the framework of the conceptual model linking the LFP to assembly activations ([Sec. 3.3](#)). In particular, our findings indicate that the parameter γ (probability of a spike to belong to an assembly) is dependent on the strength of LFP oscillations – most likely in a nonlinear fashion considering the results of [Sec. 3.1](#). To investigate such a dependency would likely require a large pool of data. Fortunately in this respect, the spatial homogeneity of cortical LFP (i.e., wave propagation with large wavelengths, cf., [Rubino et al., 2006](#)) gives reason to hope that γ does not have a significant dependence on the lateral cortical location. Relating LFP amplitude to synchronous spiking might help to better reconcile findings on the behavioral correlates of both synchronous spiking ([Riehle et al., 1997](#)) and LFP power (e.g., [Murthy and Fetz, 1992](#); [Mehring et al., 2003](#)).

Our results indicate that only part of the neuronal population is involved in assembly activity. A natural question arises: what is the mechanism that causes non-assembly spikes? It is well known that rate responses in motor cortex are correlated with movement related aspects (e.g., directional tuning [Georgopoulos et al., 1986](#); [Moran and Schwartz, 1999](#); [Georgopoulos et al., 2007](#); [van Hemmen and Schwartz, 2008](#)) and can be predictive of movement on the level of the neuronal ensembles ([Fetz, 1999](#); [Wessberg et al., 2000](#); [Moritz et al., 2008](#)). Can we confirm the hypothesis of a dual coding scheme of rate and synchrony in motor cortex (e.g., from the view point of directional tuning), and how do these processes overlap and disentangle? To this end, one hypothesis we might put forth is that if two completely separate sets of spikes code for different aspects by rate and synchrony, then the predictive performance of rate responses should increase if putative assembly spikes (after all, about one fourth according to our estimate) are not considered.

However, such an approach requires definite knowledge of whether an identified spike from the network is part of a synchronous event, or not. Even with recent advances of recording techniques that allow simultaneous

recordings from 100 or more electrodes (Nicolelis et al., 1997; Csicsvari et al., 2003; Euston et al., 2007; Fujisawa et al., 2008), the system is undersampled and recording sites are typically dispersed over a large area. Moreover, there are challenging problems in analyzing the resulting data given current methods (Grün, 2009) due to the complexity of the task and the resulting combinatorial explosion (Brown et al., 2004). Our findings on how the occurrence of excess synchrony is related to the LFP oscillation cycle offers an exciting window to circumvent the undersampling problem by independently identifying best possible candidates for spikes that are involved in assembly activations based on the LFP. To this end, we may use the knowledge of the probability of assembly activity to occur (Sec. 3.3), its locking characteristics (Sec. 3.2), and possibly its dependence on the LFP amplitude (see above). Therefore, the LFP might provide an additional measure to aid the detection of likely assembly activity where only spikes from a single neuron are directly observed.

The studies presented in this thesis make no assumption on the group size (assembly size) of individual synchronous events in the system. In order to further exploit the LFP as an indicator of hypothesized large-scale synchronous firing in the network, we must establish our model in the context of observed synchronous events of higher-order, i.e., repeated activations of patterns of more than 2 neurons. The task is quite a challenging one, as the probability of detection of significant patterns of coactivated neurons is low, and only a handful of methods exist (e.g., Baker and Gerstein, 2000; Gerstein, 2004; Shlens et al., 2006; Schneidman et al., 2006; Schrader et al., 2008; Shlens et al., 2009), most of which rely on inferring higher-order structure from the pair-wise correlations, or which are practically limited to analysis of few neurons due to the underlying computational complexity. Nevertheless, in the context of this thesis the simplest prediction we could postulate is that all synchronized events, no matter what group size, would exhibit the same phase relationship as observed in the case of pair-wise correlations. For LFP amplitudes, on the other hand, it is hard to make a prediction: one theory could be that the amplitude of LFP oscillations is indicative not only of the number of active assembly activations, but also of their respective group size. By understanding the relationship between the LFP and synchronous events consisting of a large number of neurons, we have hope to discover new estimates, such as the size of a typical assembly, from the combined knowledge

of how synchronous activity is represented on the different scales of cortical measurements.

Acknowledgements

I would like to thank Prof. Dr. Randolph Menzel for his support in supervising my thesis, and for the effort and time that he put into a fruitful collaboration in a challenging project.

My deepest gratitude extends to my supervisor Dr. Sonja Grün for all the countless hours and enormous amount of work she has invested towards my training in the neurosciences. In working with her, I profited from her diligence and her ability to keep me focused on the important questions at hand. Last not least, I also thank her for creating a work environment that provides many good times, both inside and outside of the laboratory.

Likewise, I would like to thank my supervisor Dr. Marc Timme who continuously paved the way that led to the work presented in this thesis, ever since the completion of my Diploma thesis. His sincere interest in scientific questions, productive criticism and thought-provoking suggestions always provided the right encouragement to dig a little deeper.

I thank all my collaborators in the various projects detailed in this thesis. In particular I would like to express my gratitude for the close interactions with Sébastien Roux, Henrik Lindén, Andrew Sharott, and Robert Finke. Moreover, I would like to acknowledge the overwhelming amount of support extended to me by my senior collaborators Dr. Alexa Riehle and Dr. Markus Diesmann.

I am grateful to have spent my scientific life with truly fantastic people in the groups in Berlin and Wako-shi, who would always be open for a scientific discussion. In particular, I would like to thank Antonio Paziènti, Benjamin Staude, and Denise Berger for their share in making this wonderful group become a reality in the first place.

My sincere thank goes out to my wonderful, loving parents Tillu and Manfred, whose unconditional support, openness, and encouragement has carried me over all these years. Also, I am grateful that despite the distances apart, I am fortunate for the company of reliable and caring friends in Göttingen, Berlin, Wako-shi, and elsewhere around the globe.

Anhang A

Zusammenfassung der Hauptergebnisse (Summary of Main Findings in German)

Gehirnaktivität wird mittels verschiedener Messmethoden beobachtet, die eine große Spannbreite räumlicher und zeitlicher Skalen umfassen. Im Falle extrazellulärer Ableitungen wird das resultierende Signal der Elektrode in der Regel auf zweierlei Ebenen interpretiert: während die höherfrequenten Komponenten das präzise zeitliche Auftreten von Aktionspotentialen nahe der Elektrode sichtbar machen, so liefern die niederfrequenten Anteile ein mesoskopisches Populationssignal –als lokales Feldpotenzial (LFP) bezeichnet– welches von der überlagerten Aktivität einzelner Neurone in einem größeren Umkreis dominiert wird. Obwohl im allgemeinen angenommen wird, dass die oszillatorischen Eigenschaften des LFP durch synchronisierte unter-schwellige Aktivität innerhalb der neuronalen Population hervorgerufen werden, besteht größtenteils Unklarheit darüber inwiefern die Oszillationen mit der Einzelspikeaktivität, und insbesondere mit dem Auftreten synchronisierter Spikeaktivität, zusammenhängen. Diese Dissertationsschrift enthält fünf Berichte, die unterschiedliche Aspekte beleuchten, wie der Zusammenhang zwischen der konzertierten Spikeaktivität und dem LFP Informationen über die Eigenschaften und die dynamische Organisation des zugrunde liegenden neuronalen Systems bereitstellt. Der Fokus dieser Dissertation ruht auf einer Folge von dreien dieser Berichte, die das fehlende Bindeglied zwischen präzisen Aktivitätsmustern koordinierter Aktionspotenziale und dem

LFP aufdecken.

Um die typischerweise schwache Kopplung zwischen Spikes und dem LFP zuverlässig quantifizieren zu können, wurde eine neuartige Methode entwickelt, die aus den Techniken zur Phasenanalyse kontinuierlicher Signale hervorgeht. Im Gegensatz zu herkömmlichen Maßen, die auf der Signalmittlung aufbauen (z.B. das Spike-triggered Average), analysiert diese Methode direkt die Spike-LFP Phasenkopplung auf der Basis des Einzelspikes. Alle in dieser Dissertation zusammengefassten Berichte zeigen, wie dieses methodische Werkzeug sensitiv neuronale Ensembles identifizieren und charakterisieren kann, die an die beobachteten Populationssignale gekoppelt sind.

Die erste Studie demonstriert lernabhängige Änderungen der neuronalen Antwort auf Duftstimulation im Ausgangsbereich des Antennallobus der Honigbiene während einer olfaktorischen Konditionierung. Der Lernprozess geht mit einer Änderung der Ensembleantwort auf Düfte einher (basierend auf den Ratenantworten einzelner Neurone), die für den belohnten Duft am stärksten ist. Darüber hinaus zeigen wir, dass entsprechende Änderungen in der Power des LFP in bestimmten Frequenzbändern für den belohnten Duft mit Änderungen in der Größe und der Zusammensetzung der neuronalen Subpopulation, die zu dem entsprechenden Band gekoppelt ist, korrelieren. Somit stärkt die Analyse in diesem Bericht die Hypothese, dass das LFP als Monitor für die lern-induzierte Reorganisation der zeitlichen Ensemblerepräsentation des externen Stimulus dient.

In einer zweiten Studie analysieren wir separat die Rolle vierer identifizierter neuronaler Subtypen im Striatum der anästhesierten Ratte in der Weiterleitung von Oszillationen kortikalen Ursprungs in die Basalganglien. Basierend auf der Spike-LFP und Spike-ECoG (Elektrokortikogramm) Phasenkopplung, und gestärkt durch die Analyse interneuroner Kreuzkorrelationen, identifizieren wir die Neuronenklasse, die Fast Spiking Interneurone, die zu der Erzeugung oszillatorischer LFP Komponenten des hochfrequenten Gammabereichs in Bezug gesetzt werden kann. Somit reflektiert das LFP hier die synchrone Aktivität eines speziellen, physiologisch definierten Neuronentyps.

Der größte Teil dieser Dissertation besteht aus drei aufeinanderfolgenden Arbeiten, die den Zusammenhang zwischen synchroner Spikeaktivität mit Millisekundengenauigkeit und der im LFP sichtbaren synchronen Massenaktivität aufzeigen. Die experimentelle Datenbasis für alle drei Studien

bilden Ableitungen aus dem Motorkortex des Affen in einem Experiment, in dem das Tier trainiert wird, eine zeitverzögerte Armbewegung zu vollführen (delayed pointing task). Zuerst zeigen wir, dass die Amplitude (oder Einhüllende) der LFP Oszillationen ein Maß für die Größe der phasengelockten Population darstellt. Durch eine separate Phasenanalyse von Spikes und Spikekoinzidenzen belegen wir in der darauf folgenden Arbeit direkt die langläufige Hypothese, dass synchrone Spikeaktivität im Feldpotenzial reflektiert ist. Entgegen der Intuition gilt diese Aussage jedoch nur für diejenigen präzisen Koinzidenzen, die in Zeitfenstern auftreten in denen die Anzahl der detektierten Koinzidenzen die Erwartung signifikant übertrifft. Diese überschüssigen Koinzidenzen stellen eine Signatur für koordinierte Spikemuster spezifischer neuronaler Subgruppen (Zellensembles) dar. Ein konzeptionelles Modell erklärt die experimentellen Daten im Kontext der Idee neuronaler Kodierung mittels Zellensembles. In der letzten Studie zeigen wir, wie die kombinierte Messung von Synchronizität auf der Spike- und Populationsebene einen makroskopischen Parameter schätzt, der angibt in welchem Maße aktive Zellensembles zu der gesamten Spikeaktivität des Netzwerkes beitragen. Zusammenfassend erbringen die drei Arbeiten den ersten Beweis dafür, dass LFP Oszillationen ein Abbild orchestrierter Aktivität neuronaler Ensembles darstellen, wie sie durch eine der faszinierendsten aktuellen Theorien neuronaler Informationsverarbeitung vorhergesagt werden.

Bibliography

- Abeles M (1991) *Corticons: Neural Circuits of the Cerebral Cortex* Cambridge Univ Press, NY.
- Abeles M, Hayon G, Lehmann D (2004) Modeling compositionality by dynamic binding of synfire chains. *J Comput Neurosci* 17:179–201.
- Anastassiou, Montgomery, Barahona, Buzsáki, Koch (2008) Effect of spatially inhomogeneous electric fields on single neurons. In *Frontiers in Computational Neuroscience. Conference Abstract: Bernstein Symposium 2008*. doi: 10.3389/conf.neuro.10.2008.01.023, p. doi: 10.3389/conf.neuro.10.2008.01.023.
- Baker SN, Gerstein GL (2000) Improvements to the sensitivity of gravitational clustering for multiple neuron recordings. *Neural Comput* 12:2597–2620.
- Baker SN, Lemon RN (2000) Precise spatiotemporal repeating patterns in monkey primary and supplementary motor areas occur at chance levels. *J Neurophysiol* 84:1770–1780.
- Berens P, Keliris GA, Ecker AS, Logothetis NK, Tolias AS (2008) Feature selectivity of the gamma-band of the local field potential in primate primary visual cortex. *Front Neurosci* 2:199–207.
- Berke JD, Okatan M, Skurski J, Eichenbaum HB (2004) Oscillatory entrainment of striatal neurons in freely moving rats. *Neuron* 43:883–896.
- Boashash B (1992) Estimating and interpreting the instantaneous frequency of a signal. I. Fundamentals. *Proc. IEEE* 80:520–538.
- Brown EN, Kass RE, Mitra PP (2004) Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nat Neurosci* 7:456–461.

- Buzsáki G, Bickford RG, Ponomareff G, Thal LJ, Mandel R, Gage FH (1988) Nucleus basalis and thalamic control of neocortical activity in the freely moving rat. *J Neurosci* 8:4007–4026.
- Buzsáki G (2004) Large-scale recording of neuronal ensembles. *Nat Neurosci* 7:446–451.
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929.
- Cardin JA, Carlén M, Meletis K, Knoblich U, Zhang F, Deisseroth K, Tsai LH, Moore CI (2009) Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature* 459:663–667.
- Courtemanche R, Lamarre Y (2005) Local field potential oscillations in primate cerebellar cortex: synchronization with cerebral cortex during active and passive expectancy. *J Neurophysiol* 93:2039–2052.
- Csicsvari J, Henze DA, Jamieson B, Harris KD, Sirota A, Barthó P, Wise KD, Buzsáki G (2003) Massively parallel recording of unit and local field potentials with silicon-based electrodes. *J Neurophysiol* 90:1314–1323.
- Daly KC, Christensen TA, Lei H, Smith BH, Hildebrand JG (2004) Learning modulates the ensemble representations for odors in primary olfactory networks. *Proc Natl Acad Sci U S A* 101:10476–10481.
- David F, Hugues E, Buonviso N (2007) Phase-locking of mitral cell activity during gamma lfp oscillations: experiments and modeling. In *COSYNE Conference Abstracts (online)*., p. 172.
- Denker M, Timme M, Grün S (2008) Analyzing phase synchronization between neuronal spiking activity and local field potentials. In *Neuroscience Research (Proceedings of the 31st Annual Meeting of the Japan Neuroscience Society, Neuroscience 2008)*, Vol. 61, Supplement 1, pp. S280, p3–v19.
- Denker M, Timme M, Roux S, Riehle A, Grün S (2005) Detecting transient temporal relationships between spikes and lfp by phase analysis. In *Online Abstract Viewer/Itinerary Planner. Washington, DC. Society for Neuroscience.*, p. Program No. 970.10. Society for Neuroscience 35th Annual Meeting, Washington, DC. 12.-16. November.

- Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J Neurosci* 19:4595–4608.
- Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. *Nature* 402:529–533.
- Donoghue JP, Sanes JN, Hatsopoulos NG, Gaál G (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol* 79:159–173.
- Eckhorn R, Bauer R, Jordan W, Brosch M, Kruse W, Munk M, Reitboeck HJ (1988) Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biol Cybern* 60:121–130.
- Eckhorn R, Obermueller A (1993) Single neurons are differently involved in stimulus-specific oscillations in cat visual cortex. *Exp Brain Res* 95:177–182.
- Elul R (1971) The genesis of the eeg. *Int Rev Neurobiol* 15:227–272.
- Euston DR, Tatsuno M, McNaughton BL (2007) Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. *Science* 318:1147–1150.
- Faber T, Joerges J, Menzel R (1999) Associative learning modifies neural representations of odors in the insect brain. *Nat Neurosci* 2:74–78.
- Farkhooi F, Strube-Bloss MF, Nawrot MP (2009) Serial correlation in neural spike trains: experimental evidence, stochastic modeling, and single neuron variability. *Phys Rev E Stat Nonlin Soft Matter Phys* 79:021905.
- Fetz EE (1999) Real-time control of a robotic arm by neuronal ensembles. *Nat Neurosci* 2:583–584.
- Friedrich RW, Habermann CJ, Laurent G (2004) Multiplexing using synchrony in the zebrafish olfactory bulb. *Nat Neurosci* 7:862–871.
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563.

- Fries P, Nikolić D, Singer W (2007) The gamma cycle. *Trends Neurosci* 30:309–316.
- Fries P, Schröder JH, Roelfsema PR, Singer W, Engel AK (2002) Oscillatory neuronal synchronization in primary visual cortex as a correlate of stimulus selection. *J Neurosci* 22:3739–3754.
- Fujisawa S, Amarasingham A, Harrison MT, Buzsáki G (2008) Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* 11:823–833.
- Georgopoulos AP, Schwartz AB, Kettner RE (1986) Neuronal population coding of movement direction. *Science* 233:1416–1419.
- Georgopoulos AP, Merchant H, Naselaris T, Amirkian B (2007) Mapping of the preferred direction in the motor cortex. *Proc Natl Acad Sci U S A* 104:11068–11072.
- Gerstein GL, Bedenbaugh P, Aertsen MH (1989) Neuronal assemblies. *IEEE Trans Biomed Eng* 36:4–14.
- Gerstein GL (2004) Searching for significance in spatio-temporal firing patterns. *Acta Neurobiol Exp (Wars)* 64:203–207.
- Goense JBM, Logothetis NK (2008) Neurophysiology of the bold fmri signal in awake monkeys. *Curr Biol* 18:631–640.
- Gold C, Henze DA, Koch C (2007) Using extracellular action potential recordings to constrain compartmental models. *J Comput Neurosci* 23:39–58.
- Grammont F, Riehle A (2003) Spike synchronization and firing rate in a population of motor cortical neurons in relation to movement direction and reaction time. *Biol Cybern* 88:360–373.
- Grün S (2009) Data-driven significance estimation for precise spike correlation. *J Neurophysiol* 101:1126–1140.
- Grün S, Diesmann M, Aertsen A (2002a) Unitary events in multiple single-neuron spiking activity: I. Detection and significance. *Neural Comput* 14:43–80.

- Grün S, Diesmann M, Aertsen A (2002b) Unitary events in multiple single-neuron spiking activity: II. Nonstationary data. *Neural Comput* 14:81–119.
- Hallem EA, Carlson JR (2006) Coding of odors by a receptor repertoire. *Cell* 125:143–160.
- Harris KD, Henze DA, Hirase H, Leinekugel X, Dragoi G, Czurkó A, Buzsáki G (2002) Spike train dynamics predicts theta-related phase precession in hippocampal pyramidal cells. *Nature* 417:738–741.
- Hasenstaub A, Shu Y, Haider B, Kraushaar U, Duque A, McCormick DA (2005) Inhibitory postsynaptic potentials carry synchronized frequency information in active cortical networks. *Neuron* 47:423–435.
- Hayon G, Abeles M, Lehmann D (2005) A model for representing the dynamics of a system of synfire chains. *J Comput Neurosci* 18:41–53.
- Hebb D (1949) *The Organization of Behavior: A Neuropsychological Theory* Wiley, New York.
- Hochberg LR, Serruya MD, Friehs GM, Mukand JA, Saleh M, Caplan AH, Branner A, Chen D, Penn RD, Donoghue JP (2006) Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* 442:164–171.
- Holt GR, Koch C (1999) Electrical interactions via the extracellular potential near cell bodies. *J Comput Neurosci* 6:169–184.
- Hubel DH, Wiesel TN (1959) Receptive fields of single neurones in the cat's striate cortex. *J Physiol* 148:574–591.
- Hurtado JM, Rubchinsky LL, Sigvardt KA (2004) Statistical method for detection of phase-locking episodes in neural oscillations. *J Neurophysiol* 91:1883–1898.
- Ikegaya Y, Aaron G, Cossart R, Aronov D, Lampl I, Ferster D, Yuste R (2004) Synfire chains and cortical songs: temporal modules of cortical activity. *Science* 304:559–564.
- Imas OA, Ropella KM, Wood JD, Hudetz AG (2004) Halothane augments event-related gamma oscillations in rat visual cortex. *Neuroscience* 123:269–278.

- Joerges J, Küttner A, Galizia CG, Menzel R (1997) Representations of odours and odour mixtures visualized in the honeybee brain. *Nature* 387:285–288.
- Juergens E, Guettler A, Eckhorn R (1999) Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and eeg-potentials, but not in human eeg. *Exp Brain Res* 129:247–259.
- Kamondi A, Acsády L, Wang XJ, Buzsáki G (1998) Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. *Hippocampus* 8:244–261.
- Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL, Carandini M (2009) Local origin of field potentials in visual cortex. *Neuron* 61:35–41.
- Kayser C, Montemurro MA, Logothetis NK, Panzeri S (2009) Spike-phase coding boosts and stabilizes information carried by spatial and temporal spike patterns. *Neuron* 61:597–608.
- Kilavik BE, Roux S, Ponce-Alvarez A, Confais J, Grün S, Riehle A (2009) Modifications in motor cortical spiking dynamics induced by practice. *BMC Neuroscience* 10(Suppl 1):344.
- Klausberger T, Magill PJ, Márton LF, Roberts JDB, Cobden PM, Buzsáki G, Somogyi P (2003) Brain-state- and cell-type-specific firing of hippocampal interneurons in vivo. *Nature* 421:844–848.
- Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *J Neurosci* 25:3661–3673.
- Kopell N, Ermentrout GB, Whittington MA, Traub RD (2000) Gamma rhythms and beta rhythms have different synchronization properties. *Proc Natl Acad Sci U S A* 97:1867–1872.
- Krofczik S, Menzel R, Nawrot MP (2009) Rapid odor processing in the honeybee antennal lobe network. *Front Comput Neurosci* 2:9.
- Lachaux JP, Rodriguez E, Martinerie J, Varela FJ (1999) Measuring phase synchrony in brain signals. *Hum Brain Mapp* 8:194–208.

- Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320:110–113.
- Lampl I, Reichova I, Ferster D (1999) Synchronous membrane potential fluctuations in neurons of the cat visual cortex. *Neuron* 22:361–374.
- Laurent G (1996) Dynamical representation of odors by oscillating and evolving neural assemblies. *Trends Neurosci* 19:489–496.
- Laurent G, Stopfer M, Friedrich RW, Rabinovich MI, Volkovskii A, Abarbanel HD (2001) Odor encoding as an active, dynamical process: experiments, computation, and theory. *Annu Rev Neurosci* 24:263–297.
- Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci* 3:884–895.
- Lebedev MA, Wise SP (2000) Oscillations in the premotor cortex: single-unit activity from awake, behaving monkeys. *Exp Brain Res* 130:195–215.
- Legatt AD, Arezzo J, Vaughan HG (1980) Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: effects of volume-conducted potentials. *J Neurosci Methods* 2:203–217.
- Lindén H, Pettersen KH, Tetzlaff T, Potjans T, Denker M, Diesmann M, Grün S, Einevoll GT (2009) Estimating the spatial range of local field potentials in a cortical population model. *BMC Neuroscience* 10 (Suppl I):P224.
- Liu LF, Palmer AR, Wallace MN (2006) Phase-locked responses to pure tones in the inferior colliculus. *J Neurophysiol* 95:1926–1935.
- Logothetis NK (2003) The underpinnings of the bold functional magnetic resonance imaging signal. *J Neurosci* 23:3963–3971.
- Logothetis NK, Kayser C, Oeltermann A (2007) In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron* 55:809–823.
- Logothetis NK, Wandell BA (2004) Interpreting the bold signal. *Annu Rev Physiol* 66:735–769.

- Maldonado P, Babul C, Singer W, Rodriguez E, Berger D, Grün S (2008) Synchronization of neuronal responses in primary visual cortex of monkeys viewing natural images. *J Neurophysiol* 100:1523–1532.
- Mardia KV, Jupp PE (2000) *Directional Statistics* John Wiley & Sons Ltd, Chichester.
- Martin C, Gervais R, Messaoudi B, Ravel N (2006) Learning-induced oscillatory activities correlated to odour recognition: a network activity. *Eur J Neurosci* 23:1801–1810.
- Mazzoni A, Panzeri S, Logothetis NK, Brunel N (2008) Encoding of naturalistic stimuli by local field potential spectra in networks of excitatory and inhibitory neurons. *PLoS Comput Biol* 4:e1000239.
- Mehring C, Rickert J, Vaadia E, de Oliveira SC, Aertsen A, Rotter S (2003) Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci* 6:1253–1254.
- Mitzdorf U (1985) Current source-density method and application in cat cerebral cortex: investigation of evoked potentials and eeg phenomena. *Physiol Rev* 65:37–100.
- Mitzdorf U (1987) Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. *Int J Neurosci* 33:33–59.
- Monosov IE, Trageser JC, Thompson KG (2008) Measurements of simultaneously recorded spiking activity and local field potentials suggest that spatial selection emerges in the frontal eye field. *Neuron* 57:614–625.
- Montemurro MA, Rasch MJ, Murayama Y, Logothetis NK, Panzeri S (2008) Phase-of-firing coding of natural visual stimuli in primary visual cortex. *Curr Biol* 18:375–380.
- Moran DW, Schwartz AB (1999) Motor cortical representation of speed and direction during reaching. *J Neurophysiol* 82:2676–2692.
- Moritz CT, Perlmutter SI, Fetz EE (2008) Direct control of paralysed muscles by cortical neurons. *Nature* 456:639–642.

- Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R (2005) Coupling between neuronal firing, field potentials, and fmri in human auditory cortex. *Science* 309:951–954.
- Murthy VN, Fetz EE (1992) Coherent 25- to 35-hz oscillations in the sensorimotor cortex of awake behaving monkeys. *Proc Natl Acad Sci U S A* 89:5670–5674.
- Murthy VN, Fetz EE (1996a) Oscillatory activity in sensorimotor cortex of awake monkeys: synchronization of local field potentials and relation to behavior. *J Neurophysiol* 76:3949–3967.
- Murthy VN, Fetz EE (1996b) Synchronization of neurons during local field potential oscillations in sensorimotor cortex of awake monkeys. *J Neurophysiol* 76:3968–3982.
- Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nat Neurosci* 12:70–76.
- Nicolelis MA, Ghazanfar AA, Faggin BM, Votaw S, Oliveira LM (1997) Reconstructing the engram: simultaneous, multisite, many single neuron recordings. *Neuron* 18:529–537.
- Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A, Fried I, Malach R (2007) Coupling between neuronal firing rate, gamma lfp, and bold fmri is related to interneuronal correlations. *Curr Biol* 17:1275–1285.
- Nunez PL, Shrinivasan R (2006) *Electric Fields of the Brain. The Neurophysics of EEG* Oxford University Press.
- O’Leary JG, Hatsopoulos NG (2006) Early visuomotor representations revealed from evoked local field potentials in motor and premotor cortical areas. *J Neurophysiol* 96:1492–1506.
- Pettersen KH, Devor A, Ulbert I, Dale AM, Einevoll GT (2006) Current-source density estimation based on inversion of electrostatic forward solution: effects of finite extent of neuronal activity and conductivity discontinuities. *J Neurosci Methods* 154:116–133.
- Pettersen KH, Hagen E, Einevoll GT (2008) Estimation of population firing rates and current source densities from laminar electrode recordings. *J Comput Neurosci* 24:291–313.

- Pikovski A, Kurths J, Rosenblum M (2001) *Synchronization: A Universal Concept in Nonlinear Sciences* Cambridge Univ Press.
- Poulet JFA, Petersen CCH (2008) Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. *Nature* 454:881–885.
- Quiroga RQ, Kraskov A, Kreuz T, Grassberger P (2002) Performance of different synchronization measures in real data: a case study on electroencephalographic signals. *Phys Rev E Stat Nonlin Soft Matter Phys* 65:041903.
- Quyen MLV, Foucher J, Lachaux J, Rodriguez E, Lutz A, Martinerie J, Varela FJ (2001) Comparison of hilbert transform and wavelet methods for the analysis of neuronal synchrony. *J Neurosci Methods* 111:83–98.
- Quyen MLV, Bragin A (2007) Analysis of dynamic brain oscillations: methodological advances. *Trends Neurosci* 30:365–373.
- Radman T, Su Y, An JH, Parra LC, Bikson M (2007) Spike timing amplifies the effect of electric fields on neurons: implications for endogenous field effects. *J Neurosci* 27:3030–3036.
- Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK (2008) Inferring spike trains from local field potentials. *J Neurophysiol* 99:1461–1476.
- Ravel N, Chabaud P, Martin C, Gaveau V, Hugues E, Tallon-Baudry C, Bertrand O, Gervais R (2003) Olfactory learning modifies the expression of odour-induced oscillatory responses in the gamma (60–90 hz) and beta (15–40 hz) bands in the rat olfactory bulb. *Eur J Neurosci* 17:350–358.
- Ray S, Hsiao SS, Crone NE, Franaszczuk PJ, Niebur E (2008) Effect of stimulus intensity on the spike-local field potential relationship in the secondary somatosensory cortex. *J Neurosci* 28:7334–7343.
- Riehle A, Grün S, Diesmann M, Aertsen A (1997) Spike synchronization and rate modulation differentially involved in motor cortical function. *Science* 278:1950–1953.
- Rodriguez E, George N, Lachaux JP, Martinerie J, Renault B, Varela FJ (1999) Perception's shadow: long-distance synchronization of human brain activity. *Nature* 397:430–433.

- Rosenberg JR, Amjad AM, Breeze P, Brillinger DR, Halliday DM (1989) The fourier approach to the identification of functional coupling between neuronal spike trains. *Prog Biophys Mol Biol* 53:1–31.
- Rosenblum M, Pikovsky A, Kurths J (1996) Phase synchronization of chaotic oscillators. *Phys Rev Lett* 76:1804–1807.
- Rotermund D, Taylor K, Ernst UA, Kreiter AK, Pawelzik KR (2009) Attention improves object representation in visual cortical field potentials. *J Neurosci.* 29:10120–10130.
- Roux S, Coulmance M, Riehle A (2003) Context-related representation of timing processes in monkey motor cortex. *Eur J Neurosci* 18:1011–1016.
- Roux S, Mackay WA, Riehle A (2006) The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behav Brain Res* 169:335–351.
- Rubino D, Robbins KA, Hatsopoulos NG (2006) Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci* 9:1549–1557.
- Rymar VV, Sasseville R, Luk KC, Sadikot AF (2004) Neurogenesis and stereological morphometry of calretinin-immunoreactive gabaergic interneurons of the neostriatum. *J Comp Neurol* 469:325–339.
- Sachse S, Galizia CG (2002) Role of inhibition for temporal and spatial odor representation in olfactory output neurons: a calcium imaging study. *J Neurophysiol* 87:1106–1117.
- Sandoz JC, Galizia CG, Menzel R (2003) Side-specific olfactory conditioning leads to more specific odor representation between sides but not within sides in the honeybee antennal lobes. *Neuroscience* 120:1137–1148.
- Sanes JN, Donoghue JP (1993) Oscillations in local field potentials of the primate motor cortex during voluntary movement. *Proc Natl Acad Sci U S A* 90:4470–4474.
- Schneidman E, Berry MJ, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012.

- Schrader S, Grün S, Diesmann M, Gerstein GL (2008) Detecting synfire chain activity using massively parallel spike train recording. *J Neurophysiol* 100:2165–2176.
- Schäfer C, Rosenblum MG, Abel HH, Kurths J (1999) Synchronization in the human cardiorespiratory system. *Phys Rev E Stat Phys Plasmas Fluids Relat Interdiscip Topics* 60:857–870.
- Shinomoto S, Kim H, Shimokawa T, Matsuno N, Funahashi S, Shima K, Fujita I, Tamura H, Doi T, Kawano K, Inaba N, Fukushima K, Kurkin S, Kurata K, Taira M, Tsutsui KI, Komatsu H, Ogawa T, Koida K, Tanji J, Toyama K (2009) Relating neuronal firing patterns to functional differentiation of cerebral cortex. *PLoS Comput Biol* 5:e1000433.
- Shlens J, Field GD, Gauthier JL, Greschner M, Sher A, Litke AM, Chichilnisky EJ (2009) The structure of large-scale synchronized firing in primate retina. *J Neurosci* 29:5022–5031.
- Shlens J, Field GD, Gauthier JL, Grivich MI, Petrusca D, Sher A, Litke AM, Chichilnisky EJ (2006) The structure of multi-neuron firing patterns in primate retina. *J Neurosci* 26:8254–8266.
- Singer W (1999) Neuronal synchrony: a versatile code for the definition of relations? *Neuron* 24:49–65.
- Singer W, Pipa G, Uhlhaas P, Lima B, Melloni L, Neuenschwander S, Nikolic D (2009) Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience* 3:17.
- Somogyvári Z, Zolányi L, Ulbert I, Erdi P (2005) Model-based source localization of extracellular action potentials. *J Neurosci Methods* 147:126–137.
- Stark E, Abeles M (2005) Applying resampling methods to neurophysiological data. *J Neurosci Methods* 145:133–144.
- Tetzlaff T, Rotter S, Stark E, Abeles M, Aertsen A, Diesmann M (2008) Dependence of neuronal correlations on filter characteristics and marginal spike train statistics. *Neural Comput* 20:2133–2184.
- Thiel M, Romano C, Kurths J, Meucci R, Allaria E, Arecchi FT (2002) Influence of observational noise on recurrence quantification analysis. *Physica D* 171:138–152.

- Tort ABL, Kramer MA, Thorn C, Gibson DJ, Kubota Y, Graybiel AM, Kopell NJ (2008) Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a t-maze task. *Proc Natl Acad Sci U S A* 105:20517–20522.
- Vaadia E, Haalman I, Abeles M, Bergman H, Prut Y, Slovin H, Aertsen A (1995) Dynamics of neuronal interactions in monkey cortex in relation to behavioural events. *Nature* 373:515–518.
- van Hemmen JL, Schwartz AB (2008) Population vector code: a geometric universal as actuator. *Biol Cybern* 98:509–518.
- Varela F, Lachaux JP, Rodriguez E, Martinerie J (2001) The brain-web: Phase synchronization and large-scale integration. *Nat Rev Neurosci* 2:229–239.
- Viswanathan A, Freeman RD (2007) Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nat Neurosci* 10:1308–1312.
- von der Malsburg C, Schneider W (1986) A neural cocktail-party processor. *Biol Cybern* 54:29–40.
- Wessberg J, Stambaugh CR, Kralik JD, Beck PD, Laubach M, Chapin JK, Kim J, Biggs SJ, Srinivasan MA, Nicolelis MA (2000) Real-time prediction of hand trajectory by ensembles of cortical neurons in primates. *Nature* 408:361–365.
- Witham CL, Baker SN (2007) Network oscillations and intrinsic spiking rhythmicity do not covary in monkey sensorimotor areas. *J Physiol* 580:801–814.
- Womelsdorf T, Schoffelen JM, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609–1612.
- Yu D, Ponomarev A, Davis RL (2004) Altered representation of the spatial code for odors after olfactory classical conditioning; memory trace formation by synaptic recruitment. *Neuron* 42:437–449.
- Yu S, Huang D, Singer W, Nikolic D (2008) A small world of neuronal synchrony. *Cereb Cortex* 18:2891–2901.