Single-Neuron and Network Dynamics in the Neocortical-Entorhinal-Hippocampal Circuit

by

James McFarland

B.A., Pomona College, 2006

A dissertation submitted in partial fulfillment of the requirements for the Degree of Doctor of Philosophy in the Department of Physics at Brown University

Providence, Rhode Island

May 2011

This dissertation by James McFarland is accepted in its present form by the Department of Physics as satisfying the dissertation requirement for the degree of Doctor of Philosophy.

Date _____ _____
Prof. Mayank Mehta, Advisor

Recommended to the Graduate Council

Date _____ _____
Prof. Xinsheng Sean Ling, Reader

Date _____ _____
Prof. Derek Stein, Reader

Approved by the Graduate Council

Date _____ _____
Peter M. Weber, Dean of the Graduate School

Acknowledgements

First I would like to thank my advisor and mentor, Mayank Mehta. Mayank has shown incredible patience with me since offering me the opportunity to work with him at a point when I had no experience in neuroscience and no idea what questions I wanted to ask or how. Through the years his encouragement and guidance have helped me to find my niche in research. I thank Barry Connors and Brad Marston for their generous support, Derek Stein for graciously agreeing to be my 'ghost advisor,' and Sean Ling whose service on my thesis committee is greatly appreciated. I would also like to thank my undergraduate mentor David Tanenbaum for introducing me to research.

I have been fortunate to have the opportunity to work with talented collaborators in Thomas Hahn and Evgeny Resnik. I thank my fellow graduate students and post-docs in the Mehta lab that I have had the pleasure of working with over the years: Omar Ahmed, Arvind Kumar, Lucas Santos, David Ho, David Freestone, and Zhiping Chen. I would especially like to thank Omar and Arvind for teaching me much of what I know about electrophysiology and data analysis. I also was fortunate to work with a remarkable group of undergraduates, especially Sam Reiter, Lyle Muller, Vivek Buch, Valerie Yanofsky, and Annie van Bueningen. In addition, I owe a great deal of thanks to Shawna Hollen for her friendship and support over the years.

I especially want to thank my family: my parents for their unwavering support and encouragement, and my brother John for showing me that getting a PhD doesn't have to be a miserable experience. Lastly, my deepest thanks are owed to Anne Booker for not only keeping me sane, but also making me happy.

Contents

List of Tables

List of Figures

Chapter 1: **Introduction**

Electrical Signals in the Brain

Types of Neural Data

Data within neuroscience can come in many forms, describing the structure, development, and function of the nervous system, as well as its behavioral output. Even measurements of neural activity span a range of data types including optical imaging, functional magnetic resonance imaging (fMRI), recordings of electrical activity, and measurements of neurotransmitter, neuromodulator, and ionic concentrations. The focus of this thesis will be the analysis of electrical signals in the brain, of which there are several different kinds. Macroscopic fields present at the scalp can be recorded non-invasively using scalp electrodes or magnetometers, producing the electro-encephalogram (EEG) or magneto-encephalogram (MEG) respectively. These methods are restricted to detecting the synchronous activity of large numbers of neurons, and hence are limited in their signal detection and spatial resolution capabilities. In order to record signals from local patches of neural tissue that cannot be detected with the above methods, electrodes can be inserted into the extracellular medium directly to provide the extracellular field potential (EFP). Finally, several methods exist for measuring the electrical activity of individual neurons, a signal known as the membrane potential (MP). Because the EFP (and its signal derivatives) and the MP are the signals used in the thesis, they are discussed in more detail below.

Membrane Potential

The membrane potential is defined as the electrical potential difference across the cell membrane of a neuron. As a result of ionic gradients across the cell membrane maintained by ion channels and active ion pumps, the MP of a neuron typically has a value on the order of -70 to -80mV in the resting or baseline state [1]. Fluctuations in the MP are then produced by ionic currents flowing into and out of the neuron through a wide variety of ion channels. The conductance of these ion channels can be sensitive to chemical ligands, ionic concentrations, or the MP itself (voltage-gated ion channels). Of particular importance are chemical receptors located largely in the dendrites, where chemical signals from other neurons are received at specialized structures called synapses. These receptors can control the opening and closing of ion channels both directly and indirectly, and convey the incoming chemical signals into changes in the neuron's MP. When the MP at the axon initial segment exceeds a threshold of about -55mV, a short-lasting (~1ms) high-amplitude (~100mV) electrical impulse known as an action potential is produced [1]. Action potentials are mediated by a set of specialized voltage-gated ion channels, and are actively propagated into the cell's processes, particularly down the axon where the action potential triggers the release of chemical transmitters at synapses with other neurons. Because the action potentials are seen as the output of a neuron they are often of particular interest.

Several different experimental techniques for measuring a cell's MP exist. Further, the MP can either be measured directly ("current clamp recording"), or the experimenter can maintain the cell's MP at a desired level through the use of a feedback circuit in order to measure membrane currents ("voltage clamp recording") [2]. The first method introduced for recording a cell's MP

was to insert the tip of a sharp microelectrode into the cell directly [3]. The primary disadvantage of this approach, however, is that it requires the experimenter to pierce the cell membrane and thus compromises the health of the cell. More recently, glass micropipettes have been used to create a high-resistance seal with a patch of the cell membrane [4]. After achieving such a seal, the experimenter can record currents through individual ion channels located on that patch of membrane ('cell-attached recording'). Alternatively, suction on the electrode can be used to rupture the patch of membrane and allow direct access to the intracellular medium without leaking the cell's contents into the extracellular space ('whole-cell recording'). In this case, care must be used in preparing the solution used to fill the pipette since this solution will mix with the intracellular material.

While in all of the above approaches the recording electrode is typically placed at or inside of a neuron's cell body, or soma, it is important to recognize that the MP varies both in time and as a function of position within the cell's often complex and expansive structure [1]. MP recordings from a cell's dendrites or axon are also possible, although they are typically substantially more difficult to obtain [5]. The passive propagation of dendritically generated signals to the soma is described by cable theory, and involves substantial filtering and attenuation. Thus, these additional factors, along with potential for altering a neuron's properties by the recording process itself, must be considered when interpreting MP signals.

Local Field Potential

The electric potential difference between a point in extracellular space and a remote ground electrode (the EFP) is given by a weighted sum of spatially distinct current sources and sinks created by current flowing across nearby cell membranes, with the weighting being determined

by the geometry of the extracellular space [6]. The spatial distribution of current flow across a given neuron will itself be determined by the geometry of the neuron as well as the distribution of ion channels. Thus, the orientation, arrangement, and packing of neurons can substantially alter how currents will summate to form the EFP. For instance, current flow along the oriented somato-dendritic axes of cortical pyramidal neurons is thought to provide a large contribution to the cortical EFP and EEG signals [6,7].

The EFP signal is typically divided into two distinct components based on its frequency content. By high-pass filtering the EFP above about 300-500Hz one obtains a signal reflecting the spiking activity of nearby neurons (see *Extracellular Spikes*). In contrast, by low-pass filtering the EFP below about 300Hz one obtains the so-called local field potential (LFP) [8]. It has long been thought that the LFP is reflective of the synaptic inputs received by neurons in the vicinity of the recording electrode rather than the spiking output of those neurons [6,9,10]. Later studies, however, have suggested that other factors, particularly intrinsic membrane potential oscillations [11], and spike afterpotentials [12,13] can contribute to the LFP signal as well. Direct interpretation of the LFP in terms of underlying electrophysiological processes is thus extremely difficult. Even if the LFP is taken to reflect only the synaptic inputs to the surrounding neurons, these inputs could arise from recurrent local connections, or from distant brain regions. The spatial scale over which neural activity is thought to contribute to the LFP is on the order of 100-200µm [14,15], although other studies have estimated that activity within several mm of the electrode can contribute to the LFP [16].

Even given the difficulty in their interpretation, LFP signals have proven to be useful tools in a number of settings. One important advantage is that extracellular signals are substantially

easier to acquire than recordings from individual neurons, and can be scaled to allow recordings from a large number of electrodes in parallel [17]. Such large-scale extracellular recordings can naturally provide information about local spiking output of the recorded areas as well (see *Extracellular Spikes*). In many respects the LFP is similar to the EEG signal, but with vastly superior spatial resolution and sensitivity [6]. The LFP has also been shown to correspond somewhat closely to the BOLD fMRI signal [8,18], but with much higher temporal resolution. The information provided by the LFP is often complimentary to that provided by spiking activity, being determined primarily by synchronous local synaptic activity. In fact, the information carried by LFP signals about visual stimuli [16], as well as movement planning and execution [19,20,21], has been found to be largely independent of that carried by spiking activity [22]. Such information could be particularly useful for brain-machine interface applications.

Extracellular Spikes

Since the currents generated during action potentials have a faster time course than other transmembrane currents, spiking activity can be identified by high-pass filtering (>300-500Hz) the EFP. Spikes can be detected extracellularly from neurons whose somas are up to ~150µm from the recording electrode [23], which means that with typical cell densities spiking activity of a large number of neurons can be detected (e.g. ~1000 in the hippocampal CA1 region [23]). In addition to knowing the set of spike times for the group of nearby neurons (a signal referred to as 'multi-unit activity'), it is highly advantageous to know the spike times of each constituent neuron independently. The set of spike times for a single neuron is referred to as 'single-unit activity'. Because the waveforms of extracellular action potentials depend on a number of factors including the type of neuron, the neuron's spiking history, and in particular the geometry

of the neuron-electrode configuration, different neurons will produce characteristic extracellular spike waveforms. This fact can sometimes be leveraged to identify the spikes generated by each neuron [24], a procedure known as spike-sorting. Accurate spike-sorting requires that the extracellular spike waveforms have amplitudes greater than 50μV-100μV [23,25], which restricts the set of neurons for which single-unit activity can be inferred to those within ~50μm of the electrode tip [23]. While spike-sorting can be accomplished with single-electrode recordings, it has become common to use multiple electrodes bundled together for this purpose. A group of four electrodes, known as a tetrode [26], has been shown to produce good results. The geometric relation of each neuron to the set of electrode tips results in a characteristic waveform being recorded on the set of electrodes , and allows more accurate classification of single-unit activity [25,27].

Analysis of Neural Data

Time Series Analysis

Electrical signals from the nervous system are represented mathematically as a sequence of data points ordered in time, known as a time series. Because of this, methods for analyzing time series are an invaluable tool in neuroscience, and are used extensively throughout this thesis. While dynamical systems approaches have also been employed to understand neural dynamics[28], such methods are not utilized here. Instead, electrical signals from the brain are typically modeled as stochastic, or random, processes because of the presence of noise in neural recordings and due to intrinsic randomness of neural activity. Such a probabilistic framework also offers numerous approaches for quantifying the structure of neural signals.

Because electrical signals are sampled in discrete time, a recorded signal x_t is treated as a sequence of random variables, governed by a joint probability distribution. As with other probability distributions, an effective approach to characterizing the distribution is to estimate its moments. The first moment of a time series is the mean function:

$$E[x_t] = \mu_t \tag{1-1}$$

which describes how the average value of the signal changes in time. The second (central) moment, known as the autocovariance function, determines how samples at different time points *t* and *s* covary:

$$E\left[\left(x_t - \mu_t\right)\left(x_s - \mu_s\right)\right] = K_{xx}(t,s) \tag{1-2}$$

Since estimating even these first two moments of the joint distribution requires estimating their values at each time point (or pair of time points), this is typically not feasible with neural data where only a single instance of the stochastic process is available. Thus, the assumption of stationarity is frequently made in order to simplify these calculations.

A strictly stationary process is one in which the distribution governing a sequence of samples is invariant to translations in time. The weaker requirement of wide-sense stationarity implies only that the first two moments of the distribution (the mean function and auto-covariance function) are time-translation invariant. Thus, for a wide-sense stationary process, equations (1-1) and (1-2) reduce to:

$$E[x_t] = \mu \tag{1-3}$$

$$E\left[\left(x_t - \mu\right)\left(x_s - \mu\right)\right] = K_{xx}(t,s) = K_{xx}(s-t) = K_{xx}(\tau) \tag{1-4}$$

where the mean function becomes a constant (which we will assume to be 0), and the autocovariance reduces to a function of the difference between the two time points. A related quantity known as the autocorrelation function is given by normalizing the autocovariance function by the signal variance:

$$\rho(\tau) = \frac{K_{xx}(\tau)}{K_{xx}(0)} = \frac{K_{xx}(\tau)}{\sigma^2} \quad \textbf{(1-5)}$$

Time series can be equivalently represented in the frequency domain by application of the discrete Fourier transform:

$$\tilde{X}(f) = \sum_t x_t \exp\left(-2\pi i f t\right) \quad \textbf{(1-6)}$$

where $\tilde{X}(f)$ is a complex number. The second moment of the process in the frequency domain is given by the power spectrum $S_{xx}(f)$:

$$S_{xx}(f)\delta\left(f - f'\right) = E\left[\tilde{X}^*(f)\tilde{X}(f')\right] \quad \textbf{(1-7)}$$

where the asterisk denotes complex conjugation, and the delta function indicates that $\tilde{X}(f)$ and $\tilde{X}(f')$ $f \neq f'$ are uncorrelated (and zero-mean) for a stationary process [29]. This orthogonality property of the spectral representation is particularly useful in the context of neural signal analysis because different electrophysiological processes which are inseparable in the time domain are often separable in the frequency domain.

It is worth noting that for a signal sampled in discrete time, the maximum frequency that can be resolved is called the Nyquist frequency and is equal to $f_s/2$, where f_s is the sampling frequency. Similarly, the Nyquist-Shannon sampling theorem states that a discretely sampled signal can be

perfectly reconstructed if the sampling rate is greater than or equal to 2*B*, where *B* is the highest frequency component of the original signal [30]. Thus, discrete-time sampling is only appropriate for so-called bandlimited signals.

The direct estimator of the power spectrum (known as the periodogram), obtained by simply squaring the amplitude of the Fourier transform of the signal, suffers from formidable problems of both bias and variance [31,32], and similar problems exist for its time domain equivalent (direct estimate of the autocovariance function). In fact, the Wiener-Khinchin theorem states that the autocovariance function and power spectrum are simply Fourier transform pairs [29]. For signals of finite duration, the periodogram estimate of the power at a given frequency is biased by the power at other frequencies (spectral leakage). Further, even for signals of infinite duration, the periodogram estimator is not consistent (e.g. the estimate does not converge as the time series duration goes to infinity) [31,32]. Both problems can be addressed by first multiplying the data by a window or taper function w_t, at the expense of a reduction in frequency resolution. A powerful approach is to apply multiple orthogonal taper functions to the data. The so-called multi-taper spectral estimator is then given by averaging over the individual tapered estimates [31]:

$$S_{xx}^{MT}(f) = \frac{1}{K}\sum_{k=1}^{K}\left(\left|\sum_{t=1}^{T} w_t^k x_t \exp\left(-2\pi i f t\right)\right|^2\right) \qquad \textbf{(1-8)}$$

where *K* is the number of tapers, and w_t^k is the k^{th} taper function. The discrete prolate spheroidal sequences (or Slepian sequences) provide an ideal choice for a set of orthogonal tapering functions, having minimal spectral leakage for a given frequency resolution [33]. In

addition to the reductions in bias and variance (by effectively averaging estimates of the spectrum at nearby frequencies), the multi-taper method also provides a natural means of obtaining uncertainty estimates because it produces *K* independent estimates of the spectrum [31,32].

The methods described above represent the so-called nonparametric approach to spectral analysis. An important alternative is the use of parametric models, particularly autoregressive (AR) models in which the value of the time series at a given time point is represented as a linear combination of its previous values plus Gaussian white noise. An equivalent perspective is to view the model as a filter that generates the observed data when driven by a Gaussian white noise input. The frequency response of the filter, which can be derived from the estimated AR model parameters, then gives the AR estimate of the power spectrum [34]. In addition to parametric spectral analysis, higher order moments of the spectral representation can be estimated such as with the bispectrum. These methods provide information about the interactions between components of the signal at different frequencies, and are useful for characterizing non-linearities in the signal [35].

Multivariate Time Series

When dealing with multivariate time series, such as with multiple simultaneously recorded electrical signals, the second moment is described by a matrix at each time lag (assuming wide-sense stationarity). If the *k*-dimensional vector-valued time series is: $\mathbf{z}_t = (z_t^1, z_t^2, ..., z_t^k)$, its cross-covariance function is given by:

$$E\left[\left(z_t^i-\mu^i\right)\left(z_t^j-\mu^j\right)\right]=K_{ij}(\tau) \tag{1-9}$$

where the vector $\boldsymbol{\mu}=(\mu^1,\mu^2,...,\mu^k)$ gives the mean of each component time series. It's normalized equivalent is the cross-correlation function:

$$\rho_{ij}(\tau)=\frac{K_{ij}(\tau)}{\sqrt{K_{ii}(0)K_{jj}(0)}}=\frac{K_{ij}(\tau)}{\sqrt{\sigma_i^2\sigma_j^2}} \tag{1-10}$$

As in the univariate case, the second moment of a multivariate time series can also be characterized in the frequency domain as [32]:

$$E\left[\tilde{Z}^{i*}(f)\tilde{Z}^{j}(f')\right]=S_{ij}(f)\delta(f-f') \tag{1-11}$$

The quantity $S_{ij}(f)$ is known as the cross-spectrum, and it measures the covariation of power in components z^i and z^j at frequency f. Typically, the normalized version of Equation (1-11) is used instead. This quantity, known as the coherence spectrum, is bounded between 0 and 1, and describes the fraction of power in signal z^i at frequency f that can be predicted from $S_{jj}(f)$, and is given by:

$$C_{ij}(f)=\frac{S_{ij}(f)}{\sqrt{S_{ii}(f)S_{jj}(f)}} \tag{1-12}$$

In practice, the cross-spectrum (and coherence spectrum) must be estimated from multiple instances of the signals z^i and z^j, or equivalently by dividing the data into non-overlapping segments and averaging the cross-spectrum estimates across segments. Alternatively, the coherence spectrum can be estimated from a single instance of the time series z^i and z^j using multi-taper methods [31,32]:

$$C_{ij}^{MT}(f) = \frac{\sum_{k=1}^{K}\left(\left(\sum_{t=1}^{T} w_t^k z_t^i \exp(-2\pi i f t)\right)\left(\sum_{t=1}^{T} w_t^k z_t^j \exp(-2\pi i f t)\right)\right)}{\sqrt{\sum_{k=1}^{K}\left(\left|\sum_{t=1}^{T} w_t^k z_t^i \exp(-2\pi i f t)\right|^2\right)\sum_{k=1}^{K}\left(\left|\sum_{t=1}^{T} w_t^k z_t^j \exp(-2\pi i f t)\right|^2\right)}} \quad \textbf{(1-13)}$$

such that the cross-spectrum is estimated by averaging across the *K* data tapers.

As with the univariate case, these nonparametric methods represent only one approach to characterizing the relationships between multiple time series, and numerous other methods have been employed in the analysis of multivariate neural data. Parametric methods (multivariate autoregressive models, or MAR models) can also be used to estimate the coherence spectrum between two signals. Importantly, these methods also allow for inference of causal, rather than merely correlational, relationships between time series [36]. More recently, information theoretic methods have also been developed to quantify non-linear as well as causal interactions between time series [37]. Higher moments of the multivariate time series can also be used to quantify interactions between the component time series at different frequencies. The nonparametric time- and frequency-domain analysis methods are by far the most widely used in neural data analysis, however, because of their robustness, interpretability, and well-understood statistical properties.

Time-Frequency Analysis

In the analysis described above, the signal(s) in question were assumed to satisfy the requirement of stationarity. However, neural signals are rarely stationary over substantial periods of time, and one is often interested in analyzing how the structures of neural signals evolve over time. Thus, one would like to be able to characterize the signal variance as a joint

function of time and frequency (the 'spectrogram') [32,38]. The most straightforward approach, known as the short-time Fourier transform (STFT), is to divide the signal into segments of length *L* samples and compute the power spectrum separately within each segment. Such a procedure produces a description of the signal variance in the time-frequency 'plane'. In order to achieve a higher temporal resolution, one can reduce the segment length *L*, however this results in a corresponding decrease in frequency resolution. To see this, we consider that a length *L* segment of data centered at time *t* can be obtained by multiplying the signal x_t with a rectangular window function $w(t'-t)$ centered at time *t*, and the resulting Fourier transform is given by:

$$\Im\{w(t'-t)x_t\} = \Im\{w(t'-t)\} * \Im\{x_t\} \quad \textbf{(1-14)}$$

where $w(\tau)$ is the rectangle function:

$$w(\tau) = \begin{cases} 1, & -L/2 <= \tau <= L/2 \\ 0, & \text{else} \end{cases} \quad \textbf{(1-15)}$$

$\Im\{\ \}$ denotes the Fourier transform, and the right hand side of Equation (1-14) follows from the convolution theorem. This means that the resultant power spectrum of the windowed signal will be determined by the convolution of the Fourier transforms of the signal and the window function. Indeed, this is the basis for the biases of the direct periodogram spectral estimator, since any finite duration signal can be represented by multiplication with an appropriate rectangular window function. This implies a tradeoff between the width of the window function and the resolution of the resulting spectral estimate, because a function and its Fourier transform cannot be made arbitrarily narrow. One approach is to select a window length *L*, and use the multi-taper spectral estimator as in Equation (1-8), successively translating the *K*

orthogonal tapers of length *L* to be centered on a desired series of time points at which the spectrogram will be evaluated [32]. For the set of Slepian tapers, there will be 2*LW*-1 Slepian functions well-localized in the interval [-W,W] [31,32,33], where *W* is referred to as the bandwidth.

An alternative approach which often achieves improved time-frequency resolution compared to STFT methods is to convolve the signal x_t with a series of wavelet functions which are scaled and translated versions of a 'mother wavelet' function [38,39]. The continuous wavelet transform is defined as:

$$\tilde{W}_t(s) = \sum_{t'=1}^{T} x_{t'} \tilde{\psi} * \left[\frac{(t'-t)\delta t}{s} \right] \quad \textbf{(1-16)}$$

where $\tilde{\psi}_0[\eta]$ is the mother wavelet function. A commonly used choice of mother wavelet is a (properly normalized) Gaussian-modulated plane wave ('Morlet' wavelet):

$$\tilde{\psi}(\eta) = \sqrt{\frac{\delta t}{s\sqrt{\pi}}} \exp\left(i\omega_0 \eta\right) \exp\left(\frac{-\eta^2}{2}\right) \quad \textbf{(1-17)}$$

The power of the signal x_t is then given as a function of time and scale by the squared amplitude of the wavelet transform (Equation (1-16)). Similarly, one can define the wavelet coherence spectrum between two signals x_t and y_t as [39,40]:

$$R_t^2(s) = \frac{\left| S\left(\tilde{W}_t^{xy}(s)\right) \right|^2}{S\left(\left|\tilde{W}_t^{x}(s)\right|^2\right) \cdot S\left(\left|\tilde{W}_t^{y}(s)\right|^2\right)} \quad \textbf{(1-18)}$$

where *S*() is a smoothing operator in both time and scale. While the wavelets at each time are defined by a scale parameter rather than a frequency, it is possible to define an equivalent Fourier frequency at each scale [39]:

$$\hat{f}(s) = \frac{\omega_0 + \sqrt{2+\omega_0^2}}{4\pi s} \quad \textbf{(1-19)}$$

The advantage of wavelet analysis is that the effective 'window width' is chosen optimally at each frequency, rather than using fixed window width across frequencies as with the STFT.

Both the Fourier transform and wavelet transform (when using a complex-valued mother wavelet) produce complex-valued outputs, and thus the time-frequency representation in both cases can be written in the form:

$$S(f,t) = \left|S(f,t)\right| \mathrm{e}^{i\phi(f,t)} \quad \textbf{(1-20)}$$

where $\left|S(f,t)\right|$ is the signal amplitude and $\phi(f,t)$ is the phase as a function of time and frequency. A related and frequently useful concept is that of the analytic representation $x_a(t)$ of a real-valued signal $x(t)$:

$$x_a(t) = x(t) + i\hat{x}(t) \quad \textbf{(1-21)}$$

The complex part $\hat{x}(t)$ is given by the Hilbert transform of $x(t)$:

$$\hat{x}(t) = \frac{1}{\pi} PV \int_{-\infty}^{\infty} \frac{x(t)}{t-\tau} d\tau \quad \textbf{(1-22)}$$

where PV denotes that the principal value of the integral is taken. The signal amplitude and phase are then computed as $\left|x_a(t)\right|$, and $\arg\left(x_a(t)\right)$ respectively [30]. These quantities (and

their discrete-time analogues) only have a clear physical interpretation for narrowband signals. Thus, a bandpass filter with center frequency *f* is typically first applied to produce $x_f(t)$. The analytic representation then provides the amplitude and phase of the component signal at frequency *f* as in Equation (1-20). The 'instantaneous frequency' can also be computed as:

$$\omega_f(t) = \frac{d\phi_f(t)}{dt} \quad \textbf{(1-23)}$$

Numerous other approaches to time-frequency analysis have been proposed, including time-varying AR models, the Wigner-Ville distribution, empirical mode decomposition, and the matching pursuit algorithm, although none have yet been extensively used in the analysis of neural signals.

Point Process Analysis

We are often interested in analyzing the timing of action potentials rather than continuous signals. A set of spike times ***s*** is treated mathematically as a sample from a so-called point process. Several equivalent representations can be used to describe a point process [41]. Defining N_t to be the counting process, such that N_t is the number of spikes that have occurred up to time *t*, one approach is to represent ***s*** as the time derivative of N_t (dN_t) given by a series of Dirac delta functions at each spike time. Equivalently, it is possible to represent the point process by the set of inter-spike intervals. Interestingly, much of the analysis presented above can also be applied to the counting process representation dN_t [32,41,42]. For example, the power spectrum of a point process can be computed as:

$$S_{dNdN}(f) = \left| \sum_{t}^{T} \exp(-2\pi i f t) \right|^2 \quad \textbf{(1-24)}$$

The extension to multi-taper methods, and STFT analysis is also straightforward. Furthermore, coherence spectra between pairs of point processes, or continuous and point processes, can be computed directly [32].

Neural Oscillations

Different types of oscillations

Ever since the first recordings of human scalp potentials (EEG) [43], oscillations have been a salient feature of electrical signals in the brain, and a subject of great interest to neuroscientists. A wide range of oscillation frequencies have been observed in electrical signals of the mammalian forebrain, ranging from 'infraslow' oscillations at frequencies between 0.02 and 0.2Hz [44] to 'fast ripple' oscillations at frequencies from 100-500Hz [45]. Neuronal oscillations have been historically classified into frequency bands which are presumed to represent unique neural processes and network states [46,47]. These include the infraslow (0.02-0.2Hz), slow (0.2-1Hz), delta (1-4Hz), theta (4-10Hz), alpha (8-12Hz), beta (10-30Hz), gamma (30-80Hz), fast (80-200Hz) and ultra-fast (200-600Hz) oscillations. These different oscillatory patterns have long been known to reflect different underlying brain and behavioral states [43]. For instance, the theta, alpha, beta, and gamma oscillations are closely related to different states of alertness and attention [48,49,50,51,52], as well as whether the eyes are opened or closed [43]. The spectral content of the EEG signal is also used to classify different stages of sleep [53].

Oscillation generating mechanisms

Oscillations at the macroscopic (EEG) and mesoscopic (LFP) scales are thought to arise from the synchronous and rhythmic synaptic activity of groups of neurons [6,9,10]. Thus, an important question is determining what mechanisms can generate oscillatory behavior amongst such groups of neurons. Individual neurons can exhibit intrinsic resonance properties due to a wide variety of ion channels, resulting in preferential responses to inputs at particular frequencies [54,55]. Such cellular resonance properties can also lead to self-sustaining subthreshold membrane potential oscillations [56,57,58,59,60], which can in turn be synchronized through network mechanisms [61]. By influencing the pattern of spiking outputs [60,62], these subthreshold membrane potential oscillations can underlie the synchronized oscillatory spiking of a network. Intrinsic oscillatory properties of single neurons can also involve the generation of periodic high-frequency bursts of spikes through interactions between activity-dependent and voltage-sensitive conductances [63,64,65,66,67]. Such rhythmic bursting is thought to be capable of both generating and synchronizing network oscillations [64,65]. In fact, neurons known as central pattern generators can intrinsically produce rhythmic activity that drives such motor outputs as digestion [68] and respiration [69].

Since neurons can exert both excitatory (E) and inhibitory (I) synaptic influences on one another, several mechanisms exist that can produce network synchronization. Generally speaking, these include recurrent excitation (E-E), mutual inhibition (I-I), and feedback inhibition (E-I) [70]. In fact, single neurons that spike repetitively can be treated mathematically as oscillators [71], and groups of such model neurons can be studied using the theory of coupled oscillators [72]. A key element of this approach is the phase-resetting curve, which characterizes the perturbation

responses of individual neural oscillators [71]. Recent work has even led to the surprising demonstration that networks of uncoupled neurons can be synchronized by correlated noise inputs [73,74], a phenomenon known as stochastic synchrony. Thus, a number of mechanisms have been demonstrated for generating and synchronizing oscillations in groups of neurons involving both cellular and network properties.

Cross-frequency Coupling

Extensive evidence has been acquired in recent years indicating that rather than being independent phenomena, neural oscillations at different frequencies are strongly related. Such cross-frequency relationships can occur in several unique ways. One is that fluctuations in the amplitude (or equivalently, the power) of oscillations at one frequency can be correlated with the amplitude (power) of oscillations at another frequency (amplitude-amplitude coupling). Such cross-frequency power comodulation has been described between the hippocampal theta and gamma oscillations in rodents [75], as well as between gamma and beta oscillations of prefrontal and parietal cortices in the human MEG [76]. Another form of cross-frequency coupling can exist between the phases of two oscillations. So-called phase-synchrony, or phase-phase coupling, exists when the phases of two oscillations are not independent. For oscillations at different frequencies this implies that the higher-frequency oscillation undergoes m cycles for every n cycles of the lower-frequency oscillation ($m > n$); hence this cross-frequency phase-phase coupling is also referred to as m:n phase coupling. Indeed, m:n phase coupling is a salient feature of human neocortical MEG [77,78] and EEG [79] recordings, with 1:2-1:4 synchronies present among alpha, beta, and gamma oscillations [77,78], as well as between theta and gamma oscillations [79].

By far the most well characterized form of cross-frequency coupling, however, is a dependence of the amplitude (power) of a higher-frequency oscillation on the phase of a lower-frequency oscillation (phase-amplitude coupling). Phase-amplitude coupling has been seen to occur between a large range of neural oscillations, the most commonly studied of which is between the phase and amplitude of hippocampal theta and gamma oscillations respectively [80,81,82,83,84]. Theta-gamma phase-amplitude coupling has also been seen in the human cortex [85] and rodent striatum [83], as well as between hippocampal theta and cortical gamma oscillations [86]. Phase-amplitude coupling is not unique to theta and gamma oscillations however, and has also been observed between a number of oscillation pairs including delta and theta [82], delta and low-gamma [87], delta/theta and alpha/beta [88], alpha and gamma [89,90], as well as between infraslow oscillations and a range of higher frequency oscillations [44,91]. Thus, far from being unique to any particular pair of oscillation frequencies, brain region(s), or species, cross-frequency coupling appears to be a ubiquitous organizing principle of neural activity [82,92].

Functional Role of Brain Oscillations

As described in *Oscillation generating mechanisms*, meso- and macroscopic neural oscillations are thought to reflect synchronous modulation of rhythmic synaptic inputs to large groups of neurons. Thus, it is believed that such oscillations are either a direct cause or consequence of synchronous neural activity [93]. Neural oscillations can influence information processing and transmission in several ways. Since temporally aligned inputs can more effectively drive downstream neurons [94,95], synchronous network activity, as produced by oscillations, can increase propagation to downstream structures [96]. Oscillations of downstream neurons can

also influence the effectiveness of driving inputs; and in particular, synchronization of oscillations in the driving and receiving neurons can enhance communication between the two areas [97]. Similarly, oscillatory synchrony among ensembles of neurons is thought to be important for sensory processing. For instance, oscillatory synchrony among neurons with similar stimulus responses properties [9,98] can more effectively recruit lateral inhibition of competing representations [99]. Further, coherent oscillations among distributed sensory representations are believed to provide a potential solution to the so-called 'binding problem' by linking representations of different object features [93,100,101].

Oscillations also define temporal windows within which neurons can integrate and respond to synaptic inputs. It has been suggested that the gamma oscillation could thus produce discrete frames of sensory processing [102], as well as a means to store sequences of item representations in working memory [103]. More generally, substantial evidence now supports the idea that information can be represented directly in terms of the phase of spikes relative to ongoing oscillations (also known as a 'temporal code') [104]. Recent studies have shown that visual and auditory information is represented in the phase of spikes relative to low-frequency oscillations in the LFP, and that information provided by spike phase is complimentary to that provided by spike rates [105,106]. By far the most well known example of such spike-phase coding, however, is the relationship between the phase of hippocampal neurons' spiking relative to the theta rhythm and an animal's position within an environment (see *The Temporal Code*) [26]. In addition to providing a potentially independent source of information relative to spike rates, such phase coding can produce a precise temporal ordering of spikes within an oscillations cycle, which can be critical in determining modifications of synaptic strengths through spike-

timing dependent plasticity [107,108]. In fact, it has been well-documented that induction of long-term synaptic plasticity is most effective when stimuli are delivered as theta frequency oscillations [109]. The conversion of inputs encoded by the firing rates of presynaptic neurons into a phase-coded output ('rate-phase transformation') can be accomplished in the presence of ongoing oscillatory modulation [110,111].

In addition to theoretical results demonstrating the potential functional roles of neural oscillations, numerous direct links have been made between the presence and dynamics of neural oscillations and cognitive processes. For example, attentional modulation of sensory processing through gamma oscillations has been well-established in several different sensory modalities (reviewed in [112]). This could occur through modulation of coordinated subthreshold oscillations [98] leading to enhanced responses to attended stimuli [49]. Furthermore, in line with the notion that synchronously oscillating neural ensembles may be neural correlates of object representations, maintenance of such representations in working memory has been correlated with enhanced gamma oscillations [112]. Both gamma and theta oscillations have been extensively implicated in both the encoding and retrieval of memories (reviewed in [113]). The slow oscillation during deep sleep has also been shown to be important for long term memory consolidation [114,115] (see *Functional Role of UP-DOWN states*).

Given that neural computations, and the communication between brain regions, are thought to depend on oscillations, it is perhaps not surprising that the ubiquitous interactions between oscillations at different frequencies would play a particular role in these processes. In fact, phase-amplitude coupling between theta and gamma oscillations is itself a rich and complex phenomenon, with rapid and region-specific dynamics that relate to ongoing behavior [83].

Phase-amplitude [116], phase-phase [79], and amplitude-amplitude [76] coupling of theta and gamma oscillations have been shown to be predictive of performance in working memory tasks. Cross-frequency couplings of other neural oscillations have also been related to both working memory and decision-making processes [78,87,117].

Further evidence supporting the importance of neural oscillations in normal brain function is derived from observations that aberrant oscillatory activity is linked with several neurological disorders. Reductions of the amplitudes and synchronization of beta and gamma oscillations in schizophrenic patients have been well documented [118]. Similarly, decreases in the amplitude and synchrony of alpha- and beta-band oscillations are linked with Alzheimer's disease [119]. Aberrant enhancements of beta band oscillations in the basal ganglia have been shown to strongly relate to akinesia symptoms of patients with Parkinson's disease [120]. Even while epilepsy is closely associated with neuronal hyper-excitability, more recent evidence also suggests that abnormal synchrony of local oscillations [121], as well as the generation of aberrant ultra-fast oscillations [45], may also play an important role in generating epileptic activity. Further, some evidence now exists for the presence of abnormal neural oscillations in patients with autism spectrum disorders [122]. Thus, nearly all of the major neurological disorders are correlated with changes in the presence and synchronization of different neural oscillations, emphasizing their importance in normal cognitive function.

UP-DOWN States

UP-DOWN states overview

In the early 1990's, Steriade and colleagues demonstrated the presence of a slow (<1Hz) oscillation in the cortical EEG of naturally sleeping and anesthetized animals that was separate

from the well-known delta oscillations (1-4Hz) characteristic of slow-wave sleep [123]. These and other authors further showed that these slow oscillations in the EEG corresponded to synchronous alternations of single neurons between two stable activity states [123,124], later referred to as the UP and DOWN states. The UP states were characterized by depolarization and spiking, while neurons were quiescent and hyperpolarized during DOWN states. It has since become apparent that the slow-oscillations are a ubiquitous property of neural networks under a variety of conditions.

The slow oscillation is present during all stages of human sleep other than rapid eye movement (REM) sleep[125]. In fact, the large EEG deflections known as K-complexes that are characteristic of non-REM (nREM) sleep are now believed to represent individual cycles of the slow oscillation [126]. Importantly, similar slow-oscillations are also present under several commonly used anesthetics, including urethane [123], ketamine/xylazine [123], and isoflurane [127], although they are not present under barbiturate anesthesia [123]. While the properties of the induced slow-oscillations differ for different anesthetics [123], as well as compared to natural nREM sleep [123,128], the anesthetized state bears many similarities to the natural sleep state including dynamic alternations between discrete activity states [129], functional dissociation of corticothalamic networks, and similar modulation by ascending brainstem cholinergic pathways [130]. Thus, studies of slow oscillations under anesthesia are thought to provide important insight into the mechanisms and function of natural sleep states. Since studies of anesthesia often refer to slow oscillations as UP-DOWN states (UDS), we use the terminology UDS and slow oscillation separately to distinguish the anesthetized and natural sleep states.

Since their initial discovery by Steriade et al., UDS have been shown to occur synchronously among neurons in a given cortical area [123,131,132], as well as between distant cortical regions [133]. Neocortical neurons of all types, including both excitatory and inhibitory neurons, from all cortical layers, and between hemispheres [134] are seen to participate nearly synchronously in the UDS [133,135,136,137]. In fact, transitions from the DOWN to UP state have been observed to occur as traveling waves in the cortex of rodents [138] and humans [139]. Further, rather than being confined to the neocortex, neurons from a wide range of brain regions have been shown to participate in the UDS, including thalamocortical and reticular thalamic nucleus (RTN) neurons [140,141], cerebellar [142], and basal ganglia neurons [143,144]. Even the cholinergic brainstem neurons which control the state of arousal (and whose direct activation can abolish UDS [145]) participate synchronously in UDS [146]. Interestingly, hippocampal neurons show varied participation in UDS, with CA1 R-LM interneurons [147], some dentate gyrus granule cells [148], and subicular neurons [149] exhibiting synchronized UDS, while pyramidal neurons in CA1 and CA3 [148,149], as well as some dentate gyrus granule cells [149] do not show UDS. Thus, UDS and their analogue during natural sleep the slow oscillation, represent an alternation between two activity states that is remarkably ubiquitous and synchronous.

Mechanisms of UP-DOWN states

While a number of models have been suggested to explain the generation of UDS, the topic remains highly debated. Such theories can generally be divided into two classes depending on whether the UDS are generated within the neocortex or thalamus. Two important experimental observations have been taken by many as evidence of a cortical origin of UDS. One is the

demonstration that UDS in thalamic neurons are eliminated upon removal of corticothalamic afferents [150,151]. The other is that cortical UDS can persist in a cortical slab [152] or slice preparation [135] lacking thalamocortical neurons. For cortical generation theories, the question then arises as to what intrinsic cortical mechanisms can generate these bistable network oscillations. Interestingly, cortical UP states both *in vitro* [135,136] and *in vivo* [137] are characterized by a precise and self-sustaining balance between excitation and inhibition, yet they can be initiated and terminated by synaptic inputs [136]. Within the context of such a bistable cortical network it remains unclear what cortically generated signal actually initiates and terminates the UP states. Timofeev and colleagues have suggested that the spontaneous summation of miniature excitatory post-synaptic potentials (EPSPs), particularly in large layer 5 neurons, can initiate a cascade of activity and the transition to the UP state [152,153]. Similarly, in a modeling study Holcman and Tsodyks suggested that synaptic noise could initiate transitions to the UP state [154]. Other theories argue that a subset of cortical neurons is spontaneously active even during the DOWN state [155,156]. Indeed, a subset of cortical pyramidal neurons, as well as layer 5 Martinotti interneurons, have been shown to exhibit spontaneous activity *in vitro* in the absence of synaptic inputs [157]. This intrinsically generated spontaneous activity could potentially trigger cortical UP states during UDS *in vivo* as well. Transitions to the DOWN state have been suggested to occur as a result of the buildup of activity-dependent conductances [155], depression of excitatory synapses [153], spontaneous synaptic fluctuations [154], and network mechanisms [156].

The other set of models posit that UDS are generated either within the thalamus or through thalamocortical interactions [158]. The most important evidence in favor of this possibility

comes from the observation that thalamocortical neurons consistently exhibit bursts of low-threshold calcium spikes that precede the cortical UP state transition by 20-50ms [141]. It was initially unclear how such thalamocortical activity could actually initiate the cortical UP states, since thalamocortical neurons did not exhibit UDS in the absence of cortical inputs. Hughes and colleagues provided an interesting explanation by showing that both thalamocortical [159] and RTN [160] neurons can intrinsically generate UDS *in vitro* upon activation of the mGluRa metabotropic glutamate receptors. These receptors would normally be activated by corticothalamic inputs, which accounts for the observations that UDS are absent in decorticated preparations [150,151]. Thus, the model of Crunelli et al. suggests that interactions between the intrinsically generated thalamic UDS and the synaptically mediated cortical UDS are needed to explain the full set of experimental observations [158]. Interestingly, a recent study has demonstrated that the spike timing of thalamocortical neurons relative to cortical UDS is dynamic and nucleus specific [161], emphasizing the potential importance of thalamocortical interactions during UDS.

Functional Role of UP-DOWN states

The primary electrophysiological structures present during nREM sleep and anesthesia, in addition to the slow-oscillation (UDS), are the 1-4Hz delta oscillations, short-lasting (0.5-3s) 7-14Hz cortical oscillations known as spindles, and brief (50-200ms) high-frequency (100-250Hz) hippocampal oscillations (ripples) that are accompanied by large amplitude deflections in the hippocampal LFP (sharp-waves). Importantly, all of these phenomena are related to the ongoing slow oscillations (UDS) [150,162], suggesting that the slow oscillation may serve to temporally organize the other sleep structures. As described in *Functional Role of Brain*

Oscillations, sleep, and in particular the slow oscillations, are thought to be important for memory consolidation [114]. Indeed, induction of slow oscillations (0.75Hz) through application of transcranial electric fields enhanced retention of declarative memory, while stimulation at higher frequencies characteristic of REM sleep had no effect [115]. Similarly, learning is correlated with an increased prevalence of sleep spindles during ensuing sleep [163]. Long-standing theories of memory consolidation have focused on interactions between the hippocampus and neocortex [164,165], with the idea that newly learned memories stored in the hippocampus are transferred to the neocortex through a gradual and interleaved reactivation of hippocampal memories thought to occur during sleep [165,166].

In support of such theories, Pavlides and Winson observed that correlated hippocampal spiking activity during waking behavior was reactivated during subsequent sleep [167]. More amazingly, however, it was later discovered that precise patterns of hippocampal spiking that occurred during behavior were in fact replayed during subsequent slow-wave [168,169] and REM [170] sleep. Interestingly, while replay of sequences during REM sleep occurred at the same rate as during behavior, during slow-wave sleep temporal sequences were compressed roughly 20-fold, and occurred during the transient hippocampal sharp-wave ripple complexes. Similar sleep replay has also been observed to occur in the medial prefrontal cortex [171], as well as in a coordinated fashion between visual cortical and hippocampal neurons during UP states [172]. In all cases it appears that the hippocampal sharp-wave ripples are involved in the generation of these replay events during nREM sleep, and thus could represent the transfer of information from the hippocampus to the neocortex. The alignment of hippocampal sharp-

wave ripples with cortical UDS [162] suggests that the slow oscillations indeed serve to organize the dialogue between neocortex and hippocampus during sleep.

The Entorhinal Cortex/Hippocampus System

The Hippocampus

The hippocampus is a curved structure located in the medial temporal lobe. In a highly influential study, Scoville and Milner found that a patient (known as H.M.) whose hippocampus had been surgically lesioned suffered severe anterograde amnesia, being unable to form new memories. At the same time, H.M. retained memories from long before the operation [173]. This finding led to the suggestion that the hippocampus plays an important role in the formation and consolidation of new memories. It is now widely believed that the hippocampus and anatomically related cortical structures (entorhinal, perirhinal, and parahippocampal cortices) comprise the so-called medial temporal lobe memory system, and are responsible for declarative memory (memories which can be consciously recalled) [165,174]. In addition to its suggested role in declarative memory formation, the discovery by O'Keefe and Dostrovsky of hippocampal neurons that fire selectively in certain regions of an environment (termed 'place cells') [175] revealed an additional function of the hippocampus as an integral part of the brain's spatial processing system [176]. While the spatially selective activity of hippocampal neurons is discussed in more detail below, an overview of hippocampal anatomy is given first.

The hippocampus was first divided into its now commonly used subdivisions by Lorenté de Nó in 1934 [177]. These subfields were labeled CA1 through CA4 (although CA2 and CA4 are often not distinguished), and comprise the hippocampus proper. The hippocampal formation refers to the CA subfields along with the subiculum and the dentate gyrus. The primary axis of the

hippocampus spans a substantial portion of the antero-posterial axis of the brain, arching from a dorsal/medial location more anterially to a more ventral/lateral position posterially. In cross-section, the CA subfields form a C-shaped structure with the subiculum located on the CA1 end. The dentate gyrus forms a separate C-shaped structure wrapping around the CA4 end of the hippocampus proper, creating an overall S-shaped appearance [178]. The CA subfields are composed of a large number of pyramidal neurons whose cell bodies are tightly packed into a single layer. These pyramidal neurons are 'inverted' such that their dendrites extend towards the center of the hippocampal formation while their axons run along the outside. In addition to pyramidal neurons, a wide variety of interneurons exist within the hippocampus, each thought to play a unique role in shaping the activity of hippocampal principal neurons [179]. The principal neurons of the subiculum are also pyramidal, while in the dentate gyrus granule cells provide the output.

The overall connectivity pattern of the hippocampus creates a well-defined loop with the entorhinal cortex, its primary source of both input and output (see *The Entorhinal Cortex*). The dentate gyrus, which receives its primary inputs from the entorhinal cortex, sends its outputs to CA3. CA3 receives additional direct inputs from entorhinal cortex, and projects to CA1. Similarly, CA1 receives inputs from both CA3 and entorhinal cortex, and sends its outputs to the subiculum and entorhinal cortex. Thus, the general path of signal flow is from the entorhinal cortex to the dentate gyrus, then through CA3 and CA1 before returning to the entorhinal cortex. Interestingly, while there is extensive recurrent connectivity among the pyramidal neurons of CA3, CA1 has a 'feedforward' architecture, lacking recurrent connections among the

pyramidal neurons. In both CA1 and CA3, however, the different interneuron subtypes implement both feedforward and recurrent inhibition [178].

The Entorhinal Cortex

Given the purported role of the hippocampus, and in particular the cortico-hippocampal interactions, in memory formation, it is perhaps surprising that the hippocampus is not extensively interconnected with widespread neocortical areas. Instead, the hippocampus is reciprocally connected primarily with the entorhinal cortex (EC), which itself has such extensive interconnections with neocortex. Thus, the EC is often seen as a functional and anatomical gateway between the neocortex and hippocampus, and is thought to play a role in the same memory and spatial processing systems as the hippocampus. Superficial layers of the EC (layers 2 and 3) provide the primary input to the hippocampal formation, with layer 2 EC neurons projecting to the dentate gyrus and CA3, while layer 3 EC neurons project directly to the hippocampal output structures, the CA1 and subiculum. The deep layer 5 neurons of the EC are then the primary recipients of the CA1 and subicular outputs [180,181,182]. The EC sends projections to a range of neocortical areas primarily originating from the layer 5 neurons, but certain cortical projections arise from the superficial layers of EC as well [183]. Neocortical inputs to the EC arise from a similarly diverse set of cortical regions, and primarily target neurons in layers 2 and 3 of the EC, but some projections exist to layer 5 of EC [181].

Typically the EC is divided into two primary subdivisions referred to as the medial EC (MEC) and lateral (LEC) respectively. The MEC constitutes the more dorsal, posterior, and medial portion of the EC, while the LEC is more ventral, lateral, and anterior [181]. The EC is further divided into three bands along an axis roughly perpendicular to the MEC-LEC division. These bands are the

dorsolateral, intermediate, and ventromedial bands [182]. In addition to defining the spatial organization of the EC, these divisions relate to differences in the cortical afferent [184,185] and efferent projections [183,185], as well as the pattern of projections to the hippocampal formation [181,182,185]. Specifically, the dorsolateral band of the EC projects primarily to the dorsal part of the hippocampus, while the ventromedial band projects to the more ventral hippocampus [182]. Further, projections of the layer 3 MEC neurons to CA1 are targeted nearer CA3 compared to projections from the LEC [180]. These anatomical differences between entorhinal subregions also correspond to substantial functional differences, as described below.

The Rate Code

The discovery by O'Keefe and Dostrovsky that hippocampal neurons fire spikes selectively within a certain region of an environment [175] led to the suggestion that these hippocampal neurons (termed 'place cells') could generate a cognitive map of the environment [176]. Indeed spatial selectivity is found in both CA1 and CA3 neurons [186], as well as in the dentate gyrus granule cells [187]. While the preferred locations of place cells (their 'place fields') do not show a topographic orientation [188] as is seen in sensory cortical areas, place field sizes do systematically increase along the dorso-ventral axis of the hippocampus [189]. Thus, it is believed that local ensembles of hippocampal place cells, with place fields distributed over the environment, are sufficient to represent an animal's position [188]. This representation, given by the 'instantaneous' firing rates of ensembles of hippocampal place cells, is referred to as the hippocampal rate code [190].

Since their discovery, a number of interesting properties of place cells have been revealed (see [191] for review). For instance, while place cells respond only to the position of an animal's

head within a 2-dimensional environment, they respond to both position and head direction when an animal runs on a narrow linear (1-dimensional) track, and thus represent each movement direction on the track independently [26]. The set of cells that are active (have place fields) changes from one environment to another [188]. As a consequence, if an environment is changed sufficiently a new subset of place cells becomes activated, a phenomenon referred to as remapping [192]. Consistent with the observation that the hippocampus is required for spatial learning [193], systematic changes have been observed in the firing activity of place cells with experience [188,194,195]. In particular, place fields show a rapid asymmetric expansion and backwards shift [194,195] that is predicted to occur as a result of spike-timing dependent plasticity (STDP) [107,108]. Such newly acquired 'memories', believed to be represented by the rapid modifications of hippocampal synaptic connectivity with experience, are then hypothesized to undergo 'consolidation' to form a stable and distributed representation in the neocortex. In support of this idea, ensembles of place cells which were activated in a particular sequence during recent experience are observed to become reactivated in the same spatio-temporal pattern during subsequent slow-wave sleep, a phenomenon known as hippocampal replay [168,169] (see *Functional Role of UP-DOWN states*). In addition to an animal's position, the firing rate of place cells during maze running is modulated by behavioral variables such as running speed [196,197,198,199,200,201], as well as non-spatial behaviors, stimuli, and reward conditions [197,202,203,204,205,206,207], creating potential ambiguity in the rate code for position. Thus, far from being a simple and static representation of position, the hippocampal rate code is flexible, dynamic, and multifaceted.

Beginning in 2004, a striking series of studies by the Moser group revealed that neurons in the MEC have a more important role in spatial processing than simply mediating cortical inputs to the hippocampus. These authors showed that principal neurons in all layers of the MEC, later termed 'grid cells', fired only in certain regions of an environment such that their multiple 'place fields' formed a hexagonal grid of the environment [208,209]. While the neurons in a particular region of MEC were seen to have similar grid spacing and orientation, the spatial phases of nearby grid cells were randomly distributed [210], similar to the nontopographic organization of hippocampal place cells. Further, the spatial scale of the grids increased from dorsolateral to ventromedial locations within the MEC [208], consistent with the increase in place field size seen in their respective downstream hippocampal targets [189]. EC cells that respond to the animal's head direction, as well as those that provided a conjunctive representation of position and head direction were later discovered [209], further emphasizing the role of the EC-hippocampal system in spatial processing.

The Temporal Code

During active behavior the hippocampal LFP also shows large-amplitude theta rhythmic (~8Hz) network oscillations [211]. In a groundbreaking study, O'Keefe and Recce showed that the timing of hippocampal spikes relative to the ongoing theta oscillation provided a 'temporal' code for the animal's position [26,104], in addition to the spatial information carried by their spiking rate. This temporal code was given by the progressive advancement in the timing (or phase) of a place cell's spiking relative to the theta rhythm as the animal ran through the cell's place field. When analyzed across many passes through the place field, spiking is observed to precess through a full 360 degrees of theta phase [26], while for individual passes through the

place field precession is restricted to ~180 degrees [212]. It has been suggested that precession can provide a mechanism for encoding the animal's position that is independent of behavioral variables such as running speed [26,213,214]. In fact, while most studies of phase precession have been conducted on linear tracks, recent work has shown that the temporal code can disambiguate an animal's trajectory through a two-dimensional place field [215]. Precession is also believed to be important for sequence learning [216,217,218] by producing sequential activation of place cells with adjacent place fields on the compressed timescale relevant for STDP [107,108]. Phase precession was later shown to occur in grid cells in layer 2 of the MEC as well, but interestingly was only weakly present in the layer 3 grid cells [219].

Numerous models have been proposed to explain precession, and they can roughly be divided into three general classes. A key component of these models is that the theta rhythmic oscillation of the hippocampal LFP corresponds to a rhythmic inhibition of the pyramidal neurons' somas. Indeed, this is expected given the firing patterns of most hippocampal interneurons that provide perisomatic inhibition to pyramidal neurons [179]. In the first class of models, precession is proposed to represent the 'readout' of a sequence of associated place cells within each theta cycle, arising from recurrent network connectivity [220,221]. Because of the recurrent connections assumed by such models, however, they are only appropriate for generating precession within CA3, but not CA1. A second group of models attempt to explain precession as arising from the interference of two separate oscillators [26,222]. The third set of models relies on the interaction between theta rhythmic inhibition and a spatially modulated depolarization. This spatially modulated depolarizing input is itself modulated by the theta rhythm in the 'soma-dendrite interference' models [11,223]. Alternatively, Mehta et al. have

shown that precession can arise from a combination of theta rhythmic inhibition and asymmetric ramping excitation [110]. Such asymmetric excitation has indeed been observed both in the spiking activity [110,194,195] and subthreshold membrane potential fluctuations [224] of place cells. Importantly, the latter class of models implies a relationship between the strength of excitatory inputs and the phase of spiking outputs, a phenomenon known as the rate-phase transformation [110,111].

Thesis Outline

The goal of this thesis will be to explore the dynamics of neural activity in the neocortical-entorhinal-hippocampal (Ncx-EC-Hpc) circuit from a quantitative perspective. The thesis is divided into two parts. In Part 1 (Chapters 2-4), I analyze the activity of entorhinal neurons during UDS in relation to both neocortical and hippocampal activity. These studies aim to uncover cellular and network properties of the entorhinal cortex, as well as its role in mediating cortico-hippocampal interactions during sleep. In part 2 (Chapters 5-6), I examine the activity of hippocampal neurons and their rate and temporal codes of position.

In order to provide a more quantitative description of UDS in the Ncx-EC-Hpc circuit, in chapter 2 I present a general hidden Markov model (HMM) algorithm for inferring UDS from continuous signals (LFP or MP). The HMM framework allows for robust and flexible UDS classification from any set of signals and signal features, and provides a natural means of evaluating different signal features in terms of their discrimination information. I show explicit models of the state duration distributions can also be incorporated to improve upon the implicit use of geometric state duration distributions by the standard HMM. This HMM method is validated and

quantitatively compared to standard approaches using simultaneous recordings of neocortical LFP and the MP of nearby neurons.

In Chapter 3, I apply the HMM UDS classification algorithm to study the relationship of UDS in layer 3 pyramidal neurons from lateral and medial EC (LECL3 and MECL3) to neocortical UDS. I show that MECL3, but not LECL3, neurons nearly always persistent in the UP state beyond the concurrent neocortical UP state. These MECL3 'persistent' UP states often last for several cycles of the neocortical UDS, yet MECL3 state transitions maintain precise phase relationships relative to neocortical UDS. This results in a quantized distribution of MECL3 UP state durations in units of neocortical UDS cycles. I discuss the influence of these persistent MECL3 UP states on the hippocampal LFP, along with the influence of cellular mechanisms in generating the persistent activity.

While superficial neurons of the LEC do not show persistent UP states, their unique oscillatory properties are explored in Chapter 4. In this chapter, I show that LEC neurons in layers 2 and 3 occasionally transition into states with pronounced 2-4Hz oscillations in conjunction with ongoing UDS. These oscillatory states are found to correspond with discrete changes in the global network state, and occur synchronously with similar oscillations in the neocortical and hippocampal LFP. Rather than being independent of the ongoing UDS, the 2-4Hz oscillations are phase-locked to the UDS transitions. I discuss the relation of these LEC oscillations to previously described theta and delta oscillations, as well as their potential role in controlling the relative timing of UDS in the Ncx-EC-Hpc circuit.

In Chapter 5, I investigate the activity of hippocampal CA1 'place cells' in genetically modified mice lacking the GluA1 AMPA receptor subunit. This subunit has been shown to produce

specific alterations of synaptic strength and long-term synaptic plasticity in CA1, and the KO animals are deficient at certain tests of spatial memory. I show that nearly all aspects of the hippocampal rate code of position are impaired in the KO place cells, and specific alterations of the temporal spiking properties of KO cells are also observed. In addition to the single-neuron spiking properties, I also demonstrate large differences in the network theta and gamma oscillations of KO animals.

The relationship of the hippocampal temporal code to behavioral and reward parameters is discussed in Chapter 6. In particular, I show that hippocampal phase precession is weaker for place fields nearer the reward location. Spike time autocorrelation and ISI distribution analyses are used to identify specific changes of the temporal spiking structure for these place fields. Similar changes with respect to reward proximity are also observed to occur in the hippocampal theta rhythm itself. I discuss alternative explanations of the results based on related parameters, and I present a general mechanism based on systematic increases in dopaminergic inputs.

In Chapter 7, I summarize the findings of the thesis and discuss the ways in which future research can further our understanding of the Ncx-EC-Hpc circuit and its function during sleep and behavior. Following the results of Part 1, I discuss extensions of the HMM UDS classification algorithm. I then focus on discussing experiments which can determine the mechanisms that generate the MECL3 persistent UP states, its manifestation at the network level, and its contributions in shaping cortico-hippocampal interactions during natural sleep. Similarly, I explore ways in which the 2-4Hz oscillations seen in superficial LEC neurons might be generated and synchronized, and how such oscillations could affect the relative timing of UDS in the Ncx-

EC-Hpc circuit. Future research exploring the results of Part 2 is also discussed, emphasizing general mechanistic principals underlying the rate and temporal codes. I also focus on the relationship between hippocampal activity and behavior, and propose several experiments which could clarify and extend the results of Chapter 6.

Chapter 2: **Explicit-duration hidden Markov model inference of UP-DOWN states from continuous signals**

Introduction

Different patterns of activity in the neocortex are known to correspond to different behavioral states. During slow-wave sleep, large amplitude slow (<1Hz) oscillations in the EEG and local field potential (LFP) are present reflecting synchronous fluctuations in the membrane potential (MP) and spiking activity of individual neurons. A similar state is also observed under various anesthetics whereby neurons exhibit bistability and undergo synchronous transitions between a depolarized and active UP state, and a quiescent hyperpolarized DOWN state [123,124]. Both excitatory and inhibitory neurons participate in these synchronous state transitions, and thus the active state is characterized by a balance of excitatory and inhibitory activity [135,136,137].

UP-DOWN states (UDS) present an excellent opportunity to study both cellular and network properties, and because of their global nature they can yield insight into network dynamics, both within and across brain regions. There has also been much interest in using the relatively simple discrete state dynamics of UDS to study state dependencies of neural responses to external stimuli [132,225,226,227]. Further, a number of studies have demonstrated that UDS occurring during natural sleep could serve an important role in the process of memory consolidation [114,115], thought to occur via cortico-hippocampal interactions [165]. This possibility is supported by the observations that hippocampal activity can be synchronized with

the neocortical UDS [147,149,228] and that the primary electrophysiological structures present during sleep (including sleep spindles and hippocampal sharp-wave ripples) are temporally organized by UDS [150,162,229,230]. Thus, understanding UDS will likely yield insights not only into the cellular and network dynamics at play during slow-wave sleep, but also into the role of slow-wave sleep in memory consolidation.

UDS have most often been defined in terms of the MP of individual neurons, however given that robust spiking of both excitatory and inhibitory neurons occurs exclusively during the UP state, classification of UDS based on extracellular unit activity is also possible [172,231,232,233]. Due the synchronous nature of UDS, large amplitude fluctuations in the LFP can also be used to classify them [147,148,149]. Classifying UDS from extracellular signals (multi-unit, LFP, EEG) offers the obvious advantage that the data is easier to acquire, but it also presents a natural means for estimating the collective state of an ensemble of neurons near the recording electrode. Further, due to the increasing popularity of multi-electrode recording methods, there is a strong need for a general framework for inferring UDS from extracellular signals.

Several different approaches to classifying UDS in LFP signals have been put forth which utilize different features of the LFP signal. Mukovski et al. argued that the high-frequency (20-80Hz) power measured in a suitably localized window of time provides the best signal feature for classifying UDS [234], while more recently Saleem et al. used the low-frequency (<4 Hz) phase of the LFP for classification [235]. While these works also differed in their methods of inferring UDS from the respective signal features, both approaches were based on determining a fixed threshold (or set of thresholds) and comparing the signal feature(s) to these fixed threshold(s) independently at each time point. We broadly term such approaches 'threshold-crossing'

algorithms, and argue that the framework of hidden Markov models (HMMs) provides significant advantages over threshold-crossing methods for classifying UDS. In particular, HMMs provide a consistent, flexible, and statistically principled framework for inferring UDS from different types of signals (including MPs and LFPs) that handles variations in experimental parameters such as type of anesthetic, type of electrode, precise electrode location and depth of anesthesia without the need for supervision, and can easily be adapted to accommodate non-stationarities in the data [236,237]. The efficacy of HMMs for inferring UDS from stationary point-processes (extracellular spike times) has already been demonstrated [231,232]. Here we describe an explicit-duration HMM (EDHMM) method for inferring UDS from potentially non-stationary sets of continuous signals (including MPs and LFPs). This approach also allows for a natural means of evaluating different signal features for UDS classification in terms of their discrimination information. We show that the amplitude of low-frequency LFP provides the most information about the neocortical state in our data. We then show by comparing UDS inferred from simultaneously recorded LFP and MP signals that our EDHMM procedure produces significant improvements over standard methods of classifying UDS.

Methods

Experimental Methods

Ethics Statement

All surgical procedures and experiments were conducted according to the animal welfare guidelines of the Max-Planck-Society. The protocol was approved by the responsible State Committee on the ethics of animal experiments Karlsruhe (Permit Number 35-9185.81). All efforts were made to minimize suffering.

Animals, Surgery, and Histology

Methods were similar to those described previously [147,148]. Briefly, data were obtained from 11 C57BL6 mice aged postnatal day p29 to p35 weighing between 16 and 23g. Mice were anesthetized with urethane (1.7-2.0g/kg i.p.). Body temperature was maintained at 37ºC with the help of a heating blanket. The animal was head-fixed in a stereotaxic apparatus and the skull was exposed. A metal plate was attached to the skull and a chamber formed with dental acrylic was filled with warm artificial cerebrospinal fluid. A 1 to 1.5mm diameter hole was drilled over the left hemisphere and the underlying dura mater was removed. After electrophysiological recordings, mice were transcardially perfused with 0.1M PBS followed by 4% paraformaldehyde and 150µm thick saggital brain sections were processed with the avidin–biotin–peroxidase method.

Electrophysiology and Data Acquisition

LFPs were recorded with a quartz/platinum-tungsten glass coated microelectrode (Thomas Recording GmbH, Giessen, Germany). *In vivo* intracellular membrane potential (MP) was recorded in whole-cell configuration using borosilicate glass patch pipettes with DC resistances of 4-8 MΩ, filled with a solution containing 135mM potassium gluconate, 4mM KCl, 10mM Hepes, 10mM phosphocreatinine, 4mM MgATP, 0.3mM Na3GTP (adjusted to pH 7.2 with KOH), and 0.2% biocytin for histological identification. Whole-cell recording configuration was achieved as described previously [238]. Relative to bregma, both the MP and the LFP recordings were made either around 1 to 1.5mm anterior and 1 to 1.5mm lateral (frontal) or around 3mm anterior and 1 to 1.5mm lateral (prefrontal). MP was recorded from pyramidal neurons at various depths, and LFP was recorded from upper layer 5. The recording site of the LFP was at less than 1mm distance from the neuronal soma from which the MP was obtained Both LFP and

MP were recorded continuously on an eight-channel Cheetah acquisition system (Neuralynx; Bozeman, MT) for at least 600s per experiment. That complete recording was used for subsequent analysis as described below. The MP was acquired by Axoclamp-2B (Axon Instruments, Union City, CA) and fed into a Lynx-8 amplifier (Neuralynx). LFP was sampled at 2kHz, low-pass filtered at 475Hz, and amplified 2,000 times. MP was low-pass filtered at 9kHz, sampled at 32kHz, and amplified 80–100 times. Simultaneously, the DC value of MP was recorded by an ITC18 interface (Instrutech; Mineola, NY) under the control of Pulse software (Heka; Lambrecht, Germany).

A total of 21 LFP recordings from 11 animals were analyzed. Usually two LFP recordings were performed in a single animal, the recordings being separated by 50 to 210 minutes. Statistics were computed across LFP recordings. In addition, 9 MP recordings were performed simultaneously with 9 of the LFP recordings, all from different animals.

Statistical Methods

Hidden Markov Model of UP and DOWN States

The problem of inferring the UP and DOWN state sequence given some measure of neural activity is well suited for the framework of HMMs [231,232]. Here we will focus on (discretely sampled) continuous signals such as the LFP or MP. Consider a time series of signal features **Y** (e.g. the amplitude of a filtered LFP) such that **Y** carries information about the underlying UP and DOWN states. Henceforth we will refer to **Y** as the 'observation sequence'. With a HMM we treat the instantaneous state as a discrete hidden variable that can take K possible values (in our case two, representing the UP and DOWN states). Let z_t represent the value of the hidden state variable at discrete time step t. The HMM then makes the simplifying assumption that z_t is

independent of the preceding sequence of latent variables, given the value of z_{t-1}. We can thus write

$$p(z_t \mid z_1, z_2, .., z_{t-1}) = p(z_t \mid z_{t-1}) = \mathbf{A}_{z_{t-1} z_t} \quad \textbf{(2-1)}$$

where **A** is the state transition matrix. The observation y_t at time *t* is assumed to be conditionally independent of previous values of *y* given z_t, and is determined by a state-conditional observation distribution:

$$p(y_t \mid \phi_{z_t}) \quad \textbf{(2-2)}$$

where ϕ_{z_t} is the vector of model parameters associated with state z_t .

In the general case we can define a vector-valued time series of observations $\mathbf{y}_t$. We use Gaussian observation distribution models of the form $p(\mathbf{y} \mid \phi_k) = \mathrm{N}(\mathbf{y}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, where $\phi_k = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$ are the state-conditional mean and covariance matrices. Other observation distribution models could also be considered (in particular 'mixture of Gaussians' models have been used extensively for other applications) if the state-conditional observations are determined to be highly non-Gaussian. Because some of the signal features characterizing the UP and DOWN states often change substantially over the course of the recording [236,237], we allow the mean vectors to be (slowly varying) functions of time. Time-varying covariance matrices were not considered here, however they could be included as well using a similar approach. Therefore, the Gaussian-observation HMM is fully specified by the parameter vector $\Phi = \{\mathbf{A}, \boldsymbol{\pi}, \boldsymbol{\mu}_k(t), \boldsymbol{\Sigma}_k\}$, where $\boldsymbol{\pi}$ is the distribution over the initial state z_1. Signal features which are assumed constant in time are easily handled in this framework as well.

Maximum likelihood (ML) estimates of the model parameters Φ given a sequence of observations can be determined using the expectation maximization (EM) algorithm [239,240] which proceeds using a two-step iteration. First, initial values for the parameters Φ^{old} are selected, and the posterior distribution of the latent state sequence $\mathbf{Z} = (z_1, ..., z_T)$ given the observation sequence $\mathbf{Y} = (\mathbf{y}_1, ..., \mathbf{y}_T)$, and the parameter vector is computed. The posterior distribution over the latent state sequence is then used to compute the expected complete-data log likelihood [241]:

$$Q(\Phi, \Phi^{old}) = \sum_{Z} p(\mathbf{Z} \mid \mathbf{Y}, \Phi^{old}) \log p(\mathbf{Y}, \mathbf{Z} \mid \Phi) \qquad \textbf{(2-3)}$$

Maximization of *Q* with respect to Φ gives the new estimate of the parameter vector Φ^{new}. The process is repeated iteratively until convergence. For a given parameter vector, calculation of the posterior distributions for the latent variables is achieved using a set of recursions known as the forward-backward algorithm [240].

Explicit duration HMM

An important weakness of the standard HMM is that it implicitly assumes a geometric distribution of state durations [240]. Numerous extensions of the standard HMM have been developed to address this issue, but the most frequently used is the explicit-duration HMM (EDHMM) [242,243], which is a type of hidden semi-Markov model (for review see [244]). In the discrete-time EDHMM, the state duration is assumed to be a random variable which can take integer values in the range [1, d_{max}], where d_{max} is the maximum allowable state duration. Upon transitioning to a state *k*, a sequence of conditionally independent observations of length *d* will be emitted from the observation model ϕ_k. Each state can thus be specified by the pair (*k*,*d*),

and state transitions are determined by two such pairs. In the EDHMM, the simplifying assumption is made that $p_t(k,d \mid k',d') = p(k \mid k')p_k(d) = A_{kk'}p_k(d)$ [244]. Thus, the probability of observing state *k* at time *t* depends only on the previous state *k'*, and the probability of state *k* having duration *d* depends only on the duration distribution $p_k(d)$ for state *k*. Since self-transitions are prohibited in the EDHMM, the transition matrix **A** is uniquely determined in the two state case:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \textbf{(2-4)}$$

EDHMM parameter inference:

Inference of model parameters in the EDHMM can be accomplished using a forward-backward algorithm similar to the standard HMM [242]. Yu and Kobayashi demonstrated an efficient forward-backward algorithm for the EDHMM [245], and showed how to redefine the forward and backward variables in terms of posterior probabilities to avoid numerical underflow [246]. Defining $\gamma_t(k) = P(z_t = k \mid \mathbf{Y}, \Phi)$ to be the marginal probability of state *k* at time *t* given the observation sequence **Y** and the model parameters Φ, the observation model parameters are updated in the M step of the EM algorithm in direct analogy with the estimation formulas in the standard HMM, according to:

$$\boldsymbol{\mu}_k(t) = \frac{\sum_{t'} w(t'-t)\gamma_{t'}(k)\mathbf{y}_{t'}}{\sum_{t'} w(t'-t)\gamma_{t'}(k)} \quad \textbf{(2-5)}$$

$$\boldsymbol{\Sigma}_k = \frac{\sum_{t=1}^{T} \gamma_t \left(\mathbf{y}_t - \boldsymbol{\mu}_k\right)^T \left(\mathbf{y}_t - \boldsymbol{\mu}_k\right)}{\sum_{t=1}^{T} \gamma_t} \quad \textbf{(2-6)}$$

where $w(t'-t)$ is a symmetric, time-localized window function with a time-scale chosen to reflect the time-course of fluctuations in the state-conditional means.

State duration distributions

In Ferguson's original EDHMM formulation [242], the state durations are assumed to be generated from arbitrary non-parametric distributions. Defining $D_t(k,d) = P(z_{t-d}, z_{t-d+1}, ..., z_t = k \mid \mathbf{Y})$ to be the conditional probability of state *k* starting at time *t-d* and ending at time *t* (lasting for duration *d*), Ferguson showed that ML estimates for the parameters of the non-parametric probability mass function (pmf) for state *k* are given by [242]:

$$p_k(d) = \frac{\sum_{t'=2}^{T} D_{t'}(k,d)}{\sum_{d'=1}^{d_{\max}} \sum_{t'=2}^{T} D_{t'}(k,d')} \quad \textbf{(2-7)}$$

Following Levinson's [243] use of a (continuous) parametric state duration model, Mitchell and Jamieson demonstrated how to find ML solutions for the parameters of any exponential family pmf [247]. We follow this approach and use the two-parameter gamma and inverse Gaussian exponential family discrete pmfs to model the state duration distributions. The gamma distribution is given by:

$$f(x;\alpha,\beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} \quad \textbf{(2-8)}$$

where α and β are the shape and rate parameters respectively. Correspondingly, the (censored) discrete gamma distribution is simply given by:

$$p(x;\alpha,\beta)=\begin{cases}\dfrac{x^{\alpha-1}e^{-\beta x}}{\sum_{x=x_{\min}}^{x_{\max}} x^{\alpha-1}e^{-\beta x}}, & x_{\min} <= x <= x_{\max}\\ 0, & \text{else}\end{cases} \quad \textbf{(2-9)}$$

The censored discrete inverse Gaussian distribution is given by:

$$p(x;\mu,\lambda)=\begin{cases}\dfrac{x^{-3/2}\exp\left(\dfrac{-\lambda(x-\mu)^2}{2\mu^2 x}\right)}{\sum_{x=x_{\min}}^{x_{\max}} x^{-3/2}\exp\left(\dfrac{-\lambda(x-\mu)^2}{2\mu^2 x}\right)}, & x_{\min} <= x <= x_{\max}\\ 0, & \text{else}\end{cases} \quad \textbf{(2-10)}$$

Writing these two-parameter pmfs in the exponential family form gives:

$$p(x;\boldsymbol{\Theta})=\frac{\mathrm{I}_{x_{\min},x_{\max}}}{B(\boldsymbol{\Theta})}\exp\left(-\sum_{p=1}^{2}\theta_p S_p(x)\right) \quad \textbf{(2-11)}$$

where $\mathrm{I}_{x_{\min},x_{\max}}$ is the indicator function which is 1 inside the range [x_{min},x_{max}] and 0 elsewhere, $B(\boldsymbol{\Theta})$ is a normalization constant, $\boldsymbol{\Theta}$ is the vector of natural parameters, and $\mathbf{S}$ is the vector of natural statistics. Mitchell and Jamieson showed that ML estimates for the natural parameters θ_p are given by solving the equations [247]:

$$E_{\theta_p}\left[S_p(x)\right]-E_{np}\left[S_p(x)\right]=0 \quad \textbf{(2-12)}$$

where $E_{\theta_p}\left[S_p(x)\right]$ is the expectation of the p^{th} natural statistic with respect to the exponential family distribution, and $E_{np}\left[S_p(x)\right]$ is its expectation with respect to the non-parametric

distribution. We use Matlab's nonlinear optimization routine (fminsearch) to solve for the ML estimates of the parameters $\boldsymbol{\Theta}^{ML}$. While the distribution parameters are all strictly positive, constrained optimization techniques were unnecessary because the parameters were initialized very close to their ML values, as discussed below.

Maximum likelihood state sequence

Estimation of the ML state sequence is typically achieved using the Viterbi algorithm [240], and similar algorithms have been demonstrated for the EDHMM [248]. We follow the method of Datta et al. and map the Viterbi decoding problem into that of finding the longest path in a directed acyclic graph (DAG) [249].

Discontinuous data segments

We can easily adapt the EDHMM to handle discontinuous segments of data, which can be useful if we wish to exclude certain portions of the data from analysis (for instance during periods of 'desynchronized activity' [129,130,250]). Assuming that we have N_s discontinuous segments of data for which we wash to classify UDS, we define our segmented observations at time sample t within the n^{th} segment to be y_t^n. First, we must replace π (the distribution on the initial latent state variable) with a matrix containing a distribution over the initial latent state variables for each data segment. Next we introduce a set of time-varying state-conditional mean functions, one for each data segment, which are updated according to:

$$\boldsymbol{\mu}_k^n(t) = \frac{\sum_{t'} w(t'-t)\gamma_{t'}(k)\mathbf{y}_{t'}^n}{\sum_{t'} w(t'-t)\gamma_{t'}^n(k)} \qquad \textbf{(2-13)}$$

The same forward-backward procedure can then be used to compute $\gamma_t^n(k)$ and $D_t^n(k,d)$ for each data segment independently. The time-invariant model parameters can be updated by computing expectations with sums over time samples and segments. For example, the updated estimate for the conditional variance of state *k* is given by:

$$\boldsymbol{\Sigma}_k = \frac{\sum_{n=1}^{N_s}\sum_t \gamma_t^n \left(\mathbf{y}_t^n - \boldsymbol{\mu}_k^n(t)\right)^T \left(\mathbf{y}_t - \boldsymbol{\mu}_k^n(t)\right)}{\sum_{n=1}^{N_s}\sum_t \gamma_t^n} \quad \textbf{(2-14)}$$

where N_s is the number of data segments. Computing the most likely state sequence within each data segment can be accomplished using the same Viterbi decoding algorithm to infer the ML state sequence of each data segment independently. It's important to note that this approach implicitly assumes that the first and last state within each segment start and stop at the beginning and end of those segments respectively. These assumptions can be relaxed [244], however when the data segments are long relative to the state durations, the boundary states can be excluded from analysis without substantial loss of data.

Robustness to deviations from the observation model

Large deviations of the data from the observation model can lead to errors in inferring the state sequence, and thus the results should be monitored for such deviations. While the model is not expected to be a perfect description of the data, large sudden changes in the signal properties, such as could be generated by movement artifacts [236], can create situations where the observation likelihood is very small under both state-conditional observation models. In such situations, where the observation is very far from both state-conditional means, the posterior distribution on the hidden state will tend to be dominated by the state with highest variance.

For example, this could lead to cases where the posterior probability is estimated to be higher for state k=2 than for state k=1 (where $\mu_2 > \mu_1$) at a time when $y_t < \mu_1$, if state 2 has a higher variance. Thus, such deviations from the observation model should be identified and these periods can either be excluded from analysis, or the observation likelihoods can be made more robust as follows. First, large deviations from the observation model are detected by monitoring the maximum state-conditional observation likelihood. For times when both states' observation likelihoods are smaller than a predefined threshold we can reassign the observation likelihoods for each state according to Gaussian distributions where the state-conditional variances are constrained to be equal. This produces a more robust model which will always favor the hidden state whose mean is closer to the observed data at times when the data is very far from either state's mean. We choose a threshold for the maximum state-conditional likelihood such that these instances constituted on average 3% of the LFP data.

Alignment of the decoded state transition times

In some cases, such as when the signal features are down-sampled or filtered, it is desirable to consider alignment of the initially decoded state transition times to a separate observation sequence, such as the unfiltered data. This procedure allows for initial classification of the state sequence using signal features with a lower sampling frequency f_s, while subsequent alignment of the state transition times to a signal with higher f_s prevents loss of temporal precision. This can be particularly important for the EDHMM where the complexity of the inference procedure and the Viterbi decoding algorithm both scale with f_s^2 [245,249] for a fixed maximum state duration. Alignment of state transition times to a new observation sequence can also be useful for removing the 'blurring' effects of filtering or smoothing operations on the detected state

transition times. We perform this alignment using a dynamic programming algorithm to find the set of state transition times that maximizes the likelihood of the new observation sequence given the model parameters. Let $\mathbf{Y}' = \{\mathbf{y}'_1, ..., \mathbf{y}'_T\}$ denote the 'new' observation sequence to which we want to align the state transition times. We first estimate the parameters of the state-conditional Gaussian observation distributions for **Y'** using the posterior probabilities $\gamma_t(k) = P(z_t = k \mid \mathbf{Y}, \Phi)$ computed from the EDHMM fit to the initial observation sequence **Y**. Next, let t_i be the time of the transition into the i^{th} state whose latent state variable is given by $k_i = z_{t_i}$, let δ_i be a perturbation on the time of the transition into the i^{th} state taking values in the range $(\delta_{\min}, \delta_{\max})$, and let d_i be the current estimate of the duration of the i^{th} state (such that $d_i = t_{i+1} - t_i$) (Figure 2-1). In the two state case, we can write the log likelihood as a function of the set of perturbations $\{\delta_i\}$ as:

$$\begin{aligned}\mathcal{L}(\delta_1, ..., \delta_N) = {} & \log\left(p_{k_0}(t_1 + \delta_1)\right) + \sum_{t'=t_1+\delta_{\min}}^{t_1+\delta_1-1} \log\left(p(\mathbf{y}'_{t'} \mid z_{t'} = k_0)\right) \\ & + \sum_{t'=t_1+\delta_1}^{t_1+\delta_{\max}} \log\left(p(\mathbf{y}'_{t'} \mid z_{t'} = k_1)\right) + \log\left(p_{k_1}(t_2 + \delta_2 - t_1 - \delta_1)\right) \\ & + ... + \log\left(p_{k_{N-1}}(t_N + \delta_N - t_{N-1} - \delta_{N-1})\right) + \sum_{t'=t_N+\delta_{\min}}^{t_N+\delta_N-1} \log\left(p(\mathbf{y}'_{t'} \mid z_{t'} = k_{N-1})\right) \\ & + \sum_{t'=t_N+\delta_N}^{t_N+\delta_{\max}} \log\left(p(\mathbf{y}'_{t'} \mid z_{t'} = k_N)\right)\end{aligned} \quad \textbf{(2-15)}$$

where we have not included terms which are independent of the δ_i, and where $k_0 = z_1$. If we define $\Gamma_i(\delta_i)$ to be the maximized log likelihood of perturbations $\delta_1, ..., \delta_i$ up to the i^{th} transition then we have the recursion relation:

$$\Gamma_i(\delta_i) = \max_{\delta_{i-1}} \Big(\Gamma_{i-1}(\delta_{i-1}) + \log\Big(p_{k_{i-1}}\left(t_i + \delta_i - t_{i-1} - \delta_{i-1}\right)\Big) + \sum_{t'=t_i+\delta_{\min}}^{t_i+\delta_i-1} \log\Big(p\big(\mathbf{y}'_{t'} \mid z_{t'} = k_{i-1}\big)\Big) + \sum_{t'=t_i+\delta_i}^{t_i+\delta_{\max}} \log\Big(p\big(\mathbf{y}'_{t'} \mid z_{t'} = k_i\big)\Big)\Big) \quad \textbf{(2-16)}$$

By keeping track of the values of δ_{i-1} which maximize Γ_i for each δ_i we can then backtrack the entire ML sequence of transition perturbations, as with the Viterbi algorithm. We only need to include the initialization condition:

$$\Gamma_1(\delta_1) = \log\Big(p_{k_0}\left(t_1 + \delta_1\right)\Big) + \sum_{t'=t_1+\delta_{\min}}^{t_1+\delta_1-1} \log\Big(p\big(\mathbf{y}'_{t'} \mid z_{t'} = k_0\big)\Big) + \sum_{t'=t_1+\delta_1}^{t_1+\delta_{\max}} \log\Big(p\big(\mathbf{y}'_{t'} \mid z_{t'} = k_1\big)\Big) \quad \textbf{(2-17)}$$

Seamari et al. achieved a similar realignment of state transition times by finding local maxima of the rate of change of the signal features in time[236]. Such a procedure will generally produce similar results to the ML realignment procedure presented here for univariate signal features, however additional 'thresholding' parameters may be required [236]. Further, it is unclear how such methods should be extended to multivariate signal features, and depending on the signal feature being used it may be more susceptible to noise.

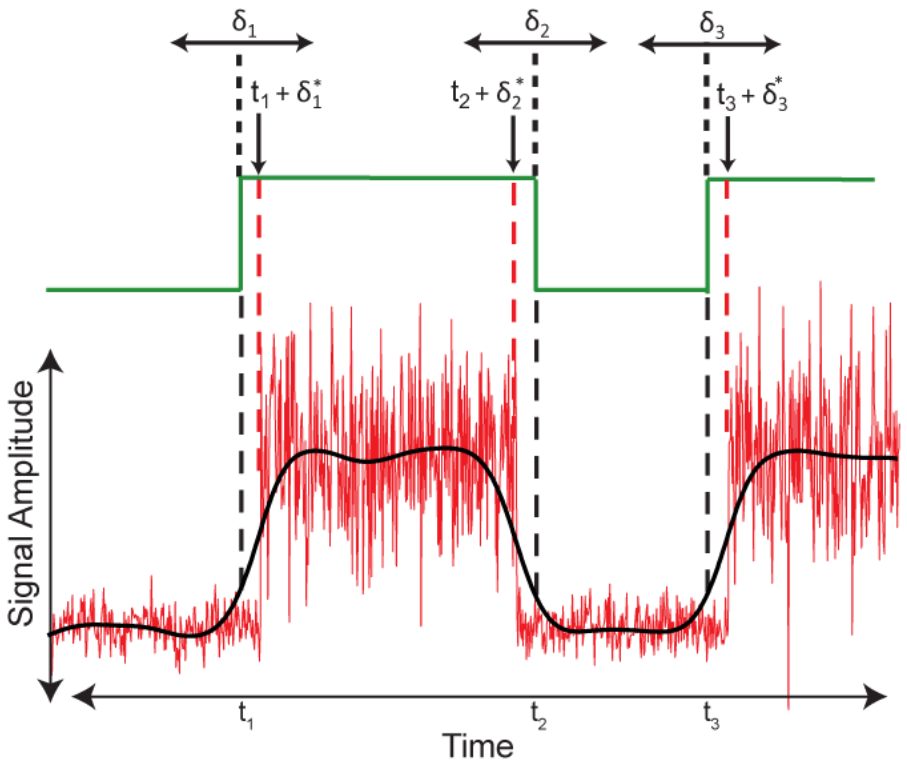

Figure 2-1: Schematic depiction of the maximum likelihood alignment procedure.

An imaginary 'two-state' signal is shown in red. The low-pass filtered (and down-sampled) observation sequence (black trace) is used to produce the initially decoded Viterbi state sequence (green trace) with state transition times t1, t2, and t3 indicated by the vertical black dashed lines. These transition times are aligned to the broadband observation sequence (red trace) by maximizing the likelihood of the broadband observation sequence with respect to a set of perturbations δ_i on the initial state transition times. The maximum likelihood values of the aligned transition times are then given by the set of times $\{t_i + \delta_i^*\}$, indicated by the red vertical dashed lines. For the signal features used in the majority of the analysis, perturbations of up to +/- 150ms were determined to be sufficient; however when exploring a large range of filtering parameters we allowed perturbations of up to +/- 300ms on the state transitions.

EDHMM parameter initialization

It is well known that proper initialization of the model parameters is important to insure that the parameter estimates of the EM algorithm represent a global maximum of the likelihood function [240]. Several steps were thus taken to initialize the model parameters near their ML values. First, the time varying state means were initialized using a sliding-window density

estimator. Specifically, Gaussian kernel density estimates were computed in overlapping, sliding windows of length *L*. The bandwidth of the density estimator was selected using Terrel's over-smoothing bandwidth selector [251]. If the density estimate was bimodal, the two modes were taken as the estimates of the UP and DOWN state means for the given interval. In cases where the density estimate was unimodal, the sign of the skewness of the density was used to determine whether the single mode represented the UP state or the DOWN state (positive skewness indicating the mode represented the DOWN state and vice versa). In such cases, the mean of the other state was taken to be:

$$\boldsymbol{\mu}_k(t) = \boldsymbol{\mu}_{k'}(t) - \frac{\sum_{t' \in \mathcal{T}} \left(\boldsymbol{\mu}_{k'}(t') - \boldsymbol{\mu}_k(t') \right)}{\left| \mathcal{T} \right|} \qquad \textbf{(2-18)}$$

where $\boldsymbol{\mu}_{k'}(t)$ is the mean estimate for the state *k′* representing the mode of the density estimate at time *t,* and the set $\mathcal{T}$ includes all times for which the density estimate was bimodal. Recordings in which the density estimate was never, or very rarely, bimodal (and which still met power spectral criteria for the presence of UP-DOWN states) were not considered here; however other initialization procedures could be introduced in such cases. Initial estimates of the state covariance matrices $\boldsymbol{\Sigma}_k$ were then determined by fitting a Gaussian mixture model to the variable $\mathbf{y}_t - \boldsymbol{\mu}_k(t)$ for each state *k*. Similar results were obtained when using sliding-window Gaussian mixture models to initialize the parameters of the state-conditional observation models, however this was more computationally expensive.

After initializing the observation model parameters as described above, the initial transition matrix **A** was set to:

$$\mathbf{A}^0 = \begin{bmatrix} 1-1/f_s & 1/f_s \\ 1/f_s & 1-1/f_s \end{bmatrix} \quad \textbf{(2-19)}$$

This initializes the expected state durations for both the UP and DOWN states to be 1s under a geometric state duration distribution. A standard HMM was then fit to the observation sequence using this set of initial parameters. Next, the ML state sequence was computed using the standard Viterbi decoding algorithm, and the set of state durations $\{d_i\}_k$ for each state was determined from the Viterbi state sequence. Initial values for the parameters of the state duration distributions were then computed from the Viterbi state durations $\{d_i\}_k$ of the HMM. For the gamma distribution, ML estimates of the state-dependent shape parameters α_k were determined numerically using the Newton-Raphson method [252]. The rate parameters β_k were then given analytically by:

$$\beta_k^{ML} = \frac{\alpha_k^{ML}}{\bar{d}_k} \quad \textbf{(2-20)}$$

where $\bar{d}_k$ is the sample mean duration of state *k*. For the inverse Gaussian distribution, ML estimates of the model parameters are given by:

$$\mu_k{}^{ML} = \bar{d}_k \quad \textbf{(2-21)}$$

$$\lambda_k^{ML} = \left(\frac{1}{N_k} \sum_{i=1}^{N_k} \left(\frac{1}{d_k^i} - \frac{1}{\mu_k^{ML}} \right) \right)^{-1} \quad \textbf{(2-22)}$$

Where N_k is the number of occurrences of state *k*. The ML estimates of the observation distribution parameters and the state duration distribution parameters determined from the HMM were used to initialize the parameters of the EDHMM.

EDHMM Method Summary

In summary, the entire EDHMM UDS inference algorithm proceeds as follows (Figure 2-2). First, all segments of data meeting the criteria for the presence of UDS were extracted. Then, signal features (the observation sequence) were calculated and down-sampled to a sampling frequency of 50Hz within each segment. Initial estimates of the time-varying state means were computed using the sliding-window kernel density estimator, and Gaussian mixture models were then used to initialize the state-conditional covariance matrices. Next, ML estimates of the parameters of a standard HMM were determined, and were used to compute the Viterbi state sequence of the HMM. The state durations of the Viterbi state sequence were used to estimate the ML parameters of the state duration distribution models. These, along with the ML estimates of the observation model parameters and the matrix of initial state probability distributions for all data segments from the HMM, were used to initialize the parameter vector for the EDHMM. Using a maximum state duration of 30s, only a few iterations of the EM algorithm were typically sufficient to achieve convergence of the log likelihood with a tolerance of $1x10^{-5}$ for the EDHMM. The most likely state sequence under the EDHMM was then determined, and subsequent alignment of the state transition times to the broadband signal sampled at 252Hz was performed by maximizing the observation likelihood with respect to a sequence of perturbations on the transition times. The full algorithm was found to take on average 5s per minute of raw data for UDS inference from scalar signal features (run using Windows XP 64-bit with an Intel Core2 Quad 2.40GHz processor with 4GB of RAM). When using the simpler HMM, without explicit state-duration modeling, the algorithm took about 1s per minute of raw data on average.

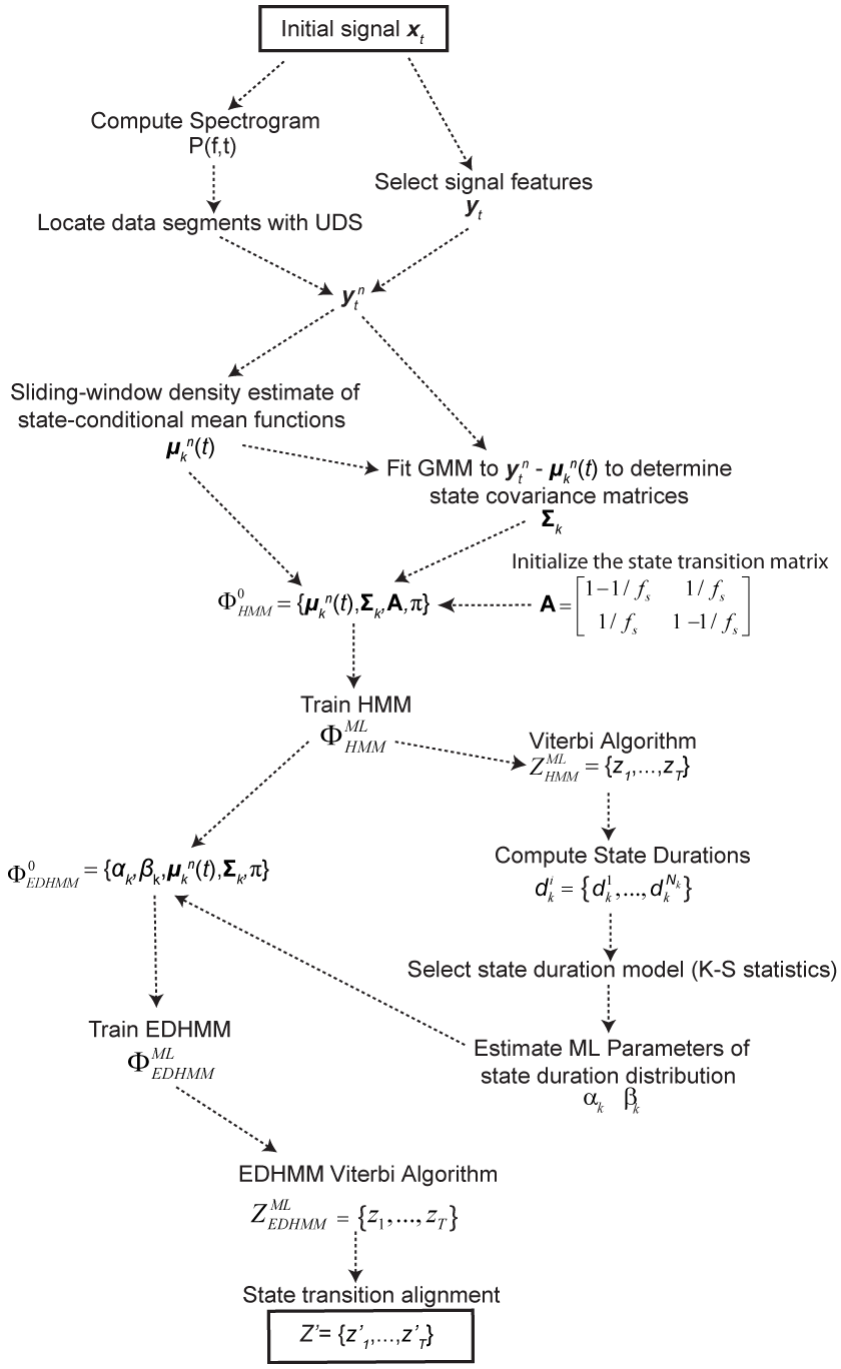

Figure 2-2: Flow diagram of the EDHMM inference algorithm.

The full procedure for UDS inference from continuous signal features is depicted schematically. The initial input to the algorithm is the raw signal $\boldsymbol{x}_t$*, while the final output of the algorithm is the Viterbi state sequence of the EDHMM* $\boldsymbol{Z'}$*, aligned to the broadband signal.*

Discrimination information (signal separability)

The Kullback-Leibler divergence between two Gaussian distributions is given by [253]:

$$\mathrm{D}_{\mathrm{KL}}\left(\mathcal{N}_1 \| \mathcal{N}_2\right)=\frac{1}{2}\left(\log\left(\frac{\det \boldsymbol{\Sigma}_2}{\det \boldsymbol{\Sigma}_1}\right)+\mathrm{tr}\left(\boldsymbol{\Sigma}_2^{-1}\boldsymbol{\Sigma}_1\right)+\left(\boldsymbol{\mu}_2-\boldsymbol{\mu}_1\right)^{\mathrm{T}}\boldsymbol{\Sigma}_2^{-1}\left(\boldsymbol{\mu}_2-\boldsymbol{\mu}_1\right)-\mathrm{p}\right) \quad \textbf{(2-23)}$$

where $\boldsymbol{\mu}_\mathbf{k}$ and $\boldsymbol{\Sigma}_\mathbf{k}$ are the mean and covariance matrix of the k^{th} Gaussian component. This quantity provides the expected 'discrimination information' for the 1st Gaussian component over the 2nd. We use the symmetrized Kullback-Leibler divergence: $\mathrm{D}_{\mathrm{KL}}\left(\mathcal{N}_1 \| \mathcal{N}_2\right)+\mathrm{D}_{\mathrm{KL}}\left(\mathcal{N}_2 \| \mathcal{N}_1\right)$ as a measure of the separability of the two component Gaussian distributions. For comparing Gaussian distributions with time-varying means we substitute the time-averaged Mahalonobis distance in the equation for the Kullback-Leibler divergence (averaging across discontinuous data segments):

$$\frac{\sum_{n=1}^{N_s}\sum_{t}\left(\boldsymbol{\mu}_2^n(t)-\boldsymbol{\mu}_1^n(t)\right)^T\boldsymbol{\Sigma}_2^{-1}\left(\boldsymbol{\mu}_2^n(t)-\boldsymbol{\mu}_1^n(t)\right)}{\sum_{n=1}^{N_s}\sum_{t}1} \quad \textbf{(2-24)}$$

Results

Non-stationarities in the UDS:

It has been well-documented that the UDS seen under anesthesia are sometimes interrupted by so-called desynchronized epochs [129,130,250], which are periods lacking clear UDS. Thus, for all our analysis we first located the epochs of data that contained clear UDS based on the spectral properties of the signal (Figure 2-3A,B). The spectrogram of the z-scored signal was computed in 15s overlapping windows using multi-taper methods (Chronux Matlab toolbox

[254]; http://chronux.org) with a time-bandwidth product of 4, and 7 tapers [31]. The maximum power in the range 0.05-2Hz ('UDS power') was then computed for each signal, along with the integral of the log power between 4 and 40Hz ('reference power'). These statistics were used to test for the presence of clear UDS in each segment. A single-threshold value was set for the 'UDS' and 'reference' power based on visual inspection across recording sessions (Figure 2-3A,B). Any data segments which had UDS power below this threshold and reference power above the threshold were excluded from analysis as desynchronized epochs. The reference power criterion insured that any segments of data which had low power in the UDS band because of very long DOWN states were not excluded from UDS analysis, since these segments would show a corresponding reduction in high-frequency power. All subsequent analysis was performed only for those segments of data containing clear UDS (see *Discontinuous data segments*).

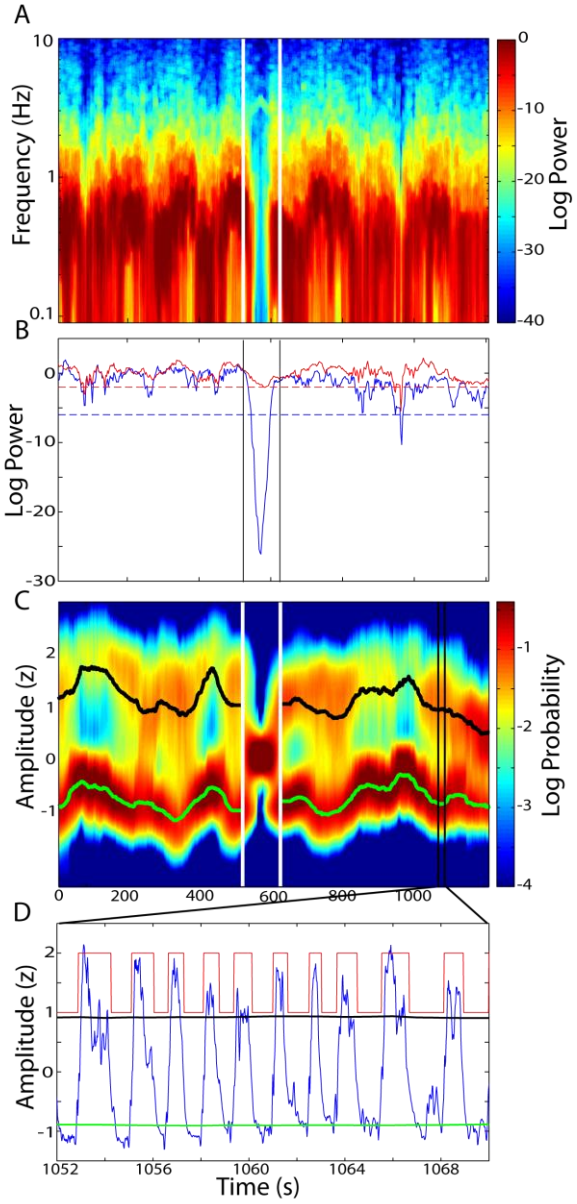

Figure 2-3: Procedure for selecting epochs with clear UDS.

A) *The spectrogram of a cortical LFP. White vertical lines indicate a period of desynchronized activity.* ***B)*** *The maximum UDS power (blue trace) and high-frequency power (red trace) are extracted from the spectrogram. Threshold values (dashed lines) for each of these statistics are used to locate desynchronized epochs (highlighted by the vertical lines).* ***C)*** *The sliding window density estimate of the low-frequency (0.05-2Hz) LFP amplitude along with the time-varying UP (black) and DOWN (green) state-conditional means estimated from the HMM.* ***D)*** *An example LFP trace taken from the region of the recording indicated by the black lines. The black and green traces again represent the time-varying UP and DOWN state-conditional means. The red trace shows the estimated UDS state sequence.*

Initially, we applied a Gaussian-observation HMM with two hidden states to classify UDS using the amplitude of the low-frequency (0.05-2Hz) LFP [147,148]. The distribution of this signal feature ('observation') within each hidden state (the UP and DOWN states) was thus modeled as Gaussian (see *Hidden Markov Model of UP and DOWN States*). When classifying UDS over long duration recordings (~20 minutes) we found that the amplitudes of the UP and DOWN states could vary substantially over the course of the recording [236,237] (Figure 2-3C). This could arise from actual changes in the amplitude of the UDS, as well as filtering artifacts of the AC-coupled amplifiers that remove slowly varying or constant signal components. To account for the variations in the UP and DOWN state amplitudes, we introduced time-varying state-conditional means into the model by using a sliding-window estimate of the state-conditional mean amplitudes (see *Hidden Markov Model of UP and DOWN States*). A window length of 50s was found to provide a good tradeoff between temporal resolution and robustness. Thus, the estimated state-conditional means of the UP and DOWN states at a given time were computed using the 50s of data surrounding the time point.

Choosing a state duration distribution model:

The distributions of UP and DOWN state durations obtained from the HMM model were very far from the geometric distribution implicitly assumed by the HMM (Figure 2-4). Thus, in order to determine an appropriate choice of duration distribution model, we computed the goodness of fit (Kolmogorov-Smirnov statistic) for several exponential family distributions including the gamma distribution and inverse Gaussian distribution, as well as the geometric distribution [231]. Both the gamma (DOWN: 0.15, 0.14-0.16 (median, inter-quartile range); UP: 0.074, 0.056-0.091; n=21 LFPs) and the inverse Gaussian (DOWN: 0.12, 0.10-0.14; UP: 0.062, 0.046-

0.071) distributions produced far better fits than the geometric distribution (DOWN: 0.28, 0.24-0.31; UP: 0.38, 0.36-0.40) for both the DOWN (gamma: p=6.9x10^{-5} (two-sided Wilcoxon signed rank test); inverse Gaussian: p=6.0x10^{-5}) and UP (gamma: p=6.0x10^{-5}; inverse Gaussian: p=6.0x10^{-5}) state duration distributions (Figure 2-4E,D). Furthermore, the inverse Gaussian distribution produced better fits than the gamma distribution for both the DOWN (p=1.2x10^{-4}), and the UP (p=0.014) state duration distributions (Figure 2-4F). Thus, to improve upon the implicit use of the geometric state duration model, we use an EDHMM with an inverse Gaussian distribution model for both the UP and DOWN state durations.

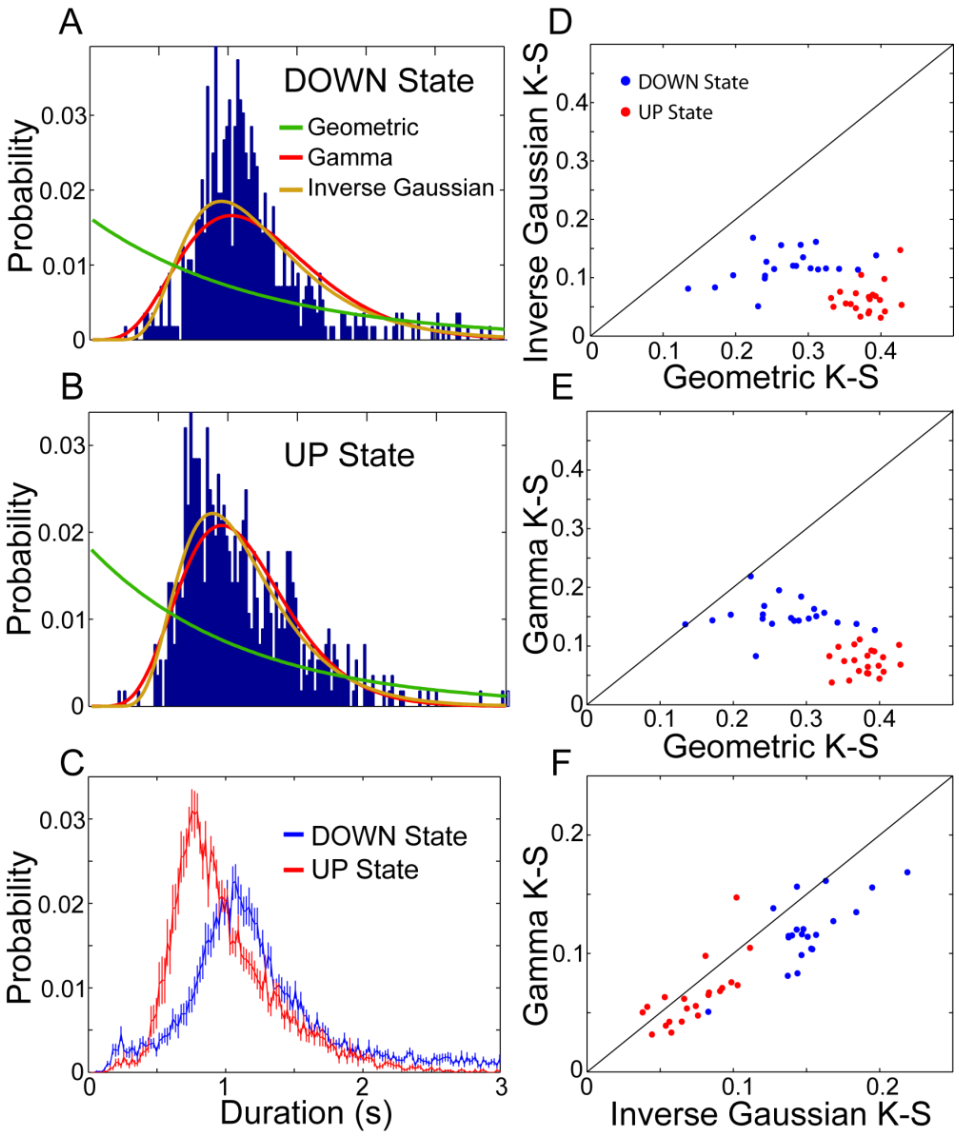

Figure 2-4: Explicit models for UP and DOWN state duration distributions.

A) *The distribution of DOWN state durations inferred by the HMM algorithm for an example LFP recording. The green trace shows the geometric distribution fit. The red and orange traces are the fits for the gamma and inverse Gaussian distributions respectively.* ***B)*** *Same as A for the UP state durations.* ***C)*** *Average DOWN (blue) and UP (red) state duration distributions across all (n=21) LFP recordings. Error bars indicate mean±SEM.* ***D)*** *The Kolmogorov-Smirnov (K-S) statistics for the geometric distribution fits are plotted against the corresponding K-S statistics for the inverse Gaussian distribution fits across all LFP recordings for the DOWN (blue dots) and UP (red dots) state duration distributions.* ***E)*** *Same as D for the gamma distribution.* ***F)*** *Comparison of gamma distribution and inverse Gaussian distribution K-S statistics for the UP and DOWN state duration distributions.*

LFP feature selection

Thus far we have only considered inferring UDS from the low-frequency amplitude of the LFP, but previous work has shown that the signal power in the 20-80Hz range can be more effective for inferring LFP UDS [234]. Thus, we sought to compare results obtained using the low-frequency amplitude (LF-amplitude) and high-frequency power (HF-power). Furthermore, classification results depend on the preprocessing steps used to compute the LF-amplitude and HF-power of the LFP. Thus, if we believe (whether because of evidence accumulated from intracellular recordings, or because we have some knowledge about the underlying mechanisms governing state transitions) that rapid (e.g. <200ms) fluctuations in the signal features are unlikely to represent transitions between states, we can apply this prior information by appropriately bandlimiting the signal features.

To obtain the 'instantaneous' high-frequency power we convolved the squared amplitude of the filtered signal with a Gaussian smoothing kernel. This smoothed power was then log-transformed in order to normalize the state-conditional distributions. We varied the high-cutoff frequency (HCF) and Gaussian smoothing sigma of the LF-amplitude and HF-power respectively to control the bandwidth of the signal features. Since such low-pass filtering of the signal features will also have a 'blurring' effect on the precise state transition times, we aligned the initially detected state transition times to a broadband signal by optimizing a sequence of perturbations on the state transition times (Figure 2-1), analogous to the procedure used by Seamari et al [236]. This also allowed for more direct comparisons of the distributions of detected state durations when using various filtering/smoothing parameters.

First, we compare the state duration distributions obtained from an HMM fit to these different signal features. Figure 2-5A,C show that as the HCF of the LF-amplitude was increased from 0.5 to 10Hz, the state duration distributions became increasingly bimodal with a peak developing at short (< 200ms) durations. These short-duration states represent a deviation from the unimodal state duration distributions that cannot be well-modeled by simple two-parameter distributions, and likely correspond to the spurious detection of state transitions. Furthermore, classification based on the HF-power produced far more short-lived states than when using the LF-amplitude, even when the smoothing sigma was as large as 100ms (Figure 2-5B,D). These results suggest that the frequency content of the signal features used for classification should be restricted to prevent excessive detection of short-lived states; however, when the signal features were excessively bandlimited we overly restricted the range of state durations we were able to detect (Figure 2-5E,F).

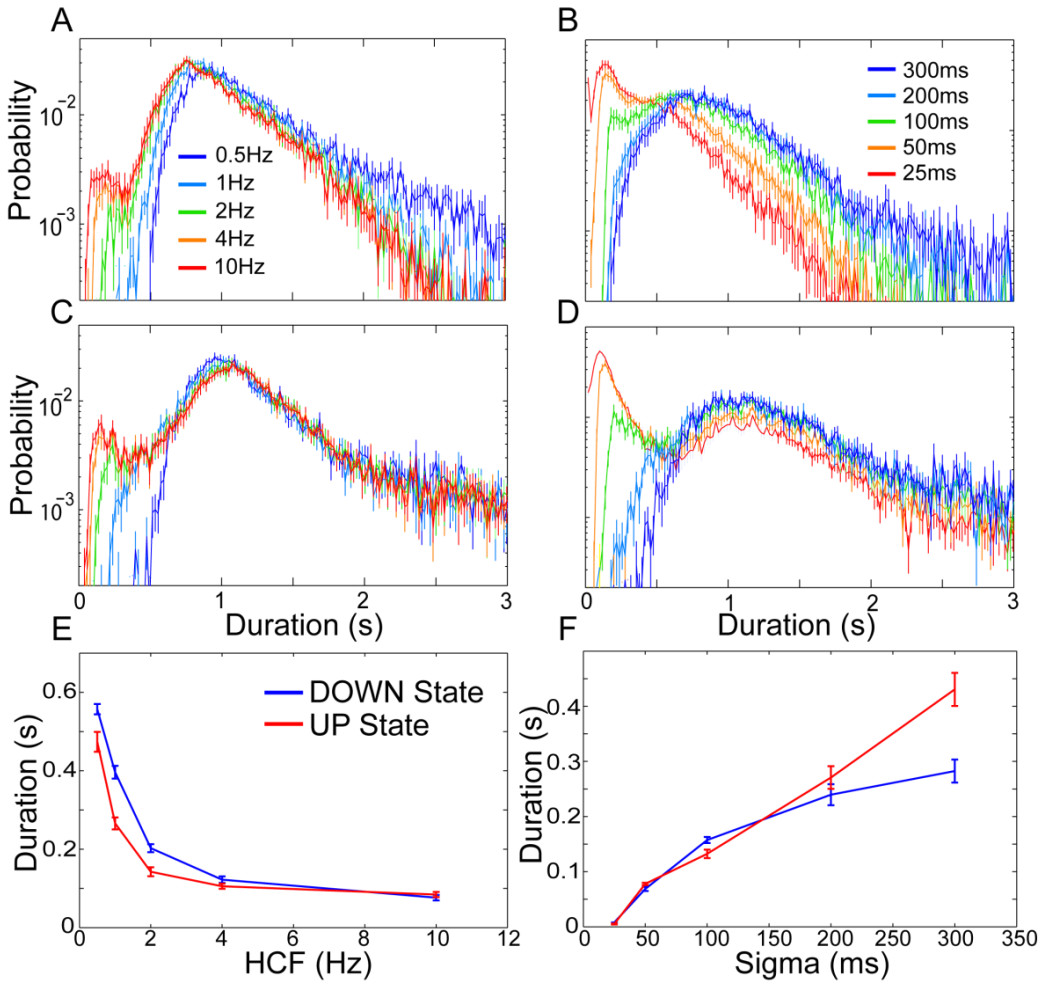

Figure 2-5: Dependence of UDS classification on signal preprocessing.

A) *The average UP state duration distribution for HMM classification based on the filtered LFP amplitude is plotted for several different values of the high-cutoff frequency (HCF). As the HCF increases from 0.5Hz to 10Hz the UP state duration distribution develops a secondary peak at short durations (< 200ms).* ***B)*** *Same as A for inference based on high-frequency LFP power using several different values of the Gaussian smoothing sigma. In this case the short-duration peak is more pronounced than for classification of the filtered LFP amplitude, and it appears for smoothing windows as large as sigma=100ms.* ***C-D)*** *Same as A-B for the DOWN state duration distributions, showing similar effects of filtering and smoothing on the detected state durations.* ***E)*** *The minimum DOWN (blue) and UP (red) state durations, averaged across recordings (n=21), is plotted as a function of the HCF. For HCFs as low as about 2Hz, UP and DOWN state durations as short as 200ms are still detected.* ***F)*** *Same as E for the high-frequency power based classification. As the smoothing sigma is increased above about 150ms the minimum state duration exceeds 200ms.*

We thus chose HCF = 2Hz and sigma = 150ms to provide a good compromise between minimizing the presence of spuriously detected states while maximizing the range of 'allowed' state durations. Other authors have used threshold minimum state durations, typically in the range 100-200ms, to avoid detecting such spurious state transitions [172,231,233,234,235,236]; however such thresholds can produce ambiguous state sequences (if multiple states with subthreshold duration occur in sequence). Further, when applied in the EDHMM framework, such thresholds can produce a large number of states whose duration is exactly equal to the minimum allowed duration. Our preprocessing of the signal features imparts a 'prior' bias against overly short state durations without imposing a hard threshold, thus avoiding these problems. It is worth noting at this point that the assumption of conditional independence of the observation sequence used in the HMM (and EDHMM) is clearly not strictly valid, particularly for the bandlimited signal features. The auto-regressive HMM (ARHMM) relaxes this assumption by modeling the signal features with state-dependent autoregressive models [255]. We found that despite this assumption, the HMM produced better results than an ARHMM for UDS classification (results not shown), likely because the UP and DOWN states are better distinguished by their state-conditional distributions than by their state-conditional autocorrelation functions.

Next, we computed the amount of discrimination information in the LFP LF-amplitude and HF-power, where discrimination information was measured by the symmetrized Kullback-Leibler divergence between the state-conditional Gaussian observation distributions computed for the EDHMM (see *Discrimination information (signal separability)*). This provides a measure of how much information each signal feature will provide on average about the underlying UDS state

sequence. Equivalently, this measure quantifies how 'separable' the state-conditional observation distributions are. We use the time-varying state-conditional means when estimating the discrimination information so that variations in the state means do not decrease the apparent discrimination information.

We found that the LF-amplitude provided significantly more discrimination information about the LFP state than the HF-power (LF: 23.8, 20.4-24.7; HF: 15.8, 11.1-22.9; p=0.013), suggesting that for our data the LFP LF-amplitude was a more effective signal feature than the HF-power for inferring UDS. The HF-power also provided significantly more variable discrimination information across LFP recordings than the LF-amplitude (standard deviations: LF-amplitude = 5.3, HF-power = 11.4; $p=1.1\text{x}10^{-3}$, two-tailed F-test for equal standard deviations). We also computed the joint discrimination information in the combined LF-amplitude and HF-power signals (LF+HF: 23.8, 21.1-26.3), and found that there was only a small (median 3.9%), but significant ($p=5.7\text{x}10^{-3}$), increase compared to the discrimination information in the LF-amplitude alone, suggesting that the two signals carry largely redundant information about the underlying state sequence.

Certainly, more 'elaborate' signal features, such as time-frequency representations (e.g. coefficients of the continuous wavelet transform) could also be used; however, because the differences in the state-conditional power spectra are largely redundant across frequencies [234], such possibilities were not explored here. We emphasize however that the EDHMM framework can be applied to any signal features which can be well modeled with state-conditional Gaussian (or mixture of Gaussians) distributions. Thus, for a given data set, various signal features should be evaluated to determine whether there is an appropriate two-state

mixture distribution, and whether there is sufficient separability between the state-conditional distributions to perform robust inference of UDS.

Comparison to threshold-crossing approaches

The results of the EDHMM UDS classification method are compared with those of a 'fixed threshold-crossing' (TC) approach, using the LF-amplitude (0.05-2Hz). The fixed threshold was selected for each LFP recording using either a 'static mixture model' (SMM) method, or a 'nonparametric' method. For the static mixture model, we selected a threshold by fitting a two-component Gaussian mixture model to the signal feature. The threshold was then chosen to be the value of the signal in the range (μ_1, μ_2) where the two mixture components had equal posterior probability. For the 'nonparametric' approach, the threshold was chosen as the location of the minimum of a Gaussian kernel density estimate of the signal in the range (μ_1, μ_2), similar to the method used by Mukovski et al [234]. In this case, the state means were again taken from the two-component mixture model fit. Given the threshold value, the threshold-crossing times were then used to identify UP and DOWN states.

While for the most part the UDS can be inferred accurately using the TC methods, non-stationarities in the data, and ambiguous, intermediate-amplitude signal features can pose substantial problems. Figure 2-6 illustrates some potential problems with the TC approach for an example LFP signal. As shown in Figure 2-6C, at around 800s there was a period with increased probability of the DOWN state which was also accompanied by an increased mean amplitude of the DOWN state, largely resulting from high-pass filtering of the long DOWN states by the AC-coupled amplifiers. The threshold value selected by the SMM method was found to

be substantially too low during this period (Figure 2-6C), such that a series of erroneous UP states were detected while the signal was apparently remaining in the DOWN state. If the threshold value was instead selected using the nonparametric method, this issue could be largely avoided, however this choice of threshold created similar problems at other times within the same recording, where fluctuations within the UP state repeatedly triggered DOWN state transitions (Figure 2-6D). Thus, the problem with the TC methods is often not one of selecting the 'appropriate' fixed threshold value, but rather that no single threshold value exists which can provide robust separation of the UP and DOWN states across the entire recording. Hence, while we found that a TC approach gave results which were mostly in agreement with those of the EDHMM method, the advantages of using the EDHMM inference procedure were particularly important when substantial non-stationarities of the signal features were present, or when the UP and DOWN state-conditional distributions were not well-separated.

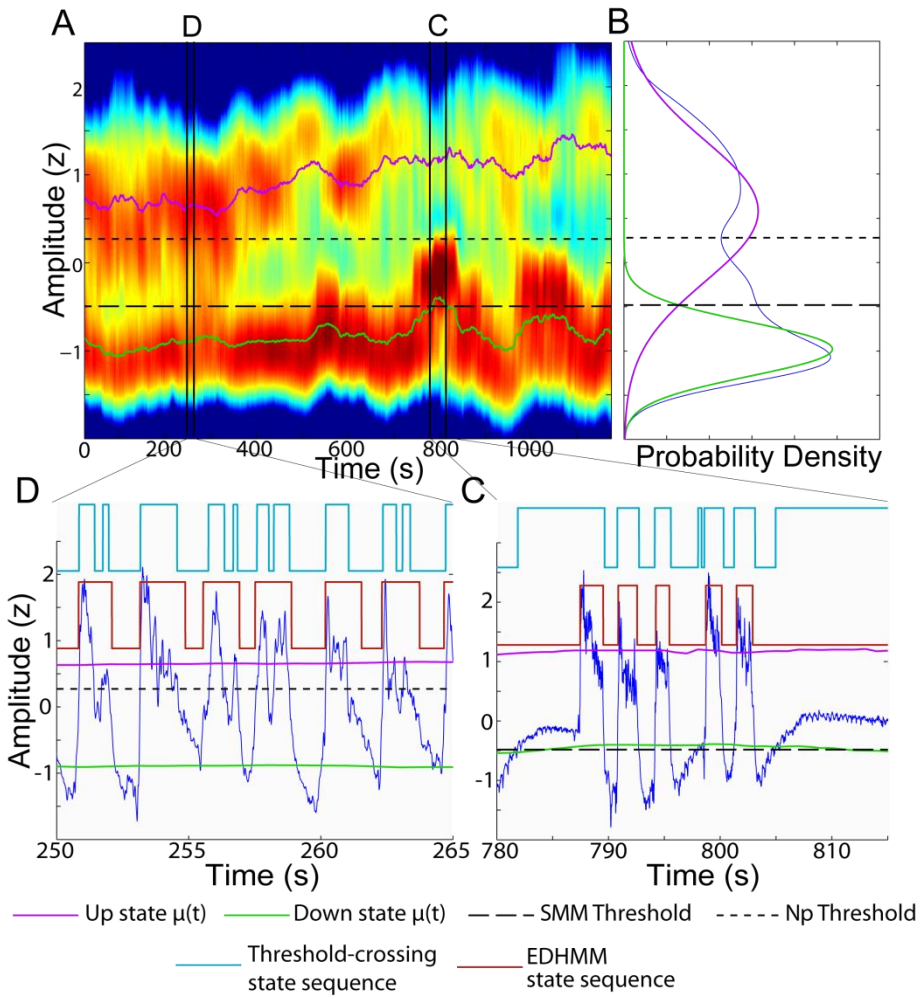

Figure 2-6: Comparison of EDHMM classification and threshold-crossing classification.

A) *The sliding window density estimate is shown for an example LFP recording. Log probability density is depicted as the color map, with the violet and green traces showing the time-varying state-conditional means. The regions 'C' and 'D' indicated by the vertical black lines show the locations within the data where the example traces in panels C and D were taken.* ***B)*** *The overall amplitude distribution, along with the UP and DOWN state component distributions, fit using a (static) Gaussian mixture model. The two values of the fixed threshold used for the 'threshold-crossing' algorithm, chosen using the non-parametric (Np) and static mixture model (SMM) approaches, are shown by the finely and coarsely dashed black horizontal lines respectively.* ***C)*** *An example LFP trace (blue) from the region indicated by the black lines is shown along with the time-varying UP (violet) and DOWN (green) state-conditional means. The Viterbi state sequence from the EDHMM is shown in brown. The state sequence classified using the static mixture model (SMM) threshold-crossing method is shown in light blue.* ***D)*** *Another example LFP trace from earlier in the same recording, comparing the EDHMM UDS classification with the nonparametric (Np) threshold-crossing method (again shown as the light blue trace).*

Evaluating classification accuracy

In order to quantitatively compare the results produced using the various methods and signal features described above, we compared the agreement between state sequences classified from the LFP to those classified from the simultaneously recorded MP of nearby cortical neurons. The state of single cortical neurons is classified relatively unambiguously from the signal amplitude [234,235], and is a reliable indicator of the cortical state since cortical neurons make nearly synchronous state transitions [133]. Thus, we computed the difference between state sequences inferred from MP/LFP pairs (n=9), using the normalized Hamming distance metric (referred to here as the 'error rate'). MP state sequences were computed from the LF-amplitude (0.05-2Hz) after removing spikes from the MP signal. When considering different algorithms (e.g. threshold-crossing vs. EDHMM), identical methods were used to infer UDS in the MP and LFP to insure fair comparisons between methods. Since the error rates themselves varied substantially across data sets (for the EDHMM based on LF-amplitude error rates ranged from 4.2% to 10.1%, with median 8.8% and IQR 6.6-9.1%), we report error rates of the various algorithms as percent changes relative to the error rate of the EDHMM with LF-amplitude (positive relative error rates thus signify poorer performance compared to the EDHMM algorithm with LF-amplitude). Two-sided Wilcoxon signed rank tests were used to test for significant differences in the median error rates.

First, we hypothesized based on the previous analysis that we would achieve better agreement between MP and LFP UDS when using the LF-amplitude of the LFP, compared to the HF-power. This was confirmed by the observation that using LFP HF-power increased the error rates substantially (+78%, +36 - +198%; p=4.0x10^{-3}) compared to using LF-amplitude. Furthermore,

using both LF-amplitude and HF-power together did not significantly improve the error rate compared to LF-amplitude alone (+0.67%, -4.2 - +3.7%; p=1.0), also in agreement with the discrimination information analysis. Next, we sought to evaluate the contribution of two key components of our method to the decoding accuracy: the explicit state duration models, and the variable state-conditional means. Thus, we compared error rates obtained using a HMM (no explicit duration models), and a 'fixed-mean EDHMM' (where the state-conditional means were constrained to be constant), with those obtained using the unconstrained EDHMM. Somewhat surprisingly, we found that neither constraining the state means to be constant (+3.4%, -1.3 - +9.3%; p=0.12), nor using the simpler HMM (+0.10%, -0.32 - +0.98%; p=0.82) produced a significant increase in the error rate compared to the unconstrained EDHMM. These results suggest in particular that the explicit state-duration models may contribute relatively little to accurate inference of UDS, even though the inverse Gaussian state-duration distributions were found to be much better fits than the geometric distributions implicitly assumed by the HMM. Thus, the much more computationally efficient HMM algorithm can be used in place of the EDHMM without a significant loss of accuracy.

Next, we compared the results of the EDHMM classification method with the threshold-crossing approach, again considering two different procedures for determining the threshold: the SMM and nonparametric methods. We found that, as predicted, TC algorithms produced higher error rates than the EDHMM method (Figure 2-7) when the threshold was selected using either the SMM method (+10%, +6.5 - +34%; p=0.012), or the nonparametric approach (+9.6%, +5.2 - +31%, p= 0.012). Even the computationally faster HMM significantly outperformed both the SMM TC (p=7.8x10^{-3}), and the nonparametric TC (p=0.012) methods. Thus, we conclude that

using the EDHMM method presented here, as well as the simpler HMM method, provides significant improvements over previously demonstrated approaches when inferring cortical UDS from LFP signals.

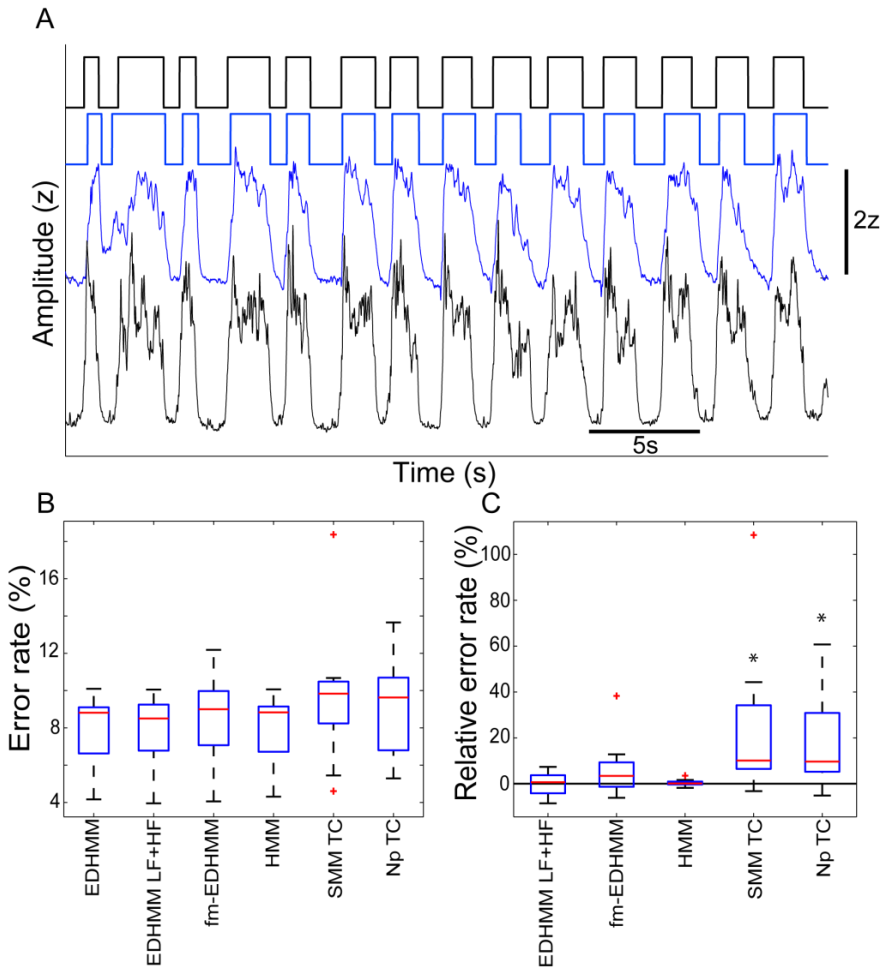

Figure 2-7: EDHMM classification produces improved agreement between simultaneously recorded LFP and MP signals.

A) *Example of an LFP (blue trace) recorded simultaneously with the MP (black trace) of a nearby cortical neuron. The corresponding LFP and MP state sequences are shown above.* ***B)*** *Box plots illustrating the distributions of error rates (n=9) for the following decoding algorithms: EDHMM; EDHMM using both LF-amplitude and HF-power (EDHMM (LF+HF)); HMM; fixed-mean EDHMM (fm-EDHMM); static mixture model threshold-crossing (SMM TC); and nonparametric threshold-crossing (Np TC).* ***C)*** *Same as B for relative error rates*

Discussion

We conclude by summarizing several of the key improvements provided by our UDS inference procedure. Firstly, the framework presented here can be used to infer UDS from different types of continuous signals (such as MPs and LFPs), so that we do not need to resort to different procedures when using different types of signals. Our method also allows for UDS inference from combinations of simultaneously recorded signals, so that multi-electrode recordings could be used to analyze global as well as region-specific UDS. Different signal features, or combinations of features (e.g. HF-power and LF-amplitude) can be used as desired in each case. Importantly, regardless of the exact experimental conditions or type of signal(s) used, estimation of the discrimination information provides a natural procedure for selecting a set of signals and/or signal features, without the need to perform calibration experiments of any kind[235]. When combining information from multiple signals and/or signal features, an important strength of a probabilistic generative model such as the EDHMM is that it will automatically account for correlations between the signal components, as well as differences in the information each component provides, when inferring the state sequence. Indeed, we found that the discrimination information provided by the LF-amplitude, and particularly the HF-power, was highly variable across recordings, suggesting a strong advantage for adaptive methods when considering multiple signal features.

The flexibility of the HMM framework presented here avoids the need for subjectively chosen rules or model parameters which are unlikely to be optimal for a broad range of experimental conditions. Our algorithm also effectively handles the non-stationarities present in large datasets by allowing the state-conditional means of the signal features to vary in time, as well as

by allowing inference of model parameters across discontinuous segments of data interrupted by periods of desynchronized activity. Another important feature of the algorithm is that it naturally provides a direct measure of uncertainty for the inference results in terms of the posterior distribution over the latent state variables. This could allow for identification of state 'transition regions', as well as the selection of particular segments of data for analysis where the estimated uncertainty about the hidden state sequence is sufficiently low (e.g. clear individual UP and DOWN states, or data epochs with particularly well-defined UP and DOWN states).

We use simultaneous LFP-MP recordings to confirm the benefits of these features, showing that our method produces small but significant decreases in the 'error rate' of the UDS inferred from the LFP. These results suggest that the EDHMM method outperforms threshold-crossing type methods regardless of the procedure used to select the fixed threshold. Interestingly, our results also indicate that the faster HMM algorithm could be used instead of the EDHMM algorithm without significant loss of accuracy, even though the EDHMM uses a more appropriate model of the UP and DOWN state durations. It is worth noting that the strength of this 'MP-LFP comparison' metric in quantifying the accuracy of a particular algorithm for inferring LFP UDS is limited by several factors, including the instantaneous variability of UDS across neurons, as well as any ambiguity in inferring UDS from the MP. In particular, previous work has shown that UDS decoded from the LFP can produce better agreement with UDS inferred from the MP of a simultaneously recorded neuron than the agreement between the UDS of two simultaneously recorded MP signals [234]. Such limitations could explain the apparent lack of improvement in the EDHMM over the HMM method, or the 'fixed-mean'

EDHMM. Thus, in addition to this quantitative measure of the accuracy of LFP UDS inference, we also emphasize other qualitative strengths of the methods presented here. Finally, we note that our EDHMM method could also be extended to include both continuous and point-process observations to allow UDS inference based on extracellular spikes as well as LFP signals, as done in [235] using a threshold-crossing type algorithm, and as suggested by [231] within the EDHMM framework.

Chapter 3: **Region specific, spontaneous persistent activity in the entorhinal cortex *in vivo***

Introduction

The entorhinal cortex (EC) provides the main source of inputs to the hippocampus and is bidirectionally connected with the neocortex (Ncx), thus functioning as a gateway between the Ncx and hippocampus [180,181]. Theoretical and experimental work suggests that cortico-hippocampal interactions are required for long-term memory consolidation [165,166,173], thought to occur during sleep [114]. In support of these ideas, lesions of the direct inputs from layer 3 EC (ECL3) neurons to the hippocampus dramatically impair consolidation [256]. The mechanisms by which ECL3 neurons shape cortico-hippocampal interactions, however, are not well understood.

To this end, we measured the MP of these neurons in urethane anesthetized mice, along with the neocortical (Ncx) local field potential (LFP) from posterior parietal cortex. The Ncx LFP shows UDS oscillations similar to those seen during SWS [123,129,132,257], and provides a robust estimate of the Ncx inputs to ECL3 neurons [181]. Layer 3 neurons from the medial and lateral EC (MECL3 and LECL3) showed UP-DOWN states (UDS) that were phase-locked with the Ncx UDS, however MECL3 neurons nearly always persisted in the UP state well beyond the Ncx UP state, often lasting for several cycles of the Ncx UDS. Despite decoupling from the afferent Ncx UDS, the duration of these MECL3 persistent UP states was quantized in units of the Ncx

UDS cycles. These persistent UP states were not seen in LECL3 neurons, and could not be evoked by current injection. Further, the decoupled UDS of MECL3 neurons was found to influence the downstream hippocampal LFP. These results provide a major dissociation between the spontaneous activity of MECL3 and LECL3 cells, and support growing evidence of their differential functions [258]. The quantized persistent activity of MECL3 neurons could play an important role in mediating cortico-hippocampal interactions during sleep, and may be critical for EC-dependent long-term memory consolidation.

Methods

Experimental Methods

Methods were as described in Chapter 2: *Experimental Methods*. Data were obtained from 32 C57BL6 mice aged postnatal day (p)26 to p42 (p31.4$\pm$0.73, mean$\pm$SEM) weighing between 13.8 and 29.2g (17.6$\pm$0.59g). Mice were anesthetized with urethane (1.2-2.0g/kg i.p.). There were no significant differences between the age, weight, or anesthetic dosage for mice where MECL3 or LECL3 cells were recorded.

Visualization of biocytin filled neurons allowed the determination of cell type, layer and recording site in medial or lateral entorhinal cortex. L3 pyramidal neurons used in this study are distinguished from other excitatory neurons and interneurons by a prominent apical dendrite and characteristic dendritic branching. Neurons that were not clearly identified, or neurons for which the layer determination was ambiguous were excluded from analysis. We further made 3D reconstructions of the dendritic tree and of the initial axonal segment from neurons with a Neurolucida system (Microbrightfield, Williston, VT), from which the insets in Figure 3-3A,B are taken.

Local field potentials (LFP) were recorded with a single-shank multisite probe (NeuroNexus Technologies, Ann Arbor, MI) inserted 2.0mm posterior to bregma and 1.5mm lateral from midline. The most superficial electrode was located in layer 2/3 of posterior parietal cortex while a deeper contact was targeted to the cell body layer of the CA1 region of the dorsal hippocampus. LFPs from frontal and prefrontal cortices (upper layer 5) were recorded with a quartz/platinum-tungsten glass coated microelectrode (Thomas Recording GmbH, Giessen, Germany) inserted 1 to 1.5mm anterior and 1 to 1.5mm lateral (frontal) or around 3 mm anterior and 1 to 1.5 mm lateral (prefrontal) to bregma.

Relative to bregma, the craniotomy for the MP recordings was made around 4.5mm posterior and 4mm lateral for MEC neurons, and 4mm posterior, 4.5mm lateral, and 4mm ventral for LEC neurons. The average initial series resistance was 49$\pm$2.9MΩ. MP values were corrected for the estimated junction potential of approximately +7mV. Both LFP and MP were recorded continuously on an eight-channel Cheetah acquisition system (Neuralynx; Bozeman, MT) for at least 600s per experiment (1160$\pm$41s). That complete recording was used for subsequent analysis as described below.

Data Analysis

Statistics and Hypothesis Testing

All central tendencies and variability are reported as mean$\pm$standard error of the mean (SEM). Unless stated otherwise, all hypothesis tests were performed using two-sided two-sample t-tests, with paired t-tests where applicable. A *p*-value of less than 0.05 was considered to be significant.

Data Preprocessing

All analysis was restricted to subthreshold fluctuations in the MP to obtain the MP in the UDS range that is not influenced by spikes, and has similar spectral content to the LFP. Hence, spikes were removed from the MP in the following way. The MP was bandpass filtered between 100Hz and 8kHz using a 4-pole Butterworth filter. The approximate derivative of this filtered signal was then calculated by differentiation of adjacent values, and converted to units of z-score. Times where the derivate of the filtered signal exceeded a fixed threshold (10z) were taken as spike times. Spikes were then removed from the MP by replacing 3ms of data following the onset of each spike by linear interpolation.

In order to determine the absolute amplitude of the MP signal we first aligned the DC-coupled data (recorded with either Instrutech or CED hardware) with the AC-coupled data (recorded with Neuralynx hardware). Such alignment was necessary because some of the DC-coupled data were recorded in discontinuous sweeps of 7 or 10s separated by 5 or 2s respectively, and thus, UDS detection was performed on the continuously recorded AC-coupled data. To align the signals, both the DC- and AC-coupled MP signals were first high-pass filtered above 2Hz using a 2-pole Butterworth filter. The DC-coupled signal was down-sampled to have the same sampling frequency as the AC-coupled signal (2016 Hz) using linear interpolation. The DC- and AC-coupled MP signals were then aligned by finding the location of the maximal value of the cross-correlogram.

Explicit-duration Hidden Markov Model Detection of UP and DOWN States

Since activity under urethane anesthesia is known to undergo periods of 'desynchronized activity' where UDS are absent (Kasanetz et al., 2002; Clement et al., 2008), we located such

epochs and excluded them from all analysis. This was accomplished by first computing the spectrogram of the signal in 15s overlapping windows using multi-taper methods (Chronux Matlab toolbox [254]; http://chronux.org) with a time-bandwidth product of 4, and 7 tapers [31]. The maximum power in the range 0.05-2Hz and the integral of the log power in the 4-40Hz range were then used to locate desynchronized epochs in the data.

UDS of both MPs and LFPs were detected as described in Chapter 2. Briefly, MP and LFP signals were first filtered in the low-frequency (0.05-2Hz) range using a 4-pole Butterworth filter. We then fit a two-state Gaussian-observation hidden Markov model (HMM) to the filtered signal. Because the amplitudes of the UP and DOWN states could vary substantially over the course of an experiment, we allowed the means of the state-conditional Gaussians to be slowly varying functions of time by introducing time-localized window functions (rectangular windows of length 50s) into the parameter re-estimation procedure. After computing the most likely 'Viterbi' state sequence from the resulting model, we fit inverse Gaussian distributions to the durations of the detected UP and DOWN states. These fits, along with the ML parameter estimates from the HMM, were used to initialize the parameters of an explicit-duration HMM (EDHMM). This model improves on the geometric state duration model implicitly assumed by the standard HMM, which was determined to produce poor agreement with the empirical distributions of both UP and DOWN state durations. After finding maximum likelihood parameter estimates of the EDHMM, the Viterbi state sequence was computed, and was used to define UDS for all analysis. In order to more accurately determine the precise times of state transitions, we aligned these initially detected state transition times to the less-filtered (0.05-20Hz) signal, by maximizing the likelihood of the broadband signal with respect to a set of

perturbations (in the range of -150 to 150ms) on the state transition times. This allowed our initial detection of state transitions to be more robust towards high-frequency fluctuations in the data, while avoiding biases due to the low-pass filtering of the signal. Based on simultaneous cortical MP-LFP recordings, we found that this algorithm produced very good agreement between neocortical MP and LFP UDS (median Hamming distance 8.8%, IQR 6.6-9.1%).

Persistent Activity

Persistent activity was defined in two ways. Type 1 persistent activity was defined as an MP UP state that lasted longer than the concurrent Ncx UP state. Type 2 persistent activity was defined as an UP state of the MP lasting longer than one entire cycle of the Ncx UDS (the concurrent Ncx UP state plus the subsequent Ncx DOWN state). The concurrent cortical UP state was defined by finding the UP transition in the LFP that was nearest in time to the MP UP transition. In order to make our definition of type 2 persistent activity more robust to false detection of brief Ncx DOWN states, we excluded MP UP states which were found to skip any Ncx DOWN state that lasted less than 0.5s (which constituted only 10% of Ncx DOWN states, Figure 3-1). All results were qualitatively very similar without imposing this criterion.

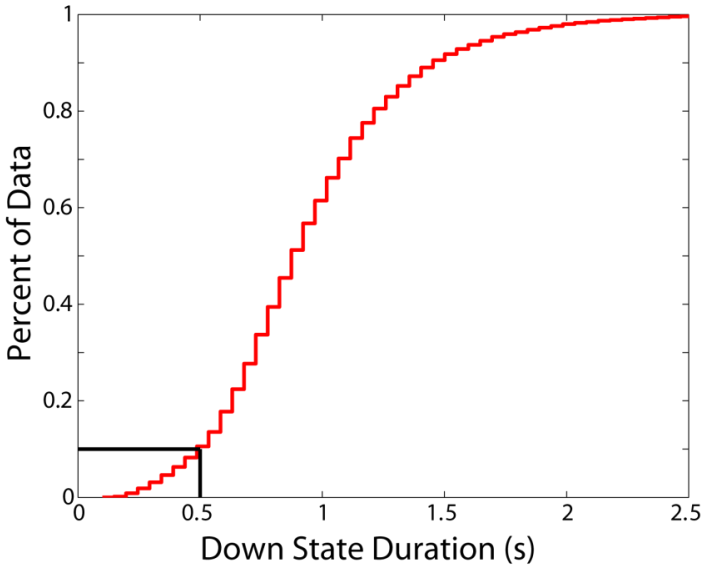

Figure 3-1: Average cumulative distribution of LFP DOWN state durations.

The threshold of 0.5s used to rule out spuriously detected type 2 persistent states was chosen to exclude the shortest 10% of LFP DOWN states.

To compute the duration of each MP UP state in units of Ncx UDS cycles, the LFP signal was first shifted so that the concurrent UP transitions in the MP and LFP were aligned. This removed any fixed temporal delay between the LFP and MP. Each complete Ncx UP or DOWN state was then treated as 0.5 cycles, and the fraction of the next Ncx state completed before the MP DOWN transition was added to get the duration of the MP UP state (Figure 3-2).

MP UP transition lags were computed relative to the concurrent Ncx LFP UP transition, as defined above. MP DOWN transition lags were computed relative to whichever occurred later between the previous LFP DOWN transition, or the first LFP DOWN transition after the concurrent UP transition.

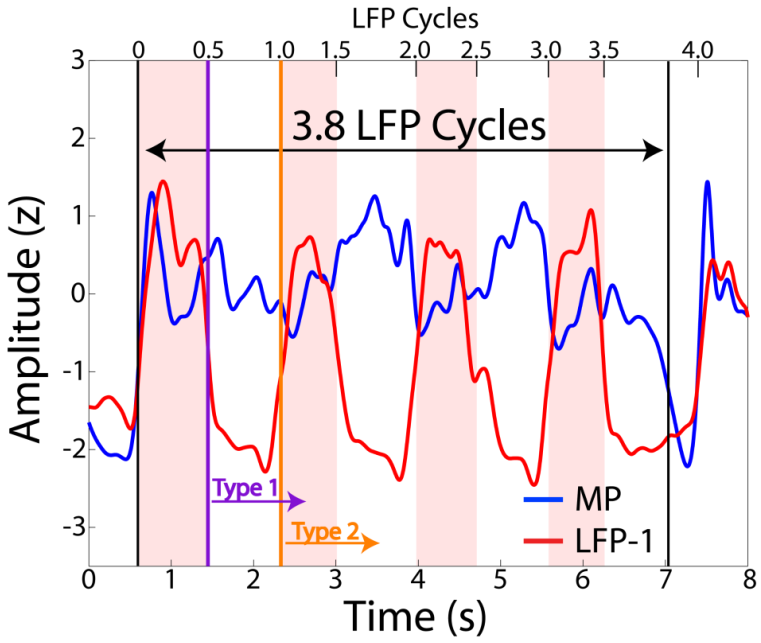

Figure 3-2: Determination of MP UP state duration in units of Ncx UDS cycles.

A single MP UP state is depicted (blue trace) along with the corresponding LFP signal (red) aligned to the MP UP transition (and with 1 standard deviation subtracted for ease of visual comparison). The beginning and end of the MP UP state are denoted by vertical black lines. Individual Ncx UP states are depicted as red shaded regions. In this case, the total duration of the MP UP state was 3.8 NCX UDS cycles. The threshold durations for type 1 and type 2 persistent activity are indicated by the lavender and orange vertical lines respectively.

Spectral Analysis

Power spectra and coherence spectra were estimated using multi-taper methods (Chronux Matlab toolbox [254]; http://chronux.org). Power spectra were computed in non-overlapping 50s windows using a time-bandwidth product of 3 and 5 tapers [31], and the average across time windows and tapers was taken. The coherence spectrum between two signals *x* and *y* was given by:

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)} \quad \textbf{(3-1)}$$

When computing coherence spectra a time-bandwidth product of 7 and 13 tapers was used. Power spectra and coherence spectra estimates were corrected for finite sample size biases according to [259]:

$$\log(\tilde{S}(f)) = \log(S(f)) - \psi(\upsilon/2) + \ln(\upsilon/2) \quad \textbf{(3-2)}$$

$$\tanh^{-1}(\tilde{C}(f)) = \tanh^{-1}(C(f)) - \frac{1}{\upsilon - 2} \quad \textbf{(3-3)}$$

Where $\tilde{S}(f)$ and $\tilde{C}(f)$ are the bias-corrected power spectrum and coherence spectrum, and $\upsilon = 2NK$ (where *N* is the number of data windows, and *K* is the number of tapers) is the degrees of freedom, equal to twice the number of independent estimates of the spectrum [42]. Partial coherence spectra between signals *x* and *y*, controlling for signal *z*, were computed according to:

$$C_{xy|z}(f) = \frac{\left|C_{xy}(f) - C_{zx}(f)C_{yz}(f)\right|^2}{\left(1 - \left|C_{zx}(f)\right|^2\right)\left(1 - \left|C_{yz}(f)\right|^2\right)} \quad \textbf{(3-4)}$$

Results

A recent study showed UDS oscillations in all layers of EC in the urethane-anesthetized rat, although it did not differentiate between lateral and medial EC, and did not document persistent activity [149]. Consistent with this study, the MP of MECL3 (n=22) (Figure 3-3A) and LECL3 (n=14)(Figure 3-3B) pyramidal neurons showed clear, spontaneous UDS oscillations (Table 3-1). While the LECL3 MP showed UDS which were very similar to the Ncx UDS, the MP of MECL3 neurons often persisted in the UP state while the afferent Ncx activity continued to fluctuate between the UP and DOWN states (Figure 3-3C,D).

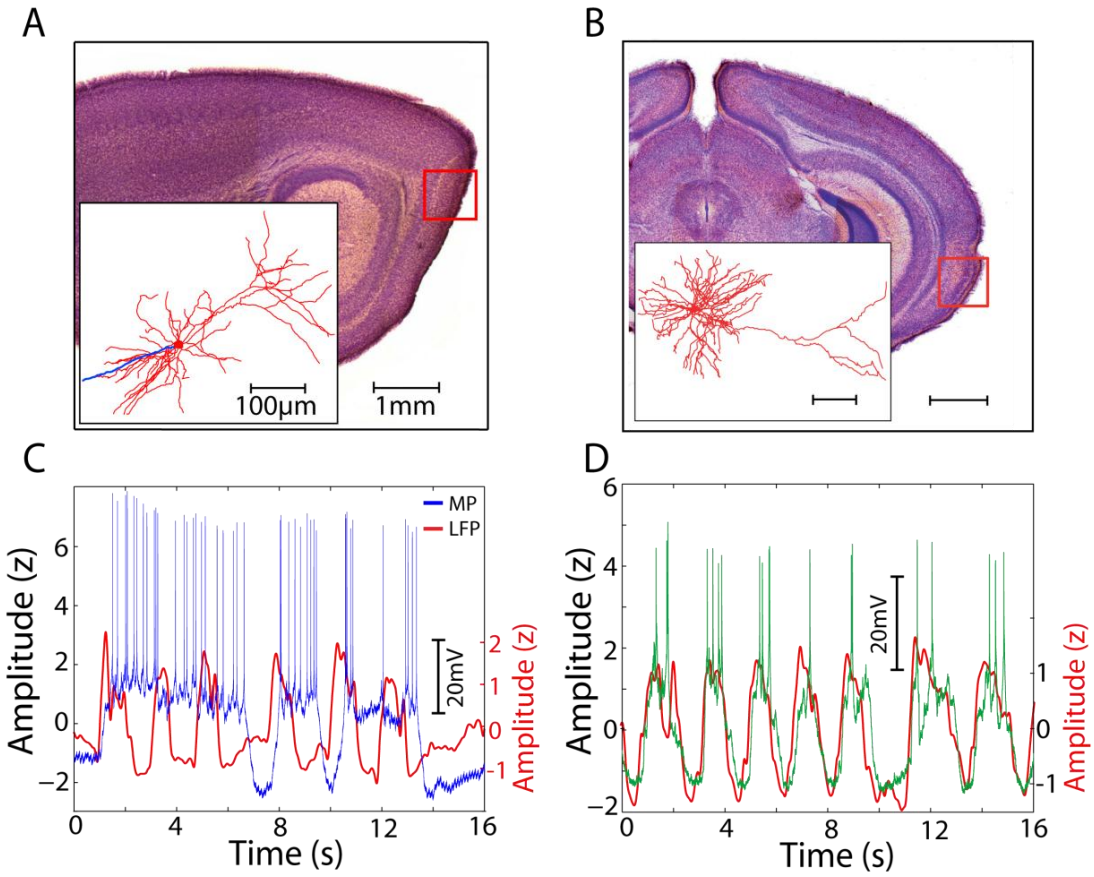

Figure 3-3 UP-DOWN states in example MECL3 and LECL3 pyramidal cells.

A) *Sagittal brain slice showing the region where the example MECL3 neuron was recorded (red box). The white inset shows a 3D reconstruction of the example neuron (dendrites in red, axon in blue)* ***B)*** *Same as A, but with a coronal section showing an LECL3 neuron.* ***C)*** *UDS in the membrane potential of the example MECL3 neuron. Simultaneously recorded Ncx LFP is shown in red. Amplitudes are in units of z-score to allow comparison across cells and with the LFP.* ***D)*** *Same as C, but for the LECL3 cell.*

To allow comparison across different neurons, and between the MP and simultaneously recorded local field potential (LFP), all data were converted to units of z-score. Due to the relative stability of the UDS, and the rapid transitions between states, the average distribution of the MP for all MECL3 (and to a lesser extent LECL3) pyramidal cells was bimodal (Figure 3-4A).

Averaged across all the data, the UP states in the MECL3 neurons, but not LECL3 neurons, lasted significantly longer than the UP states of the Ncx LFP (Figure 3-4B).

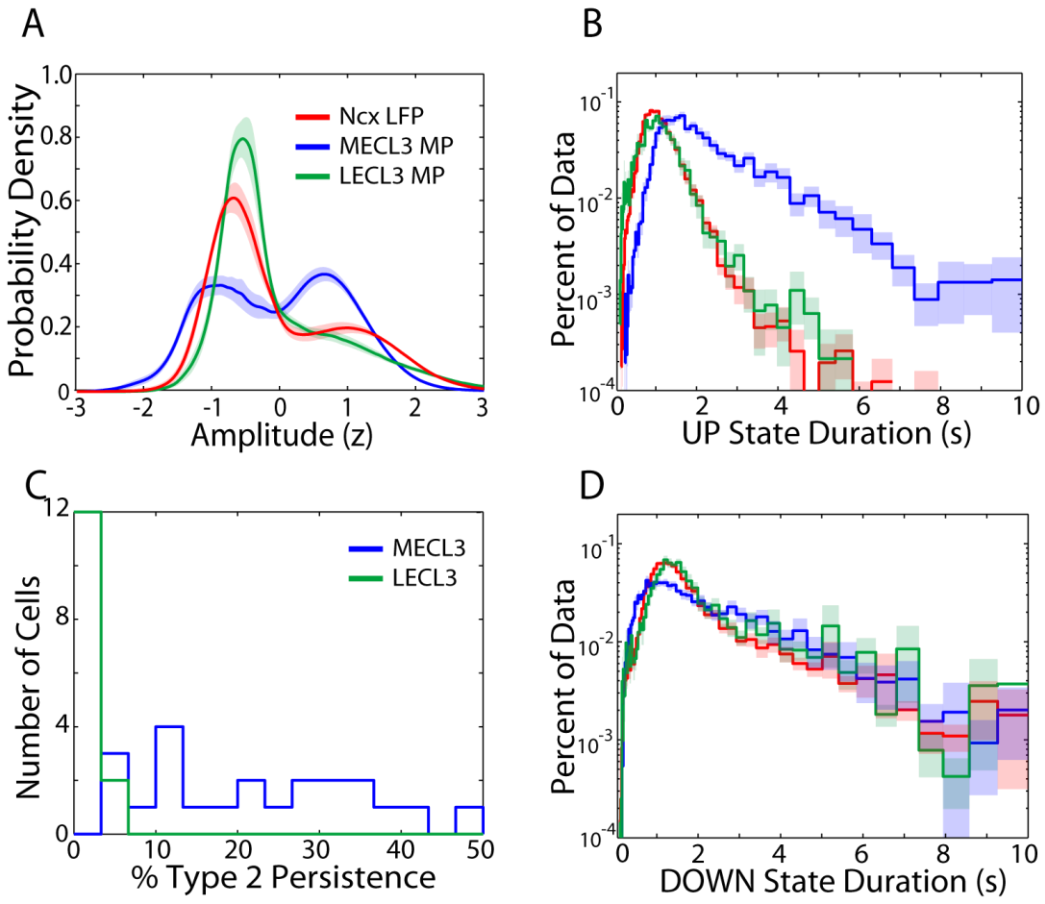

Figure 3-4: Persistent UP states in MECL3, but not LECL3 cells.

***A**) Average of all amplitude distributions showing bimodality of the Ncx LFP (red), MECL3 MP (blue), and to a lesser extent the LECL3 MP (green). Uncertainty is reported as the interval (mean±SEM), depicted as a shaded region around the mean in all subsequent figures. **B**) Distribution of UP state durations averaged across all recording sessions. The average duration of UP states was larger in the MECL3 MP (2.0±0.14s) than in the Ncx LFP (0.98±0.023s, $p=1.0x10^{-7}$, paired t-test), but not between the LECL3 MP (0.92±0.055s) and Ncx LFP (p = 0.070). Furthermore, the MECL3 MP UP states lasted significantly longer than the LECL3 MP UP states ($p = 1.3x10^{-6}$). **C)** Histogram of the percentage of MP UP states showing type 2 persistent activity. Averaged across all neurons, MECL3 UP states were 12-times more likely to have type 2 persistent UP states (23±2.8%) than LECL3 UP states (1.7±0.39%, $p=7.2x10^{-7}$). **D**) Same as B for DOWN state durations. Both MECL3 (1.6±0.16s; p=0.46) and LECL3 (1.8±0.21s; p=0.18) DOWN state durations were not significantly different than Ncx DOWN states (1.6±0.11s).*

We defined type 1 persistent UP states as those EC UP states that lasted longer than the concurrent Ncx UP state (see *Persistent Activity*). Nearly all of the MECL3 neurons' UP states (90±1.5%) showed type 1 persistent activity, which was 2.4 times more than the LECL3 neurons' (37+4.9%, p=4.4x10^{-14}). Our more restrictive definition of type 2 persistent activity required that the EC UP state last longer than the concurrent Ncx UP state and the subsequent Ncx DOWN state, thus lasting longer than an entire cycle of the Ncx UDS. A large majority of MECL3 cells (18/22) showed substantial amounts (>10% of up states) of type 2 persistent activity, but none of the LECL3 neurons did (Figure 3-3C). By definition, all type 2 persistent UP states are also type 1 persistent, but not the converse. Interestingly, even those MECL3 UP states that did not show type 2 persistence were far more likely to show type 1 persistence than LECL3 UP states. This was not simply because the MECL3 neurons underwent UDS oscillations at a lower frequency, since there was no significant difference between the average duration of Ncx and MECL3 DOWN states (Figure 3-3D). Importantly, while we have defined persistence in terms of the subthreshold state of the MP, persistent UP states also corresponded to persistent spiking activity such that there was only a small difference between the MECL3 UP state firing rate during Ncx DOWN states (5.7±0.40Hz) compared to Ncx UP states (6.3±0.52Hz; p=7.1x10^{-3}).

	MECL3 (n=22)	LECL3 (n=14)	p-value (two-sample t-test)
MP during UP state (mV)	-52.0±0.90	-68±2.2	4.4×10^{-9}
MP during DOWN state (mV)	-72±1.3	-79±1.7	0.0026
Firing rate during UP state (Hz)	6.0±0.45	1.3±0.41	2.7×10^{-8}
Overall firing rate (Hz)	2.9±0.28	0.4±0.13	4.6×10^{-8}
UP state duration (s)	2.0±0.14	0.92±0.055	1.3×10^{-6}
DOWN state duration (s)	1.6±0.16	1.8±0.21	0.47
UP transition lag from cortical UP transition (ms)	240±17	130±25	3.7×10^{-4}
DOWN transition lag from cortical DOWN transition (ms)	750±45	-30±46	3.9×10^{-13}
Percent of UP states showing type 1 persistent activity	90±1.5	37±4.9	4.4×10^{-14}
Percent of UP states showing type 2 persistent activity	23±2.8	1.7±0.39	7.2×10^{-7}

$p^* = 0.14$ (MP during UP state vs. DOWN state duration)

$p^* = 4.3\times10^{-3}$ (MP during DOWN state vs. DOWN state duration)

$p^* = 3.4\times10^{-9}$ (UP state duration vs. UP transition lag from cortical UP transition)

$p^* = 5.9\times10^{-3}$ (DOWN state duration vs. DOWN transition lag from cortical DOWN transition)

*p-values represent paired t-tests

Table 3-1: Summary comparison of MECL3 and LECL3 UDS properties.

The temporal relationship between ECL3 neurons' MP and Ncx activity was characterized by computing the MP-LFP cross-correlogram. Both MECL3 and LECL3 neurons were strongly correlated with the afferent Ncx LFP (Figure 3-5A), yet the peak cross-correlation with the Ncx LFP was significantly weaker for MECL3 neurons (0.35±0.025) than LECL3 neurons (0.64±0.022; $p=1.6\times10^{-9}$), and occurred at a longer latency (MEC: +550±35ms; LEC: +90±26ms; $p=6.2\times10^{-11}$). For comparison, we also recorded simultaneous LFPs from posterior parietal and frontal (n=8) or prefrontal (n=16) cortices and found that Ncx UDS were much more strongly coupled, even across these substantial cortical distances, than between the parietal cortical LFP and EC MP (peak correlation with frontal: 0.936±0.0064; prefrontal: 0.901±0.0074; Figure 3-5A).

This shows that the UDS measured by the parietal cortical LFP are indeed representative of the Ncx UDS.

As expected from the cross-correlation analysis, the maximal coherence between MECL3 MP and Ncx LFP (0.69±0.021) was significantly less than between LECL3 MP and the Ncx LFP (0.88±0.014; $p=2.3x10^{-9}$; Figure 3-5B). To determine the influence of the EC neurons' on the downstream hippocampus (Hpc) we computed the coherence spectrum between the EC MP and the LFP from the CA1 region of dorsal Hpc. We used partial coherence analysis to isolate effects that could not be explained by mutual correlation with the Ncx LFP signal and to remove any effects of volume conduction (see *Spectral Analysis*). Interestingly, the peak partial coherence between MECL3 MP and Hpc LFP (0.55±0.049) was significantly larger than for the LECL3 MP and Hpc LFP (0.31±0.036; $p=1.2x10^{-3}$) (Figure 3-5B). Further, the frequency at which maximal (partial) coherence occurred between MECL3 MP and Hpc LFP (0.29±0.043Hz) was significantly lower than for Ncx LFP (0.47±0.018Hz; $p=9.8x10^{-4}$). The lower frequency range indeed corresponded to the frequency of MECL3 MP (0.25±0.021Hz) and Hpc LFP (0.31±0.029Hz) UDS, while the higher frequencies were representative of LECL3 MP (0.41±0.027Hz) and Ncx (0.43±0.014Hz) UDS (Figure 3-5C). In this regard the Ncx and Hpc UDS can be viewed as separate but coupled oscillations [172,228] occurring at two different frequencies. Since the lower frequencies of UDS in the MECL3 MP resulted from the persistent UP states, these observations suggest that the persistent UP states indeed exert a direct influence on the hippocampus.

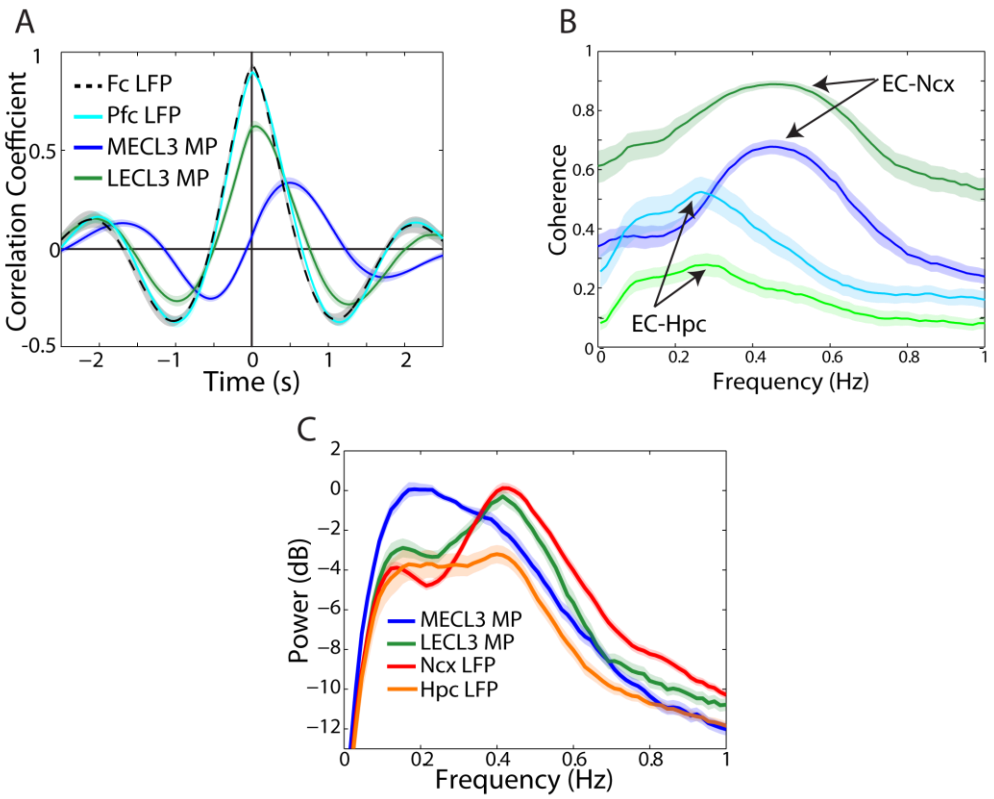

Figure 3-5: MECL3 neurons are less coupled to Ncx UDS, but more coupled to the Hpc LFP than LECL3 neurons.

A) *Average cross-correlograms with the parietal cortical LFP for the simultaneously recorded frontal cortical LFP (dashed black trace), prefrontal cortical LFP (light blue trace), MECL3 MP (dark blue trace), and LECL3 MP (dark green trace).* ***B)*** *Average coherence spectra computed between the parietal cortical LFP and the MECL3 MP (dark blue), and LECL3 MP (dark green). The average partial coherence spectra with the hippocampal LFP are shown in light blue and light green for the MECL3 MP and LECL3 MP respectively.* ***C)*** *Average power spectra for the MECL3 MP (dark blue), LECL3 MP (dark green), parietal cortical LFP (red), and hippocampal LFP (orange).*

The coupling between Ncx activity and MECL3 and LECL3 neurons' MP could be different for UP and DOWN state transitions. Hence, the phase-locking of ECL3 state transitions to the Ncx activity was directly quantified by computing the timing of MP UP and DOWN transitions relative to Ncx UP and DOWN transitions respectively (Figure 3-6A,B). Consistent with previous

work [149], both MECL3 (240±17ms) and LECL3 (130±25ms) neurons made UP transitions significantly after the Ncx UP transition. Further, although UP transitions of MECL3 neurons were delayed nearly twice as much as LECL3 neurons relative to the Ncx UP transition ($p=3.7x10^{-4}$), the differences in the timing of MECL3 and LECL3 neurons' DOWN transitions were even more striking. While the DOWN transitions of LECL3 neurons occurred synchronously with the Ncx DOWN transition (-30±46ms; p=0.60), MECL3 neurons' DOWN transitions occurred 750±45ms after the Ncx DOWN transition, nearly half the duration of the average Ncx DOWN state. This long delay in the MECL3 neurons' DOWN transitions was present even for the UP states that were not type 2 persistent (760±51ms), reflecting the observation that type1 persistent activity occurred consistently, even in the absence of type 2 persistent activity. Thus, while for LECL3 neurons DOWN transitions were significantly less delayed relative to the Ncx LFP than for UP transitions ($p=5.9x10^{-3}$), for MECL3 neurons the DOWN transitions were significantly more delayed ($p=3.4x10^{-9}$). Further, cells that made a DOWN transition at a longer latency with respect to the Ncx DOWN transition were more likely to persist beyond an entire concurrent Ncx UDS cycle, i.e. show type 2 persistent activity (Pearson's correlation coefficient MECL3: r=0.64, $p=1.2x10^{-3}$; LECL3: r=0.75, $p=1.8x10^{-3}$; Figure 3-6C). This shows that type 1 and type 2 persistence are related phenomena, resulting in a systematic relationship in the coupling between the Ncx LFP and EC MP: the maximum value of the cross-correlation between the Ncx LFP and EC neuron's MP was inversely related to the latency at which this maximal value occurred (r=-0.89, $p=4.5x10^{-13}$; Figure 3-6D).

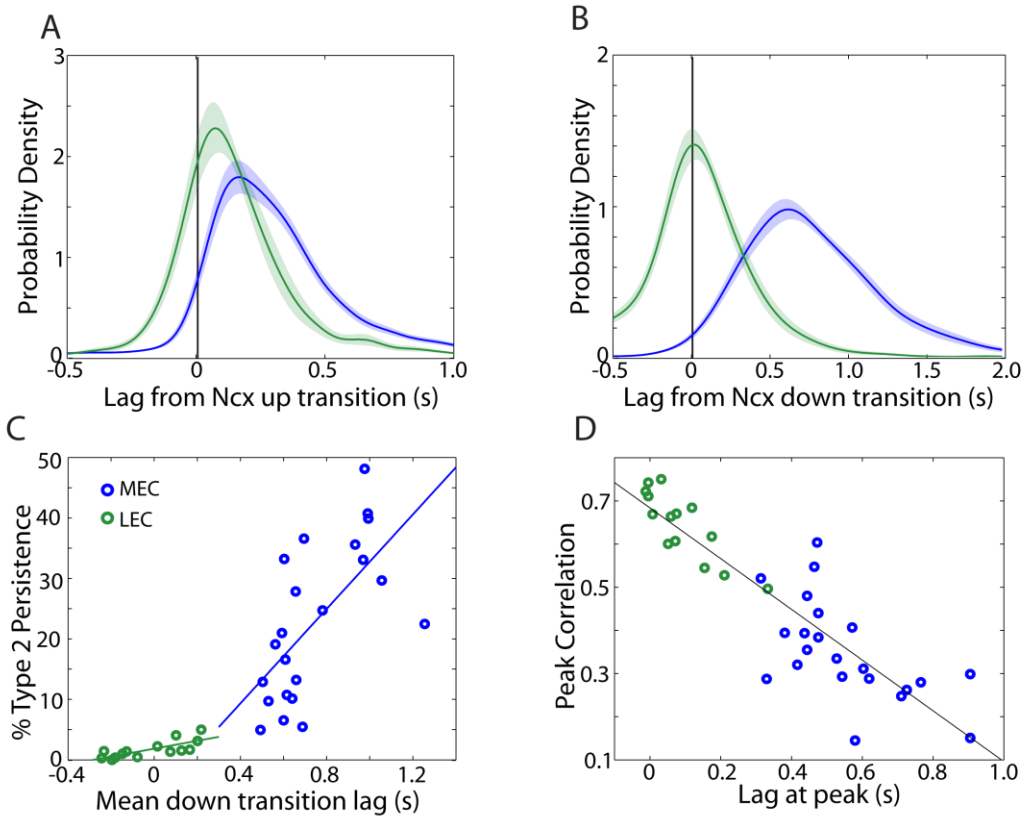

Figure 3-6: Differential phase-locking of MECL3 and LECL3 neurons' MP to Ncx UDS.

A) *Average distributions of the time lag between the MP UP transition and the concurrent Ncx UP transition.* ***B)*** *Average distributions of the time lag between the MP DOWN transition and corresponding Ncx DOWN transition.* ***C)*** *The percent of UP states that are persistent (type 2) is plotted against the average DOWN transition delay. MECL3 neurons are indicated by blue circles while LECL3 neurons are shown as green circles.* ***D)*** *The maximal value of the cross-correlogram between MP and Ncx LFP is plotted against the latency of the cross-correlation peak for MECL3 (blue circles) and LECL3 cells (green circles).*

To understand the precise temporal relationship of the MECL3 persistent states to the Ncx UDS, we measured the duration of MP UP states in units of the corresponding Ncx UDS cycles. Segments of the MP were extracted around every transition to the UP state, and these segments were sorted according to the duration of the ensuing UP state. The sorted segments were centered on the UP transition and assembled into a single matrix, shown for the MECL3

cell depicted in Figure 3-3A,C (Figure 3-7A). This cell shows UP states with durations ranging from 0.3s to 12s (3.1±0.13s). Corresponding segments of the Ncx LFP (aligned to the UP transitions in the MP) were also assembled into a matrix (Figure 3-7B). This procedure shows that the example neuron transitions to the DOWN state almost exclusively during the Ncx DOWN state. It further shows that the phase-locking of the UP and DOWN transitions was independent of the duration of the MP UP state. Thus the MECL3 UP states lasted for integer multiples of Ncx UDS cycles, with a constant offset. Figure 3-7C shows examples of MP UP states that persisted for about 0.8, 1.8, 2.8 and 3.5 times the underlying Ncx UDS period respectively.

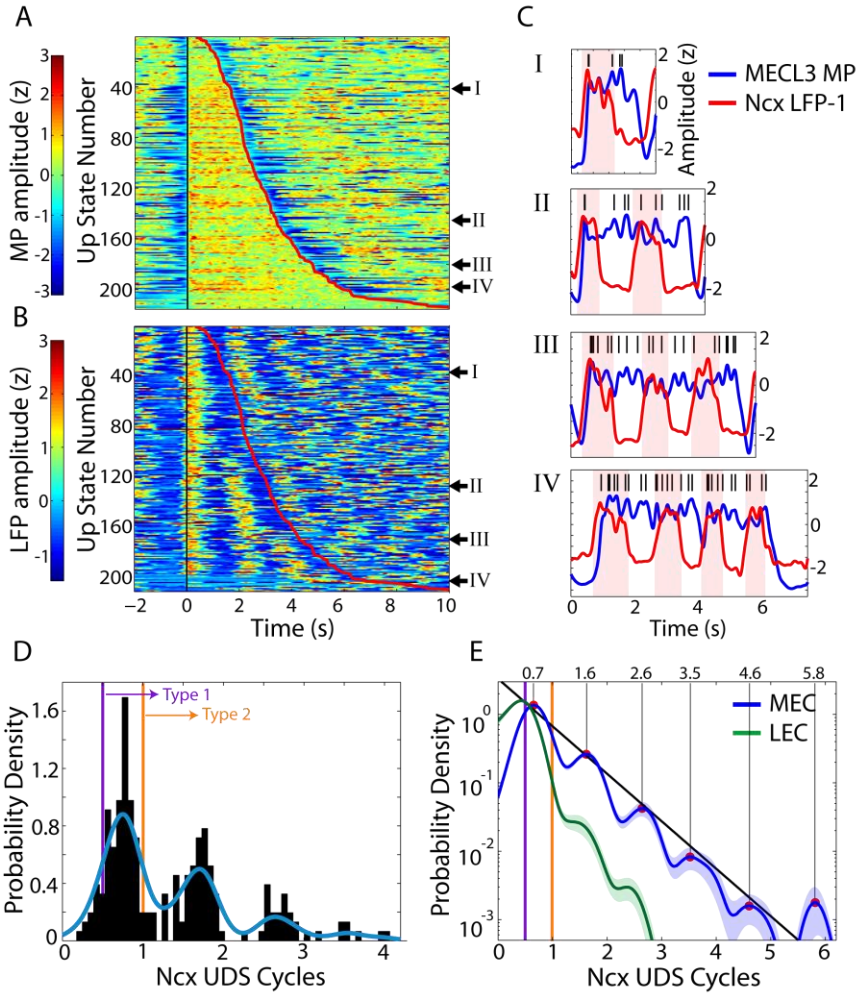

Figure 3-7: Quantization of the MECL3 persistent UP states.

A) *Matrix composed of all 218 MP UP states from an example MECL3 neuron, aligned to the onset of the UP state, and ordered from top to bottom according to increasing UP state duration. Colors depict the amplitude of the MP (in units of z-score). The vertical black line shows the onset of the MP UP states (at t=0), and the red trace shows the time of termination of each MP UP state.* ***B)*** *Matrix of corresponding Ncx LFP traces during each of the MP UP states shown in A. Black and red traces show the time of onset and offset of the corresponding MP UP states respectively (as in A).* ***C) I-IV*** *Example traces from the MP and Ncx LFP where the MP UP state persists for 1, 2, 3 and 4 Ncx UP states respectively. In these figures, action potentials are indicated by tick marks and 1z is subtracted from the LFP signal to provide better visual comparison of LFP and MP. Arrows on the right side of panels A and B show the locations of each of the individual traces within the matrices.* ***D)*** *Histogram of MP UP state durations in units of Ncx UDS cycles. The cyan trace is the smoothed histogram. Threshold durations for type 1 and type 2 persistent activity are indicated by the lavender and orange vertical lines respectively.* ***E)*** *Distribution of MP UP state durations averaged across all cells, in units of the corresponding Ncx UDS cycles for MECL3 (blue) and LECL3 (green) neurons.*

The distribution of UP state durations for this example MECL3 neuron (Figure 3-3) was calculated in the units of the simultaneously measured Ncx UDS periods (see *Persistent Activity*). In these units the distribution was multimodal, with the modes all having integer spacing (Figure 3-7D). Significant quantization of the UP state duration in units of the Ncx UDS cycles can be seen in this cell for up to 4 Ncx UDS cycles. Additional examples of such quantization are shown in Figure 3-8. Averaged across all cells, quantization of the persistent activity is clearly visible in the distribution of MECL3 UP state durations (Figure 3-7E), indicating that MECL3 cells showing different levels of type 1 and type 2 persistent activity have a similar degree of phase-locking with the Ncx input. Importantly, the first six peaks of the average MP UP state duration distribution have a nearly constant spacing of about one UDS cycle, illustrating quantization in the ensemble of cells that was apparent even in the very long MECL3 UP states. The probability of observing a MECL3 UP state lasting for *k* cycles of the Ncx UDS was also found to have a nearly geometric distribution (Figure 3-7E), suggesting that MECL3 UP states would terminate during a given Ncx DOWN state with a fixed probability, independent of the number of Ncx UDS cycles already skipped.

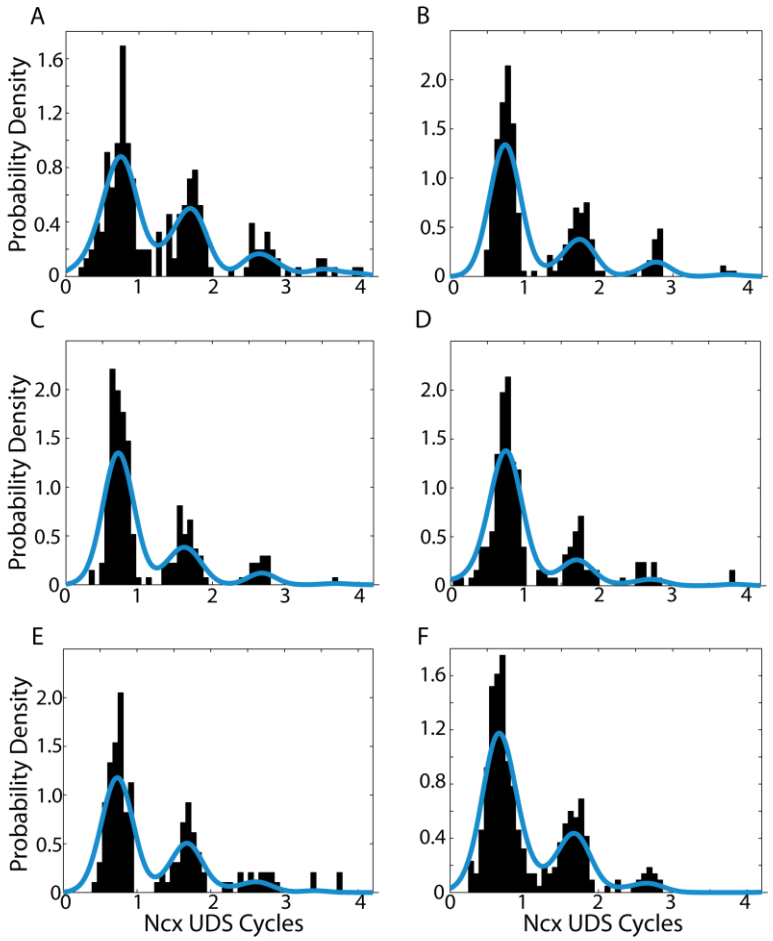

Figure 3-8: Examples of quantized MP UP state duration in six different MECL3 cells.

Histograms of MP UP state duration in units of the Ncx UDS period for six MECL3 cells. The cyan traces are the smoothed histograms. ***A)*** *Example histogram from the same cell shown in Figure 3-7D.* ***B-F)*** *Similar figures for five other MECL3 cells. All examples show multimodal distributions with integer spacing between the modes, indicating that the MP UP state durations are quantized in the units of Ncx UDS period.*

Previous *in vitro* studies have shown that EC neurons can persist in a depolarized and spiking state after a period of depolarizing stimulation[260,261,262]. However, *in vivo*, the network is spontaneously active and this spontaneous activity could interfere with any stimulation-induced persistent activity. Hence we asked if the persistent states in MECL3 neurons can be generated

by current injection *in vivo*. We injected depolarizing current for 1s or 12s, durations comparable to those used for *in vitro* studies (Figure 3-9). This current injection put the MECL3 neurons at their spiking threshold, with 220±61% and 340±86% increases from the spontaneous firing rate for 1s and 12s current injections respectively. Yet, the MP did not persist in the depolarized state upon the termination of the current injection lasting either 1s (Figure 3-9A,B) or 12s (Figure 3-9C,D). Thus, current injection failed to generate either type 1 or type 2 persistent activity.

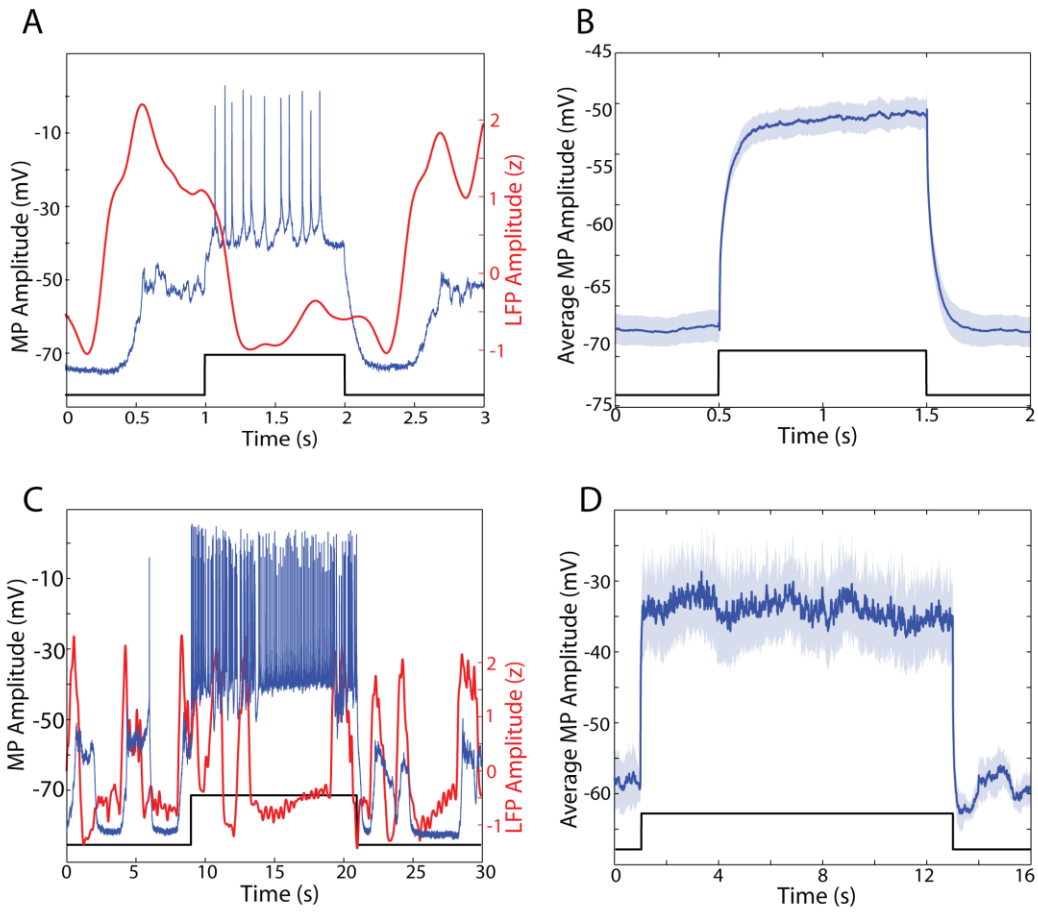

Figure 3-9: MECL3 persistent UP states cannot be elicited with depolarizing current injections.

A) *Example trace showing the MP of an MECL3 neuron during a 1s depolarizing current injection (indicated by the black trace). The simultaneous Ncx LFP is shown in red. In this example the cell returns to the DOWN state following the termination of current injection.* ***B)*** *Average across 658 1s current injections over 11 MECL3 cells, showing that the neurons were depolarized by over 20mV, but did not generate persistent UP states.* ***C)*** *Example trace for the same neuron as in A, during 12s depolarizing current injection. Again, the cell returns to the DOWN state immediately upon termination of the current injection.* ***D)*** *Same as B, but for 31 12s current injections in 5 neurons. In this case the average depolarization is nearly 30mV, yet the cells return to the pre-stimulus value immediately after the termination of current injection.*

Discussion

Persistent changes in spiking activity have been observed during working memory tasks *in vivo* in the prefrontal cortex [263] as well as the LEC [264,265], and those changes are thought to be a neural correlate of working memory. Persistent depolarization and spiking has also been induced *in vitro* in both LEC[261] and MEC [260,262] neurons by electrical, synaptic, and chemical stimulation. Further, while some studies have referred to the Ncx UP states themselves as persistent states [266], here we define persistent UP states of the EC neurons as those states that outlast the concurrent Ncx UP states. Our results demonstrate that persistent activity can occur spontaneously in pyramidal neurons in layer 3 of the MEC *in vivo*. This is the first demonstration of persistent activity in MEC neurons *in vivo*. The phase-locking of MECL3 activity to spontaneous Ncx activity with a short delay suggests that UP states in the MECL3 MP are triggered by Ncx inputs, yet can decouple from these inputs to produce prolonged, persistent UP states. Further, partial coherence analysis suggests that these MECL3 persistent UP states exert a direct influence on the downstream Hpc LFP. Surprisingly, these persistent UP states were not observed in LECL3 neurons. This difference might be explained by the less-depolarized MP and reduced firing rates of the LEC neurons' UP states (Table 3-1). However, since UDS in the LEC neurons' MP were more tightly locked to the Ncx LFP than the MEC neurons, this suggests there could be differences in the local balance of excitation and inhibition, as well as the relative contributions of Ncx, Hpc, and local inputs in these regions. There may be other behavioral or neuromodulatory conditions where persistent activity can occur in LECL3 neurons.

Both single-cell and network mechanisms have been proposed to explain several forms of persistent activity (for review see [267]). *In vitro* studies of persistent activity in EC have focused on the cellular mechanisms [268]. However, we were unable to generate persistent activity in MECL3 using similar protocols as those used *in vitro*. This suggests that in addition to these cellular mechanisms, network mechanisms are needed to generate persistent activity *in vivo*. For example, cholinergic inputs, intracellular calcium concentration [268], and the reverberation of activity among groups of MECL3 cells could determine the initiation and termination of MECL3 persistent activity. Recent work suggests that MECL3 neurons may receive a unique pattern of convergent local network inputs that could allow such network reverberation [269]. The same mechanisms leading to MECL3 persistent activity might also play a role in mediating the learning of associations between stimuli and the formation of the unique spatially periodic and conjunctive firing patterns observed in MECL3 [209], but not LECL3 [258], neurons.

Since memory consolidation is thought to occur via cortico-hippocampal interactions during slow-wave sleep (SWS) [114,165,166], the EC is anatomically ideally suited to mediate such interactions. Indeed, inputs from MECL3 neurons to the hippocampus [180,185,270] are required for spatial memory consolidation [256]. Since SWS is highly analogous to UDS [123,129,132,257], this raises the question of whether the MECL3 persistent UP states demonstrated here could play an important role in these consolidation processes. In support of this possibility, we found that the lower-frequency MECL3 UDS exerted a direct influence on the Hpc LFP, which may result in the activation of CA1/CA3 pyramidal neurons during the Ncx down state [148,149]. Such activation of hippocampal (and entorhinal) neurons during the Ncx down state could produce the interleaved activation of old and new memories in the cortico-

entorhinal-hippocampal circuit that is required by theories of memory consolidation [166], thus transferring recently learned spatial information [194,195] from the hippocampus to the rest of the circuit. On the other hand, by decoupling the efferent hippocampal pyramidal neuron's activity from neocortical inputs, MECL3 persistent activity could facilitate erasure of memories from the hippocampus [271]

Chapter 4: **Delta (2-4Hz) oscillations in lateral entorhinal cortex phase-locked to UP-DOWN states**

Introduction

The UP-DOWN states (UDS) which occur under different forms of anesthesia, are thought to be analogous to slow oscillations (<1Hz) present during natural sleep [123,129,132,149,257]. Cortico-hippocampal interactions during these slow oscillations have been hypothesized to play an important role in memory consolidation [114,165,173,256]. The entorhinal cortex (EC) is bidirectionally connected to many cortical areas as well as the hippocampus, and thus is uniquely situated to mediate these cortico-hippocampal interactions [180]. During conditions where neocortical (Ncx) neurons show UDS oscillations similar to those seen during SWS [123,129,132,257], EC neurons undergo synchronized UDS [149], with the notable exception of the layer 3 pyramidal neurons of the medial EC (MEC) which exhibit persistent UP states that last for integer numbers of Ncx UDS (see *Chapter 3*).

In addition to their ability to produce self-sustaining 'persistent' active states [260,261,262], there has been substantial interest in the oscillatory properties of EC neurons, particularly in the theta range (4-12Hz). Alonso and colleagues showed that EC stellate cells in layer 2 [57], but not superficial EC pyramidal cells [272,273], produce intrinsic membrane potential theta oscillations through a combination of hyperpolarization-activated cation and persistent sodium currents [274,275]. Layer 2 MEC stellate cells were later shown to have intrinsic oscillation frequencies

that vary systematically with their dorso-ventral location within the EC [276], in a similar fashion to the spacing of the spatial firing 'grids' seen by these neurons during navigation [210]. Quilichini et al. [269] further demonstrated that layer 2 stellate neurons produced larger theta oscillations *in vivo*, under anesthesia, than neurons in either layers 3 or 5. In contrast, such intrinsic resonance properties have not been reported in neurons of the lateral EC (LEC). Neurons in the LEC do not show grid-like spatial firing patterns, and are thought to play a functionally distinct role from those in the MEC [258]. Indeed, recent work has shown that theta oscillations during natural behavior are substantially larger in the MEC than the LEC [277]. Given that the LEC neurons have much weaker theta-oscillatory behavior than MEC neurons, this raises the question of whether there are different conditions *in vivo* under which LEC neurons do oscillate, potentially relating to their differential functional role.

To investigate this question, we measured the membrane potential (MP) of superficial neurons (layers 2 and 3) from the MEC and LEC in urethane-anesthetized mice. To compare EC neurons' behavior to the network UDS activity seen under these conditions [123], we simultaneously recorded the local field potentials (LFP) from posterior parietal cortex and hippocampus (Hpc). As demonstrated in Chapter 3, the posterior parietal LFP provides a robust measure of the Ncx UDS. The Hpc LFP is thought to reflect a separate, but related, Hpc slow oscillation [172,228] that also reflects the lower-frequency UDS associated with MECL3 persistent UP states (see Chapter 3).

Methods

Experimental Methods

For experimental methods see Chapter 3 *Methods.*

Data Analysis

Preprocessing of the MP and LFP signals was performed as in Chapter 3 (see *Data Preprocessing*). Similarly, UP-DOWN states were detected as described in Chapter 2: . *I also focus on the relationship between hippocampal activity* and behavior, and propose several experiments which could clarify and extend the results of Chapter 6.

Explicit-duration hidden Markov model inference of UP-DOWN states from continuous signals.

Identification of Neurons with Delta Oscillations

In order to determine which neurons had significant MP delta oscillations we first computed the prewhitened power spectrum [278]. Prewhitening of the signal was used to equalize power across frequencies, and was achieved by fitting a second order autoregressive model to the MP signal (after spike subtraction) using the Levinson Durbin recursion. The coefficients of the autoregressive model were used to filter the signal, and the model residuals were then used to compute the prewhitened power spectrum. Power spectral estimates were computed using multi-taper methods (Chronux Matlab toolbox [254]; http://chronux.org), averaged across non-overlapping 20s time windows, with a time-bandwidth product of 3 and 5 tapers [31]. 95% confidence intervals of the power spectrum were computed using a jackknife procedure [278]. The peak of the prewhitened spectrum in the delta band (1.5-4.5Hz) was found, and significant delta-band peaks were defined as those for which the lower 95% confidence bound of the peak power exceeded the upper 95% confidence bounds 3 bandwidths (0.45Hz) on either side of the peak. This definition provided good separation of datasets for which the MP power spectrum had a visually distinguishable peak in the delta range.

Classification of Delta Epochs

For those neurons which were determined to have significant delta power in the average power spectrum, we located epochs of data containing delta oscillations ('delta epochs') in the following way. First, we bandpass filtered the MP signal in the delta band (2-4Hz) using a 4-pole Butterworth filter. The analytic signal amplitude (derived from the Hilbert transform) then provided a measure of the amplitude envelope of delta oscillations. The running average of this

delta amplitude signal was computed in 20s windows (with 1s resolution). This running average delta amplitude was then used to fit a two-state Gaussian-emissions hidden Markov model (HMM) (Figure 4-1). The state-conditional mean and variance were initialized by fitting a static two-state Gaussian mixture model. The expectation-maximization algorithm [240] was then used to iteratively fit the parameters of the HMM, with iterations terminating upon convergence of the log likelihood with a tolerance of $1x10^{-5}$. After fitting the parameters of the HMM, the Viterbi algorithm was used to find the most likely hidden state sequence [240]. This state sequence then defined the epochs with and without delta oscillations. The main advantage of fitting a probabilistic model to define delta epochs for each MP signal was that there was substantial variability in the spectral properties (of both UDS and delta oscillations) which made selection of a single threshold for defining delta epochs across all neurons difficult.

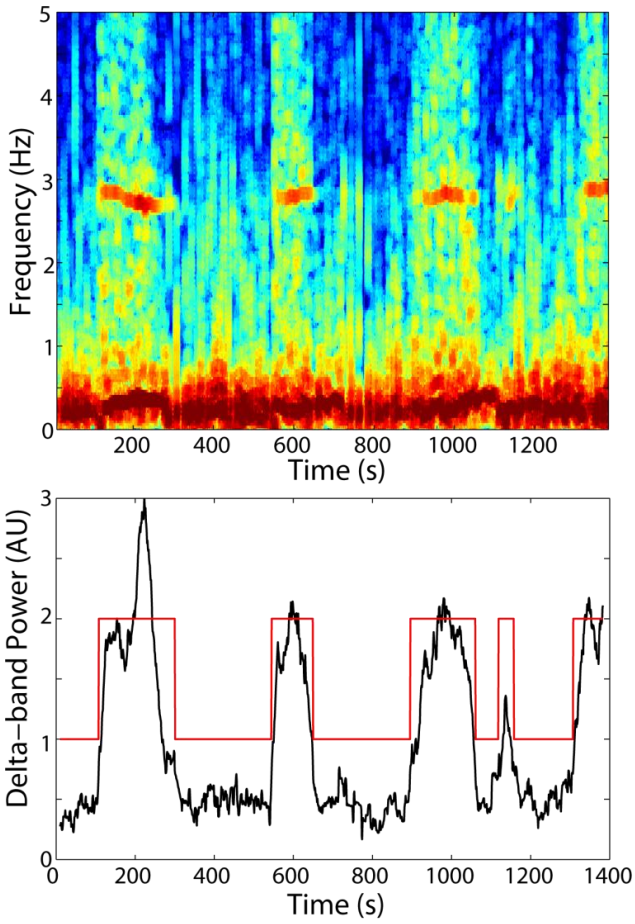

Figure 4-1: Classification of delta epochs.

(Upper panel) *Spectrogram of the same example neuron shown in Figure 4-2.* ***(Lower panel)*** *The amplitude of the Hilbert-transform of the MP signal filtered between 2 and 4Hz is computed and averaged within each of the 20s windows used to compute the spectrogram (black trace). A Gaussian-emissions hidden Markov model is then fit to this signal, and the resulting Viterbi state sequence (red trace) was used to define delta epochs.*

Spectral Analysis

Power spectra and coherence spectra of MP and LFP signals were computed separately during delta and non-delta epochs using multi-taper methods. Spectra were computed across discontinuous data segments in non-overlapping 20s windows, with a time-bandwidth product of 5 and 9 tapers [31]. In order to compare power spectra and coherence spectra between

groups with unequal number of data windows (e.g. between delta epochs and non-delta epochs), we corrected spectra for biases due to finite sample sizes [259] (see Equations (3-2)(3-3)).

In order to compute power spectra separately during UP and DOWN states, we used wavelet analysis. The continuous wavelet transform (CWT) $W_t(s)$ maps a signal into the time-scale (or time-frequency) plane. We computed the CWT using a complex Morlet wavelet with nondimensional frequency parameter ω_0=6 [39] (see Time-Frequency Analysis). Wavelet software was provided by C. Torrence and G. Compo, and is available at URL: http://paos.colorado.edu/research/wavelets/. Wavelet coherence spectra were given by Equation (1-18). Temporal smoothing was chosen to match the Gaussian envelope of the Morlet wavelet:

$$S_{time}(W)\big|_s = \left(W_t(s) * c_1^{\frac{-t^2}{2s^2}} \right)\Bigg|_s \qquad \textbf{(4-1)}$$

and smoothing in scale was chosen to match its scale decorrelation length:

$$S_{scale}(W)\big|_n = \left(W_n(s) * c_2 \Pi(0.6s)\right)\big|_n \qquad \textbf{(4-2)}$$

where c1 and c2 are normalization constants and Π is the rectangle function [39,40].

Delta-UDS Coupling

To quantify phase-locking of the MP and LFP delta oscillations to the UDS transitions we first computed the transition-triggered averages of signals filtered in the delta (2-4Hz) range (using only state transitions that occurred during delta epochs). We filtered MP and LFP signals using a causal 4-pole Butterworth filters so that the state transitions themselves could not interfere

with estimation of the delta-band signal prior to the state transitions. Similar results were obtained using a zero-phase filter. Further, when computing transition-triggered averages, we only included data up to the preceding and following state transitions. Thus, for the DOWN-UP triggered ensemble, all data at negative time lags were from the DOWN state, while all data at positive time lags were from the UP state. A similar procedure was used for the UP-DOWN triggered averages. This allowed us to dissociate the delta-UDS coupling at various time lags relative to the state transitions from the distribution of UP and DOWN state durations themselves. In order to compute the transition-triggered average delta signals across different delta epochs, as well as across different neurons, where the precise frequency of the delta oscillations varied between 1.5 and 4.5Hz, we used a time-rescaling procedure as follows. For each delta epoch, the local frequency of the delta oscillation was first determined by finding the peak of the prewhitened power spectrum (computed from the direct periodogram estimator) in the range 1.5-4.5Hz. Time was then locally rescaled into units of delta oscillation cycles using linear interpolation. UP and DOWN states which were shorter than 0.5 seconds were excluded from this analysis.

To test for significant delta-UDS cross-frequency coupling we first computed the ensemble of delta phases at each time lag relative to the DOWN-UP (UP-DOWN) transition using the procedure described above. Delta-phase was given by the complex argument of the delta-band analytic signal (derived from the Hilbert transform). The mean resultant vector length across the ensemble of delta phases at each time lag was then computed, and was taken as a measure of the delta phase-locking to the UDS transition at each relative time lag. A bootstrapping procedure was then used to generate a distribution of vector lengths at each time lag under the

null hypothesis of no delta-UDS phase-locking. Importantly, samples from the null distribution at a given time lag were computed using a random sample of *N* delta phases, matched to the empirical sample size at that lag. 500 samples from this null distribution were generated at each time lag, for each neuron, and the upper 95th percentile of samples was taken as the upper 95% confidence interval which we used to test for significant phase-locking.

Results

Superficial neurons in both the LEC and MEC, showed UDS throughout the recordings, only rarely being interrupted by desynchronized epochs and loss of UDS [129,250]. LEC neurons, but not MEC neurons, frequently transitioned into a state with prominent 2-4Hz oscillations accompanying the UDS. The spectrogram of an example layer 2 LEC neuron is shown in Figure 4-2A. While a spectral peak in the UDS range (~0.5Hz) is present throughout the recording, the power spectrum of the MP spontaneously alternates between a state characterized by an additional spectral peak well-localized at 3Hz, along with a broadband increase in power. Interestingly, while epochs of oscillations in the 2-5Hz frequency range in anesthetized animals have been reported in a number of studies [11,129,269,279], these oscillations (generally classified as theta oscillations) were accompanied by a loss of UDS (occurring as part of the 'desynchronized', or 'activated' state) [129,269]. In the current study however, the 2-4Hz oscillations occurred along with clear UDS. For these reasons, and others discussed below (see *Discussion*), we refer to the 2-4Hz oscillations seen here as 'delta oscillations'.

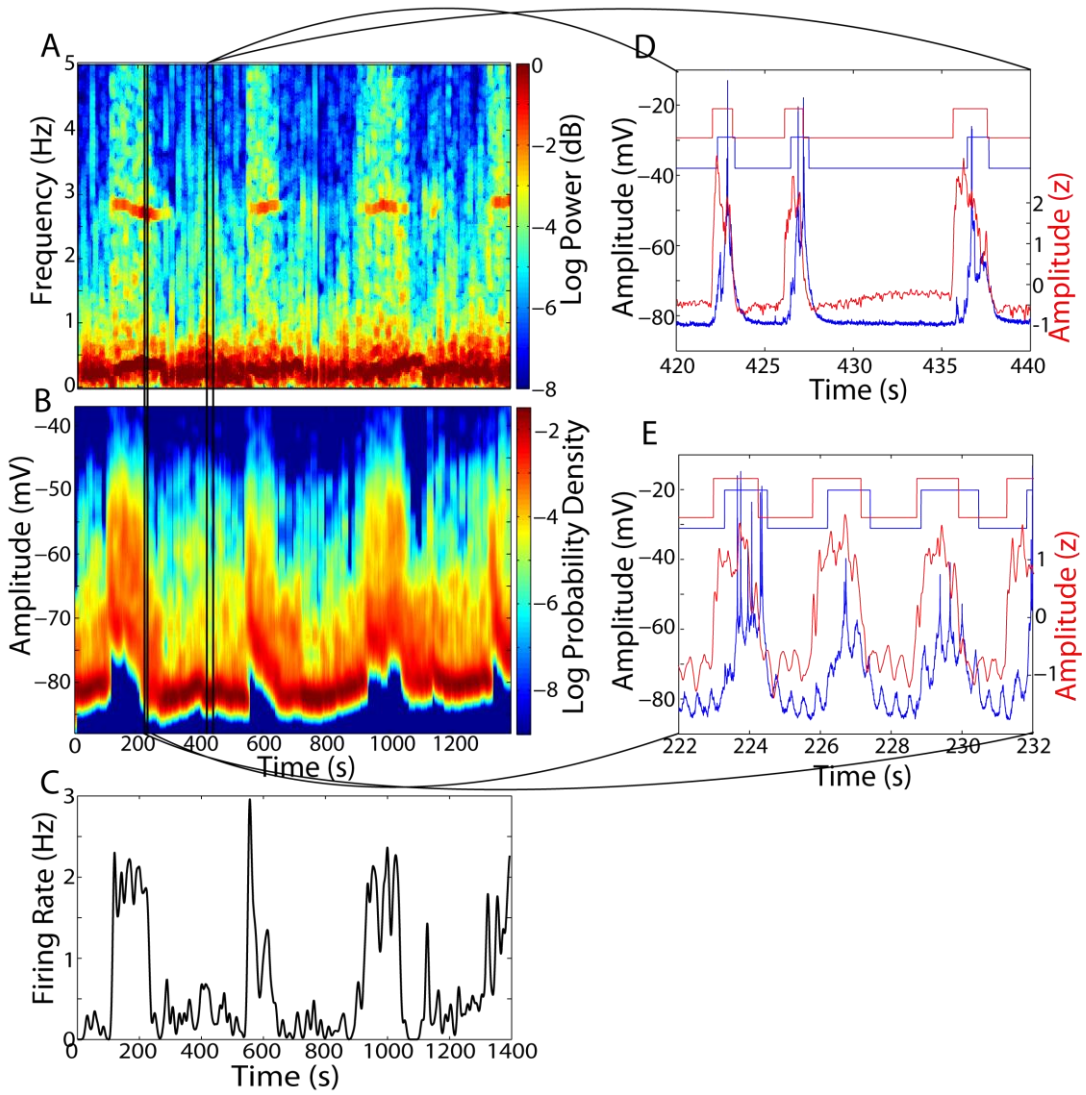

Figure 4-2: Delta oscillation epochs in an example layer 2 LEC neuron.

A) *Spectrogram across the entire 1400s recording showing four discrete epochs with a sharp peak at 3Hz.* ***B)*** *Sliding window amplitude distribution for the same neuron showing that the cell becomes depolarized by as much as 10mV within both the UP and DOWN states during the delta epochs.* ***C)*** *The firing rate increased dramatically during the delta epochs for this neuron, increasing from an average of 0.23Hz during non-delta epochs to 1.4Hz during delta epochs.* ***D)*** *An example trace of the MP (blue trace) and simultaneously recorded cortical LFP (red trace) from the non-delta epoch indicated by the black lines. The red and blue lines above indicate the detected UDS for each signal. Note the lack of signal variance in either the MP or LFP during the DOWN states.* ***E)*** *An example trace from the delta epoch indicated by the black lines. Large, extremely periodic 3Hz oscillations are clearly apparent in the MP and even the LFP, particularly during the DOWN states.*

The example cell in Figure 4-2 entered the delta oscillation state 5 times during the1400s recording, with the delta epochs lasting between 40s and 190s. The sliding-window amplitude distribution for this cell shows that the MP became substantially more depolarized (~10mV) within both the UP and DOWN states during the delta epochs (Figure 4-2B). The firing rate of the cell also increased dramatically during these periods (Figure 4-2C). Example traces from this cell during a non-delta epoch (Figure 4-2D) and a delta epoch (Figure 4-2E) show that UDS were clearly present during the delta epochs, and that the delta oscillations persisted throughout both the UP and DOWN states. The simultaneously recorded Ncx LFP also showed delta-oscillations within the Ncx DOWN states during the delta epochs that appeared to be synchronous with the MP delta oscillations (Figure 4-2E). Finally, we note the markedly shorter DOWN states apparent during the example delta epoch compared to the example non-delta epoch.

In order to quantify the properties of these delta epochs across neurons, we first analyzed the overall power spectrum of each neuron's MP and tested for the presence of a significant spectral peak in the 1.5-4.5Hz range. To this end we first prewhitened the power spectrum by fitting a low-order autoregressive model to the signal in order to remove the ~1/f slope (Figure 4-3A,B) [278], and allow more accurate detection of spectral peaks (see *Identification of Neurons with Delta Oscillations*). Across the ensemble of superficial LEC cells (n=35) there was a clear 2-4Hz peak in the average power spectrum that was not present for superficial MEC cells (n=50) (Figure 4-3C). Of the 35 LEC cells, 21 were determined to have a significant spectral peak in the delta band, while much fewer MEC cells did (3 of 50; $p<0.0001$, Pearson's chi-square test). Further, a significantly larger proportion of LECL2 cells (12/15) showed delta oscillations than

LECL3 cells (9/20) (p=0.036, Figure 4-3D). Layer 2 cells in the LEC had primarily either 'fan' or 'multipolar' morphologies [280], while layer 3 LEC cells were mostly pyramidal, with a few multipolar cells. In contrast, layer 2 MEC cells had almost exclusively stellate morphology, while layer 3 MEC cells were also primarily pyramidal with a few multipolar cells [269,272,280]. In all cases the delta oscillations were found to be extremely periodic (frequency-localized), with the peak frequencies of those LEC cells that had significant delta oscillations ranging between 2.1 and 3.6 Hz (2.86$\pm$0.088 Hz, mean$\pm$SEM). Only the superficial LEC neurons determined to have significant delta oscillations (n=21) were used for the subsequent analysis.

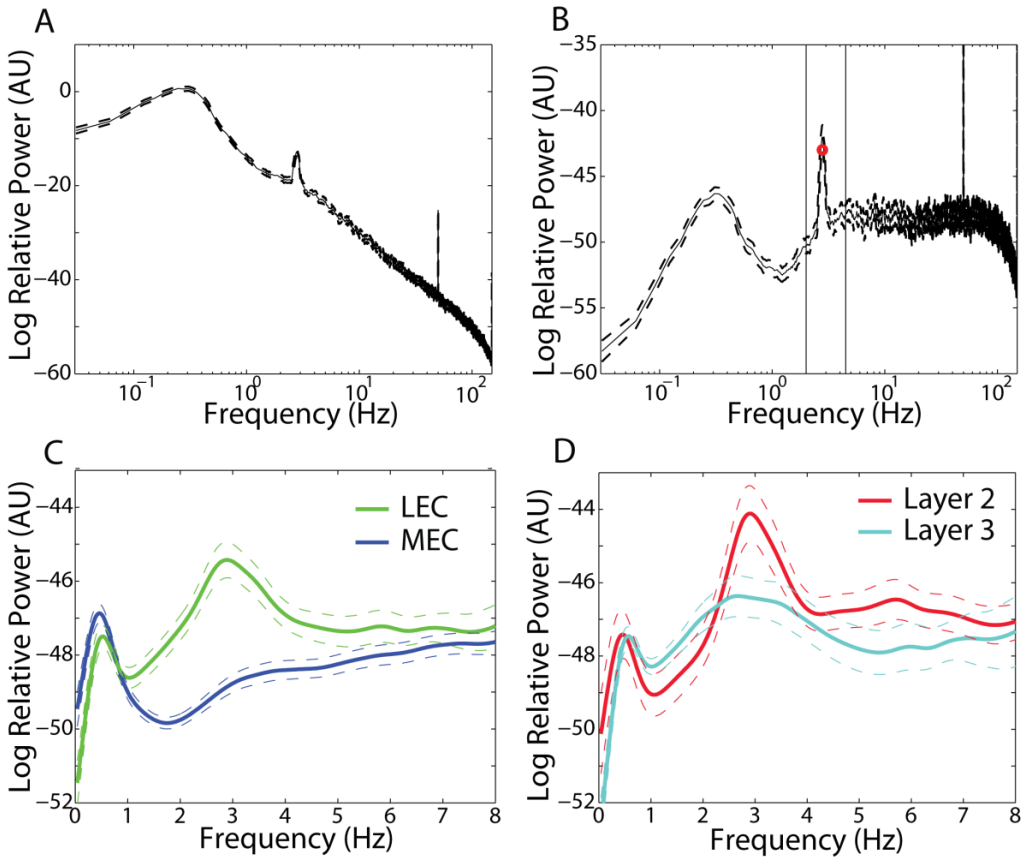

Figure 4-3: Superficial LEC, but not MEC, neurons exhibit delta oscillations.

A) *Power spectrum of the MP from the example neuron shown in Figure 4-2. Dashed traces indicate the 95% confidence interval. Note the sharp peak at 3Hz, as well as the ~1/f slope in the power spectrum making identification of the delta peak more difficult.* ***B)*** *Power spectrum of the residuals of a second-order autoregressive model for the same example neuron. This 'prewhitening' procedure removes the overall slope of the spectrum, making identification of significant delta-band peaks easier. The vertical black lines show the 'delta-band'. The red circle shows the spectral peak in the delta band.* ***C)*** *Average prewhitened MP spectrum for superficial LEC (green) and MEC (blue) neurons. Dashed traces indicate the region of one s.e.m. surrounding the mean. Note the complete lack of a delta-band peak in the ensemble average MEC spectrum, compared to the prominent 2-4Hz peak in the LEC average spectrum.* ***D)*** *Average MP spectra for layer 2 LEC (red) and layer 3 LEC (light blue) neurons. While the layer 3 LEC average has a broad peak in the delta band, the average layer 2 LEC spectrum has a significantly larger peak at 3Hz (p=0.029, Wilcoxon rank sum test).*

After identifying the LEC neurons which had significant delta oscillations, we located the delta epochs within these recordings using a Gaussian-emissions hidden Markov model of the delta-band power (see *Classification of Delta Epochs*; Figure 4-1). Across the 21 LEC neurons with significant delta oscillations, these neurons were determined to be in the delta state 35±2.9% of the time. As seen for the example cell in Figure 4-2, the MP was significantly more depolarized during delta epochs than non-delta epochs for both the UP (delta: -65±1.6mV; non-delta: -68±1.7mV; $p=1.2x10^{-4}$, Wilcoxon signed rank test) and the DOWN (delta: -74±1.5mV; non-delta: -77±1.6mV; $p=2.8x10^{-4}$) states (Figure 4-4A). These cells also fired at a significantly higher rate during delta epochs within both the UP (delta: 2.0±0.41Hz, non-delta: 0.80±0.26Hz; $p=4.8x10^{-4}$) and DOWN (delta: 0.10±0.047Hz; non-delta: 0.0030±0.0010Hz; $p=8.5x10^{-4}$) states (Figure 4-4B). Further, cells spent a significantly greater percentage of time in the UP state during delta epochs (36±2.0%) compared to non-delta epochs (27±1.8%; $p=1.4x10^{-4}$). This resulted from both an increase in the duration of LEC UP states (delta: 0.93±0.052s; non-delta: 0.82±0.051s; $p=1.2x10^{-3}$), as well as a decrease in the duration of LEC DOWN states (delta: 1.7±0.12s; non-delta: 2.3±0.16s; $p=1.9x10^{-3}$) (Figure 4-4C). Interestingly, the Ncx UDS also spent a significantly larger percentage of time in the UP state during delta epochs (40±1.3%) compared to non-delta epochs (35±1.1%; $p=5.4x10^{-4}$). While the duration of Ncx UP states did not increase significantly during delta epochs (delta: 1.01±0.032s; 0.98±0.032s; p=0.24), the duration of Ncx DOWN states did show a significant decrease (delta: 1.57±0.074s; non-delta: 1.9±0.11s; $p=1.7x10^{-3}$) (Figure 4-4D). Even the time lag between Ncx UP transitions and LEC UP transitions was significantly decreased during delta epochs (delta: 160±32ms; non-delta: 210±29ms; p=0.014; data not shown). Thus, delta epochs were characterized by an increased level of activation in both the

LEC MP as well as the Ncx LFP, with more depolarized and spiking LEC neurons and more prominent LEC and Ncx UP states.

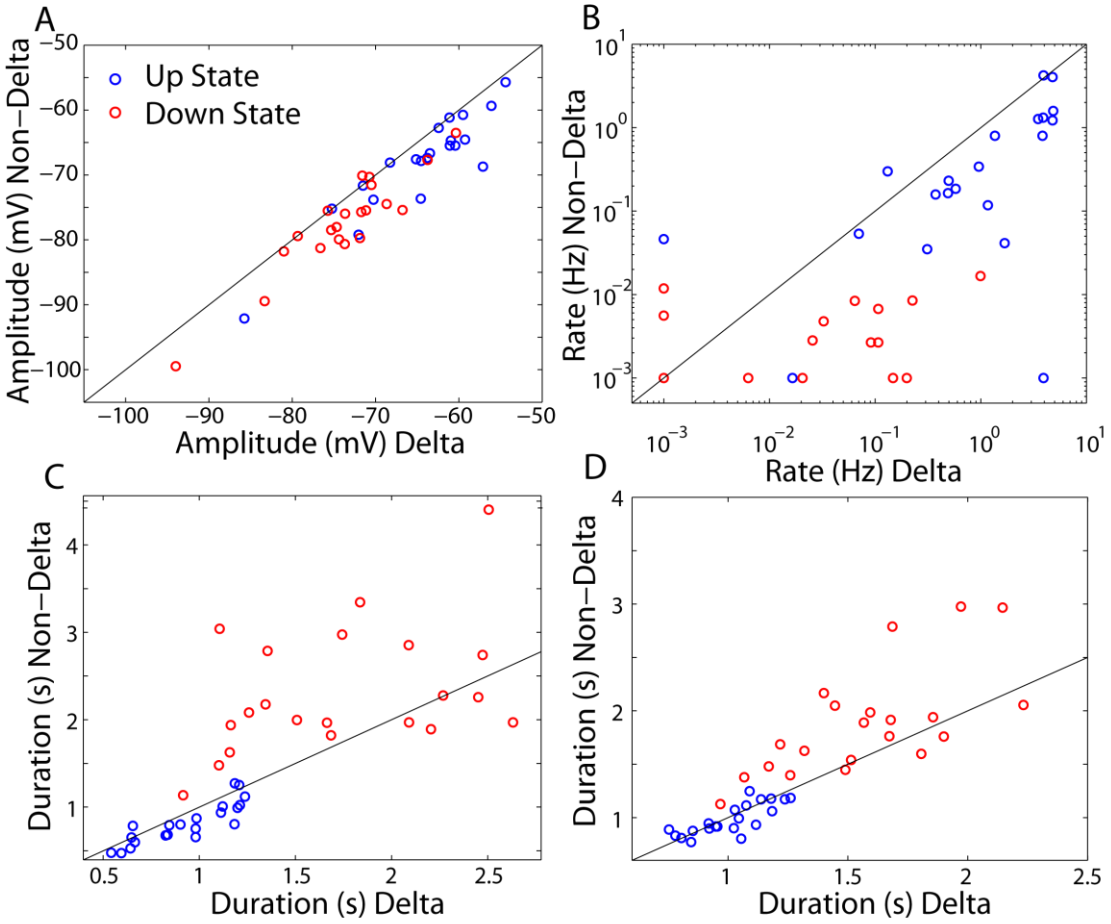

Figure 4-4: Superficial LEC neurons, and the Ncx LFP, are more active during delta epochs.

A) *Scatter plot of the average UP (blue circles) and DOWN (red circles) state amplitudes in non-delta epochs vs. delta epochs. The LEC neurons were significantly more depolarized during delta epochs in both the UP and DOWN states.* ***B)*** *Average UP and DOWN state firing rates during non-delta epochs vs. non-delta epochs. Firing rates increased by 0.6Hz on average during delta epochs.* ***C)*** *Average MP UP and DOWN state durations during non-delta epochs are plotted against the corresponding state durations during delta epochs. MP UP states were significantly longer during delta epochs while MP DOWN states were significantly shorter.* ***D)*** *Similar to C, but for Ncx UP and DOWN state durations. While the duration of Ncx UP states was not significantly different during delta epochs, Ncx DOWN states were significantly shorter, and hence the Ncx LFP also spent a significantly larger portion of time in the UP state during delta epochs.*

The fact that Ncx UDS also had a greater probability of occupying the UP state during delta epochs suggests that the delta epochs are reflective of a global change in network state. To better understand the relationship between LEC delta oscillations and global network activity, we compared spectral properties of the LEC MP, as well as Ncx and Hpc LFP during delta and non-delta epochs. For the LEC MP, the difference in the log power spectrum during delta and non-delta epochs indeed revealed a peak at 2-4Hz, as expected (Figure 4-5A). The increased power during delta epochs was not confined to the delta-band however, with substantial increases in power present across a broad range of frequencies (>4Hz). Consistent with the observed increases in firing rate and depolarization during delta epochs, this suggests that there is increased synaptic activity during the delta epochs. For the Ncx LFP, the differences in the delta and non-delta power spectra were somewhat smaller, however a clear peak at 2-4Hz, as well as a slight increase in higher frequency power were observed (Figure 4-5B). Thus, the delta oscillations themselves were present in the Ncx LFP, although their amplitude was smaller than in the LEC MP. The Hpc LFP also showed a similar pattern of increased power in the delta-band, and across a range of higher frequencies during delta epochs (Figure 4-5C).

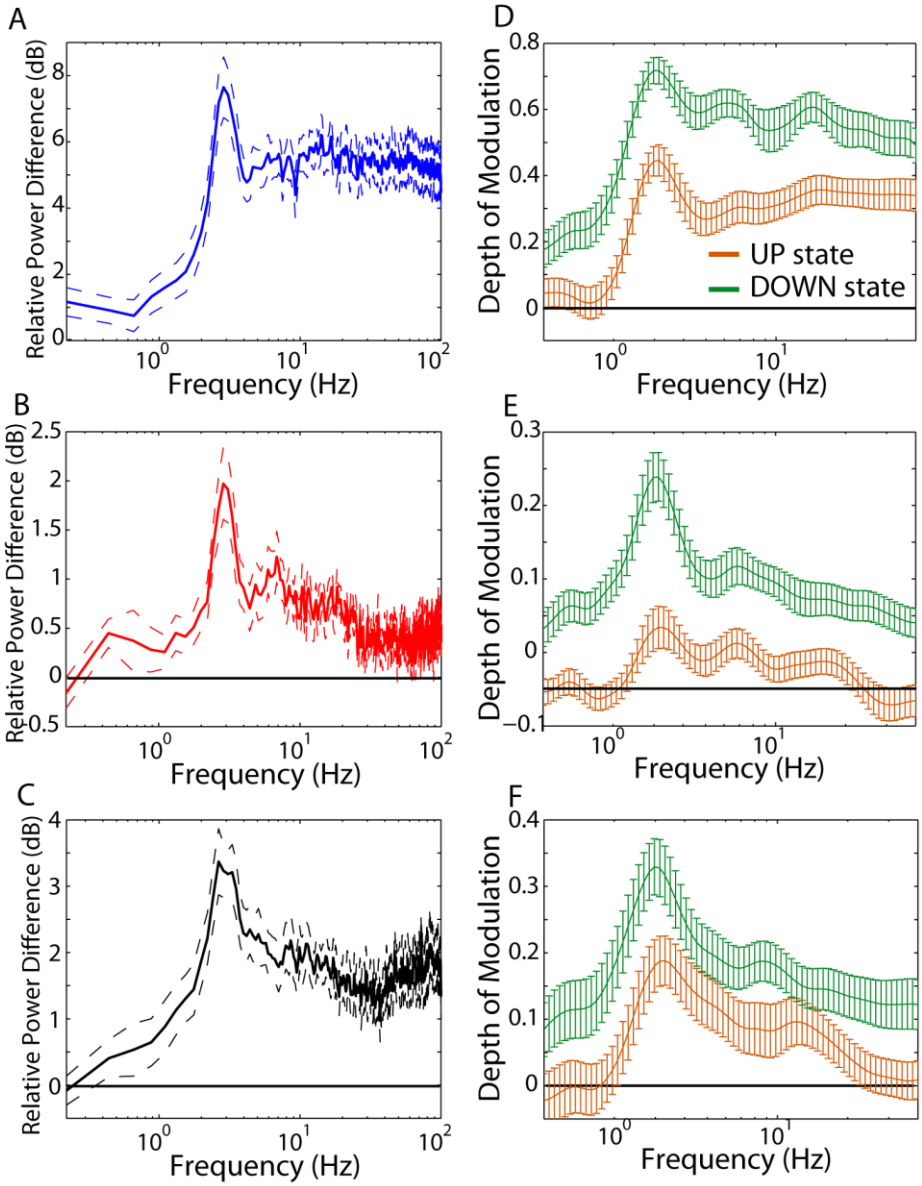

Figure 4-5: Power spectral properties during delta and non-delta epochs.

A) *Difference of the MP log power spectrum during delta and non-delta epochs.* ***B)*** *Same as A for the Ncx LFP.* ***C)*** *Same as A-B for the Hpc LFP.* ***D)*** *The depth of modulation of the amplitude spectrum (computed from the continuous wavelet transform) during delta vs. non-delta epochs is shown separately for the MP UP (orange trace) and DOWN (green trace) states.* ***E)*** *Same as D for the Ncx LFP.* ***F)*** *Same as D-E for the Hpc LFP.*

To determine whether these changes in the spectral properties of the MP and LFP signals were present within both the UP and DOWN states, we computed the continuous wavelet transform (CWT, see *Spectral Analysis*). The CWT provides optimal simultaneous time and frequency resolution, and thus allowed us to estimate the changes in spectral properties during delta epochs separately within the LEC DOWN and UP states. While the amplitude spectrum of the LEC MP during delta epochs was increased in the delta and high-frequency bands within both the UP and the DOWN states, the increases were far more pronounced within the DOWN states (Figure 4-5D). This is consistent with the visual impression of the example trace in Figure 4-2D, where the delta oscillations in the MP appear to persist with a similar amplitude across the UP and DOWN states. Since there is very little signal variance during the DOWN states, the delta oscillations contribute more to the total variance within the DOWN states. This contrast between the UP and DOWN states was even more apparent for the Ncx LFP, where there was virtually no difference in the UP state signal amplitude, even at delta frequencies (Figure 4-5E), while there were clear increases across a broad range of frequencies, particularly delta-frequencies, within the DOWN states. For the Hpc LFP, the differential increases in signal amplitude within the UP and DOWN states were also apparent, although to a lesser degree than for the Ncx LFP (Figure 4-5F).

We next sought to quantify the relationship between delta oscillations in the LEC MP and those in the Ncx and Hpc LFP by computing the coherence spectra between these signals (*see Spectral Analysis*). The average coherence between the LEC MP and the Ncx (Figure 4-6A) and Hpc (Figure 4-6B) LFP were indeed higher in the delta band during delta epochs compared with non-delta epochs. Comparing the coherence values at the corresponding frequency of MP delta

oscillations, MP-LFP coherence was more than a factor of two higher during delta epochs for both the Ncx (delta: 0.46±0.038; non-delta: 0.20±0.021; $p=1.2x10^{-4}$), and Hpc LFP (delta: 0.45±0.041; non-delta: 0.20±0.020; $p=9.2x10^{-5}$). Interestingly, MP-LFP coherence was actually significantly decreased in the lower-frequency UDS range (<2Hz) during delta epochs, particularly for the Ncx LFP (Figure 4-6A). This suggests a slight decoupling of LEC and Ncx UDS during delta epochs, even though the signals remained strongly coherent at UDS frequencies during delta epochs. Consistent with the power spectral analysis, we also found that the (wavelet-based) coherence spectra between LEC MP and both Ncx (Figure 4-6C) and Hpc (Figure 4-6D) LFPs showed much higher values within the DOWN states compared to UP states, particularly in the delta band and its first harmonic (4.5-8Hz).

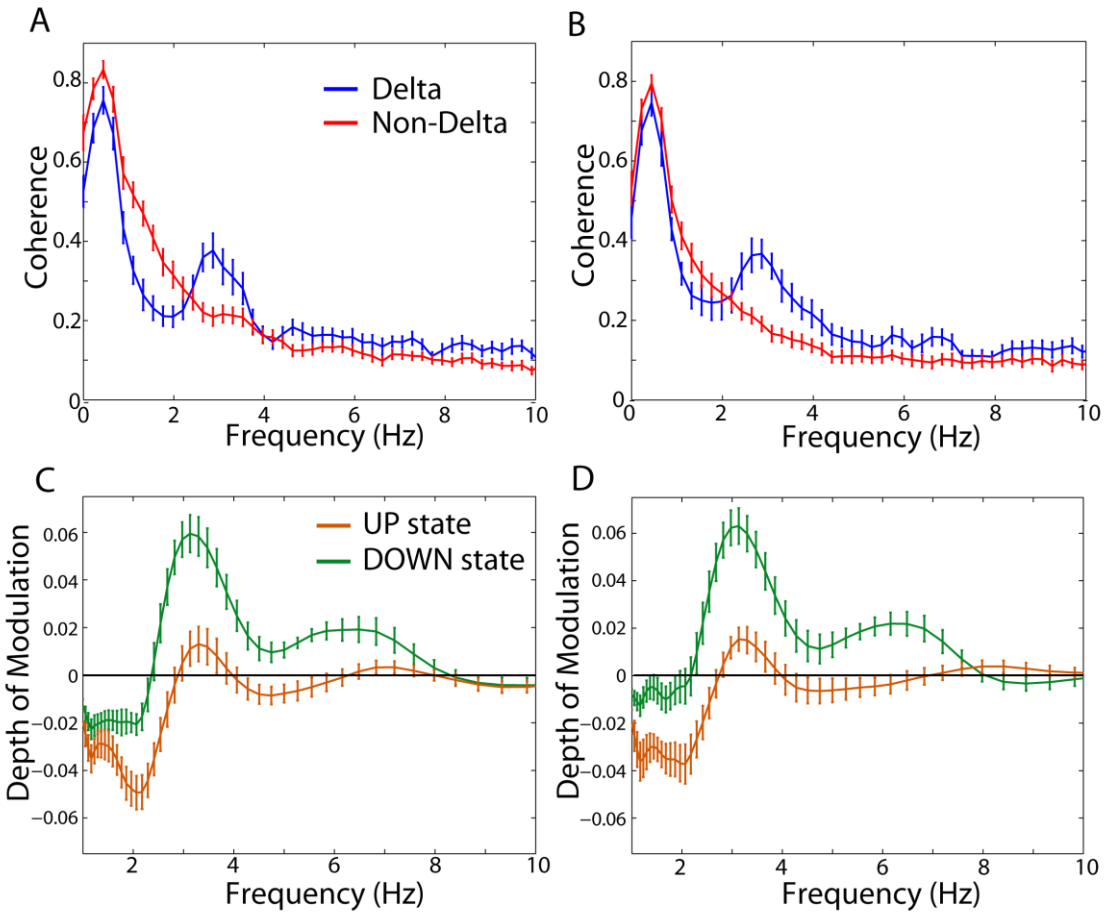

Figure 4-6: Delta oscillations in LEC MP and the LFP are coherent, particularly during DOWN states.

A) *The coherence spectrum between LEC MP and Ncx LFP during delta epochs (blue trace) and non-delta epochs (red trace). While the MP-Ncx LFP coherence in the UDS-band (< 2Hz) is higher during non-delta epochs, delta-band coherence is significantly higher during delta epochs.* ***B)*** *Same as A for the MP-Hpc LFP coherence spectrum, showing similar results.* ***C)*** *The depth of modulation of the MP- Ncx LFP coherence spectrum during delta vs. non-delta epochs is computed separately during the MP UP (orange trace) and DOWN (green trace) states. While delta-band coherence is unchanged within the MP UP states, there were clear increases in delta-band coherence within the DOWN States.* ***D)*** *Same as C but for the MP- Hpc LFP coherence showing similar results.*

Because the delta and UDS oscillations were present concurrently, we wanted to determine whether the two oscillations were related to one another, or whether they were independent phenomena. Figure 4-7A shows the same example trace as in Figure 4-2D with the MP filtered

separately in the delta (2-4Hz) and UDS (0.05-1.5Hz) bands. In this example it appears that the delta oscillations are phase-locked to the DOWN-UP and UP-DOWN state transitions, with DOWN-UP transitions occurring at peaks of the delta oscillation and UP-DOWN transitions occurring during valleys of the delta oscillations. Figure 4-7B shows the same example trace along with the simultaneously recorded Ncx LFP and the delta-band filtered Ncx LFP. Since the Ncx LFP delta oscillations are closely related to the LEC MP delta, Ncx LFP delta also appear to be phase-locked to the MP state transitions. It is important to note that this phase-phase coupling is a different form of cross-frequency coupling than the phase-amplitude coupling often described [80,85,86], since it does not depend on the amplitude of either component signal, but only on the relative timing between them.

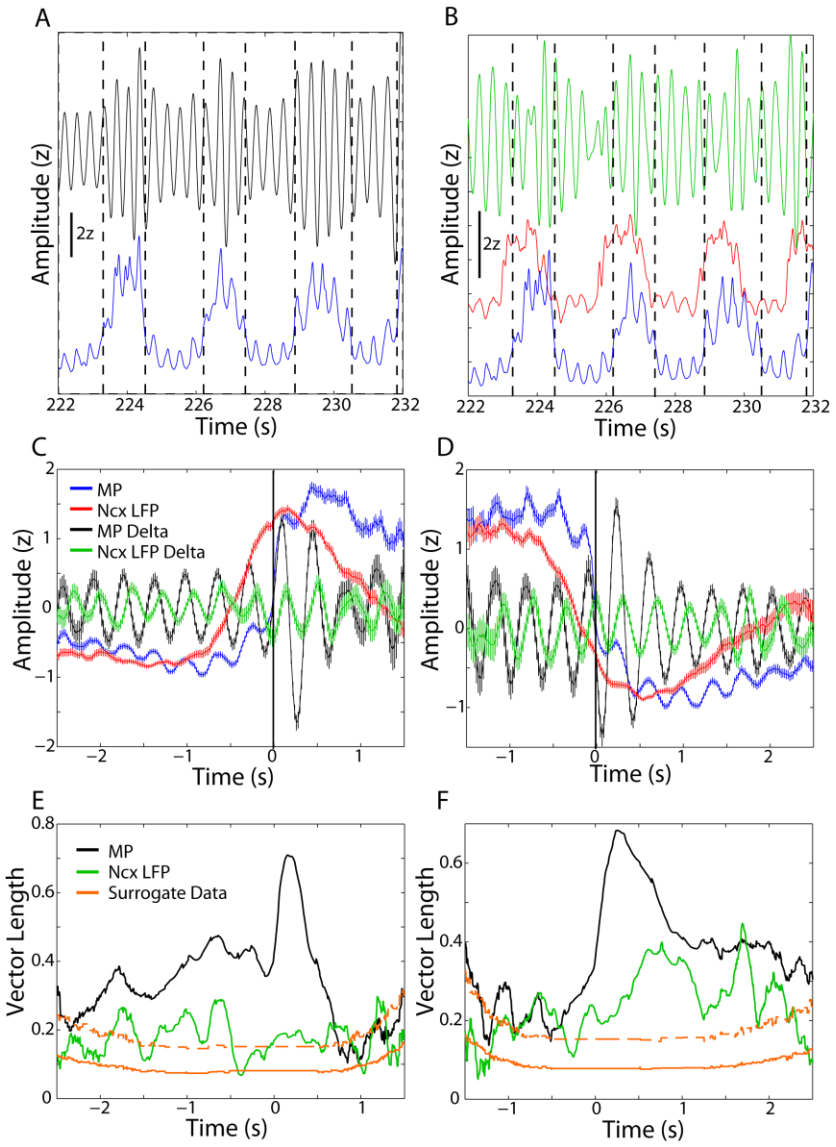

Figure 4-7: Phase-locking of MP and LFP delta oscillations to MP state transitions in an example recording.

A) *The same example MP trace as Figure 4-2D is shown (blue trace) along with the MP signal filtered in the delta band (2-4Hz; black trace). MP state transitions are shown as vertical dashed lines.* ***B)*** *Same example trace as in A, but showing the simultaneous Ncx LFP (red trace), and the delta-band filtered LFP (green trace).* ***C)*** *MP DOWN-UP triggered averages of the MP (blue), delta-band MP (black), Ncx LFP (red), and delta-band LFP (green) for the example neuron.* ***D)*** *Same as C for the MP UP-DOWN transition.* ***E)*** *The mean resultant vector length of the set of MP (black) and Ncx LFP (green) delta phases at each time lag relative to the MP DOWN-UP transition. The solid and dashed orange traces indicate the mean and upper 95% confidence interval of the null distribution of vector lengths generated using a boot-strapping procedure.* ***F)*** *Same as E for the MP UP-DOWN transition.*

To quantify this cross-frequency coupling, we computed the UP-DOWN and DOWN-UP triggered averages of the delta-band MP and LFP signals, using only the state transitions occurring during delta epochs (see *Delta-UDS Coupling*). For the example LEC neuron, both the DOWN-UP triggered average (Figure 4-7C) and the UP-DOWN triggered average (Figure 4-7D) of the delta-band MP showed clear delta oscillations, indicating strong phase-locking of delta oscillations to UDS for this cell. Consistent with the visual impression from the example traces (Figure 4-7A,B), the delta-band Ncx LFP was also locked to both the MP DOWN-UP (Figure 4-7C) and UP-DOWN (Figure 4-7D) transitions. We determined the significance of the delta-UDS coupling by computing the mean resultant vector length of delta phases across the ensemble of DOWN-UP or UP-DOWN transitions as a function of time relative to the state transition. These empirical vector lengths were compared to null distributions generated using a bootstrapping procedure (see *Delta-UDS Coupling*). Strong delta-UDS phase-locking would result in a narrow distribution of delta phases, and a large resultant vector length. Indeed, for the example LEC neuron, significant UDS-delta coupling was present for more than 2s during the DOWN states preceding the DOWN-UP transition, as well as for several delta cycles within the ensuing UP states (Figure 4-7E). Similarly, significant phase-locking of delta oscillations was observed to the MP UP-DOWN transitions for ~1s preceding the UP-DOWN transition, and for well over 2s during the following DOWN state (Figure 4-7F). The Ncx LFP delta oscillations which occurred primarily during DOWN states, had similarly significant phase-locking to the MP DOWN-UP transition and UP-DOWN transitions, showing particularly clear locking to the UP-DOWN transitions within the ensuing DOWN states (Figure 4-7E,F).

In order to compare delta-UDS coupling across the ensemble of LEC neurons which had significant delta oscillations we performed a local time-rescaling procedure, measuring time in units of the local delta oscillations (see *Delta-UDS Coupling*). This allowed for comparisons across different delta epochs within a recording, as well as across different LEC neurons, where the precise frequency of the delta oscillations could vary substantially. Clear delta-UDS coupling was apparent in the ensemble average MP DOWN-UP (Figure 4-8A) and UP-DOWN (Figure 4-8B) transition-triggered averages of the delta-band MP, Ncx LFP, and Hpc LFP. The strongest phase-locking was observed in the delta-band MP, however the Ncx and Hpc LFPs still showed substantial delta-locking to the MP UDS. The fraction of datasets (of those with significant MP delta oscillations) which had significant phase-locking of delta-band MP and LFP signals to MP DOWN-UP and UP-DOWN state transitions are shown in Figure 4-8C,D respectively. It is worth noting that the statistical power of this test is limited by the number of state transitions observed during delta epochs, as well as by the limited durations of the UP and DOWN states themselves. Despite this limitation, more than half of the data showed significant phase-locking of the DOWN state MP delta oscillations to the MP DOWN-UP (Figure 4-8C) and UP-DOWN (Figure 4-8D) transitions for more than 4 cycles preceding and following the transitions respectively. Phase-locking of delta during the UP states was more difficult to ascertain given the shorter UP state durations, but UP state delta oscillations in the MP were phase-locked to the MP UP-DOWN transition for several delta cycles preceding the transition, suggesting that the delta oscillations may serve to time the termination as well as initiation of UP states. As apparent in the transition-triggered averages, Ncx and Hpc LFP delta oscillations also showed significant phase-locking to MP DOWN-UP and UP-DOWN transitions lasting well into the preceding and ensuing DOWN states (Figure 4-8C,D). We also computed the strength of delta-

locking to the DOWN-UP and UP-DOWN transitions of the Ncx LFP but found that phase-locking was much weaker than to the MP UDS (Figure 4-8E,F). This may indicate that the timing of LEC, but not Ncx, UDS are influenced by the delta oscillations. However, this result could also arise because the Ncx UDS were measured from the LFP signal, and it was thus more difficult to determine the precise timing of Ncx UDS.

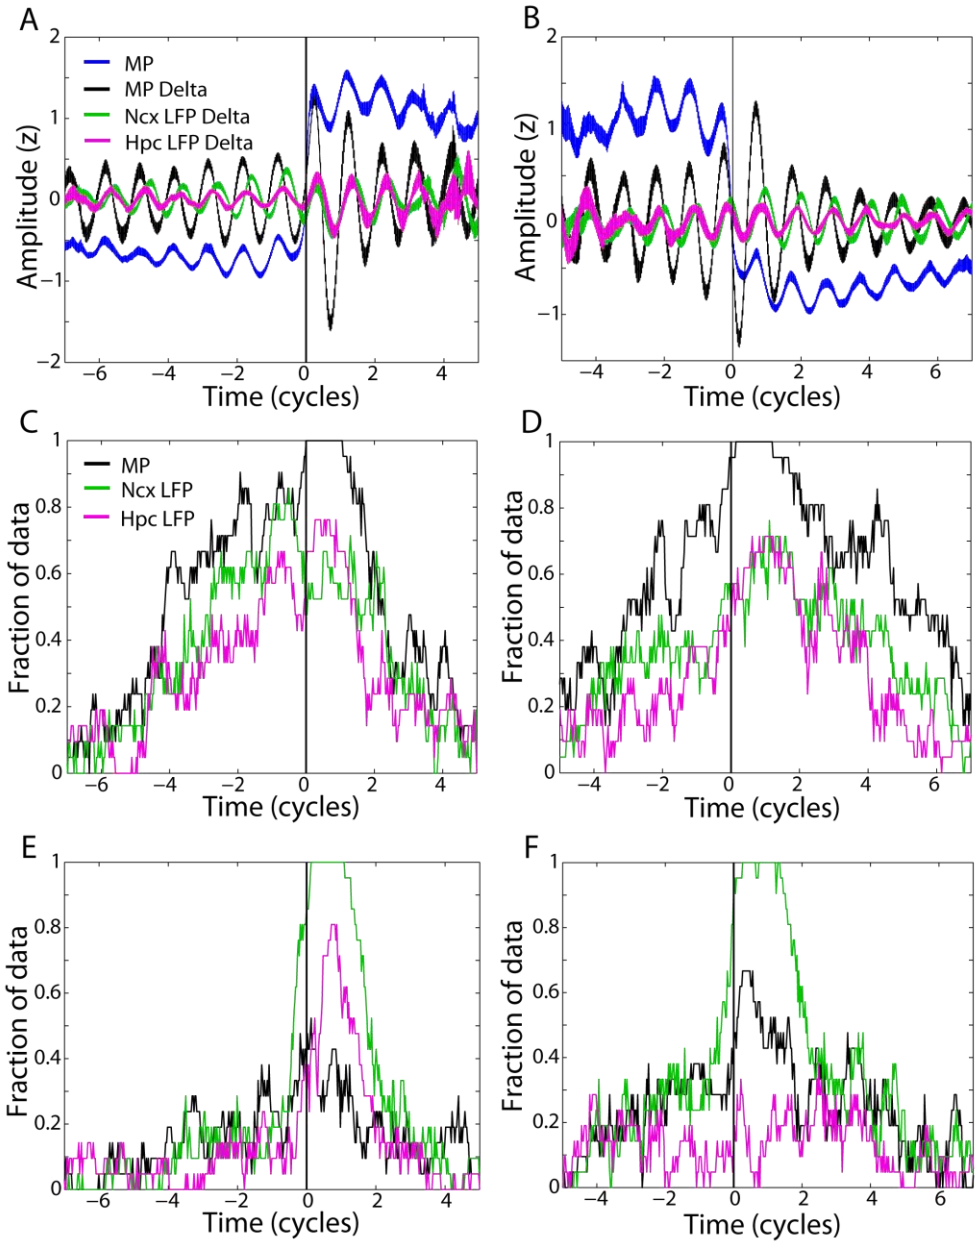

Figure 4-8: Significant MP and LFP delta-UDS coupling across the ensemble of LEC neurons

A) *Ensemble-average MP DOWN-UP triggered averages of the MP (blue), delta-band MP (black), delta-band Ncx LFP (green), and delta-band Hpc LFP (purple).* ***B)*** *Same as A for the MP UP-DOWN transition.* ***C)*** *The fraction of the data which was found to have significant delta phase-locking to the MP DOWN-UP transition as a function of time relative to the transition is plotted for the MP (black), Ncx LFP (green), and Hpc LFP (purple).* ***D)*** *Same as C for the MP UP-DOWN transition.* ***E)*** *Same as C for the Ncx DOWN-UP transition. Phase-locking of delta oscillations in the MP and LFP was substantially weaker to the Ncx DOWN-UP transition compared to the MP DOWN-UP transition.* ***F)*** *Same as E for the Ncx UP-DOWN transition.*

Discussion

Here, we show that superficial neurons in the LEC, but not the MEC, of urethane-anesthetized mice undergo spontaneous transitions into states with prominent 2-4Hz delta oscillations that occur simultaneously with the ongoing UDS. These states occurred in a majority of superficial LEC neurons, including several different cell types in both layers 2 and 3, although delta oscillations were more prevalent in layer 2 than layer 3 LEC neurons. For cells which exhibited delta oscillations, the delta activity was confined to discrete epochs lasting for ~100s and constituting about 1/3 of those datasets. During these epochs the delta oscillations occurred throughout both the UP and DOWN states, appearing more prominently during the quiescent DOWN states. While the delta epochs were accompanied by significant changes in the properties of the LEC neurons' UDS, including more depolarized UP and DOWN state amplitudes, increased firing rates, and increased time spent in the UP state, delta epochs were not restricted to the LEC. Rather, the delta epochs were accompanied by global changes in the UDS and power spectral properties of the Ncx and Hpc LFP; in particular, an increased probability of Ncx UP states. The delta oscillations themselves were also present in both the Ncx and Hpc LFPs, most prominently during DOWN states, and these LFP delta oscillations were significantly coherent with LEC MP delta. Finally, we show that the delta oscillations were phase-locked to the state transitions of the ongoing LEC UDS.

It is important to note the similarity of the delta oscillations reported here to theta oscillations that have been reported in anesthetized and naturally sleeping animals [11,129,269,279]. While theta oscillations occur at a similar, though somewhat higher (3-5Hz), frequency, and are also observed to appear during changes of network state [129], theta oscillations exhibit several

properties which distinguish them from the oscillations reported here. Most importantly, theta oscillations in anesthetized animals have been shown to occur in the absence of UDS during 'desynchronized' or 'activated' states, such as elicited by tail-pinch [281], stimulation of brainstem cholinergic nuclei [129,145,250], or cholinergic agonist injection [129]. Similar results are also observed in naturally sleeping animals, whereby theta oscillations occur during REM sleep, in the absence of the slow oscillations. Further, theta oscillations during activated network states are present in the Hpc [11,129,269,279] and EC [269] LFPs, but have not been shown to occur in the Ncx LFP as we describe here. Instead, we believe that the 2-4Hz oscillations shown here are analogous to delta oscillations reported in thalamocortical neurons [65,67] as well as the Ncx EEG [150] of anesthetized and naturally sleeping animals. These thalamocortical delta oscillations occur simultaneously with UDS during deep sleep and anesthesia [65,130,140], and appear most prominently during the DOWN states, similar to the oscillations described here.

There has been substantial interest in the oscillatory properties of entorhinal neurons in the 4-12Hz frequency range, first demonstrated by Alonso et al. to occur in MEC layer 2 stellate neurons [57], but not superficial MEC pyramidal neurons [272,273]. These intrinsic subthreshold oscillations were subsequently shown to involve persistent sodium and hyperpolarization-activated cation currents [274,275]. Indeed, layer 2 MEC stellate cells have been shown to produce much stronger theta oscillations than pyramidal neurons in layers 3 or 5 during activated states in anesthetized animals [269]. However, little is known about the theta oscillatory properties of the LEC. Recently, theta oscillations in both the MEC LFP and single unit activity during behavior were shown to be substantially stronger than those in the LEC [277].

Taken together with our results, this suggests that superficial neurons in both the LEC and MEC can oscillate in the 2-12Hz frequency range, but do so under different conditions. Intrinsic properties of MEC layer 2 stellate neurons are likely to produce theta oscillations during activated network states (such as during desynchronized epochs of anesthesia, and REM sleep), as well as during behavior. On the other hand, superficial LEC neurons (particularly the fan and multipolar cells in layer 2) are predicted to participate in the lower-frequency (1-4Hz) delta oscillations that occur concurrently with UDS (slow-oscillations) during anesthesia and deep sleep. These differences in the oscillatory behaviors of MEC and LEC neurons could also be related to their differential functional roles during behavior [258].

A number of studies have shown that UDS organize the higher-frequency structures present during anesthesia and non-REM sleep, including thalamocortical delta oscillations and spindles [150], as well as hippocampal ripples [162]. Coupling between the delta and theta oscillations [82], as well as between theta and gamma oscillations have also been demonstrated in awake animals [80,85,86]. This has lead to the suggestion that a hierarchy of interrelated oscillations exists, both during behavior as well as during anesthesia and natural sleep [82]. Interestingly, the majority of cross-frequency coupling that has been reported is of a different nature than that described here, with the amplitude of the higher-frequency oscillations being coupled to the phase of the lower-frequency oscillation. For the delta-UDS coupling, the amplitude of delta oscillations does not change throughout the UP or DOWN states, but rather UDS state-transitions are phase-locked to the delta oscillations, resulting in a phase-phase coupling between the two oscillations.

The fact that delta oscillations observed in the LEC MP are coherent with Ncx and Hpc LFP delta suggests that these oscillations may arise from, or be synchronized by, a common input such as diffusely projecting thalamic nuclei. Interestingly, the reuniens nucleus indeed provides strong inputs to the hippocampal CA1 region in addition to diverse thalamocortical projections, and has even been shown to project more strongly to superficial LEC than MEC [282], making it a possibility for such a common driving force. On the other hand, the LEC itself projects to the hippocampus [180], and has bidirectional connections with diverse Ncx areas [183]. Further, the intrinsic subthreshold oscillations of EC neurons have been shown to synchronize upon cholinergic activation [283], suggesting that such synchronization might occur spontaneously within the EC and be passed along to the neocortex and hippocampus. Thus, the network of LEC neurons could themselves be the generators of globally synchronous delta oscillation. It will be important to determine whether other LEC neurons (in particular the principal neurons of layer 5) generate similar delta oscillations as well.

If the LEC delta epochs shown here are indeed analogous to, or reflective of, the thalamocortical delta oscillations reported by Steriade et al., it is perhaps surprising that the network is in a more activate state during delta epochs compared to non-delta epochs, with LEC neurons being more depolarized and having increased spike rates, and both LEC MP and Ncx LFP spending more time in the UP state. Considering that the thalamocortical delta oscillations have been shown to occur only during times of decreased cholinergic tone and correspondingly reduced corticothalamic input [65,130], one might instead expect the delta epochs to occur during similar periods of increased hyperpolarization and longer DOWN states [130,145]. It is possible that the non-delta UDS epochs reflect a cholinergic tone which is too low even for generating

sustained delta oscillations. This could arise for instance because the thalamocortical delta oscillations can be elicited and sustained by corticothalamic volleys [65]. The relationship between the LEC delta oscillations shown here and thalamocortical delta oscillations reported elsewhere will thus be an important avenue for further study.

Regardless of the underlying cause of the LEC MP delta oscillations, their coupling to ongoing UDS and their widespread coherence with Ncx and Hpc LFP delta oscillations suggest that they could serve an important role in controlling the relative timing of UDS in different brain regions. Specifically, because delta oscillations during the DOWN state are phase-locked to the LEC DOWN-UP transition, peaks of the delta oscillation could serve to initiate LEC DOWN-UP transitions, particularly following the afferent Ncx DOWN-UP transition. Similarly, because the UP-DOWN transitions are phase-locked to delta oscillations within the preceding UP state, delta oscillations during the UP state could trigger the termination of UP states, and establish a clock signal for timing the ensuing DOWN-UP transition. It will be important to determine the extent of delta synchronization amongst LEC neurons, and the participation of LEC interneurons in the delta oscillations. An interplay between local excitatory and inhibitory delta oscillations might create narrow windows during which the LEC network could transition into the UP state, resulting in precise timing of LEC UDS relative to Ncx UDS. This could be particularly important considering that precise spatiotemporal patterns of spiking have been observed following the Ncx DOWN-UP transition [138]. UDS transition timing by LEC delta oscillations would then be conveyed to hippocampal neurons, and could serve to mediate cortico-hippocampal interactions, which have been suggested to play an important role in memory consolidation processes during natural sleep [114,165,173]. The observations that delta oscillations are also

found in the Hpc LFP, and that these delta oscillations are coherent with LEC MP delta, further show that LEC delta oscillations indeed have an effect on downstream hippocampal activity, and thus could play an important role in these processes.

Chapter 5: **Impaired hippocampal spatial selectivity and network oscillations in GluA1 KO mice**

Introduction

Since the discovery by O'Keefe and Dostrovsky of hippocampal cells which fire in a spatially selective fashion ('place cells'), the hippocampus has been hypothesized to provide a cognitive map of an animal's environment [175,176]. Consistent with this theory, lesions of the hippocampus have been shown to impair performance on several tests of spatial processing, including learning of a fixed location (spatial reference memory (SRM)) [193,284], and generation of spatial behavior contingent on short-term memory (spatial working memory (SWM)) [284]. In addition to spatially selective spiking of individual neurons, the hippocampus also shows theta and gamma rhythmic network oscillations during exploration [47,211]. The mechanisms that produce these two aspects of hippocampal activity, as well as their contributions to hippocampus-dependent spatial behavior, are not well understood.

Genetic deletion of the GluA1(GluR1, GluR-A) AMPA receptor subunit has been shown to produce specific alterations of synaptic plasticity in the hippocampus. Hippocampal neurons from these GluA1 knockout (KO) mice show reduced tetanically-induced long-term potentiation (LTP) [285], as well as an impaired short-term component of theta-burst plasticity [286,287]. However, the long-term component of theta-burst plasticity, as well as spike-timing dependent plasticity (STDP), are unaltered [286,287]. Additionally, GluA1 KO hippocampal neurons do not

show the large distance-dependent dendritic scaling seen in wild-type (WT) neurons [288]. Behaviorally, the GluA1 KO mice display profound impairments in several different SWM tasks, but perform as well as WT animals on SRM tasks [284,285,289]. Thus, the GluA1 KO animals present an interesting opportunity to study the relationship between these different forms of spatial behavior and hippocampal activity, as well as the role of different forms of synaptic plasticity and dendritic-scaling in generating these phenomena.

To this end, we measured hippocampal dorsal CA1 place cells' activity in GluA1 KO mice. Compared to WT neurons, there was a large reduction in all measures of hippocampal spatial selectivity in the KO neurons. Across trials, the firing rate maps of KO neurons also showed more random variability, and less systematic change than their WT counterparts. Further, the KO mice had greatly reduced hippocampal theta and gamma oscillations, and the coupling between theta and gamma oscillations was similarly impaired. These results suggest that while hippocampal integrity is required for SRM and SWM, spatial selectivity and plasticity of CA1 neurons, as well as the theta and gamma rhythmic hippocampal network oscillations, are not required for SRM. These phenomena may, however, be needed for SWM.

Methods

All experiments were carried out in 11 wild-type (C57/BL6) and 9 GluA1 knockout [285] age-matched male mice. The mice were housed individually, maintained on a 12-h light/dark cycle and had *ad libitum* access to food and water unless otherwise specified. All experiments were conducted in a double-blind fashion with respect to mouse genotype. All experiments were conducted in accordance with the animal welfare guidelines of the Max Planck Society.

Electrophysiology

Mice were implanted with hyperdrives above the right dorsal CA1 region of the hippocampus (AP-2.0 mm, L-1.5mm with respect to bregma). Each hyperdrive contained up to 4 independently movable tetrodes and a reference electrode [188,290]. Tetrodes were made of four 15-μm nichrome wires (Kanthal, Palm Coast FL), which were twisted and heat-fused together. Two stainless steel screws were implanted over the opposite frontal cortex (for anchoring) and the cerebellum (for ground).

Behavioral procedures

One week after surgery, the mice were trained to run back and forth on a linear track for food reward (sweetened milk) located at the opposite ends of the track. The track (length 140cm, width 3cm) was made of wood and painted black. Such a narrow track ensured that the mice walked through the same region of space when moving in both directions of the journey and across trials. Throughout the entire training and further recording the mice were food-deprived and maintained at ~85% of their postoperative *ad libitum* body weights.

Recording methods

Throughout the training period, tetrodes were advanced gradually over the course of many days to place them in the CA1 pyramidal cell layer. The arrival of each tetrode into the hippocampus was recognized by several criteria, including the presence of 100-300 Hz 'ripples' in the local field potential (LFP) [291,292,293], the polarity of 'sharp waves' in the LFP [294], and the sudden appearance of multiple simultaneously recorded cells with a complex spike discharge. During recording, the mice were attached to a unitary gain head-stage preamplifier (HS-16; Neuralynx; Bozeman, MT) via a cable suspended on the supporting metal string to minimize an additional

load to the animal's head. Signals were filtered and either amplified with respect to ground, or differentially amplified against the reference electrode, by Lynx-8 programmable amplifiers (Neuralynx). Whenever the amplitude of the spike signal exceeded a predetermined threshold, each tetrode channel acquired a 1-ms sample of data at a rate of 32kHz. These spike samples were time-stamped, amplified by a factor of 5,000-10,000 and stored on a personal computer running Cheetah data acquisition software (Neuralynx). The animal's position and head direction were obtained by overhead video tracking (Cohu iDome, USA) of two light-emitting diodes (red and blue) attached to the headstage. Video recording was made with spatial resolution of 0.25cm/pixel and a sampling rate of 50Hz. Daily recording sessions consisted of 8-40 laps on the linear track surrounded by rest/sleep sessions (20min) on an elevated platform. The platform sessions were used to confirm stability of recordings made during track running. Following completion of the experiments the mice were deeply anesthetized and electrolesion (200µA for 4 sec via one channel of each tetrode with respect to the ground screw) was performed to reconstruct the position of the tetrodes.

Data analysis

The data were analyzed offline using custom software written in Matlab (MathWorks). Since the data were often not distributed normally statistics are reported as medians and inter-quartile ranges (IQR), and hypothesis testing was done using the Wilcoxon rank sum test. For all the statistical tests a p-value <0.05 was considered to be statistically significant.

Unit isolation

Single units were manually isolated offline using graphical cluster-cutting software MClust-3.4 (A. David Redish; available at: http://redishlab.neuroscience.umn.edu/). The procedure

consisted of displaying all two-dimensional projections of the multidimensional parameter space (peak, energy, valley) for all spike waveforms and applying boundaries to each apparent unit cluster. Examples of two out of six possible projections of spike amplitudes on a tetrode are shown in Figure 5-1A(D) for WT (KO) mice. Based on visual inspection of the unit clusters, all isolated units were split into 3 categories. Units that were clearly isolated from other units on at least 3 (out of 6) projections of spike amplitudes on a tetrode were category 1 (18 in WT, 21 in KO). Cells that were well isolated on at least 2 projections were category 2 (52 in WT, 36 in KO). Cells that were isolated on at least 1 projection were category 3 (67 in WT, 31 in KO). The visual impression of cell isolation was confirmed by computing the isolation distance [295]. Cells in the KO mice were at least as well isolated as those in WT (WT: median=31, IQR=24-39; KO: median=41, IQR=34-59, p=1.5x10^{-8}). All detected genotype-related differences in place field characteristics were found to be consistently independent from the cell category. Putative pyramidal cells in the CA1 region were distinguished from putative interneurons using standard criteria [202,290,296,297]. The classification criteria for inclusion of a unit into the pyramidal cell category was that the unit must fire at least a small number of complex spike bursts during the recording session, have spike width (peak to valley) of at least 250ms, and an overall mean spiking rate of <2.5Hz while the animal ran on the track. Goal locations were defined as the regions within 7.5cm of the reward wells. The subsequent analysis of spiking rate maps was restricted to only those cells that had a clear refractory period in the autocorrelogram, a mean spiking rate >0.1Hz and fired >100 spikes during track-running (excluding data at the goal locations and during periods of immobility). Eventually, 86 of 137 WT and 75 of 88 KO initially isolated putative pyramidal cells met these criteria.

Exclusion of periods of immobility:

The position data was first linearized to represent the position of the animal along the length of the track. The linearized running speed of the animal along the track was estimated by smoothing the linearized position data using a Gaussian smoothing kernel (sigma=100ms) before computing its approximate temporal derivative. Data when the running speed of the animal was less than 5cm/s were excluded from all analysis to ensure that 'non-spatial' behaviors such as grooming or rearing did not contribute to the results. Furthermore, trials in which the mouse did not run all the way from one goal location to the other were excluded from analysis.

Quantification of spatial spiking properties

Since most place cells on linear tracks are directional, the left-to-right and right-to-left trials were treated separately [26,194,195,298]. The total number of spikes of a cell in each spatial bin (0.5cm) and the time spent by the animal in that bin were computed. These spike count and occupancy maps were smoothed slightly using a Gaussian smoothing kernel (sigma = 1cm). The spiking rate map was then computed by dividing the spike count in each bin by the occupancy in the corresponding bin. These spiking rate maps were then smoothed again with a Gaussian kernel (sigma = 2cm).

Several quantitative measures were used to describe and compare properties of the spiking rate maps in the WT and KO mice:

(1) Peak spiking rate over the entire environment was defined as the maximum spiking rate across all spatial bins.

(2) Spatial information content per spike was defined by the formula [299]:

$$I = \frac{1}{\bar{\lambda}} \sum_{i=1}^{2L} P_i \lambda_i \log_2 \left(\frac{\lambda_i}{\bar{\lambda}} \right) \quad \textbf{(5-1)}$$

where λ_i is the averaged spiking rate of a cell in position bin *i*; P_i P_i is the probability of occupancy in spatial bin *i*; and $\bar{\lambda}$ $\bar{\lambda}$ is the spiking rate averaged across all times during track-running. Because the spiking rates in the opposite directions of the journey were treated separately, the spiking rate map vector λ was of length 2*L*, where *L* is the length of the track excluding the goal locations. We also estimated Shannon's mutual information between the position and spike count variables (referred to as "spike-position mutual information") as shown by Olypher et al [300]. This was defined as:

$$I(X,K) = \sum_{i=1}^{N} \sum_{k \geq 0} P_{k,x_i} \log \left(\frac{P_{x_i,k}}{P_{x_i} P_k} \right) \quad \textbf{(5-2)}$$

where *X* is the sequence of mouse positions and *K* is the corresponding sequence of binned spike counts, $P_{x_i,k}$ is the joint probability of observing *k* spikes and the mouse being in position bin *i*, P_{x_i} is the marginal probability of the mouse occupying position bin *i*, and P_k is the marginal probability of observing *k* spikes in a time bin. For this calculation we used time bins of 100ms and spatial bins of 4.5cm. To avoid undetermined quantities appearing in the sum in Equation (5-2) as a result of probability estimates equaling 0, we used the 'K-T estimate' for the probabilities [300] defined by:

$$\hat{p}_i^{KT} = \frac{n_i + 0.5}{N + 0.5J} \quad \textbf{(5-3)}$$

where n_i is the number of times outcome *i* is observed, *N* is the total number of observations, and *J* is the number of unique outcomes.

(3) Direction selectivity was defined as:

$$d = \left| \frac{\bar{\lambda}_L - \bar{\lambda}_R}{\bar{\lambda}_L + \bar{\lambda}_R} \right| \quad \textbf{(5-4)}$$

where $\bar{\lambda}_L$ is the average value of the spiking rate map for right-to-left movements and $\bar{\lambda}_R$ is the average for left-to-right movements. The direction selectivity defined in this way is bounded between 0 and 1 and measures the difference between the average spiking rates as the mouse moved in the two directions along the track, with 1 being a maximal difference and 0 being no difference.

(4) Coherence was used to estimate the local smoothness of the spiking profile. It was defined analogously to the definition used in [301]. First, we computed the spike count and occupancy maps using coarser 2cm wide spatial bins. Spike rate maps were then computed by dividing the spike count by the total occupancy in each bin without any smoothing. A separate, 'smoothed' version of the same spiking rate map was then computed by averaging with a uniform 4-point (nearest + next-nearest neighbor) kernel. The coherence was then defined as the z-transform of the correlation coefficient between the smoothed and un-smoothed rate maps.

(5) 'Trial-shuffled correlation', a measure of the stability of the spatial spiking of a cell across trials, was calculated in several steps. Each trial was defined as a complete journey executed by the mouse running from one end of the track to the other and back. First, the set of all trials from a session on the linear track was randomly split into two groups, so that both groups

contained equal numbers of trials. The z-transformed Pearson's correlation coefficient between the two spiking rate maps, computed by averaging across the trials in each group, was calculated. This procedure was repeated 500 times and the mean and standard deviation of the Fisher z-transformed Pearson's correlation coefficients were computed. These correlation coefficients, termed the 'trial-shuffled correlations', were used as an estimate of the amount of trial-to-trial variability in spatial spiking that did not occur systematically across trials: the larger the trial-shuffled correlations the less trial-to-trial variability the cell exhibited.

(6) 'Trial-ordered correlation' was defined similarly to the trial-shuffled correlation. In this case, the set of all trials from a session on the linear track was split into two groups, so that these groups constituted the first and the second halves of the trials from the session. The correlation coefficient (z-transformed) between the spiking rate maps estimated from the first and second halves of trials was taken as the 'trial-ordered correlation.'

(7) 'Systematic variation' in spatial spiking was then defined as the difference between the mean trial-shuffled correlations and the trial-ordered correlation, normalized by the standard deviation of the trial-shuffled correlations. This measure gave an estimate of the amount of systematic change, or the plasticity, of the spiking rate maps beyond that expected by chance. Higher values of the systematic variation indicate that more of the trial-to-trial variability in the cell's spatial spiking is due to changes that were systematic across trials.

(8) Place fields were defined as regions where the mean spiking rate crossed above a certain threshold for at least 10 contiguous bins (5cm). Two different threshold values were used: either a fixed 1Hz threshold, or an adaptive threshold of 20% of the overall peak spiking rate. Two

place fields were considered distinct if the average spiking rate crossed below the threshold for at least 10 bins (5cm) in between the field peaks.

(9) Place field width was defined as the distance between the outermost two bins where the spiking rate profile crossed the threshold.

(10) The number of place fields for a given cell was the sum of the number of fields in its left-to-right and right-to-left spiking rate maps.

(11) The signal-to-noise ratio (SNR) of a place cell was defined by the formula:

$$SNR = \frac{\langle \lambda_s \rangle - \langle \lambda_n \rangle}{\langle \lambda_s \rangle + \langle \lambda_n \rangle} \quad \textbf{(5-5)}$$

where $\langle \lambda_s \rangle$ is the mean spiking rate of the cell within all its place fields (as defined above), the so-called 'signal', and $\langle \lambda_n \rangle$ is the mean spiking rate outside all the place fields, the so-called 'noise'. The SNR is a dimensionless number bounded between -1 and +1 with higher values indicating more spatially restricted spiking. The SNR was calculated using place fields detected by both the 1Hz and the 20% thresholds, giving qualitatively similar results. When using the 1Hz threshold two of the KO cells had no place fields on the track.

(12) The rate map sparsity Ψ of a cell was defined as:

$$\psi = \frac{1}{2L} \frac{\left(\sum_{i=1}^{2L} \lambda_i \right)^2}{\sum_{i=1}^{2L} (\lambda_i)^2} x100 \quad \textbf{(5-6)}$$

where λ_i is the averaged spiking rate of the cell in a position bin *i*; and *L* is the length of the track excluding goal locations. Thus the spiking rate map sparsity for a given cell is the average of the left-to-right rate map sparsity and the right-to-left rate map sparsity. Sparsity is a positive dimensionless number between 0 and 100, representing the percentage of the environment over which each cell was active.

Quantification of temporal spiking properties

The empirical ISI distribution was computed for ISIs less than 1s using 100 logarithmically spaced bins ranging from 1ms to 1s. The average estimated probability of an ISI being in the range defined by each bin was then computed separately for WT and KO cells. The depth of modulation of the difference between these average WT and KO ISI probabilities was computed as:

$$\frac{\bar{P}_{WT}(ISI) - \bar{P}_{KO}(ISI)}{\bar{P}_{WT}(ISI) + \bar{P}_{KO}(ISI)} \quad \textbf{(5-7)}$$

and was smoothed slightly using a Gaussian kernel. The 3 ISI values where this smoothed depth of modulation function crossed 0 were calculated, and used to define 4 non-overlapping ISI intervals in the range 1ms to 1s. The fraction of ISIs occurring within each of these ISI ranges was then computed for each cell. Only spikes fired during track-running were used for analysis of temporal spiking properties.

LFP analysis

For the analysis of LFP signals we used only those LFPs for which at least one place cell was recorded in the same session. Since some of the LFPs were recorded differentially with respect to a local reference, we added the reference signal to the differential signal in these cases to get

an estimate of the LFP with respect to ground for the purposes of comparing power spectral properties. Multi-taper spectral analysis was performed using the Chronux Matlab toolbox [254] (http://chronux.org) using a time-bandwidth product of 3, and 5 tapers [31]. Spectra were estimated in all segments of data during trial-running, for which the animal's speed did not drop below 5cm/s. Segments of duration less than 1s were discarded. Estimates for the mean and standard error of the log power spectrum were then given by:

$$\log \hat{P}(f) = \sum_i w_i \log P_i(f) \quad \textbf{(5-8)}$$

$$\hat{\sigma}(f) = \sqrt{\sum_i w_i^2 \sigma_i^2(f)} \quad \textbf{(5-9)}$$

where the weight w_i given to each segment is the duration of the segment divided by the sum of the duration across all segments.

For theta-gamma coupling analysis, theta phase was estimated by first filtering the LFP between 5 and 11Hz using a zero-phase 4-pole Butterworth filter. The peaks and troughs of the filtered signal were then detected and defined as 0/360 and 180 degrees respectively. The remaining LFP samples were assigned a theta phase using linear interpolation between the peaks and troughs. The continuous wavelet transform of the signal was then computed using a Morlet wavelet with non-dimensional frequency parameter ω0=6 [39] (see Time-Frequency Analysis). Wavelet software was provided by C. Torrence and G. Compo, and is available at URL: http://paos.colorado.edu/research/wavelets/. The log of the wavelet power at each scale was z-scored, and the average of the normalized log power was computed as a function of theta phase at each scale (frequency) using 30 phase bins. The strength of theta-gamma coupling was

then given as the difference between the maximal and minimal values of this normalized power across theta phases.

Results

Place cell activity was recorded from 11 WT (C57/BL6) and 9 GluA1 KO mice, as they ran back and forth on a linear track for food rewards at the ends. In total, 137 and 88 putative CA1 pyramidal cells were respectively recorded from the WT and KO mice. Of these, 86 WT and 75 GluA1 KO cells had sufficient spiking activity during running on the track between the goal locations, and all analysis was restricted to these cells (see *Unit isolation*).

While most WT cells had clear, focused place fields in a small region of space, the KO neurons did not show clear spatial selectivity (Figure 5-1). Most place fields in WT were directional, spiking in only one direction of the journey, whereas the KO cells showed poor directional selectivity (Figure 5-1). These impressions were quantitatively confirmed (Figure 5-2).

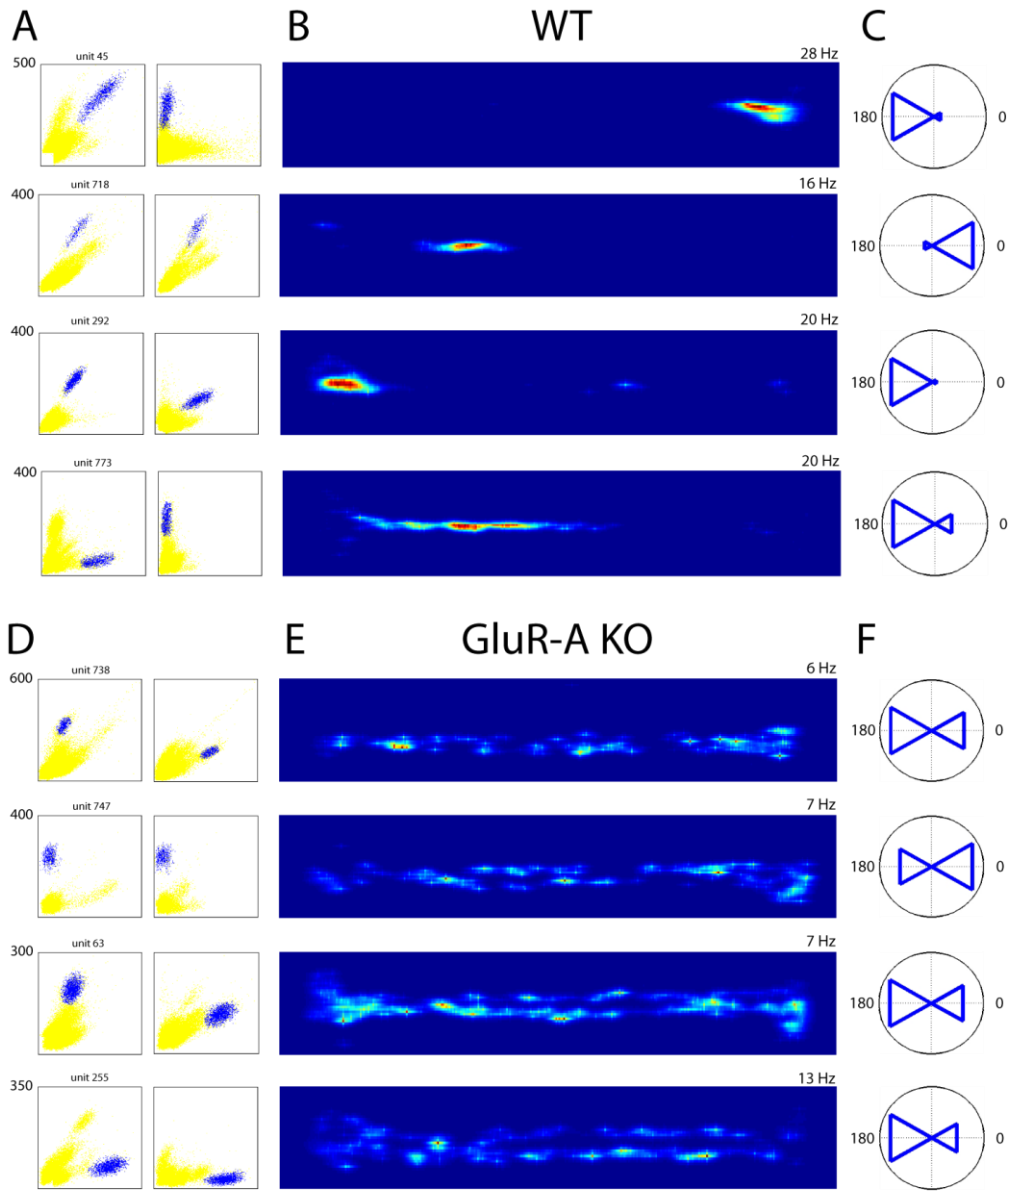

Figure 5-1: Comparison of spatial selectivity in WT and KO place cells.

*Comparison of unit isolation and spatial activity of four WT and four KO neurons. Data from WT are shown in panels **A-C**, and from KO in **D-F**. **A)** Each yellow dot represents one spike. Each blue dot represents a spike from the isolated neuron. Axes refer to the amplitude of spikes along each tetrode channel (mV). Two adjacent panels show the same tetrode data projected onto two different peak amplitude projections. **B)** Spiking rate of the isolated neurons in A as a function of the mouse's position. The x- and y-axes refer to the position of the mouse, and the color indicates the spiking rate (Hz). **C)** The spiking rate of each neuron is shown as a function of the head direction of the mouse in polar coordinates where the polar angle depicts the head direction and the radial distance depicts the spiking rate of the neuron for that head direction.*

The direction selectivity of the hippocampal neurons (see *Quantification of spatial spiking properties*) showed a 2.4-fold reduction in the KO mice (Table 5-1). The rate code for neurons in the WT mice was more robust than that in the KO mice, as evidenced by a 3.0-fold reduction in the peak spiking rate of KO neurons compared to WT (Figure 5-2A). The KO neurons were active in a 2.7-times larger fraction of the environment [302] (Table 5-1) compared to the WT neurons (Figure 5-2B). Further, the spatial information content per spike [299,302] was 3.1-fold lower for KO than WT (Figure 5-2C). The spike-position mutual information [300] also showed a 1.9-fold decrease in the KO (Table 5-1). The spike rate maps of KO cells were also less locally smooth, as measured by a 2.3-fold decrease in their rate map coherence (Figure 5-2D).

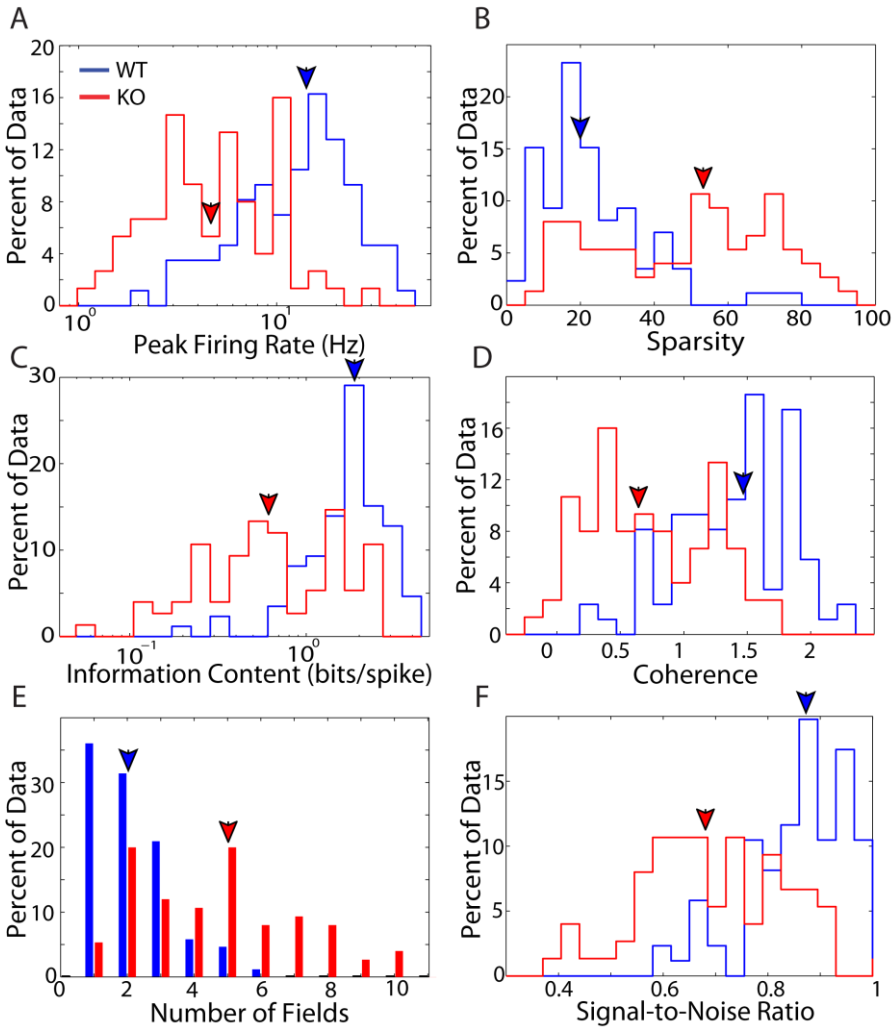

Figure 5-2: The KO place fields are fragmented and have weak spatial and directional tuning.

A) *Distributions of the peak spiking rate for WT (blue line) and KO (red line) cells. WT cells had significantly higher peak spiking rates than KO cells (WT: 14, 7.9-20Hz; KO: 4.6, 2.9 – 8.5Hz; $p=1.7x10^{-13}$). Median values for WT and KO cells are indicated by the blue and red arrows, respectively.* ***B)*** *Distributions of spiking rate map sparsity for the WT (blue) and KO (red) neurons. Rate map sparsity was 2.7-fold higher in KO (53%, 26%-68%) than in WT (20%, 14%-31%) ($p=1.1x10^{-11}$). Thus, the KO neurons were active over larger regions of space.* ***C)*** *The median spatial information per spike for WT neurons was 1.9bits/spike (1.3–2.4) while KO cells had 0.61bits/spike (0.34–1.4) ($p=5.3x10^{-12}$).* ***D)*** *WT cells also had smoother spiking rate maps than KO cells, as quantified by the rate map coherence (WT: 1.5, 1.1–1.8; KO: 0.65, 0.35–1.2; $p=2.3x10^{-13}$).* ***E)*** *There were 2.5-times as many place fields per neuron in KO than WT mice using a spiking rate threshold of 20% of the cell's peak rate (WT: 2, 1-3. KO: 5, 2.25–6; $p=1.9\cdot10^{-12}$). When a 1 Hz threshold was used the results were similar (WT: 2.5, 2-4. KO: 4, 2–5; $p=2.4x10^{-3}$).* ***F)*** *The signal-to-noise ratio was higher in WT cells (0.87, 0.80–0.93) than in KO cells (0.68, 0.60-0.81) ($p=4.0x10^{-13}$).*

The poor spatial selectivity of the KO rate maps often made it difficult to define clear place field boundaries. A commonly used definition of place fields [303] showed that there were 2.5-times as many place fields per cell for the KO neurons than the WT (Figure 5-2E), but individual place fields had similar widths in the two genotypes (p=0.84) (Table 5-1). Furthermore, KO place cells had a greater tendency to spike outside of their place fields, as quantified by a 22% reduction in the 'signal-to-noise ratio' of KO place cells (Figure 5-2F). These methods of quantifying the place fields of each cell are parameter sensitive, yet the results remain qualitatively true for a range of parameters tested.

	WT	**KO**	**p-value**
rate map avg (Hz)	1.3 (0.77-2.3)	0.91 (0.60-1.5)	8.9×10^{-3}
rate map peak (Hz)	14 (7.9-20)	4.6 (2.9-8.5)	1.7×10^{-13}
direction selectivity	0.62 (0.34-0.81)	0.26 (0.13-0.69)	1.1×10^{-3}
Info. content (bits/spike)	1.9 (1.3-2.4)	0.61 (0.34-1.4)	5.3×10^{-12}
spike-position info (bits)	0.15 (0.11-0.19)	0.081(0.057-0.11)	4.7×10^{-12}
rate map coherence	1.5 (1.1-1.8)	0.65 (0.35-1.2)	2.3×10^{-13}
rate map sparsity (%)	20 (14-31)	53 (26-68)	1.1×10^{-11}
# place fields/cell (1Hz)	2.5 (2-4)	4 (2-5)	2.4×10^{-3}
# place fields/cell (20%)	2 (1-3)	5 (2.25-6)	1.9×10^{-12}
place field size (cm) (1Hz)	24 (18-32)	20 (14-34)	0.29
place field size (cm) (20%)	21 (16-33)	221 (16-32)	0.84
SNR (1Hz threshold)	0.93 (0.87-0.97)	0.64 (0.51-0.83)	2.3×10^{-16}
SNR (20% threshold)	0.87 (0.80 0.93)	0.68 (0.60-0.81)	4.0×10^{-13}
trial-trial rate map variability	1.58 (1.15-1.99)	0.63 (0.28-1.22)	2.0×10^{-14}
rate map systematic change	1.3 (0.29-2.4)	0.05 (-0.68-0.93)	4.3×10^{-8}

Table 5-1: Summary of WT and KO spatial firing properties.

Stability [304,305] and plasticity [194,195,217] of hippocampal place fields are thought to be important for SRM. However, similar place field dynamics cannot be compared directly between WT and KO mice, because most KO neurons did not have clearly defined place fields. Instead, rate maps were computed for an entire trial (defined as the mouse traveling from one goal location to the other and back), each trial lasting less than a minute. The changes in rate maps

across trials can provide an estimate of both the stability and plasticity of rate maps on this short time scale of the order of seconds [110,195,298,306,307,308]. Specifically, the random variability of rate maps was quantified by computing the trial-by-trial fluctuations in the rate map. On the other hand, potential plasticity-related changes in the rate maps were estimated by computing the amount of systematic change in the rate map across trials (see *Quantification of spatial spiking properties*). The rate maps for an example WT and KO cell computed separately from either the first and second halves of trials, or the set of even-numbered and odd-numbered trials are shown in Figure 5-3. In this example it is clear that the WT neuron has a sharply different firing rate function in the first half and second half of trials, yet nearly identical firing rate maps when computed separately for even and odd trials. This suggests the presence of systematic across-trail changes in the rate map that were much larger than any random fluctuations in the rate map from trial to trial. In contrast, the KO neuron shows equally dissimilar rate maps when computed from the first and second halves of trials or the even and odd trials, suggesting that this neuron's rate map showed no systematic change across trials, and a comparatively large amount of random trial-by-trial variability.

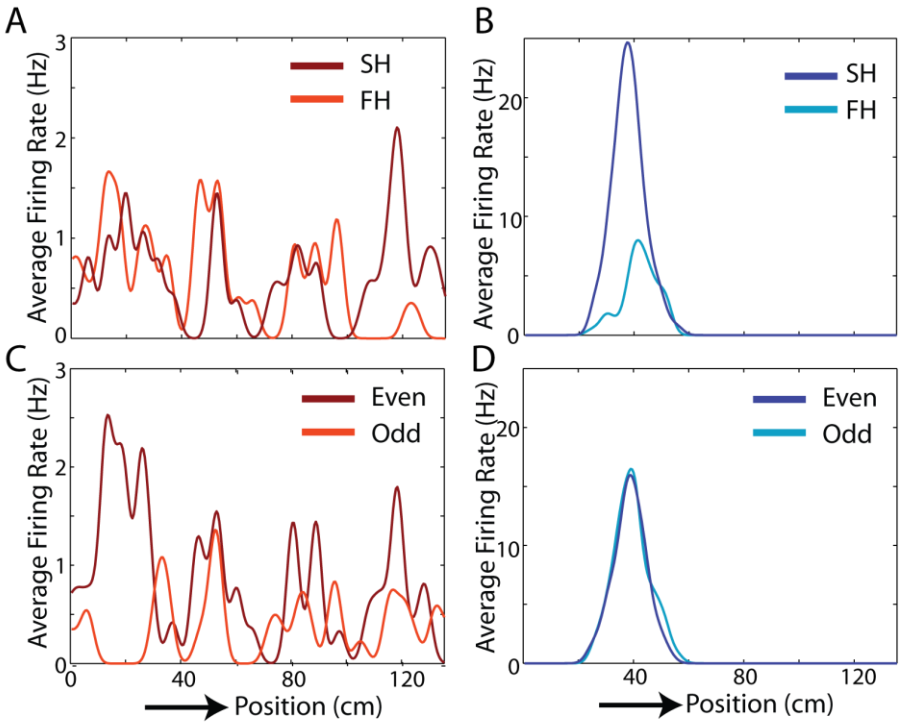

Figure 5-3: Systematic change and trial-trial variability in example WT and KO neuron's ratemaps.

A) *The rate maps from the first half (FH) and second half (SH) of trials are plotted separately for a well-isolated example KO neuron.* ***B)*** *Same as A, but for an example WT neuron.* ***C-D)*** *Same as A-B except rate maps are computed separately for the set of even-numbered and odd-numbered trials.*

Strikingly, across all data the trial-shuffled correlations for KO cells were 2.4-fold smaller than those for the WT neurons (Figure 5-4). Thus, the rate maps of KO neurons showed far more random trial-by-trial variability than those of the WT neurons. The trial-ordered correlations, measuring the similarity in the rate maps during the first and second halves of trials, were 2.0-fold smaller in KO than WT neurons. The z-scored difference between the trial-ordered correlations and the trial-shuffled correlations, termed the 'systematic variation', provides an estimate of the amount of systematic change or the plasticity of the rate maps beyond that expected by chance. More positive values of the systematic variation indicate that more of the

trial-by-trial variability in the cell's spatial spiking is due to changes that are systematic across trials. For the ensemble of WT neurons, the systematic variation was significantly positive, indicating that there were indeed systematic changes in the rate maps (Figure 5-4). In contrast, for the ensemble of KO neurons, the systematic variation was not significantly different from zero, suggesting that there were no systematic changes in the KO rate maps.

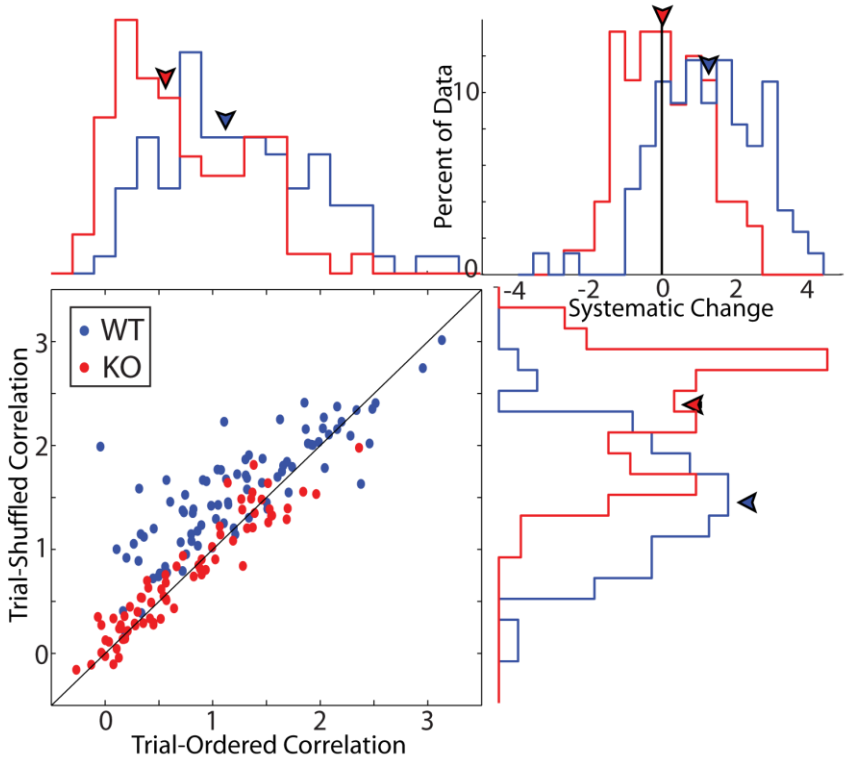

Figure 5-4: WT neurons had more systematic change and less random rate map variability than KO neurons.

(Bottom-left) *Scatter plot of the correlations between rate maps calculated by randomly partitioning trials for each cell (mean trial-shuffled correlations) and by dividing trials into the first half and second half of the recording (trial-ordered correlations). Each dot represents a cell, and the color indicates the genotype.* ***(Bottom-right)*** *Distribution of average trial-shuffled correlations for each genotype (arrows indicate median values). WT trial-shuffled correlations (1.6, 1.2–2.0) were much higher than KO (0.63, 0.28-1.22) ($p=2.0x10^{-14}$), indicating increased trial-by-trial variability in KO cells.* ***(Top-left)*** *Trial-ordered correlations were also higher for WT (1.13, 0.72-1.74) than KO (0.56, 0.19-1.27) ($p=6.8\ x10^{-6}$), although the difference was much smaller.* ***(Top-right)*** *The systematic variation is given by the normalized difference between the trial-ordered correlations and the trial-shuffled correlations for each genotype. The systematic variation was significantly positive in WT (1.31, 0.29-2.44; $p=5.5x10^{-11}$; Wilcoxon signed rank test), but not in KO (0.56, 0.19-1.27; p=0.53), indicating that there were systematic changes in WT rate maps during the recording that were not present in KO cells.*

Since the firing properties of hippocampal neurons are known to depend on an animal's running speed [196,197,199,200], it is possible that the above results could be due to behavioral

differences between WT and KO mice. Across all sessions for which a recorded cell was used for the analysis, the average running speeds during epochs used for analysis (during trial-running) were slightly higher in the WT (*n*=52 sessions, 20cm/s, 18-26cm/s) than KO mice (*n*=33 sessions, 18cm/s, 12-20cm/s) ($p=7.5x10^{-3}$). However, this small behavioral difference could not explain the much larger differences in the spatial firing properties of WT and KO neurons. An analysis of covariance (ANCOVA) revealed that even after accounting for the differences in running speed, the differences between WT and KO in all measures of spatial selectivity were still highly significant.

We next compared the temporal spiking properties of WT and KO neurons by computing the distribution of inter-spike intervals (ISIs) during trial-running for each cell. The ISI distributions for the same example WT and KO neurons from Figure 5-3 are shown in Figure 5-5A. Notably, the KO neuron showed an increased fraction of bursting (< 15ms) ISIs compared to the WT neuron, as well as a decrease in the fraction of intra-theta-cycle ISIs (15ms < ISI < 125ms). Averaging the ISI distributions across the populations of WT and KO cells revealed that these differences were consistent in the population (Figure 5-5B). In order to quantify the differences in ISI distributions we divided the range of ISIs (< 1s) into 4 non-overlapping ISI ranges (see *Quantification of temporal spiking properties*; Figure 5-5C). While the WT neurons had a significantly larger fraction of very short (<4ms) ISIs, the KO neurons actually had a substantially greater fraction of ISIs in the range (4ms < ISI < 13ms). Thus, KO neurons were more bursty overall, while WT neurons were able to achieve higher-frequency bursting. Additionally, WT neurons were significantly more likely to have ISIs in the intra-theta cycle range (13ms < ISI < 160ms), while KO neurons were significantly more likely to show very long ISIs (>160ms).

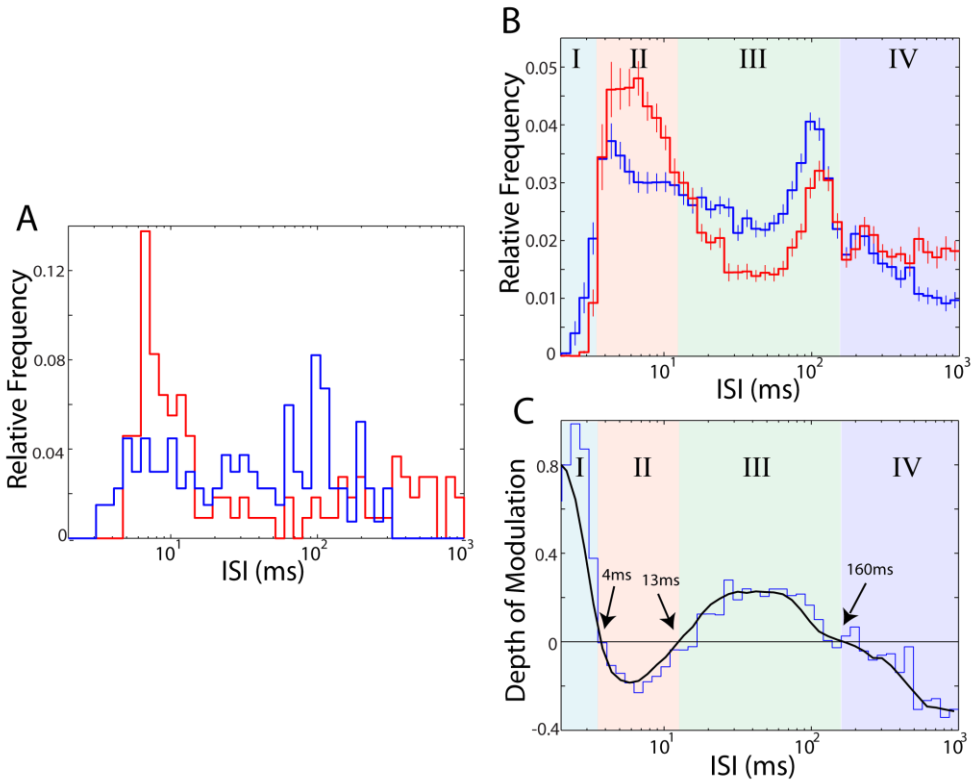

Figure 5-5: Differential nature of WT and KO ISI distributions.

A) *Example ISI distribution for the same WT (blue) and KO (red) neurons as in Figure 5-3 showing altered burst structure and theta modulation.* ***B)*** *The population average ISI distribution (mean ± s.e.m.) for WT (blue) and KO (red) neurons, showing significant differences between WT and KO mice.* ***C)*** *The relative difference between the average WT and KO ISI probability ($P_{WT}(ISI) – P_{KO}(ISI))/(P_{WT}(ISI) + P_{KO}(ISI))$ as a function of ISI. The ISI 'depth of modulation' is smoothed slightly (solid black trace) to reliably detect the ISI values where the relative difference crosses 0 (occurring at 4ms, 13ms, and 160ms). These threshold crossings define 4 ISI ranges which are depicted as regions labeled (I-IV) shaded in different colors. The fraction of ISIs in region I was larger in WT (4.8%, 1.4-7.7%) than KO (1.1%, 0-5.2%; $p=2.6x10^{-4}$). The fraction of ISIs in region II was larger in KO (37%, 27-44%) than in WT (27%, 21-32%; $p=4.4x10^{-7}$). The fraction of ISIs in region III was larger in WT (49%, 39-58%) than in KO (36%, 28-48%; $p=2.4x10^{-8}$), and the fraction in region IV was larger in KO (22%, 17-29%) than in WT (16%, 9.5-23%; $p=2.1x10^{-4}$).*

Given the substantial differences observed in the spatial and temporal spiking properties of KO cells, we next investigated whether changes were also present in the hippocampal local field potential (LFP) of KO mice. Example traces of a WT and KO LFP are shown in Figure 5-6A,B,

illustrating that the amplitude of theta oscillations was much smaller in the KO. Further, the presence of higher-frequency (25-80Hz) gamma oscillations was much reduced in the KO LFP as well. The time-frequency energy distributions for the example LFPs confirmed these visual impressions, and also revealed a comparative lack of theta-rhythmic modulation of gamma power in the KO LFP (Figure 5-6C,D).

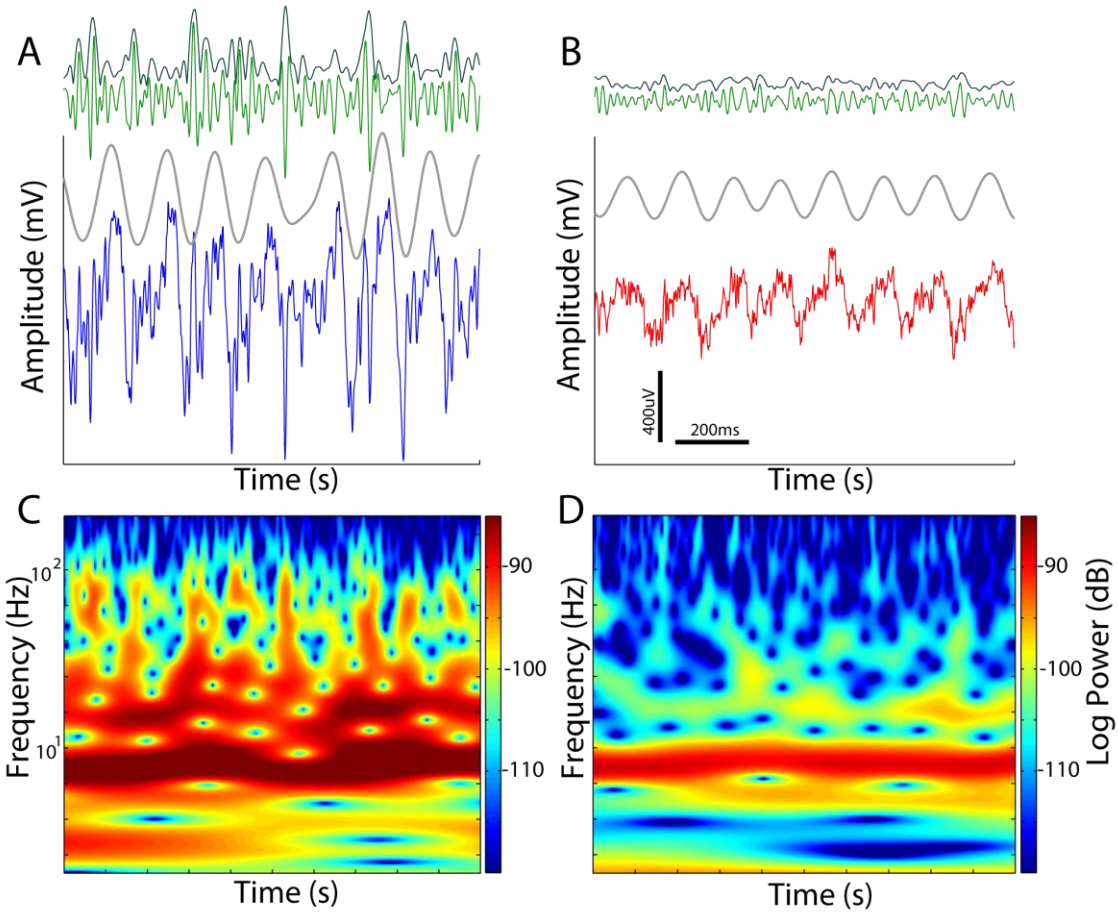

Figure 5-6: Spectral properties of example WT and KO LFPs.

A) *Trace from an example WT LFP recording (blue). The theta-band filtered signal is shown above (gray trace), along with the gamma-band filtered signal (green trace), and the envelope of the gamma-band signal (dark green trace).* ***B)*** *Same as A for an example KO LFP.* ***C)*** *The time-frequency energy distribution of the LFP trace from A is computed from the continuous wavelet transform.* ***D)*** *Same as C for the KO LFP trace in B.*

The power spectra for these example WT and KO LFPs were computed across all epochs during trial-running (see *LFP analysis*). Indeed, there were dramatic reductions of power in the KO LFP compared to the WT LFP across a broad range of frequencies, notably at theta and gamma frequencies (Figure 5-7A). The average LFP power spectra across all 82 WT and 44 KO LFPs showed similarly dramatic reductions at theta and gamma frequencies (Figure 5-7B). Indeed, there was 1.5-fold as much power in the theta band (5-11Hz) in the WT LFPs (5.1, 4.2-8.1x$10^{-3}$$mV^2$) compared to the KO LFPs (3.4, 2.4-5.4x$10^{-3}$$mV^2$; p=3.0x$10^{-5}$). The frequency of KO LFP theta oscillations (7.7, 7.4-8.3Hz) was also significantly lower than WT (8.1, 8.1-8.4Hz; p=2.8x10^{-5}). The difference in gamma-band power (25-80Hz) was even more pronounced, with WT LFPs (3.3, 2.6-5.3x$10^{-3}$$mV^2$) having 2.7-fold more gamma power than KO LFPs (1.2, 0.90-1.8x$10^{-3}$$mV^2$; p=9.5x$10^{-15}$).

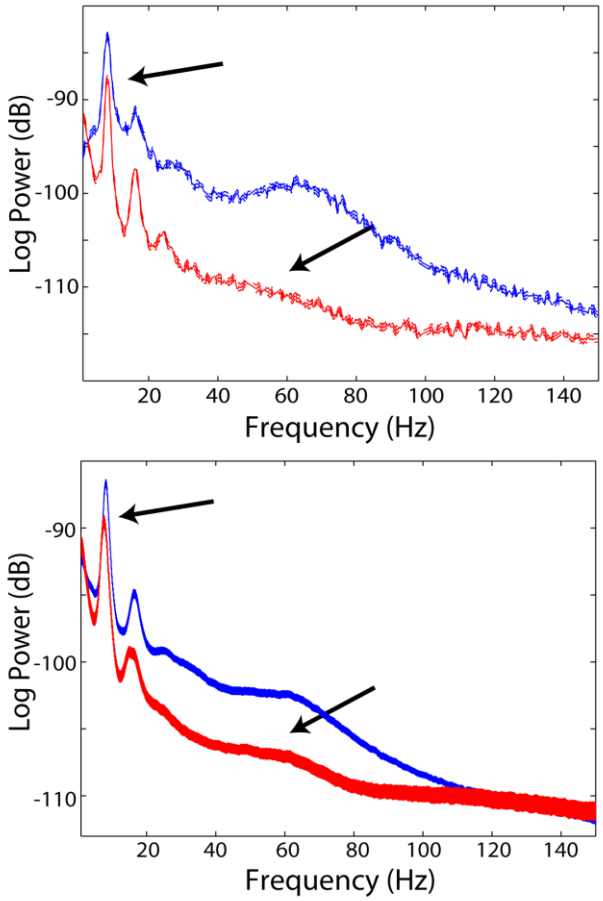

Figure 5-7: Decreased theta and gamma power in the KO LFP.

(Upper panel) *The power spectrum of the same example WT and KO LFPs from Figure 5-6 are computed across all data epochs during trial-running. The dashed lines indicate the 95% confidence interval. Note the reduced power in the KO LFP at theta and gamma frequencies.* ***(Lower panel)*** *Population average (n=82 WT LFPs, n=44 KO LFPs) LFP power spectra.*

Numerous studies have demonstrated a strong relationship between the amplitude of hippocampal gamma oscillations and the phase of the theta oscillation [80,85,86]. Further, the reduction of theta and gamma power in the KO LFP does not imply that phase-amplitude coupling between the two rhythms will necessarily be impaired as well [309]. Thus, we quantified the relationship between gamma power and theta phase by computing the average

high-frequency power as a function of frequency and theta phase using the continuous wavelet transform (Figure 5-8A,B; see *LFP analysis*). While the shape of the theta oscillations were similar in the WT and KO LFP, the WT gamma power was much more modulated by theta phase than KO. Indeed, theta modulation of gamma power was dramatically reduced in the KO across the range of gamma frequencies (Figure 5-8C). Maximal theta-gamma coupling was reduced by 1.7-fold in the KO (0.34, 0.25-0.44) compared to the WT (0.60, 0.55-0.66; $p=2.5x10^{-14}$).

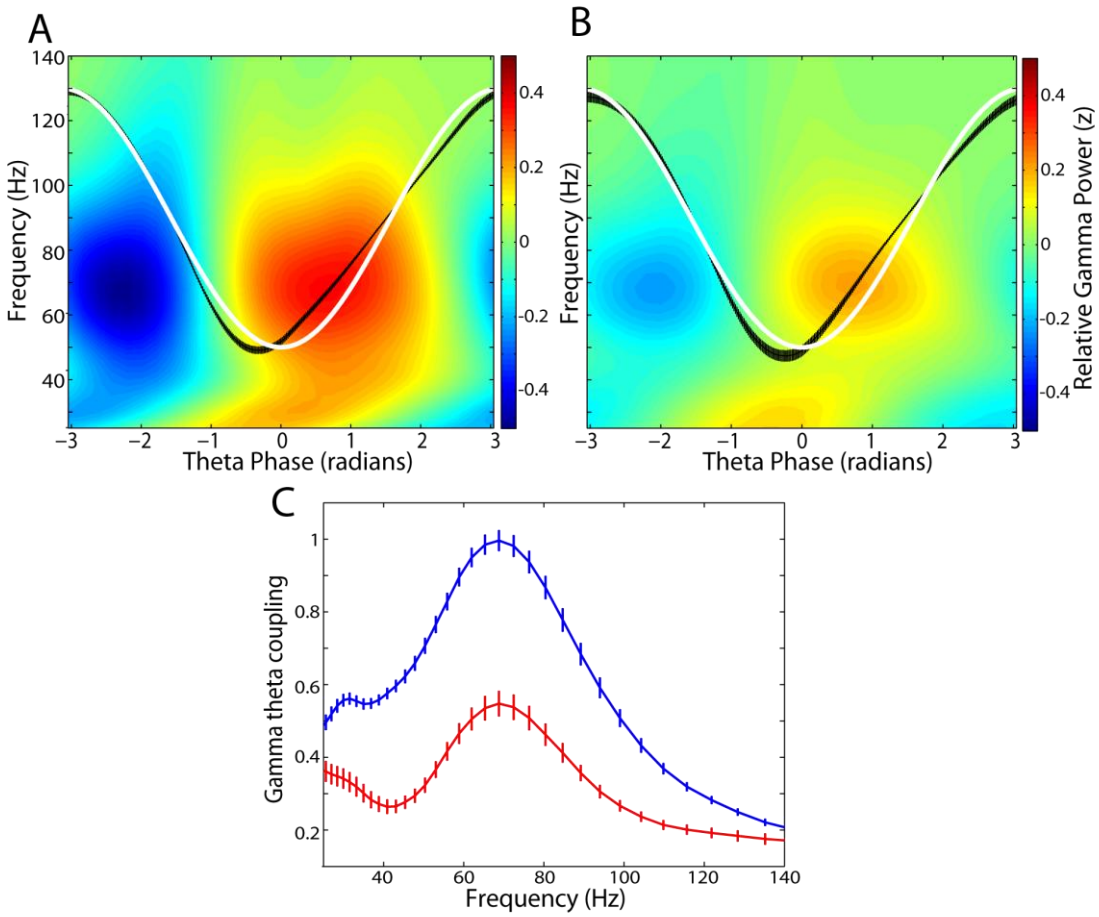

Figure 5-8: Impaired theta-gamma coupling in the KO LFP.

A) *Average gamma power at each frequency is plotted as a function of theta phase. Gamma power at each frequency is normalized to zero mean and unit variance. The black trace shows the average relative amplitude of the LFP (filtered above 5Hz) as a function of theta phase, illustrating the corresponding average shape of the theta wave. The white trace shows a corresponding sinusoid.* ***B)*** *Same as A for KO.* ***C)*** *The average strength of gamma-theta coupling (see LFP analysis) is plotted as a function of frequency for WT (blue) and KO (red) LFPs.*

Discussion

These results demonstrate that the lack of GluA1 containing AMPA receptors dramatically reduced all measures of the spatial selectivity of CA1 neurons' spiking activity. Place cells in the KO mice not only had reduced spatial and directional selectivity, but also showed increased

random trial-by-trial fluctuations in their spatial firing rate maps. Further, the firing rate maps of KO place cells did not show any systematic changes across trials, while WT neurons did. We also found that KO place cells' spiking had an altered temporal structure. Namely, KO neurons were more likely to burst, but at lower frequencies than WT neurons, and KO neurons were less likely than WT to have intra-theta-cycle ISIs. Finally, we tested whether the altered properties of KO place cells would result in changes to the hippocampal network behavior by comparing the hippocampal LFPs in WT and KO mice. Interestingly, we found dramatic reductions in both the theta and gamma frequency oscillations in the KO mice. Phase-amplitude coupling between theta and gamma oscillations was also greatly reduced in the KO LFPs. Thus, deletion of the GluA1 subunit resulted in profound alterations in both the spiking properties of CA1 place cells, and the oscillatory properties of the hippocampal network.

GluA1 subunits are thought to be important for the induction of some forms of plasticity. In particular, insertion of GluA1 containing AMPARs into the postsynaptic membrane in response to certain stimuli is believed to be a physiological substrate for long term plasticity (LTP) [310]. Selective insertion of GluA1 containing AMPARs also appears to be responsible for the large distance-dependent scaling of synaptic strength, such that larger AMPAR-mediated synaptic currents are present at more distal synapses resulting from an increased density of GluA1 containing AMPARs [288]. Although CA1 neurons of the GluA1 KO mice do not show LTP in response to 100Hz tetanic stimulation [285], other forms of plasticity such as the long-term component of theta-burst plasticity [286,287], and spike-timing dependent plasticity [287] remain intact. The latter protocols for induction of LTP are more similar to physiological inputs received by CA1 neurons during navigation. Thus, removal of GluA1 AMPAR subunits eliminates

some forms of synaptic plasticity but leaves others intact. CA1 neurons lacking GluA1 AMPAR subunits also lack the distance-dependent dendritic scaling which may result in a substantial reduction of the direct perforant path inputs from the entorhinal cortex which arrive in the distal SLM layer [311].

One mechanism that could generate the observed results is the impairment of an inhibition-mediated winner-take-all mechanism. This could result from the reduced excitatory drive onto distal dendrites producing a weaker activation of lateral inhibition within CA1, and thus a reduced sparsity of CA1 place cell activity. Interestingly, in addition to explaining the impaired spatial selectivity in the KO place cells, the reduced excitatory drive could account for the lower frequency bursting of KO neurons. Impaired recurrent inhibition could also account for the increased overall burst probability of KO neurons. In fact, similar reductions in the spatial selectivity of hippocampal place cells were observed in mice lacking NMDA receptors in PV-expressing hippocampal interneurons [309], a manipulation which reduced the excitatory drive onto inhibitory interneurons. Thus, impaired inhibition-mediated 'competition' within CA1 in both KO animals could result in the reduced spatial selectivity. It is less clear how the lack of systematic changes across trials in the firing rate maps of GluA1 KO place cells would result from such a mechanism. These results are especially surprising given that experience-dependent changes in place cell firing are thought to be mediated by STDP [312], which is intact in the GluA1 KO mice [287]. It is possible, however, that other GluA1-dependent forms of synaptic plasticity are also required for such experience-dependent changes, or that the altered temporal structure of KO place cell activity itself results in a lack of experience-dependent changes even with STDP intact.

Impairments of both hippocampal theta and gamma oscillations in the GluA1 KO mice could also result from a reduced excitatory drive and impairment of inhibition within CA1. Indeed, in hippocampal slices in which the GluA1 subunit was deleted in PV-expressing interneurons the amplitude of kainite-induced gamma oscillations was reduced [313]. In contrast, gamma oscillations were actually enhanced in mice lacking the NMDARs in PV-expressing interneurons [309], suggesting that the simplistic lateral inhibition based model alone cannot account for the observed alterations in network oscillation patterns. While the interneuron NMDAR KO mice had enhanced gamma oscillations, the theta oscillations were substantially weakened, as was the strength of theta-gamma coupling [309]. Thus, different mechanisms may lead to the reduced theta and gamma oscillations in the GluA1 KO mice.

We also note that the GluA1 subunit is knocked out in the entire forebrain in these mice, and thus the observed changes in the activity of CA1 neurons and hippocampal network oscillations could arise in afferent structures such as the entorhinal cortex. However, we focus our attention on CA1 because alterations of synaptic scaling and LTP have been demonstrated only in CA1 [285,287,288], and the GluA1 subunit has highest expression within the hippocampus [314]. In fact, the forms of plasticity which were impaired in the CA1 neurons of GluA1 KO mice were found to be unaltered in neocortical pyramidal neurons [287], suggesting minimal disruption of extra-hippocampal circuits. Further, juvenile mice display GluA1 independent long-term plasticity (LTP) which could generate intact wiring of hippocampal circuits in KO mice during development [315]. Accordingly, conditional restoration of GluA1 receptors in adult KO mice rescues both synaptic and behavioral deficits [316].

Previous studies have found that systemic blockade [317] and genetic silencing of NMDA receptors [318] within CA1 abolished long-term synaptic plasticity and experiential plasticity of place cells, but did not have as large of an impact on the spatial selectivity of place cells. In contrast, despite the relatively intact physiological forms of synaptic plasticity in the hippocampus and neocortex of the GluA1 KO mice, there was a large reduction in both spatial selectivity and experiential plasticity. This suggests that the NMDARs are required for experiential modification of place fields but are not required for the formation of focused place fields, and that spatial selectivity in CA1 arises via other mechanisms such as recurrent inhibition and AMPAR-mediated excitation.

The GluA1 KO mice show profound and specific behavioral deficits in spatial memory tasks known to depend on the hippocampus. While the mice perform as well as controls on spatial reference memory (SRM) tasks such as the Morris water maze task, they show profound impairments on a variety of spatial working memory (SWM) tasks [284,285,289]. This pattern of behavioral impairments is similar to those seen in mice lacking NMDARs in PV-expressing interneurons [309]. The fact that GluA1 KO mice are unimpaired on the SRM tasks while spatial selectivity and experiential plasticity of CA1 place cells show such profound impairments implies that these aspects of hippocampal activity are not required for SRM tasks. On the other hand, hippocampal-lesioned animals, as well as the GluA1 KO mice, were severely impaired in SWM such as delayed alteration on a T-maze, suggesting that spatial selectivity, stability, and experiential plasticity of CA1 neurons may be required to perform these tasks. Interestingly, animals lacking CB1 receptors show a similar pattern of SWM, but not SRM, impairments [319]. Yet administration of CB1 antagonist was found to leave hippocampal spatial selectivity intact

while impairing the temporal organization of CA1 activity [320]. Thus, alteration of the temporal structure of CA1 neurons in GluA1 KO mice, rather than the impaired spatial selectivity, may be the underlying cause of the SWM impairments.

Chapter 6: **Behavioral modulation of the hippocampal temporal code**

Introduction

The spiking of hippocampal pyramidal cells ('place cells') is thought to encode an animal's position in two distinct ways. First, place cells fire spikes selectively within certain regions of an environment (the 'place fields'), presenting a rate code for the animal's position across the ensemble of such cells. Secondly, a temporal code of position is given by the advancement of spike times with respect to the theta rhythm as the animal runs through the place field. In addition to an animal's position, the firing rate of place cells during maze running has been shown to be modulated by behavioral variables such as running speed [196,197,199,200], as well as non-spatial behaviors, stimuli, and reward conditions [197,202,203,204,205,206,207]. In contrast, the temporal code is thought to be largely independent of behavioral variables such as stopping [26] and running speed [26,213,214]. While such a robust representation could be advantageous for maintaining consistent position information across various conditions, phase precession is widely believed to be important for sequence learning [216,217,218] by producing sequential activation of cells with adjacent place fields on the fast timescale relevant for spike-timing dependent plasticity (STDP) [107,108]. Hence, modulation of precession by task conditions, stimuli, reward, and/or behavior could be important for providing relevant context for sequences being learned. While a slight dependence of phase precession on running speed within single trials has been previously reported [212], it is unknown whether the hippocampal temporal code shows a clear relation to any behavioral or task parameters. Further, most

precession studies have focused on rats. With the availability of a range of transgenic mice, it is important to have a systematic study of precession in mice, and comparison to results reported in rats.

Here, we report that precession in mice is significantly weaker than in rats. However, these differences can be accounted for by behavioral variables such as stopping within the place field and lower running speeds. Further, we demonstrate that the relative position of the place field on the track strongly influences precession, with precession becoming progressively weaker as the animal approaches the reward location. Such a dependence of phase precession on the location of the field relative to reward could be important for goal-directed sequence learning via STDP.

Methods

Experimental Methods

Experimental methods were as described in Chapter 5 *Methods*.

Data Analysis

The data were analyzed offline using custom software written in Matlab (MathWorks). Since most place cells on linear tracks are directional, the left-to-right and right-to-left trials were treated separately [26,194,195,298] by first 'linearizing' the position data. The running speed of the animal along the track was estimated by smoothing the linearized position data using a Gaussian smoothing kernel (sigma = 50ms) before computing the approximate temporal derivative of position. The running speed was then smoothed with a Gaussian kernel (sigma = 300ms) before computing its temporal derivative to give the

acceleration. Trials were defined by the animal running from one goal region to another, and the initiation (termination) of each trial was determined by the times when the animal's running speed first (last) exceeded 2cm/s. Trials in which the mouse did not run all the way from one goal location to the other were excluded from analysis. Further, times when the animal's running speed towards the target goal location was less than 2cm/s were excluded from all analysis. The total number of spikes of a cell at each spatial bin (0.5cm) and the time spent by the animal in that bin were computed. These spike count and occupancy maps were smoothed slightly using a Gaussian smoothing kernel (sigma = 1cm). The spiking rate map was then computed by dividing the spike count in each bin by the occupancy in the corresponding bin. These spiking rate maps were then smoothed again with a Gaussian kernel (sigma = 2cm).

Place fields were defined as regions of the track where the firing rate map exceeded 2Hz. The boundaries of each place field were taken as the points where the rate map dropped below 5% of the peak rate within the field. Trials where the animal stopped or paused were defined as trials in which the minimum running speed within the place field was below 2cm/s. Only place fields which had at least 50 spikes, 5 laps, and 5 spikes per lap on average were included in analysis. Further, place fields which were significantly overlapping with the goal locations (had a firing rate exceeding 25% of the peak rate of the field within the goal location) were excluded from analysis. Before excluding trials with stops there were 59 place fields meeting this criteria, while this number was reduced to 48 place fields after excluding stop trials. We also excluded two place fields that were wider than ½ the length of the track due to a lack of spatial selectivity. The results remained qualitatively similar for a range of place field selection criteria.

Theta phase was estimated by first filtering the LFP between 5 and 11Hz using a zero-phase 4-pole Butterworth filter. The peaks and troughs of the filtered signal were then detected and defined as 0/360 and 180 degrees respectively. The remaining LFP samples were assigned a theta phase using linear interpolation between the peaks and troughs. Since some of the LFPs were recorded differentially with respect to a local reference electrode, we added the reference signal to the differential signal in these cases to get an estimate of the LFP with respect to ground for consistency across recordings.

The phase-position correlation was computed using several commonly used methods. Unless otherwise stated, we used the linear Pearson's correlation coefficient between spike phase and position after rotating all spike phases by an amount minimizing the phase-position correlation (rotating phase in increments of 1 degree to find the minimum) [110,212,320]. We also computed the linear phase-position correlation after implementing the phase rotation maximizing the square of the linear correlation coefficient R^2 [219,224]. Finally, we also used the circular-linear correlation between phase and position [212,321]. Circular statistics were computed using the CircStats toolbox for Matlab [322]. When computing the phase-position correlation for single trials we maintained a common phase rotation (minimizing R) across all trials. Single-trial phase-position correlations were only included in analysis for trials with at least 5 spikes. Similarly, when computing correlations across trials between single-trial precession and other single-trial properties for each place field, only place fields with at least 5 useable trials (>= 5 spikes) were included.

To compute spatio-temporal receptive fields (STRF), we divided the number of spikes fired within each spatio-temporal bin (spatial width=1cm, temporal width=8 degrees) by the total

time spent by the mouse in that bin [110,219]. After imposing periodic boundary conditions for phase, the result was then smoothed by convolution with a two-dimensional Gaussian (spatial width=2cm, temporal width=15 degrees). When computing STRFs for single place fields as examples (rather than computing ensemble average STRFs), we used a slightly broader Gaussian smoothing kernel (spatial width=3cm, temporal width=22.5 degrees) for ease of visualization.

The place field spike time autocorrelation functions were computed by first estimating the spike time autocorrelation within each trial (only for trials with at least 5 spikes) for the spikes fired within a given place field. The average of the autocorrelation function was then computed across trials, and the across-trial average autocorrelation function was then smoothed slightly (Gaussian smoothing kernel, sigma=6ms). The inverse Fourier transform of the average autocorrelation function was used to estimate the (normalized) power spectrum of spike times.

For analysis of behavioral modulation of theta oscillations, we first computed the amplitude envelope of the theta-band (5-11Hz) filtered signal (4-pole Butterworth filter) using the Hilbert transform. Theta amplitude was then normalized across all epochs of trial running (speed >= 2cm/s) to have zero mean and unit variance. After excluding data in the goal regions, theta amplitude was similarly normalized within the restricted data set. To analyze the effects of relative position and acceleration on theta amplitude beyond those expected given the relationship between running speed and theta amplitude, we first performed a linear regression analysis of theta amplitude on running speed. The dependence of the residual theta amplitude on relative position and acceleration were then computed. We used 30 equally spaced bins to compute the nonparametric dependence of theta amplitude on running speed and acceleration, while using 40 equally spaced bins to compute its dependence on relative position.

Results

We recorded activity from place cells in dorsal CA1 as mice ran back and forth on a linear track to receive a reward at both ends of the track. A total of 137 well isolated complex spiking cells were recorded from eleven mice. Since mice engaged in non-spatial behavior at the goal locations between trials, these data were excluded from further analysis. Only the cells with a mean firing rate exceeding 0.1Hz and having greater than 100 spikes during trial running were used. A total of 107 cells met these criteria. After computing linearized firing rate maps for each of these track-active pyramidal cells, we defined place fields as regions of the linearized track where the average firing rate exceeded 2Hz, excluding any place fields which overlapped significantly with the goal locations (see *Methods*). A total of 59 place fields met these criteria and were used for the subsequent analysis. We next computed the correlation between spike theta phase and position, henceforth referred to as precession, for each place field (see *Methods*). The firing rate maps for two representative place cells are shown in Figure 6-1A,B along with the spike phase and position across all spikes (Figure 6-1C,D). While the place field depicted in Figure 6-1A shows clear precession (phase-position correlation r:-0.50, phase-position slope: -4.9 deg/cm), the place field in Figure 6-1B had much weaker precession (r: -0.20, slope: -2.0 deg/cm). The spatio-temporal receptive fields (STRFs; Figure 6-1E,F) [110,219] for these example place fields further illustrate a large difference in precession between the two fields. This suggests that a substantial amount of variability in precession is present across place fields in mice.

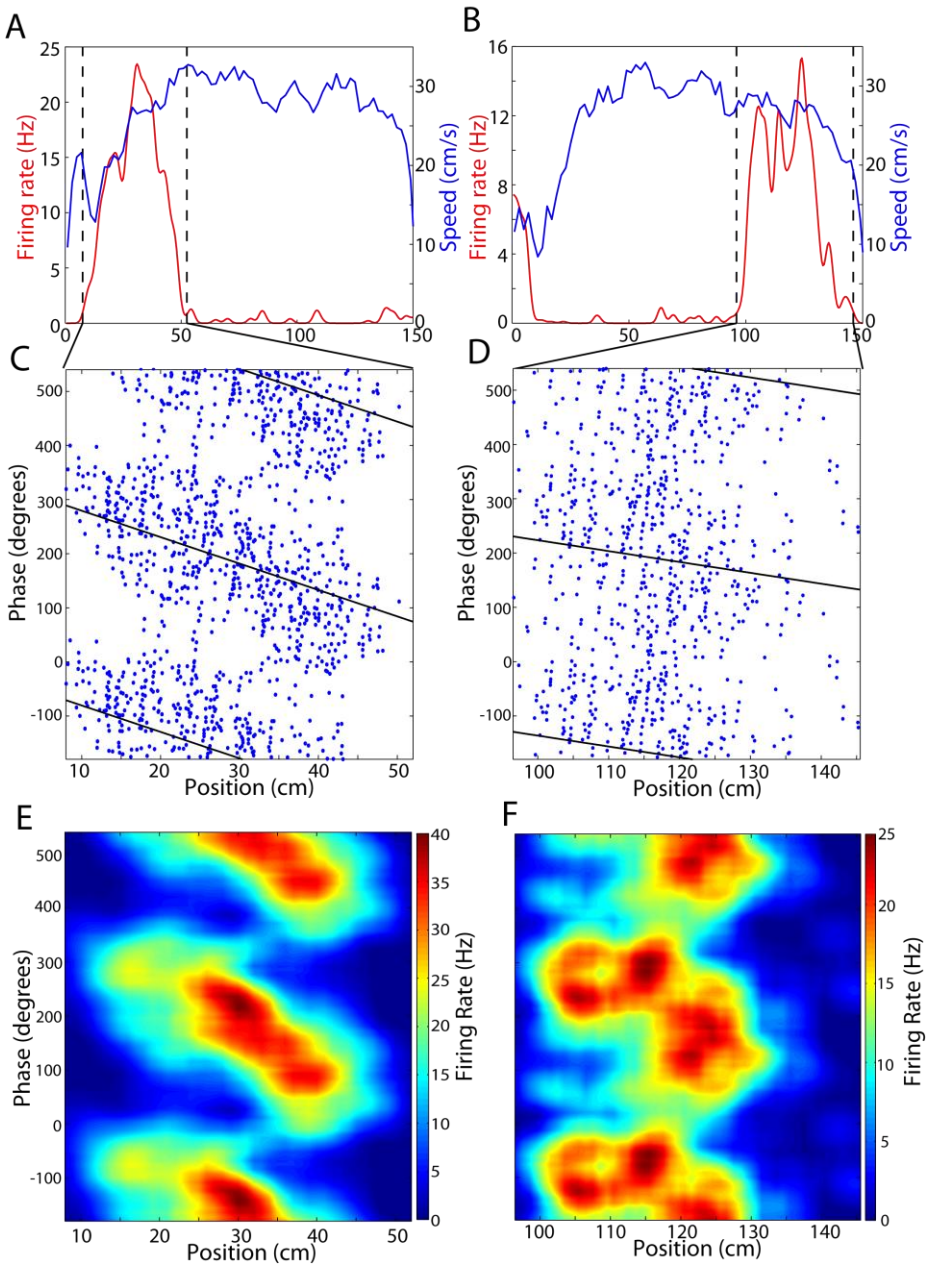

Figure 6-1: Precession analysis for two example place fields showing large variability.

A) *The average firing rate as a function of position (for the left-to-right movement direction) for an example place cell (red trace) is plotted along with the average running speed of the animal as a function of position (blue trace). The boundaries of the place field are indicated by the vertical dashed lines.* ***B)*** *Same as A for a second example place cell.* ***C)*** *Theta phase is plotted against the position of the animal at the time of each spike fired within the place field indicated in A. The linear fit of spike phase against position is shown in black.* ***D)*** *Same as C for the example field in B.* ***E)*** *Average firing rate is plotted as a function of theta phase and position (STRF) for the example field in A.* ***F)*** *Same as E for the example field in B.*

Across the ensemble of place fields the average correlation between spike phase and position was -0.37±0.015 (mean±SEM here and in all subsequent results; Figure 6-2A) when phase was rotated to minimize the phase-position correlation [110,212,320] (see *Methods*). Since this method creates a bias towards finding a negative phase-position correlation we also computed precession by rotating phase to maximize the square of the phase-position correlation [219,224]. When computed this way, the average phase-position correlation was -0.22±0.052 (Figure 6-2B), which is comparable to precession reported for mice navigating in a virtual maze [224]. The population averaged STRFs corresponding to these two methods are shown in Figure 6-2C,D.

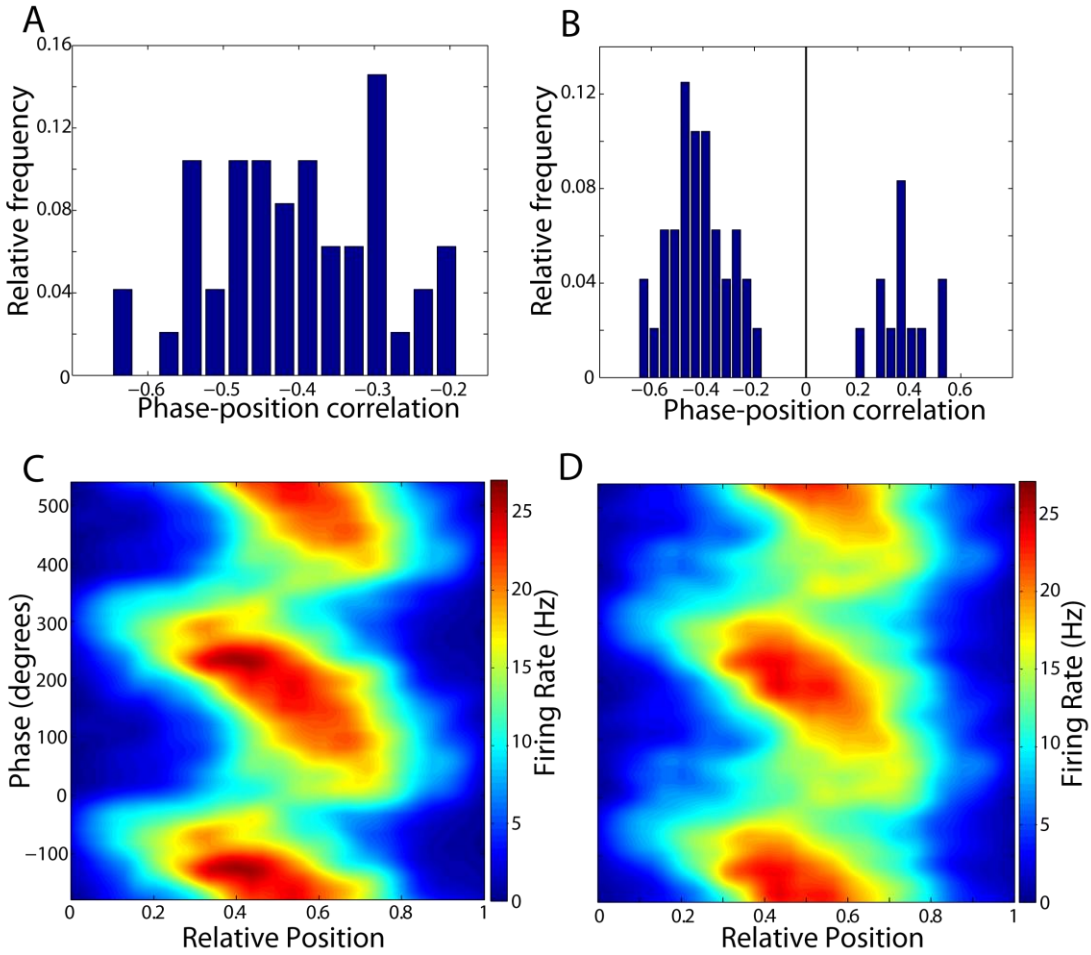

Figure 6-2: Phase precession in mice is weaker than reported in rats performing a similar task.

A) *Distribution of phase-position correlations across 59 place fields, with phase rotated to minimize the phase-position correlation (see Methods).* ***B)*** *Same as A, except with phase rotated to maximize the square of the phase-position correlation.* ***C)*** *Population-average spatio-temporal receptive field (STRF) showing firing rate as a joint function of relative position within the field and LFP theta phase, with phase rotated as in A.* ***D)*** *Same as C, except with phase rotated to maximize the square of the phase-position correlation.*

Notably, we find that precession is weaker in mice than has been reported in rats performing a similar task [26,110,212,219,320]. This could arise either due to inter-species difference in the temporal code itself, or as a result of differences in other parameters related to precession. Specifically, we tested the possibility that behavioral variability could account for some of the

impaired precession in mice. To this end we first excluded all trials in which the animal paused or stopped within the place field, and compared phase precession with and without inclusion of such trials. There were 48 place fields which still met our criteria for analysis after excluding such trials, and precession strength significantly improved for this set of place fields after excluding trials with stops and pauses (with stop trials: -0.37+0.016; without stop trials: -0.40$\pm$0.016; $p=2.6\times10^{-3}$, paired t-test). Hence, consistent with previous studies [212,213,224,323], all subsequent analyses were restricted to only the trials without stops. We next computed precession in this set of 48 place fields for each individual trial [212], along with the corresponding within-field running speed and acceleration (see *Methods*). Single-trial precession was significantly correlated with running speed when pooling trials across all place fields (n=479 trials, r=-0.13, $p= 4.3\times10^{-3}$; Figure 6-3A), or when computed separately for each place field (r=-0.13+0.047; n=45 fields, $p=8.6\times10^{-3}$, t-test). This is consistent with similar results reported in rats [212]. Interestingly, this weak but significant relationship between speed and single-trial precession is predicted to account for the difference in average precession between mice and rats, given the substantially higher running speeds observed in rats (Figure 6-3A) [212].

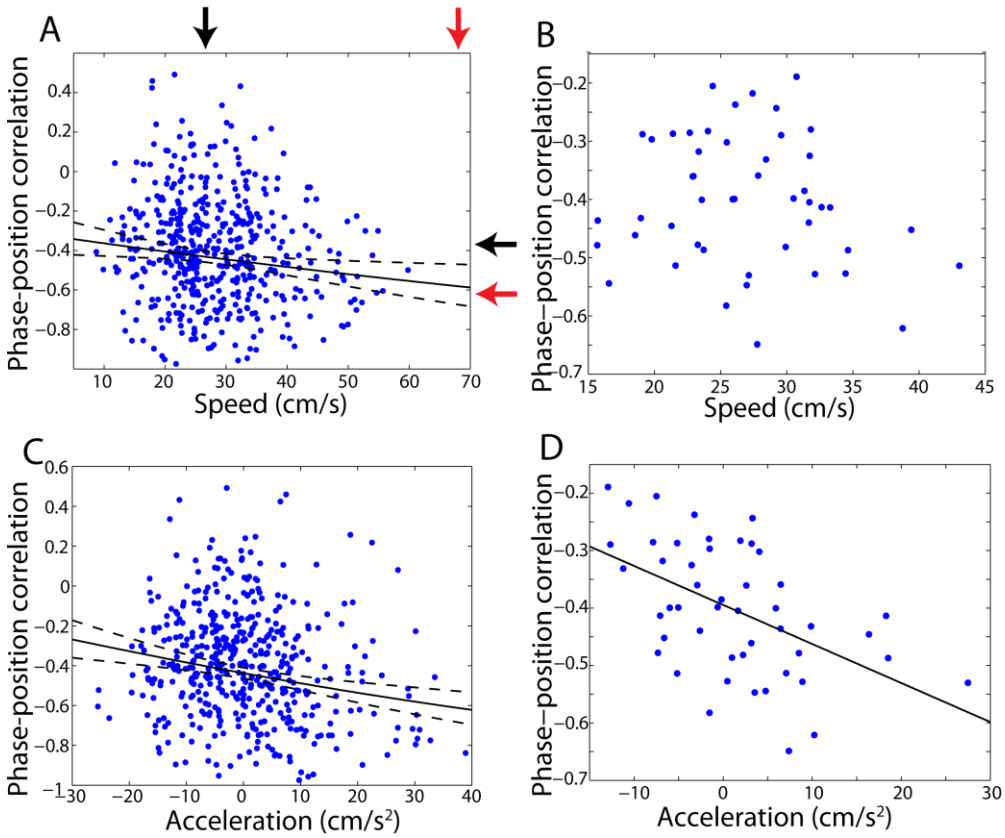

Figure 6-3: Single-trial precession is correlated with single-trial running speed, while overall precession is correlated with average acceleration.

A) *Single-trial phase-position correlation is plotted against single-trial running speed for all 479 trials pooled across place fields. The black trace shows a linear fit to the z-transformed correlation values (dashed traces indicating the 95% confidence interval of the fit). This linear fit predicts that the average single-trial precession in mice would be equivalent to that reported in rats (-0.61) when evaluated at the higher running speeds of rats (median 68cm/s). The black arrows indicate the median running speed and average phase-position correlations across all trials in mice, while the red arrows indicate the corresponding values reported by Schmidt et al.* ***B)*** *In contrast to single-trial precession, the overall precession of a place field was not correlated with the average running speed of the animal within the field.* ***C)*** *Pooled across place fields, single-trial precession was significantly correlated with single-trial acceleration.* ***D)*** *Overall precession was strongly correlated with the average acceleration of the animal within the field.*

Although single-trial speed was significantly correlated with single-trial precession, place fields where mice consistently ran at higher speeds did not have stronger precession compared to

place fields with low average speeds (correlation between average running speed and overall precession: r=-0.15, p=0.32; Figure 6-3B). As a further test, we performed an analysis of covariance (ANCOVA) to model the effects of between-field variability (group effects) and running speed (covariate effects) on single-trial precession. This analysis showed that both running speed ($p=6.2x10^{-4}$) and place field identify ($p=6.5x10^{-5}$) were significant predictors of single-trial precession. Thus, speed and precession were significantly correlated on a trial-by-trial basis within each place field, but this relationship could not account for the variability in average precession strength across place fields.

In addition to speed, we tested whether acceleration influences precession, a relationship that has not been previously investigated. When pooled across the ensemble of place fields, single-trial acceleration was weakly but significantly correlated with single-trial precession (r=-0.16, $p=3.4x10^{-4}$; Figure 6-3C). This dependence of single-trial precession on acceleration could not be explained by the speed-precession relationship since the partial correlation between precession and acceleration controlling for speed was still highly significant (r=-0.20, $p=1.6x10^{-5}$). Interestingly, unlike speed which could account for a significant portion of the within-field variability in single-trial precession, acceleration and precession were not significantly correlated across trials within a given field ($r=0.014\pm0.055$, p=0.80). Further, an ANCOVA showed that the overall correlation between pooled single-trial acceleration and precession could be entirely accounted for by between-field differences in precession strength (group effect: p=0.028, covariate effect: p=0.67). Thus, while across-trial changes in speed within a given place field influenced precession, acceleration did not exert such an influence. Instead, these results suggest that the variability in average precession across place fields must be related to the

average acceleration within the field. Indeed, the average within-field acceleration was strongly predictive of the overall precession for a given place field (Figure 6-3D), such that place fields where the mouse consistently accelerated had significantly better precession than those where the mouse consistently decelerated.

Since acceleration was only a significant predictor of the overall precession strength, and not of the single-trial precession strength, these results suggest that the acceleration-precession relationship may result from other parameter(s) that were systematically related to acceleration. In fact, on linear tracks the mice consistently accelerated as they left one goal location at the beginning of the track and consistently decelerated as they approached the target goal location at the other end of the track (Figure 6-4A). This raises the possibility that the average precession within a place field was dependent on the distance between the field and the reward location, rather than the within-field acceleration *per se*. Indeed, precession was significantly correlated with the distance between the end of the place field and the target reward location (Figure 6-4B), henceforth referred to as the 'relative position' of the place field. This result was true when using the overall precession strength ($r=-0.56$, $p=4.2\times10^{-5}$), or the average single-trial precession for each field ($r=-0.51$, $p=2.1\times10^{-4}$). Further, a similar relationship was found using other methods of estimating the phase-position correlation (maximizing R^2: $r=-0.40$, $p=5.5\times10^{-3}$ (Spearman's rho), circular-linear correlation: $r=-0.51$, $p=2.5\times10^{-4}$), and when considering both movement directions independently (left-to-right: $n=21$, $r=-0.64$, $p=1.8\times10^{-3}$; right-to-left: $n=27$, $r=-0.49$, $p=9.3\times10^{-3}$). In fact, the example place field that showed strong precession (Figure 6-1A) was at the beginning of the track, whereas the example place field with poor precession was at the end of the track near the reward (Figure 6-1B). This relationship

between precession and relative position was further confirmed by computing the population-average STRF across the set of 16 fields (1/3 of all fields) furthest from the reward location (Figure 6-4C) versus the set of 16 fields located closest to the reward (Figure 6-4D). Fields at the beginning of the track had 1.4-fold stronger precession than those at the end (beginning: r=-0.46$\pm$0.018; end: r=-0.32$\pm$0.022; $p=4.7\times10^{-5}$, two-sample t-test). Differences in the population-average STRFs for fields at the beginning and end of the track were even more pronounced after rotating the phase to maximize the square of the phase-position correlation (Figure 6-4E,F; beginning: r=-0.29$\pm$0.095; end: r=-0.10$\pm$0.086; p=0.025, Wilcoxon rank sum test).

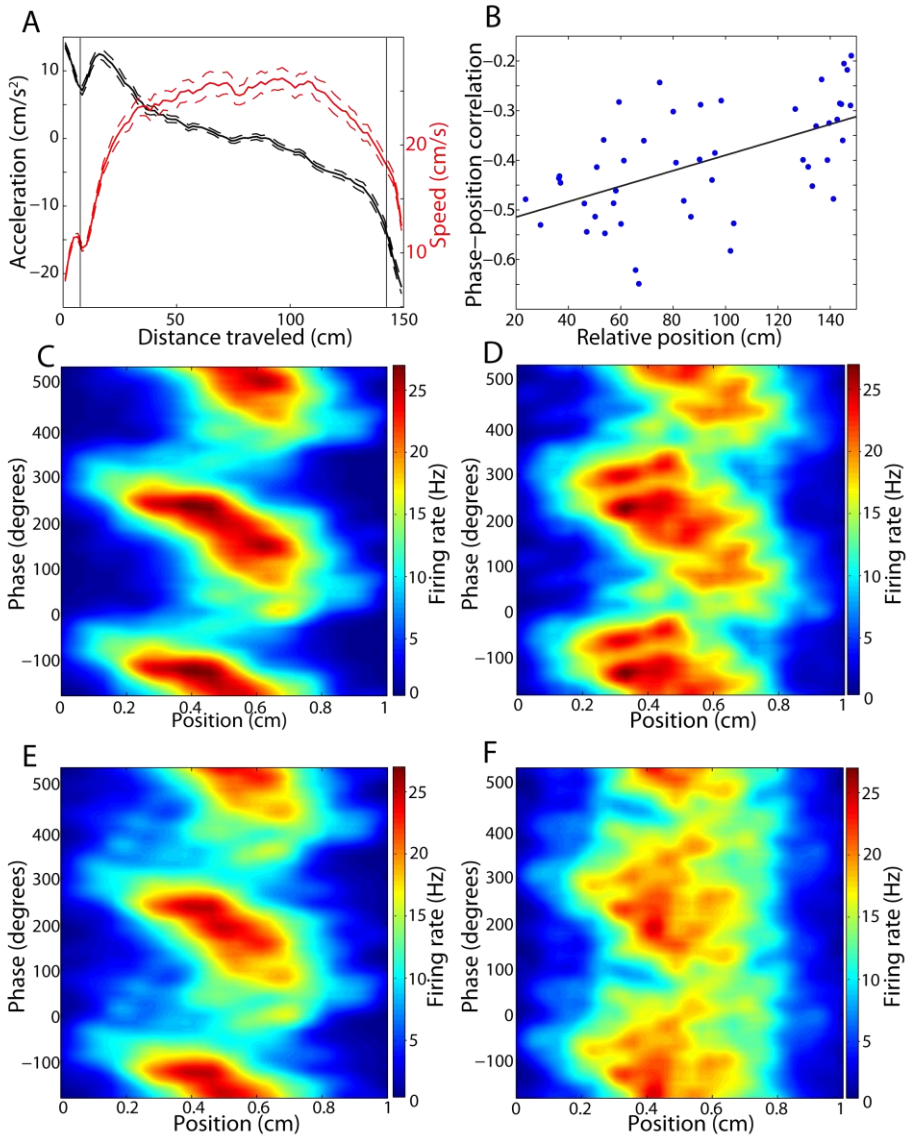

Figure 6-4: Precession is dependent on the relative position of the place field on the track.

A) *Both average acceleration (black trace) and average speed (red trace) were strongly dependent on the relative position of the animal on the track. While the average speed of the animal had a roughly parabolic profile, the average acceleration decreased nearly monotonically with distance traveled. Dashed lines indicate the interval (mean±SEM).* ***B)*** *Across-trial precession strength was strongly dependent on the relative position of the place field. The black line indicates the linear fit.* ***C)*** *Ensemble average STRF for the set of 16 place fields furthest from the reward location, with phase rotated to minimize the phase-position correlation,.* ***D)*** *Same as in C for the 16 place fields nearest the reward location showing poorer precession.* ***E)*** *Same as C, except with phase rotated to maximize the square of the phase-position correlation.* ***F)*** *Same as E for the set of place fields nearest the end of the track, showing almost no precession at all.*

Differences in running speed could not account for the observed relationship between precession and relative position, since the partial correlation between precession and relative position after controlling for average running speed was still highly significant ($r=-0.60$, $p=8.3\times10^{-6}$). Similarly, while the partial correlation between average acceleration and precession controlling for relative position was very small ($r=-0.075$, $p=0.62$), the partial correlation between precession and relative position controlling for acceleration was still considerable ($r=-0.27$, $p=0.071$); although neither partial correlation was found to be significant. Further, while the correlation with precession was strongest when defining place field 'relative position' as the distance between the end of the place field and the reward location, qualitatively similar results were obtained when using the distance between the reward location and the center of mass of the place field ($r=0.53$, $p=9.3\times10^{-5}$) or the beginning of the place field ($r=0.51$, $p=2.4\times10^{-4}$).

To better understand the nature of the impaired precession for place fields nearer the reward locations, we estimated the autocorrelation function of spike times within each place field (see *Data Analysis*). Interestingly, the minimum value of the autocorrelation function for lags < ~110ms (within one theta cycle) was significantly correlated with the relative position of the place field ($r=0.40$, $p=5.1\times10^{-3}$), such that the place fields at the end of the track had less well-defined valleys in the autocorrelation function that those at the beginning of the track (Figure 6-5A). To further test this, we computed the inter-spike interval (ISI) distributions for place fields at the beginning and end of the track. Indeed, place fields at the end of the track had fewer ISIs in the theta range (~100ms) than those at the beginning of the track (Figure 6-5B). We defined intra-theta-cycle (intra-cycle) ISIs as ISIs in the 15-80ms range. The fraction of intra-

cycle ISIs was significantly correlated with the relative location of the place field ($r=0.42$, $p=3.2\times10^{-3}$; Figure 6-5C) and also with its precession ($r=0.40$, $p=4.7\times10^{-3}$), but was not significantly correlated with acceleration ($r=-0.19$, $p=0.21$). This increase in the number of intra-cycle ISIs for place fields at the end of the track could thus disrupt precession and result in the observed relationship between precession and relative position. Interestingly, while most other place field properties, such as the average rate and peak rate, were not correlated with the relative location of the place field, place field width showed a significant increase with relative position on the track ($r=0.50$, $p=3.4\times10^{-4}$), such that place fields nearer the goal location were wider (Figure 6-5D).

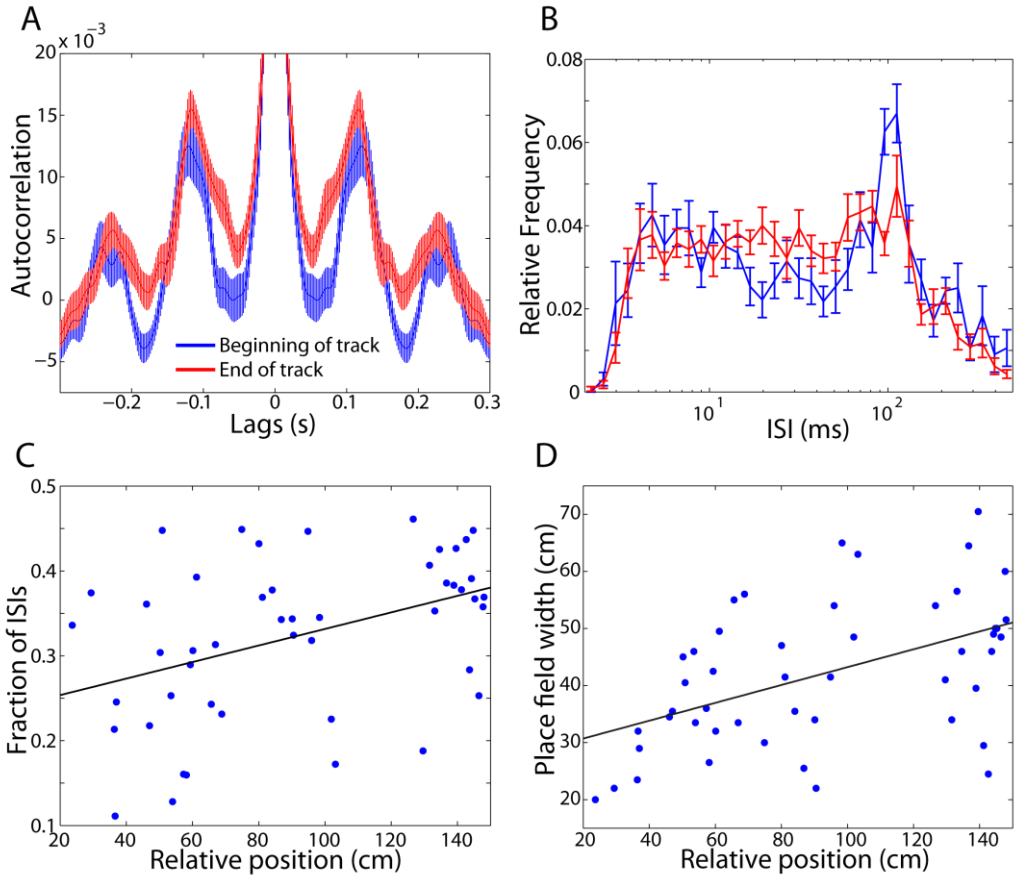

Figure 6-5: Increase in the number of intra-cycle ISIs for place fields at the end of the track.

A) *Average autocorrelation functions for the set of place fields (n=16) at the beginning (blue trace) vs. the end (red trace) of the track showing an increased correlation at intra-cycle lags (< 100ms) for place fields at the end of the track.* ***B)*** *Average ISI distribution for fields at the beginning (blue trace) and end (red trace) of the track. Note the reduction of theta-range (~100ms) ISIs and the increase in intra-cycle ISIs (15-80ms) for place fields near the end of the track.* ***C)*** *The fraction of intra-cycle ISIs was significantly correlated with the relative location of the place field.* ***D)*** *Place field width increased significantly as a function of relative position.*

Given that the precession of single-unit spike times relative to the LFP theta rhythm varies systematically with the relative position of the place field, we next asked whether the LFP theta rhythm itself also depends on the relative position of the animal. To this end, we computed the amplitude of theta oscillations as a function of relative position on the track. Theta amplitude increased (decreased) dramatically as the animal initiated (terminated) trial-running in the 'goal

regions' of the track (Figure 6-6A). Further, the parabolic profile of the average theta amplitude as a function of relative position suggested that the results could be reflective of a relationship between running speed and theta amplitude. Indeed, theta amplitude significantly increased with speed across the ensemble of recordings (slope: $9\pm1.6\times10^{-3}$ z·s/cm; $p=6.0\times10^{-7}$; Figure 6-6B). In order to more directly compare the LFP results with the precession analysis, and to separate the effects of transient changes in theta amplitude at the beginning and end of trials, we excluded data within the goal locations. After excluding the goal regions, theta amplitude still showed a significant dependence on running speed (slope: $3\pm1.6\times10^{-3}$ z·s/cm; p=0.040), although the relationship was substantially weaker (Figure 6-6B). Thus, in order to determine the direct dependence of theta amplitude on relative position and acceleration we first removed the linear dependence of theta amplitude on speed for each LFP recording. Interestingly, after excluding the large transient change in theta amplitude within the goal regions, there was a clear dependence of theta amplitude on relative position (slope: $-10\pm3.1\times10^{-4}$ z/cm; $p=2.1\times10^{-3}$; Figure 6-6C). Theta amplitude was also found to have a significant dependence on acceleration (slope: $13\pm5.7\times10^{-4}$ z·s^2/cm; p=0.022; Figure 6-6D). However, we found that relative position explained a significantly larger portion of the (residual) variance in theta amplitude (4.0±0.38%)

than acceleration (3.1±0.32%; $p=8.6\times10^{-5}$; paired t-test).

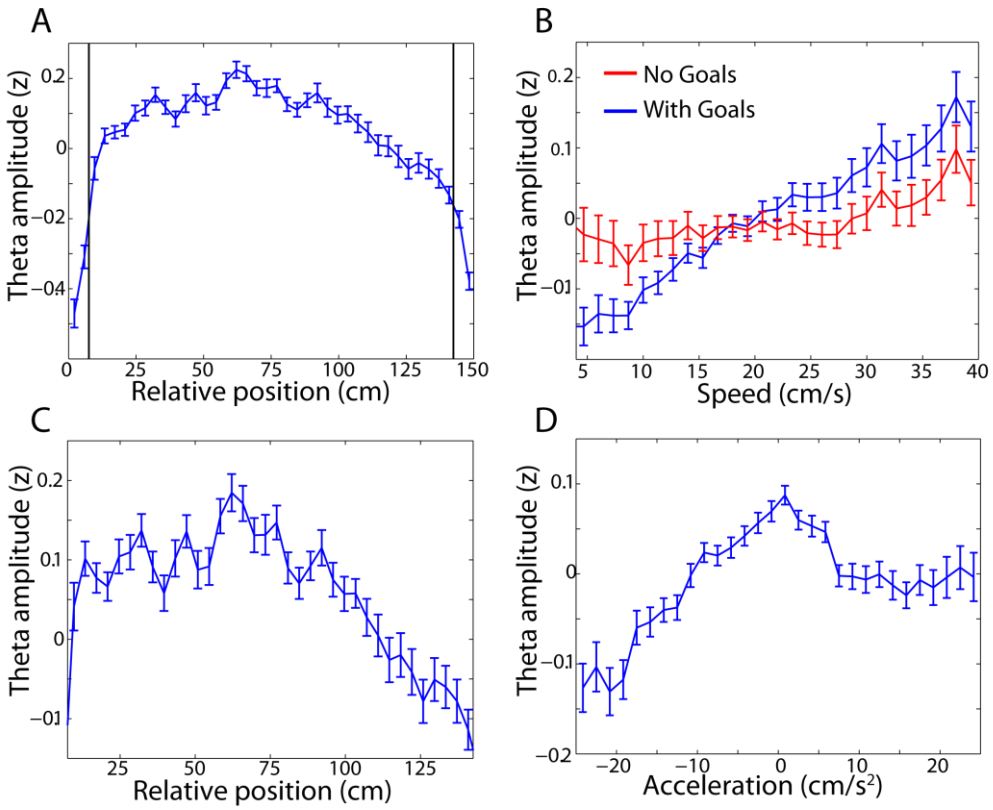

Figure 6-6: LFP theta oscillations decrease in amplitude with relative position on the track.

A) *Average LFP theta amplitude as a function of relative position on the track. The vertical black lines indicate the boundaries of the 'goal regions.' Note the dramatic decrease in the amplitude of theta oscillations within the goal regions where the mice initiated and terminated trial-running.* ***B)*** *Theta amplitude as a function of running speed using all data during trials (blue trace) or after exclusion of data within the goal regions (red trace). While theta amplitude showed a significant increase with running speed even after exclusion of goal regions, much of the dependence of theta amplitude on running speed could be accounted for by the lower theta amplitudes within the goal regions where running speeds were lowest.* ***C)*** *Similar to A, theta amplitude is plotted as a function of relative position on the track. In this case, any linear relationship between speed and theta amplitude is removed, and goal regions are excluded from analysis.* ***D)*** *Same as C, except for theta amplitude as a function of acceleration, showing a slight, but non-monotonic, increase in theta amplitude with increasing acceleration.*

Thus, not only did the precession of spikes relative to the theta rhythm change systematically with relative position, but the LFP theta rhythm itself decreased in amplitude as the animal

approached the goal location. We next asked whether the decrease in precession strength with relative position on the track could be explained by an increased uncertainty in assigning spike phases during low-amplitude theta oscillations. While the amplitude of LFP theta oscillations within the place field was indeed significantly correlated with the relative position of the field ($r=-0.42$, $p=3.0\times10^{-3}$), this relationship could not account for the dependence of precession on relative position since the partial correlation between precession and relative position was still highly significant after accounting for average LFP theta amplitude ($r=0.53$,$p=1.3\times10^{-4}$). Further, within-field LFP theta amplitude was not significantly correlated with precession strength ($r=-0.090$, $p=0.54$) on a trial-by-trial basis. Thus, the decrease in precession strength could not be explained by the decrease in LFP theta amplitude. We cannot determine whether the decrease in theta amplitude with relative position is a result of the systematic changes to phase precession across the ensemble of hippocampal neurons, or whether the two phenomena are independent of one another, potentially sharing a common cause.

Discussion

Here we quantified hippocampal phase precession in mice performing a simple linear track task, and found that phase precession depended on several behavioral parameters. In agreement with a previous report in rats [212], we found that phase precession in individual trials was weakly but significantly correlated with the animal's running speed such that stronger precession was observed at higher speeds. Furthermore, our results suggest that this relationship between speed and precession can account for the substantially reduced precession observed in mice compared to rats, such that the temporal code is likely similar across these species. Even though trial-by-trial fluctuations in speed and precession within a

place field were correlated, this did not result in any difference in the overall precession for place fields where the animal systematically ran faster, in agreement with previous reports in rats [26,213,214]. Thus, the dependence of single-trial precession on speed is much weaker than the variability in precession across place fields.

In addition to speed, we also found a significant correlation between precession and acceleration across trials that has not been previously reported. However, because within-field fluctuations in precession strength across trials were not correlated with variation in acceleration, it is more likely that precession strength is determined by the location of the place field relative to the reward, rather than the acceleration of the animal. After computing the spike time autocorrelation functions and ISI distributions for place fields at the beginning and end of the track, we found that the decrease in precession towards the end of the track likely resulted from an increase in the probability of intra-theta-cycle ISIs. Interestingly, place fields nearer to the reward location were also found to be significantly wider than those that were further from the reward.

Finally, we show that in addition to the systematic change in precession of single neurons with relative position, the amplitude of LFP theta oscillations also systematically decreased as the animal approached the reward location. In agreement with the precession analysis, the dependence of LFP theta amplitude on acceleration was weaker than the corresponding dependence on relative position. Further, this decrease in LFP theta amplitude could not explain away the observed relationship between precession and relative position. While a number of studies have reported a dependence of the hippocampal rate code on behavioral and

task-dependent variables [196,197,202,203,204,205,206,207], our results provide the first evidence for a clear dependence of the hippocampal temporal code on behavior.

What mechanisms could explain the strong relationship between phase precession and the relative position of the place field? One possibility is that dopaminergic levels in the hippocampus increase as the mouse gets closer to the goal location either in anticipation of reward, in response to reward-predicting stimuli, or in conjunction with the production of reward-seeking behavior [324,325,326,327,328,329,330,331,332,333]. Recent studies show that dopamine levels profoundly modulate the induction of long term plasticity [334,335,336], and in particular spike-timing-dependent plasticity (STDP) [337]. Zhang et al. showed specifically that application of physiological levels of dopamine abolished long-term depression (LTD) and resulted in a temporally symmetric STDP 'kernel', in contrast to the well-known temporally asymmetric STDP with both LTP and LTD [337]. Previous studies suggest that phase precession could be generated by spatially asymmetric synaptic inputs to place cells, which in turn arise from temporally asymmetric STDP [110,195]. Thus, the more symmetric STDP lacking LTD that arises in the presence of dopamine could result in impaired phase precession. The lack of LTD might also produce the wider place fields observed at the end of the track.

It is possible that impaired precession resulting from systematic increases in dopamine could in turn produce the reduced amplitude of theta oscillations also observed near the reward locations. However, the converse could be true as well, such that a dopamine-dependent reduction of theta amplitude could trigger a reduction in precession. Models of phase precession require theta rhythmic inhibitory modulation of CA1 place cells [110,222,223,338], impairment of which is expected to result in such a reduction of phase precession. Indeed,

administration of cannabinoid receptor agonist was shown to reduce both theta oscillation amplitude and the strength of phase precession, and also produced profound working memory impairments [320,339,340]. Interestingly, a recent study by Benchenane et al. showed that theta-band coherence between the hippocampus and prefrontal cortex (Pfc) was highest at the choice point in a Y-maze task, and that Pfc neurons' spike phases with respect to hippocampal theta also varied systematically at these times. These authors hypothesized that such alterations to theta oscillations and theta organization of spike timing could arise due to increased dopamine levels at the choice point, and confirmed that injection of dopamine reproduced similar effects under anesthesia [341].

While to our knowledge hippocampal dopamine levels have not been measured in freely behaving animals performing a similar spatial task, there is substantial evidence suggesting that dopaminergic input to the hippocampus is indeed likely to increase as the animal approaches the reward location. Increased activity of midbrain dopaminergic neurons, which project to the hippocampus, have been shown in response to reward-predicting stimuli [325,326,331,342], and during periods of reward expectation [329,330,331]. Further, direct measurements of dopamine levels in rodents [327,332,333,343] and humans [330] have demonstrated increases in dopamine levels during generation of reward seeking behaviors, or periods of reward expectation. Modulation of hippocampal activity by reward anticipation is also consistent with the findings that there are more place cells around reward regions [204,303,344] (but see [197,345]), and that the firing rates of place cells can increase with reward expectancy [197,203,206,346]. On the other hand, reward-related activity of dopaminergic neurons has also been shown to involve largely transient, sub-second responses to reward or salient

conditioned stimuli [325,326], and has been shown to decrease substantially for predictable rewards [326,331]. Thus, direct measurement of hippocampal dopamine levels in animals performing a similar task will be necessary to confirm this prediction.

Other mechanisms could potentially contribute to these results as well. In particular, even though both precession and LFP theta amplitude showed stronger dependencies on relative position than acceleration, we cannot rule out the possibility that acceleration was the determining factor in both cases. Such a mechanism would predict, however, that fluctuations in acceleration across trials would be correlated with changes in single-trial precession, an effect which was not observed. Differences in acceleration were also unable to account for the dependence of place field width on relative position. We note that kinematic variables, such as speed and acceleration, will be difficult to distinguish from the effects of reward anticipation, since the behavior of the animal depends on its motivational state, and because the dopamine system is thought to directly control the generation of reward-seeking behaviors [332,347]. Recent recordings in freely behaving rodents from the ventral tegmental area (VTA), which provides dopaminergic input to the hippocampus, have even shown that in addition to reward-related activity, these neurons can respond to the animal's velocity and acceleration directly [348]. Thus, further studies will be required to fully understand the contributions of behavior and reward expectancy to the hippocampal temporal code.

Our findings about the dependence of precession and LFP theta oscillation amplitude on relative position could have important functional consequences. Precession is hypothesized to be crucial for sequence learning [216,217,218]. Thus, stronger precession of place fields at the beginning of the track, as the mouse embarks on a journey, will facilitate predictive coding of the temporal

sequence of places [110,215,349] leading to the reward location. On the other hand, as the mouse reaches the reward zone, such predictive coding of a planned trajectory becomes less important. In addition to providing the context for this temporal code, the hippocampal theta rhythm has been shown to synchronize Pfc neurons [350], and is thought to play a role in working memory [103,340,351,352,353]. It will be important to explore these relationships more fully in order to better understand the functional role of the hippocampal theta rhythm and phase precession in behavior.

Chapter 7: **Conclusions**

Part 1: UDS in the Cortical-Entorhinal-Hippocampal Circuit

Summary of the results

In the first part of the thesis we explored the relationship between entorhinal and neocortical UP-DOWN states. In chapter 2, we presented an algorithm for inferring UDS from any set of continuous signals using hidden Markov models (HMMs). The HMM framework allows for any combination of signal features to be incorporated, and provides a natural means of evaluating different signal features in terms of their discrimination information. The method is highly robust to differences in the statistical properties of UDS across experimental conditions, and handles large non-stationarities within individual recordings by using time-varying model parameters and by inferring UDS only during epochs with identified UDS oscillations. Explicit state duration models were integrated into the UDS inference algorithm to avoid the implicit assumption of geometric state durations distributions, which were found to produce poor fits to the empirical state duration distributions. Finally we quantitatively evaluated the performance of the proposed algorithm in relation to standard methods by comparing the agreement between UDS inferred in the neocortical (Ncx) LFP and the MP of nearby Ncx neurons. Small, but consistent and significant, improvements on the order of 10% were observed in the agreement produced between LFP and MP UDS.

In chapter 3, we applied the EDHMM UDS inference algorithm to analyze the relationship between Ncx UDS and UDS in layer 3 pyramidal neurons of the medial and lateral entorhinal cortex (MECL3 and LECL3). We found that both MECL3 and LECL3 neurons exhibited UDS that

were phase-locked to Ncx UDS with a short delay, suggesting EC UDS were strongly influenced by upstream Ncx UDS. MECL3 neurons, but not LECL3 neurons, nearly always persisted in the UP state after the termination of the Ncx UP state, and nearly 25% of MECL3 UP states skipped the ensuing Ncx DOWN state entirely. Even when MECL3 UDS decoupled from Ncx UDS in the persistent UP state, the MECL3 UP-DOWN transitions occurred at precise times during the Ncx DOWN state. This led to a quantized distribution of Ncx UDS cycles spanned by MECL3 UP states. As a result of the persistent UP states, MECL3 neurons' UDS were less correlated with Ncx UDS, and had substantially lower frequency content compared to L3LEC neurons. Interestingly, a clear influence of these lower-frequency MECL3 UDS on the downstream hippocampal (Hpc) LFP was identified using partial coherence analysis. Finally, we found that while a number of studies have demonstrated the ability of EC neurons in all layers to generate self-sustaining depolarized and spiking states in response to electrical and/or synaptic stimulation, we were unable to elicit the MECL3 persistent UP states *in vivo* using similar depolarizing stimuli. Thus, network mechanisms, as well the cellular mechanisms described *in vitro*, are likely to contribute to the observed MECL3 persistent UP states *in vivo*.

While LECL3 pyramidal neurons did not exhibit persistent UP states as seen in MECL3 neurons, in Chapter 4 we showed that superficial LEC neurons instead showed periods of pronounced 2-4Hz oscillations that occurred in conjunction with ongoing UDS, and were not seen in superficial MEC neurons. These 'delta' oscillations occurred in more than half of the superficial LEC neurons tested, and appeared exclusively during discrete epochs typically lasting on the order of 100s. While individual LEC neurons were substantially more depolarized and active at these times, the delta epochs were also found to correspond to clear changes in the global network

state. In fact, 2-4Hz oscillations were also observed in the Ncx and Hpc LFP, particularly during the DOWN states, and the network LFP delta oscillations were synchronized with the MP oscillations of individual LEC neurons. Finally, we showed that rather than being independent of ongoing UDS, the delta oscillations in both the Ncx LFP and LEC MP were phase-locked to LEC DOWN-UP and UP-DOWN state transitions.

Future Research

An important part of a quantitative brain-wide understanding of UDS is the ability to reliably and accurately determine the timing of both single-cell and network-level UDS. The EDHMM approach developed in this thesis handles both LFP and MP signals in the same general framework, thus allowing quantification and direct comparison of both types of UDS. Because of the flexibility and robustness of the algorithm, it can also be applied directly to data recorded during non-REM (NREM) sleep where similar slow oscillations are present [123,128,129]. The advantages of the EDHMM algorithm are likely to be even greater in such cases, where the UP and DOWN states are more difficult to identify [128,172,231]. In addition to allowing the study of UDS (slow-oscillations) under natural conditions, careful quantification of UDS properties across the different stages of sleep and depths of anesthesia will be important for testing models of UDS.

Several important generalizations of our method should also be considered. One such generalization will be to incorporate point process data in the UDS inference process. Point-continuous process hybrid inference is achievable within the HMM framework [231], and would allow for more robust and precise UDS inference. Another interesting possibility would be to perform simultaneous inference of UDS across multiple brain regions from multi-channel LFP or

EEG recordings. This could be accomplished within the more general framework of dynamic Bayesian networks by modeling the UDS of each brain region as a set of interdependent latent random variables, each associated with a unique set of observable signal features. In this way, simultaneous inference of UDS across multiple brain regions could be achieved by learning/utilizing the relationships between UDS in the different regions. Unbiased classification of UDS across different brain regions and neuron types, where UDS properties can vary widely, also provides an opportunity to study network dynamics and functional interactions between different brain regions. Of particular interest is the characterization of the complete neocortical-entorhinal-hippocampal (Ncx-EC-Hpc) circuit, since this circuit is believed to mediate long-term memory consolidation processes during slow-wave sleep [114,164,165]. To this end, we applied our algorithm to classify and compare UDS from superficial neurons in the MEC and LEC and their relation to Ncx UDS.

The demonstration that MECL3 pyramidal neurons selectively decouple from the Ncx UDS and persist in the UP state represents an important step towards understanding the role of the EC in mediating cortico-hippocampal interactions during sleep. Future work should aim to uncover the precise nature of the MECL3 persistent activity, as well as the underlying mechanisms which produce it. One key question will be whether the MECL3 persistent UP states occur in an isolated fashion among individual MECL3 neurons, or whether groups of MECL3 neurons or even the entire MECL3 network exhibit persistent UP states synchronously. Several experimental approaches, including dual whole-cell recordings of MECL3 neurons, as well as single-unit (tetrode) and LFP recordings from MECL3 might provide answers. The presence of a high degree of synchrony amongst MECL3 UDS would indicate a strong contribution of network mechanisms

to the generation of MECL3 persistent UP states. Indeed, our finding that depolarizing current injection was unable to elicit persistent UP states *in vivo*, unlike in previous *in vitro* studies [260,261,262], already suggests an important role of such network mechanisms. However, these *in vitro* studies also demonstrated the importance of muscarinic acetylcholine receptor (mAChR) activation in producing persistent activity in EC neurons. Thus, pharmacological manipulations (such as the direct application of mAChR agonists and antagonists) could be used to determine whether such receptors also contribute to persistent UP states *in vivo*.

A better understanding of the persistent UP states at the level of the entire MECL3 network will also provide insight into its functional role. We hypothesize that MECL3 persistent activity during natural slow-wave sleep could control cortico-hippocampal interactions thought to be required for long-term memory consolidation [114,165,173]. Specifically, by activating hippocampal neurons during the Ncx DOWN state, MECL3 persistent UP states could produce the interleaved activation of old and new memories required by theories of memory consolidation [166]. To this end, it will be important to determine whether MECL3 persistent activity indeed occurs during natural sleep, and whether its properties are similar to those we observed under anesthesia. Further, while we showed that the MECL3 persistent activity was reflected directly in the downstream Hpc LFP, a better understanding of the precise influence of MECL3 persistence on Hpc activity will be required. For instance, the MECL3 persistent UP states could explain the activation of CA1/CA3 pyramidal neurons during the Ncx DOWN states [148,149], and might also initiate the replay of Hpc activity patterns produced during previous behavior [169,172]. To this end, simultaneous recordings of hippocampal single-unit activity will be important. Another approach may be to determine the effects of selective lesions of MECL3

and LECL3 projections to Hpc on UDS in the Ncx, Hpc, and EC. Indeed, general (not distinguished between MEC and LEC) lesions of these ECL3 inputs have been shown to result in profound impairments of spatial memory consolidation, even while spatial learning and short-term memory remained intact [256].

The fact that persistent UP states were distinctly absent in the LECL3 pyramidal neurons adds further support to the idea that the LEC and MEC have distinct functional roles. The MEC is known to receive stronger visuo-spatial Ncx inputs than the LEC [184,185], and has been shown to exhibit spatially periodic grid cell activity [208] not found in LEC neurons [258]. Thus, the MEC is thought to be involved in spatial processing while the LEC mediates non-spatial processing. It will be important to determine whether the LEC and MEC serve similarly segregated roles in memory consolidation. For example, it is possible that persistent UP states could be required for consolidating spatial information, but not other types of information.

As a further dissociation of LEC and MEC activity during UDS, we found that superficial (layer 2 and 3) neurons in the LEC, but not the MEC showed periods of pronounced 2-4Hz oscillations that occurred along with the ongoing UDS. Similar work in anesthetized rodents has shown the presence of oscillations in a slightly higher-frequency range (3-5Hz) occurring during 'desynchronized' or 'activated' network states, where the hyperpolarizing phases of the UDS are absent. These 'theta' oscillations are thought to be analogous to the 6-10Hz theta present during REM sleep and during natural behavior [129,352], which are important for spatial processing [26]. Interestingly, a recent study showed that 6-10Hz theta oscillations during behavior were much stronger in MEC neurons compared to LEC neurons [277], consistent with the purported functional roles of these regions in spatial and non-spatial processing

respectively. In contrast, the 2-4Hz oscillations described in Chapter 4 were only present in LEC neurons, suggesting they may be of a qualitatively different nature than the theta oscillations.

We noted several other differences between the LEC oscillations and previously described theta rhythms in anesthetized animals. Most importantly, instead of appearing only during the activated state, the LEC oscillations were present simultaneously with ongoing EC and Ncx UDS, and in fact were most prominent during the DOWN states. Concurrent and synchronized 2-4Hz oscillations were also present in both the Ncx and Hpc LFPs, whereas theta oscillations have not been described in the Ncx LFP. In this regard, we believe the LEC 2-4Hz oscillations are instead analogous to delta oscillations previously described in thalamocortical (TC) neurons and the Ncx EEG [65,150]. Thus, our results indicate that while superficial LEC and MEC neurons are both capable of producing oscillations in the 2-10Hz frequency range, different network states result in the presence of either 2-4Hz delta oscillations in the LEC or 3-5Hz (6-10Hz in non-anesthetized preparations) theta oscillations in the MEC. It will be important for future studies to compare the activity of superficial LEC and MEC neurons during different stages of sleep to determine whether these different oscillations are present in natural sleep conditions. Indeed, a recent report showed that urethane anesthesia produces much of the brain-state and physiological dynamics of natural sleep [129]. Thus, we hypothesize that the LEC neurons participate in the 1-4Hz delta oscillations that occur during deep stages of nREM sleep [354], while the MEC neurons exhibit 6-10Hz oscillations during REM sleep.

Because we observed LEC delta oscillations during particular network states in which synchronized delta oscillations were also apparent in the Ncx and Hpc LFP, it is highly likely that delta oscillations would be present synchronously among superficial LEC neurons at these times.

Extracellular (LFP/MUA) recordings from LEC in conjunction with simultaneous recordings of MP from LEC neurons could determine the prevalence and degree of synchronization of LEC delta oscillations. Another important question for future experiments will be whether the observed delta oscillations are generated within the LEC, within the thalamus (by the intrinsic cellular mechanisms of TC neurons [67]), or by a separate generator entirely. One clear direction for distinguishing these possibilities would be to record thalamic LFP/MUA or MP signals simultaneously with the LEC MP and Ncx/Hpc LFP. Appropriate analysis techniques (e.g. the Granger causality spectrum) could then be used to determine the causal relationships between delta rhythms present in the thalamus, Ncx, Hpc, and LEC. Another approach may be to selectively lesion either the LEC or thalamus and measure the effects on delta oscillations at these various locations. Either of these approaches, however, could have difficulty resolving a mechanism based on cortico-thalamic interactions, or the presence of multiple independent, but interacting, generators of delta oscillations. In fact, it is believed that there are at least two independent generators of delta oscillations during natural sleep, one based on intrinsic properties of TC neurons [67], and another, less well-understood, delta oscillation thought to be generated intracortically [150].

To date, the functional role of delta oscillations during natural sleep is unknown. Steriade and colleagues have claimed that the thalamic delta oscillations are temporally organized, or 'grouped', by ongoing UDS [150]. This is likely analogous to the phase-locking of LEC delta oscillations to both the DOWN-UP and UP-DOWN transitions that we described in Chapter 4. While the delta oscillations were found to be clearly phase-locked to LEC UDS, and only weakly phase-locked to Ncx UDS as measured by the Ncx LFP, a similar comparison of delta oscillations

with the UDS of individual Ncx neurons will be needed to determine whether the Ncx and EC UDS are actually differentially phase-locked to delta oscillations. Similarly, experiments relating hippocampal single- and multi-unit activity to delta oscillations will be critical towards determining the effects of delta oscillations on UDS in the Ncx-EC-Hpc circuit.

Since delta oscillations in both TC and LEC neurons are most prominent during the DOWN states, synchronized delta-mediated spikes within the DOWN states could serve to generate the DOWN-UP transition. Since UDS also occur in the absence of delta oscillations this is not likely to be the only mechanism generating DOWN-UP transitions, however delta oscillations could control the relative timing of UDS between different regions within particular network states. Specifically, delta oscillations during both the DOWN and UP states could define temporal windows within which DOWN-UP and UP-DOWN state transitions can occur. Varying degrees of delta synchronization between TC and LEC delta would in turn determine the relative timing of Ncx and Hpc UDS. In fact, the observation that precise spatiotemporal patterns of Ncx activity occur following the DOWN-UP transition [138](but see [355]) suggests that the precise timing of such state transitions may be critical.

Interestingly, both the MECL3 persistent UP states described in Chapter 3 and the delta oscillations of superficial LEC neurons shown in Chapter 4 exhibit a form of cross-frequency coupling to the Ncx (or LEC) UDS. The quantization of the number of Ncx UDS cycles spanned by MECL3 UP states can equivalently be viewed as *m*:1 cross-frequency phase-phase coupling of Ncx and MECL3 UDS, with a variable frequency ratio *m*. Similarly, the phase-locking of LEC delta oscillations to the DOWN-UP and UP-DOWN transitions can be viewed as *m*:1 coupling of delta and UDS oscillations. Thus, both phenomena can be fit into the context of a growing body of

research indicating the striking prevalence of cross-frequency coupling between different neural oscillations. These studies also support and extend the notion that the UDS (slow-oscillations) temporally organize the other electrophysiological structures present during sleep.

This work thus represents an important contribution towards fully characterizing UDS in the Ncx-EC-Hpc circuit, however, several key questions still remain. Work by Hahn et al. and Isomura et al. have shown that neurons in the different Hpc subfields have varying participation in UDS. CA1 and CA3 pyramidal neurons were found to lack clear UDS and were weakly or negatively related to Ncx UDS [148,149], while R-LM interneurons in CA1 showed clear UDS phase-locked to Ncx UDS [147]. Dentate gyrus (DG) granule cells were also found to have somewhat weaker UDS modulation [148,149]. Interestingly, while UDS observed in DG granule cells and CA1 R-LM interneurons were delayed relative to Ncx UDS it is unclear whether these delays are fully consistent with the substantial (100-250ms) delay in the afferent superficial LEC and MEC UDS observed both in our work and that of Isomura et al [149]. Further, Isomura et al. showed that layer 5 pyramidal neurons in the EC (ECL5), as well as subicular neurons, lack such a delay entirely. This is particularly surprising considering that these neurons are the recipients of hippocampal output from CA1 pyramidal neurons. One possibility is that the Ncx DOWN-UP transition could produce rapid activation of ECL5 neurons via direct cortical afferents [181], creating both synchronous DOWN-UP transitions in ECL5 neurons, and triggering a wave of IPSPs onto superficial EC neurons [356]. This wave of inhibition would then produce the substantial delays in UDS observed in superficial EC neurons, as well as in downstream Hpc structures. To this end, a systematic characterization of EC interneuron behavior during UDS would be important. Since simultaneous recording from all components of the Ncx-EC-Hpc circuit is not

feasible, careful statistical analysis of the behavior of each component independently in relation to the Ncx UDS and Hpc activity, as demonstrated here, may be the most effective approach.

Although the behaviors of Ncx neurons during UDS and natural sleep have been relatively well-documented, much less is known about the relation of neurons in the Hpc and EC to Ncx UDS. Understanding the Ncx-EC-Hpc circuit during slow-oscillations will be critical for understanding how memories are consolidated into the Ncx during sleep, and may shed light on the functions of these neurons during behavior. The work presented in this thesis adds to an evolving picture of UDS in the Ncx-EC-Hpc circuit in which a variety of inter-related phenomena create a far more complex picture than previously thought.

Part 2: The hippocampal rate and temporal codes for position

Summary of the Results

In the second part of the thesis we investigated the hippocampal rate and temporal codes of position using genetically modified mice and novel analysis approaches. In Chapter 5, we showed that hippocampal CA1 neurons ('place cells') in KO mice lacking the GluA1 AMPAR subunit had dramatic reductions in all forms of spatial and directional selectivity compared to WT neurons. KO place cells also showed an increase in random trial-by-trial fluctuations in their spatial firing rate maps, as well as a decrease in systematic changes across trials. The temporal structure of KO spiking was also found to be altered, showing an overall increase in bursting, as well as lower-frequency bursts and more intra-theta-cycle ISIs. These changes in the activity of

individual KO neurons were also paralleled by dramatic reductions in theta and gamma network oscillations in the LFP, as well as a reduction of theta-gamma coupling.

In Chapter 6, we shifted focus to the temporal code of position given by the advancement ('precession') of hippocampal neurons' spike phase relative to the theta rhythm as the animal runs through the neuron's place field. Instead of providing a behavior-independent code for position as previously thought, we found that precession within a given place field was strongly correlated with the location of the place field relative to the reward, such that precession progressively got weaker as the animal approached the reward. Spike time autocorrelation and ISI distribution analyses revealed that this was likely the result of an increased probability of intra-theta-cycle ISIs nearer the reward location. We also found a similar decrease in the amplitude of hippocampal theta oscillations as the animal approached the reward.

Future Research

A growing body of literature using genetic, pharmacological, and lesion manipulations is uncovering the mechanisms that generate the activity patterns of CA1 neurons and their relationship to behavior. Of particular importance are the various forms of synaptic plasticity, as well as the different synaptic inputs from CA3, ECL3, and a plethora of inhibitory interneuron types. In this vein our results lend support and new insights to our current understanding. CA1 pyramidal neurons from GluA1 KO mice show specific deficits in tetanically induced long-term plasticity (LTP), as well as the transient component of theta-burst induced plasticity. In addition, the lack of GluA1 containing AMPARs results in a profound reduction in the strength of distal synapses, but not proximal ones. Both the impairments of synaptic plasticity and synaptic strength could potentially result in the observed deficits in CA1 spatial selectivity. Indeed,

manipulations which produce more dramatic alterations of LTP, such as genetic activation of CaMKII [304], application of an NMDA blocker [317], or genetic silencing of NMDARs within CA1 [318] have been shown to produce comparable, or even less pronounced, effects on CA1 spatial selectivity. Similarly, manipulations which produce much larger perturbations of synaptic inputs to CA1, such as through complete lesion of CA3 inputs [357], genetic inactivation of CA3 [358], or lesion of direct ECL3 inputs [359], produce relatively mild impairments of CA1 spatial selectivity. Thus, it is possible that both factors contribute to the large impairments of spatial selectivity in GluA1 KO mice.

Interestingly, these widely different manipulations all seem to produce a qualitatively similar 'failure mode' whereby the CA1 rate code for position becomes less sparse. We propose that this could result from the impairment of inhibition-mediated competition among CA1 place cells. From this perspective, any manipulation that decreases the overall excitatory drive onto CA1 would result in a reduced recruitment of feedforward and lateral inhibition, thereby reducing competitive interactions between CA1 place cells and enabling more diffuse spatial firing. Recent work by Korotkova et al. showing similar impairments of CA1 spatial selectivity following genetic deletion of NMDARs from PV-expressing interneurons also supports this model [309], since a decrease in the E-I driving force would also impair competition within CA1. Another important prediction of this model that remains to be tested is whether there are fewer CA1 pyramidal neurons which do not spike within a given environment ('silent cells') [360] in the GluA1 KO animals, or indeed following any manipulation reducing competitive interactions within CA1. This could be determined by comparing spiking activity during track-running to that recorded during sleeping/resting both before and after track-running.

The relationships between CA1 activity and behavior appear similarly complex. While lesion experiments have shown that the hippocampus is required for a number of tests of spatial memory [193], previous behavioral experiments on the GluA1 KO mice have shown a stark dissociation between two forms of spatial memory known as spatial reference memory (SRM) and spatial working memory (SWM) [284,285,361]. Since the GluA1 KO mice are severely impaired on various tests of SWM but not SRM, our results imply that spatial selectivity of any form is not likely to be required for spatial SRM tasks. This is consistent with earlier experiments showing SRM deficits corresponding to very mild alterations of CA1 spatial selectivity [317], and normal SRM function despite reductions in CA1 spatial selectivity [358]. Our results do indicate that CA1 spatial selectivity is required for normal SWM, however it is unclear which aspect(s) of CA1 activity in particular are necessary. Evidence from manipulations of CB1 receptors shows that while CB1 KO animals have similar impairments to SWM and not SRM [319], application of CB1 antagonist leaves the rate code unaffected while severely impairing the temporal organization of CA1 activity [320]. Thus, it will be important for future work to test whether similar alterations of the hippocampal temporal code are also observed in the GluA1 KO mice.

Future experiments should also aim to uncover the relationship of hippocampal theta and gamma network oscillations to both CA1 spatial selectivity and spatial memory. Numerous studies have implicated both theta and gamma oscillations [113], as well as the theta-gamma coupling [116] in memory function, yet the role of these oscillations in shaping CA1 activity (and vice versa) remains less clear. Models of gamma oscillations depend on reverberation among networks of inhibitory interneurons [362] or interactions between excitation and recurrent

inhibition [363]. Hence, the observed deficit in gamma oscillations in the GluA1 KO would fit with the impaired inhibition model. However, this cannot explain the increase in gamma oscillations observed by Korotkova et al. in the PV-interneuron NMDAR KO animals [309]. Careful comparison of the theta and gamma oscillations, as well as their coupling, following each of the genetic, pharmacological, or lesion manipulations discussed will likely be required to formulate a viable model.

In Chapter 6 we show that the hippocampal temporal code, as well as the theta rhythm that defines it, depend on the proximity of the animal to the reward location. Although this is the first demonstration of any behavioral or task-related modulation of the temporal code, our results are still consistent with the hypothesis that the temporal code, in contrast to the rate code, provides a representation of position that is invariant to trial-by-trial variability in kinematic parameters such as running speed. While we found that the animal's relative proximity to the reward location was a more likely explanation for the large differences in precession and theta oscillation amplitude, we cannot rule out the possibility that strongly related parameters such as the average acceleration of the animal, or the time relative to movement initiation, could instead be the underlying cause. Future experiments should be able to confirm or deny such possibilities. For instance, the effects of time relative to movement initiation could be dissociated from reward proximity by designing tasks to manipulate the running speed of the animals. Similarly, tasks in which the animal is required to systematically accelerate/decelerate at various locations relative to reward delivery could clearly distinguish the effects of acceleration and reward proximity.

Based on the interpretation that the observed variation in precession is indeed related to the animal's proximity to reward, we suggest that this could arise due to systematic changes in dopaminergic inputs to the hippocampus. Numerous studies have shown that dopaminergic neurons respond to reward receipt, reward-predicting stimuli, or in conjunction with reward-seeking behavior [324,325,326,327,328,329,330,331,332,333]. Further, mounting evidence indicates that dopamine modulates long-term synaptic plasticity [334,335,336], and particularly spike-timing dependent plasticity (STDP), by changing long-term depression into potentiation and removing the temporal asymmetry of STDP [337]. Such changes to STDP could indeed result in impaired precession [110,312], and could also explain the observed increase in place field size near the reward locations. Other effects of dopamine, such as changes to the inhibitory theta rhythmic modulation of pyramidal neurons, could also play a role in altering phase precession. A recent study by Benchenane et al. argued that increased dopamine levels at the 'choice point' in a Y-maze task produced an increase in theta coherence between hippocampus and prefrontal cortex (Pfc), as well as an alteration of Pfc spike timing relative to hippocampal theta [341]. It would thus be interesting to determine whether similar changes to Pfc theta also occur near the reward location in our linear-track task.

Direct measurement of dopamine levels in the hippocampus in behaving animals with the required temporal resolution has recently become possible using fast-scan cyclic voltammetry [364]. Such measurements of fluctuations in hippocampal dopamine concentration and their relation to precession and theta oscillations would thus be the most direct test of our hypothesis. Simultaneous recordings of the activity of neurons in the ventral tegmental area (VTA), which provide dopaminergic input to the hippocampus, would be another approach,

although distinguishing dopaminergic and non-dopaminergic VTA neurons from extracellular recordings alone is difficult [348]. Alternatively, systematic manipulations of reward magnitude and probability have been shown to correlate strongly with the activity of dopaminergic neurons [326,331], and such manipulations could be correlated with hippocampal phase precession in a straightforward manner.

Several key questions remain that underlie both our study of the hippocampal rate code in GluA1 KO mice, and the hypothesized reward-dependence of the temporal code. One is whether our results reported in mice can be generalized to rats, with which most studies of hippocampal activity have been conducted. Several differences in rat and mouse hippocampal anatomy have been noted, although the extent of similarities in the rate and temporal codes for position are less well understood [190]. Another important step will be to study the activity of simultaneously recorded place fields. In both of our studies we analyzed the activity of place cells independently, however it remains to be seen whether the effects of GluA1 deletion or the proposed reward-modulation of precession will have additional consequences for the precise spike timing of sequences of place cells with overlapping place fields. Finally, it will be interesting to investigate the nature of hippocampal replay during sleep, both in GluA1 KO mice, and as a function of the location of place fields relative to reward. Such future studies, combined with the work presented in this thesis, will significantly advance our understanding of the mechanisms that generate the hippocampal rate and temporal codes, and the relations between this activity and behavior.

Bibliography

1. Hodgkin AL, Huxley AF (1952) A quantitative description of membrane current and its application to conduction and excitation in nerve. J Physiol 117: 500-544.

2. Hodgkin AL, Huxley AF, Katz B (1952) Measurement of current-voltage relations in the membrane of the giant axon of Loligo. J Physiol 116: 424-448.

3. Ling G, Gerard RW (1949) The normal membrane potential of frog sartorius fibers. J Cell Physiol 34: 383-396.

4. Hamill OP, Marty A, Neher E, Sakmann B, Sigworth FJ (1981) Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches. Pflugers Arch 391: 85-100.

5. Stuart GJ, Dodt HU, Sakmann B (1993) Patch-clamp recordings from the soma and dendrites of neurons in brain slices using infrared video microscopy. Pflugers Arch 423: 511-518.

6. 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.

7. Creutzfeldt OD, Watanabe S, Lux HD (1966) Relations between EEG phenomena and potentials of single cortical cells. I. Evoked responses after thalamic and epicortical stimulation. Electroencephalogr Clin Neurophysiol 20: 1-18.

8. Logothetis NK (2003) The underpinnings of the BOLD functional magnetic resonance imaging signal. J Neurosci 23: 3963-3971.

9. Gray CM, Singer W (1989) Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc Natl Acad Sci U S A 86: 1698-1702.

10. Avitan L, Teicher M, Abeles M (2009) EEG generator--a model of potentials in a volume conductor. J Neurophysiol 102: 3046-3059.

11. Kamondi A, Acsady L, Wang XJ, Buzsaki 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.

12. Gustafsson B (1984) Afterpotentials and transduction properties in different types of central neurones. Arch Ital Biol 122: 17-30.

13. Buzsaki G, Kandel A (1998) Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat. J Neurophysiol 79: 1587-1591.

14. 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.

15. Xing D, Yeh CI, Shapley RM (2009) Spatial spread of the local field potential and its laminar variation in visual cortex. J Neurosci 29: 11540-11549.

16. Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, et al. (2006) Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. Neuron 49: 433-445.

17. Csicsvari J, Henze DA, Jamieson B, Harris KD, Sirota A, et al. (2003) Massively parallel recording of unit and local field potentials with silicon-based electrodes. J Neurophysiol 90: 1314-1323.

18. Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. Nature 412: 150-157.

19. Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat Neurosci 5: 805-811.

20. Mehring C, Rickert J, Vaadia E, Cardosa de Oliveira S, Aertsen A, et al. (2003) Inference of hand movements from local field potentials in monkey motor cortex. Nat Neurosci 6: 1253-1254.

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

22. Belitski A, Gretton A, Magri C, Murayama Y, Montemurro MA, et al. (2008) Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. J Neurosci 28: 5696-5709.

23. Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, et al. (2000) Intracellular features predicted by extracellular recordings in the hippocampus in vivo. J Neurophysiol 84: 390-400.

24. Abeles M, Goldstein M (1977) Multispike train analysis. Proceedings of the IEEE 65: 762-773.

25. Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsaki G (2000) Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. J Neurophysiol 84: 401-414.

26. O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. Hippocampus 3: 317-330.

27. Gray CM, Maldonado PE, Wilson M, McNaughton B (1995) Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. J Neurosci Methods 63: 43-54.

28. Izhikevich EM (2007) Dynamical systems in neuroscience: the geometry and excitability of bursting. Cambridge: MIT Press.

29. Priestley MB (1981) Spectral Analysis and Time Series. London: Academic Press.

30. Lyons R (2004) Understanding digital signal processing: Prentice Hall.

31. Thomson DJ (1982) Spectrum Estimation and Harmonic-Analysis. Proceedings of the IEEE 70: 1055-1096.

32. Mitra P, Bokil H (2008) Observed Brain Dynamics: Oxford University Press.

33. Slepian D, Pollack H (1961) Prolate spheroidal wavefunctions: Fourier analysis and uncertainty I. Bell System Tech Journal 40: 43-63.

34. Pardey J, Roberts S, Tarassenko L (1996) A review of parametric modelling techniques for EEG analysis. Med Eng Phys 18: 2-11.

35. Barnett TP, Johnson LC, Naitoh P, Hicks N, Nute C (1971) Bispectrum analysis of electroencephalogram signals during waking and sleeping. Science 172: 401-402.

36. Granger CWJ (1969) Investigating causal relations by econometric models and cross-spectral methods. Econometrica 37: 424-438.

37. Schreiber T (2000) Measuring information transfer. Physical Review Letters 85: 461-464.

38. Mallat S (1998) A wavelet tour of signal processing. San Diego: Academic Press.

39. Torrence C, Compo G (1998) A Practical Guide to Wavelet Analysis. Bulletin of the American Meteorological Society 79: 61-78.

40. Grinsted A, Moore J, Jevrejeva S (2004) Application of the cross wavelet transform and wavelet coherence to geophysical time series. Nonlinear Processes in Geophysics 11: 561-566.

41. Cox D, P L (1966) The statistical analysis of series of events: Chapman and Hall.

42. Jarvis MR, Mitra PP (2001) Sampling properties of the spectrum and coherency of sequences of action potentials. Neural Comput 13: 717-749.

43. Berger H (1929) Uber das elektroenkephalogramm des menschen Archiv fur Psychiatrie und Nervenkrankheiten 87: 527-570.

44. Vanhatalo S, Palva JM, Holmes MD, Miller JW, Voipio J, et al. (2004) Infraslow oscillations modulate excitability and interictal epileptic activity in the human cortex during sleep. Proc Natl Acad Sci U S A 101: 5053-5057.

45. Bragin A, Engel J, Jr., Wilson CL, Fried I, Mathern GW (1999) Hippocampal and entorhinal cortex high-frequency oscillations (100--500 Hz) in human epileptic brain and in kainic acid--treated rats with chronic seizures. Epilepsia 40: 127-137.

46. Buzsaki G, Draguhn A (2004) Neuronal oscillations in cortical networks. Science 304: 1926-1929.

47. Buzsaki G (2006) Rhythms of the brain: Oxford University Press.

48. Bouyer JJ, Montaron MF, Rougeul A (1981) Fast fronto-parietal rhythms during combined focused attentive behaviour and immobility in cat: cortical and thalamic localizations. Electroencephalogr Clin Neurophysiol 51: 244-252.

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

50. Thut G, Nietzel A, Brandt SA, Pascual-Leone A (2006) Alpha-band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. J Neurosci 26: 9494-9502.

51. Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science 315: 1860-1862.

52. Capotosto P, Babiloni C, Romani GL, Corbetta M (2009) Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. J Neurosci 29: 5863-5872.

53. Rechtschaffen A, Kales A, editors (1968) A manual of standardized terminology, techniques and scoring system for sleep stages of human subjects. Washington, D.C.: Government Printing Office.

54. Gutfreund Y, yarom Y, Segev I (1995) Subthreshold oscillations and resonant frequency in guinea-pig cortical neurons: physiology and modelling. J Physiol 483 (Pt 3): 621-640.

55. Hutcheon B, Yarom Y (2000) Resonance, oscillation and the intrinsic frequency preferences of neurons. Trends Neurosci 23: 216-222.

56. Llinás R, Yarom Y (1986) Oscillatory properties of guinea-pig inferior olivary neurones and their pharmacological modulation: an in vitro study. J Physiol 376: 163-182.

57. Alonso A, Llinás RR (1989) Subthreshold Na+-dependent theta-like rhythmicity in stellate cells of entorhinal cortex layer II. Nature 342: 175-177.

58. Llinás RR, Grace AA, Yarom Y (1991) In vitro neurons in mammalian cortical layer 4 exhibit intrinsic oscillatory activity in the 10- to 50-Hz frequency range. Proc Natl Acad Sci U S A 88: 897-901.

59. Leung LW, Yim CY (1991) Intrinsic membrane potential oscillations in hippocampal neurons in vitro. Brain Res 553: 261-274.

60. Desmaisons D, Vincent JD, Lledo PM (1999) Control of action potential timing by intrinsic subthreshold oscillations in olfactory bulb output neurons. J Neurosci 19: 10727-10737.

61. Cobb SR, Buhl EH, Halasy K, Paulsen O, Somogyi P (1995) Synchronization of neuronal activity in hippocampus by individual GABAergic interneurons. Nature 378: 75-78.

62. Amitai Y (1994) Membrane potential oscillations underlying firing patterns in neocortical neurons. Neuroscience 63: 151-161.

63. Steriade M, Domich L, Oakson G, Deschenes M (1987) The deafferented reticular thalamic nucleus generates spindle rhythmicity. J Neurophysiol 57: 260-273.

64. Silva LR, Amitai Y, Connors BW (1991) Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. Science 251: 432-435.

65. Steriade M, Dossi RC, Nunez A (1991) Network modulation of a slow intrinsic oscillation of cat thalamocortical neurons implicated in sleep delta waves: cortically induced synchronization and brainstem cholinergic suppression. J Neurosci 11: 3200-3217.

66. Gray CM, McCormick DA (1996) Chattering cells: superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. Science 274: 109-113.

67. McCormick DA, Pape HC (1990) Properties of a hyperpolarization-activated cation current and its role in rhythmic oscillation in thalamic relay neurones. J Physiol 431: 291-318.

68. Marder E, Eisen JS (1984) Electrically coupled pacemaker neurons respond differently to same physiological inputs and neurotransmitters. J Neurophysiol 51: 1362-1374.

69. Smith JC, Ellenberger HH, Ballanyi K, Richter DW, Feldman JL (1991) Pre-Botzinger complex: a brainstem region that may generate respiratory rhythm in mammals. Science 254: 726-729.

70. Wang XJ (2010) Neurophysiological and computational principles of cortical rhythms in cognition. Physiol Rev 90: 1195-1268.

71. Gutkin BS, Ermentrout GB, Reyes AD (2005) Phase-response curves give the responses of neurons to transient inputs. J Neurophysiol 94: 1623-1635.

72. Marder E, Calabrese RL (1996) Principles of rhythmic motor pattern generation. Physiol Rev 76: 687-717.

73. Teramae JN, Tanaka D (2004) Robustness of the noise-induced phase synchronization in a general class of limit cycle oscillators. Phys Rev Lett 93: 204103.

74. Ermentrout GB, Galan RF, Urban NN (2008) Reliability, synchrony and noise. Trends Neurosci 31: 428-434.

75. Friston KJ (1997) Another neural code? Neuroimage 5: 213-220.

76. Shirvalkar PR, Rapp PR, Shapiro ML (2010) Bidirectional changes to hippocampal theta-gamma comodulation predict memory for recent spatial episodes. Proc Natl Acad Sci U S A 107: 7054-7059.

77. Palva JM, Palva S, Kaila K (2005) Phase synchrony among neuronal oscillations in the human cortex. J Neurosci 25: 3962-3972.

78. Palva JM, Monto S, Kulashekhar S, Palva S (2010) Neuronal synchrony reveals working memory networks and predicts individual memory capacity. Proc Natl Acad Sci U S A 107: 7580-7585.

79. Sauseng P, Klimesch W, Gruber WR, Birbaumer N (2008) Cross-frequency phase synchronization: a brain mechanism of memory matching and attention. Neuroimage 40: 308-317.

80. Bragin A, Jando G, Nadasdy Z, Hetke J, Wise K, et al. (1995) Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. J Neurosci 15: 47-60.

81. Buzsaki G, Buhl DL, Harris KD, Csicsvari J, Czeh B, et al. (2003) Hippocampal network patterns of activity in the mouse. Neuroscience 116: 201-211.

82. Lakatos P, Shah AS, Knuth KH, Ulbert I, Karmos G, et al. (2005) An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. J Neurophysiol 94: 1904-1911.

83. Tort AB, Kramer MA, Thorn C, Gibson DJ, Kubota Y, et al. (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.

84. Wulff P, Ponomarenko AA, Bartos M, Korotkova TM, Fuchs EC, et al. (2009) Hippocampal theta rhythm and its coupling with gamma oscillations require fast inhibition onto parvalbumin-positive interneurons. Proc Natl Acad Sci U S A 106: 3561-3566.

85. Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, et al. (2006) High gamma power is phase-locked to theta oscillations in human neocortex. Science 313: 1626-1628.

86. Sirota A, Montgomery S, Fujisawa S, Isomura Y, Zugaro M, et al. (2008) Entrainment of neocortical neurons and gamma oscillations by the hippocampal theta rhythm. Neuron 60: 683-697.

87. Siegel M, Warden MR, Miller EK (2009) Phase-dependent neuronal coding of objects in short-term memory. Proc Natl Acad Sci U S A 106: 21341-21346.

88. Cohen MX, Elger CE, Fell J (2009) Oscillatory activity and phase-amplitude coupling in the human medial frontal cortex during decision making. J Cogn Neurosci 21: 390-402.

89. Osipova D, Hermes D, Jensen O (2008) Gamma power is phase-locked to posterior alpha activity. PLoS One 3: e3990.

90. Cohen MX, Axmacher N, Lenartz D, Elger CE, Sturm V, et al. (2009) Good vibrations: cross-frequency coupling in the human nucleus accumbens during reward processing. J Cogn Neurosci 21: 875-889.

91. Monto S, Palva S, Voipio J, Palva JM (2008) Very slow EEG fluctuations predict the dynamics of stimulus detection and oscillation amplitudes in humans. J Neurosci 28: 8268-8272.

92. He BJ, Zempel JM, Snyder AZ, Raichle ME (2010) The temporal structures and functional significance of scale-free brain activity. Neuron 66: 353-369.

93. Singer W (1993) Synchronization of cortical activity and its putative role in information processing and learning. Annu Rev Physiol 55: 349-374.

94. Magee JC (2000) Dendritic integration of excitatory synaptic input. Nat Rev Neurosci 1: 181-190.

95. Azouz R, Gray CM (2003) Adaptive coincidence detection and dynamic gain control in visual cortical neurons in vivo. Neuron 37: 513-523.

96. Diesmann M, Gewaltig MO, Aertsen A (1999) Stable propagation of synchronous spiking in cortical neural networks. Nature 402: 529-533.

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

98. Lampl I, Reichova I, Ferster D (1999) Synchronous membrane potential fluctuations in neurons of the cat visual cortex. Neuron 22: 361-374.

99. Reynolds JH, Chelazzi L, Desimone R (1999) Competitive mechanisms subserve attention in macaque areas V2 and V4. J Neurosci 19: 1736-1753.

100. Mishra J, Martinez A, Sejnowski TJ, Hillyard SA (2007) Early cross-modal interactions in auditory and visual cortex underlie a sound-induced visual illusion. J Neurosci 27: 4120-4131.

101. Lakatos P, Chen CM, O'Connell MN, Mills A, Schroeder CE (2007) Neuronal oscillations and multisensory interaction in primary auditory cortex. Neuron 53: 279-292.

102. Fries P, Nikolic D, Singer W (2007) The gamma cycle. Trends Neurosci 30: 309-316.

103. Lisman JE, Idiart MA (1995) Storage of 7 +/- 2 short-term memories in oscillatory subcycles. Science 267: 1512-1515.

104. Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. Nature 376: 33-36.

105. 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.

106. 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.

107. Bi G, Poo M-M (1998) Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic sell type. J Neurosci 18: 10464-10472.

108. Markram H, Lubke J, Frotscher M, Sakmann B (1997) Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. Science 275: 213-215.

109. Larson J, Wong D, Lynch G (1986) Patterned stimulation at the theta frequency is optimal for the induction of hippocampal long-term potentiation. Brain Res 368: 347-350.

110. Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. Nature 417: 741-746.

111. 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.

112. Jensen O, Kaiser J, Lachaux JP (2007) Human gamma-frequency oscillations associated with attention and memory. Trends Neurosci 30: 317-324.

113. Nyhus E, Curran T (2010) Functional role of gamma and theta oscillations in episodic memory. Neurosci Biobehav Rev 34: 1023-1035.

114. Walker MP, Stickgold R (2004) Sleep-dependent learning and memory consolidation. Neuron 44: 121-133.

115. Marshall L, Helgadottir H, Molle M, Born J (2006) Boosting slow oscillations during sleep potentiates memory. Nature 444: 610-613.

116. Tort AB, Komorowski RW, Manns JR, Kopell NJ, Eichenbaum H (2009) Theta-gamma coupling increases during the learning of item-context associations. Proc Natl Acad Sci U S A 106: 20942-20947.

117. Cohen MX, Axmacher N, Lenartz D, Elger CE, Sturm V, et al. (2009) Nuclei accumbens phase synchrony predicts decision-making reversals following negative feedback. J Neurosci 29: 7591-7598.

118. Uhlhaas PJ, Singer W (2010) Abnormal neural oscillations and synchrony in schizophrenia. Nat Rev Neurosci 11: 100-113.

119. Jeong J (2004) EEG dynamics in patients with Alzheimer's disease. Clin Neurophysiol 115: 1490-1505.

120. Brown P (2003) Oscillatory nature of human basal ganglia activity: relationship to the pathophysiology of Parkinson's disease. Mov Disord 18: 357-363.

121. Garcia Dominguez L, Wennberg RA, Gaetz W, Cheyne D, Snead OC, 3rd, et al. (2005) Enhanced synchrony in epileptiform activity? Local versus distant phase synchronization in generalized seizures. J Neurosci 25: 8077-8084.

122. Uhlhaas PJ, Singer W (2007) What do disturbances in neural synchrony tell us about autism? Biol Psychiatry 62: 190-191.

123. Steriade M, Nunez A, Amzica F (1993) A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components. J Neurosci 13: 3252-3265.

124. Cowan R, Wilson C (1994) Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. J Neurophysiol 71: 17-32.

125. Achermann P, Borbely AA (1997) Low-frequency (< 1 Hz) oscillations in the human sleep electroencephalogram. Neuroscience 81: 213-222.

126. Cash SS, Halgren E, Dehghani N, Rossetti AO, Thesen T, et al. (2009) The human K-complex represents an isolated cortical down-state. Science 324: 1084-1087.

127. Doi A, Mizuno M, Katafuchi T, Furue H, Koga K, et al. (2007) Slow oscillation of membrane currents mediated by glutamatergic inputs of rat somatosensory cortical neurons: in vivo patch-clamp analysis. Eur J Neurosci 26: 2565-2575.

128. Steriade M, Timofeev I, Grenier F (2001) Natural waking and sleep states: a view from inside neocortical neurons J Neurophysiol 85: 1969-1985

129. Clement EA, Richard A, Thwaites M, Ailon J, Peters S, et al. (2008) Cyclic and sleep-like spontaneous alternations of brain state under urethane anaesthesia. PLoS One 3: e2004.

130. Steriade M (2000) Corticothalamic resonance, states of vigilance and mentation. Neuroscience 101: 243-276.

131. Cossart R, Aronov D, Yuste R (2003) Attractor dynamics of network UP states in the neocortex. Nature 423: 283-288.

132. Petersen C, Hahn T, Mehta M, Grinvald A, Sakmann B (2003) Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. PNAS 100: 13638-13643.

133. Volgushev M, Chauvette S, Mukovski M, Timofeev I (2006) Precise long-range synchronization of activity and silence in neocortical neurons during slow-wave oscillations [corrected]. J Neurosci 26: 5665-5672.

134. Mohajerani MH, McVea DA, Fingas M, Murphy TH (2010) Mirrored bilateral slow-wave cortical activity within local circuits revealed by fast bihemispheric voltage-sensitive dye imaging in anesthetized and awake mice. J Neurosci 30: 3745-3751.

135. Sanchez-Vives MV, McCormick DA (2000) Cellular and network mechanisms of rhythmic recurrent activity in neocortex. Nat Neurosci 3: 1027-1034.

136. Shu Y, Hausenstaub A, McCormick DA (2003) Turning on and off recurrent balanced cortical activity. Nature 423: 288-293.

137. Haider B, Duque A, Hasenstaub AR, McCormick DA (2006) Neocortical network activity in vivo is generated through a dynamic balance of excitation and inhibition. Journal of Neuroscience 26: 4535-4545.

138. Luczak A, Bartho P, Marguet SL, Buzsaki G, Harris KD (2007) Sequential structure of neocortical spontaneous activity in vivo. Proc Natl Acad Sci U S A 104: 347-352.

139. Massimini M, Huber R, Ferrarelli F, Hill S, Tononi G (2004) The sleep slow oscillation as a traveling wave. J Neurosci 24: 6862-6870.

140. Steriade M, Contreras D, Curro Dossi R, Nunez A (1993) The slow (< 1 Hz) oscillation in reticular thalamic and thalamocortical neurons: scenario of sleep rhythm generation in interacting thalamic and neocortical networks. J Neurosci 13: 3284-3299.

141. Contreras D, Steriade M (1995) Cellular basis of EEG slow rhythms: a study of dynamic corticothalamic relationships. J Neurosci 15: 604-622.

142. Ros H, Sachdev RN, Yu Y, Sestan N, McCormick DA (2009) Neocortical networks entrain neuronal circuits in cerebellar cortex. J Neurosci 29: 10309-10320.

143. Magill PJ, Bolam JP, Bevan MD (2000) Relationship of activity in the subthalamic nucleus-globus pallidus network to cortical electroencephalogram. J Neurosci 20: 820-833.

144. Stern EA, Jaeger D, Wilson CJ (1998) Membrane potential synchrony of simultaneously recorded striatal spiny neurons in vivo. Nature 394: 475-478.

145. Steriade M, Amzica F, Nunez A (1993) Cholinergic and noradrenergic modulation of the slow (approximately 0.3 Hz) oscillation in neocortical cells. J Neurophysiol 70: 1385-1400.

146. Mena-Segovia J, Sims HM, Magill PJ, Bolam JP (2008) Cholinergic brainstem neurons modulate cortical gamma activity during slow oscillations. J Physiol 586: 2947-2960.

147. Hahn T, Sakmann B, Mehta M (2006) Phase-locking of hippocampal interneurons' membrane potential to neocortical up-down states. Nature Neuroscience 9: 1359-1361.

148. Hahn T, Sakmann B, Mehta M (2007) Differential responses of hippocampal subfields to cortical up-down states. PNAS 104: 5169-5174.

149. Isomura Y, Sirota A, Ozen S, Montgomery S, Mizuseki K, et al. (2006) Integration and segregation of activity in entorhinal-hippocampal subregions by neocortical slow oscillations. Neuron 52: 871-882.

150. Steriade M, Nunez A, Amzica F (1993) Intracellular analysis of relations between the slow (< 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram. J Neurosci 13: 3266-3283.

151. Timofeev I, Steriade M (1996) Low-frequency rhythms in the thalamus of intact-cortex and decorticated cats. J Neurophysiol 76: 4152-4168.

152. Timofeev I, Grenier F, Bazhenov M, Sejnowski TJ, Steriade M (2000) Origin of slow cortical oscillations in deafferented cortical slabs. Cereb Cortex 10: 1185-1199.

153. Bazhenov M, Timofeev I, Steriade M, Sejnowski TJ (2002) Model of thalamocortical slow-wave sleep oscillations and transitions to activated States. J Neurosci 22: 8691-8704.

154. Holcman D, Tsodyks M (2006) The emergence of Up and Down states in cortical networks. PLoS Comput Biol 2: e23.

155. Compte A, Sanchez-Vives MV, McCormick DA, Wang XJ (2003) Cellular and network mechanisms of slow oscillatory activity (<1 Hz) and wave propagations in a cortical network model. J Neurophysiol 89: 2707-2725.

156. Parga N, Abbott LF (2007) Network model of spontaneous activity exhibiting synchronous transitions between up and down States. Front Neurosci 1: 57-66.

157. Le Bon-Jego M, Yuste R (2007) Persistently active, pacemaker-like neurons in neocortex. Front Neurosci 1: 123-129.

158. Crunelli V, Hughes SW (2010) The slow (<1 Hz) rhythm of non-REM sleep: a dialogue between three cardinal oscillators. Nat Neurosci 13: 9-17.

159. Hughes SW, Cope DW, Blethyn KL, Crunelli V (2002) Cellular mechanisms of the slow (<1 Hz) oscillation in thalamocortical neurons in vitro. Neuron 33: 947-958.

160. Blethyn KL, Hughes SW, Toth TI, Cope DW, Crunelli V (2006) Neuronal basis of the slow (<1 Hz) oscillation in neurons of the nucleus reticularis thalami in vitro. J Neurosci 26: 2474-2486.

161. Slezia A, Hangya B, Ulbert I, Acsady L (2011) Phase advancement and nucleus-specific timing of thalamocortical activity during slow cortical oscillation. The Journal of Neuroscience 31: 607-617.

162. Battaglia FP, Sutherland GR, McNaughton BL (2004) Hippocampal sharp wave bursts coincide with neocortical "up-state" transitions. Learn Mem 11: 697-704.

163. Gais S, Molle M, Helms K, Born J (2002) Learning-dependent increases in sleep spindle density. J Neurosci 22: 6830-6834.

164. Marr D (1971) Simple memory: a theory for archicortex. Philos Trans R Soc Lond B Biol Sci 262: 23-81.

165. Squire L (1992) Memory and the hippocampus: a synthesis from findings with rats, monkeys, and humans. Psychol Rev 99: 195-231.

166. McClelland J, McNaughton B, O'Reilly R (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102: 419-457.

167. Pavlides C, Winson J (1989) Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. J Neurosci 9: 2907-2918.

168. Nadasdy Z, Hirase H, Czurko A, Csicsvari J, Buzsaki G (1999) Replay and time compression of recurring spike sequences in the hippocampus. J Neurosci 19: 9497-9507.

169. Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. Neuron 36: 1183-1194.

170. Louie K, Wilson MA (2001) Temporally structured replay of awake hippocampal ensemble activity during rapid eye movement sleep. Neuron 29: 145-156.

171. Euston DR, Tatsuno M, McNaughton BL (2007) Fast-forward playback of recent memory sequences in prefrontal cortex during sleep. Science 318: 1147-1150.

172. Ji D, Wilson MA (2007) Coordinated memory replay in the visual cortex and hippocampus during sleep. Nat Neurosci 10: 100-107.

173. Scoville WB, Milner B (1957) Loss of recent memory after bilateral hippocampal lesions. J Neurol Neurosurg Psychiatry 20: 11-21.

174. Cohen NJ, Eichenbaum H (1993) Memory, amnesia, and the hippocampal system. Cambridge, MA: MIT Press.

175. O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. Brain Res 34: 171-175.

176. O'Keefe J, Nadal L (1978) The hippocampus as a cognitive map. Oxford: Clarendon Press.

177. Lorenté de Nó R (1934) Studies on the structure of the cerebral cortex. II. Continuation of the study of the ammonic system. Journal for Psychologie und Neurologiey 46: 113-177.

178. Amaral DG, Witter MP (1989) The three-dimensional organization of the hippocampal formation: a review of anatomical data. Neuroscience 31: 571-591.

179. Klausberger T, Somogyi P (2008) Neuronal diversity and temporal dynamics: the unity of hippocampal circuit operations. Science 321: 53-57.

180. Witter M, Groenewegen H, Silva HLd, Lohman A (1989) Functional organization of the extrinisic and intrinsic circuitry of the parahippocampal region. Prog Neurobiol 33: 161-254.

181. Witter MP (1993) Organization of the entorhinal-hippocampal system: a review of current anatomical data. Hippocampus 3 Spec No: 33-44.

182. Dolorfo CL, Amaral DG (1998) Entorhinal cortex of the rat: topographic organization of the cells of origin of the perforant path projection to the dentate gyrus. J Comp Neurol 398: 25-48.

183. Insausti R, Herrero MT, Witter MP (1997) Entorhinal cortex of the rat: cytoarchitectonic subdivisions and the origin and distribution of cortical efferents. Hippocampus 7: 146-183.

184. Burwell R, Amaral D (1998) Cortical afferents of the perirhinal, postrhinal, and entorhinal cortices of the rat. J Comp Neurol 398: 179-205.

185. Kerr K, Agster K, Furtak S, Burwell R (2007) Functional neuroanatomy of the parahippocampal region: the lateral and medial entorhinal areas. Hippocampus 17: 697-708.

186. O'Keefe J (1979) A review of the hippocampal place cells. Prog Neurobiol 13: 419-439.

187. Jung MW, McNaughton BL (1993) Spatial selectivity of unit activity in the hippocampal granular layer. Hippocampus 3: 165-182.

188. Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. Science 261: 1055-1058.

189. Jung MW, Wiener SI, McNaughton BL (1994) Comparison of spatial firing characteristics of units in dorsal and ventral hippocampus of the rat. J Neurosci 14: 7347-7356.

190. Ahmed O, Mehta M (2009) The hippocampal rate code: anatomy, physiology and theory. TINS 32: 329-338.

191. Muller R (1996) A quarter of a century of place cells. Neuron 17: 813-822.

192. Kubie JL, Muller RU (1991) Multiple representations in the hippocampus. Hippocampus 1: 240-242.

193. Morris RG, Garrud P, Rawlins JN, O'Keefe J (1982) Place navigation impaired in rats with hippocampal lesions. Nature 297: 681-683.

194. Mehta M, Barnes C, McNaughton B (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. PNAS USA 94: 8918-8921.

195. Mehta M, Quirk M, Wilson M (2000) Experience-dependent, asymmetric shape of hippocampal receptive fields. Neuron 25: 707-715.

196. McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. Exp Brain Res 52: 41-49.

197. Wiener SI, Paul CA, Eichenbaum H (1989) Spatial and behavioral correlates of hippocampal neuronal activity. J Neurosci 9: 2737-2763.

198. Sharp PE, Kubie JL, Muller RU (1990) Firing properties of hippocampal neurons in a visually symmetrical environment: contributions of multiple sensory cues and mnemonic processes. J Neurosci 10: 3093-3105.

199. Zhang K, Ginzburg I, McNaughton BL, Sejnowski TJ (1998) Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. J Neurophysiol 79: 1017-1044.

200. Muir GM, Bilkey DK (2003) Theta- and movement velocity-related firing of hippocampal neurons is disrupted by lesions centered on the perirhinal cortex. Hippocampus 13: 93-108.

201. Lu X, Bilkey DK (2010) The velocity-related firing property of hippocampal place cells is dependent on self-movement. Hippocampus 20: 573-583.

202. Ranck JB, Jr. (1973) Studies on single neurons in dorsal hippocampal formation and septum in unrestrained rats. I. Behavioral correlates and firing repertoires. Exp Neurol 41: 461-531.

203. Eichenbaum H, Kuperstein M, Fagan A, Nagode J (1987) Cue-sampling and goal-approach correlates of hippocampal unit activity in rats performing an odor-discrimination task. J Neurosci 7: 716-732.

204. Kobayashi T, Nishijo H, Fukuda M, Bures J, Ono T (1997) Task-dependent representations in rat hippocampal place neurons. J Neurophysiol 78: 597-613.

205. Wood ER, Dudchenko PA, Eichenbaum H (1999) The global record of memory in hippocampal neuronal activity. Nature 397: 613-616.

206. Holscher C, Jacob W, Mallot HA (2003) Reward modulates neuronal activity in the hippocampus of the rat. Behav Brain Res 142: 181-191.

207. Ho SA, Hori E, Kobayashi T, Umeno K, Tran AH, et al. (2008) Hippocampal place cell activity during chasing of a moving object associated with reward in rats. Neuroscience 157: 254-270.

208. Fyhn M, Molden S, Witter M, Moser E, Moser M-B (2004) Spatial representation in the entorhinal cortex. Science 305.

209. Sargolini F, Fyhn M, Hafting T, McNaughton B, Witter M, et al. (2006) Conjunctive representation of position, direction, and velocity in entorhinal cortex. Science 312: 758-762.

210. Hafting T, Fyhn M, Molden S, Moser M, Moser E (2005) Microstructure of a spatial map in the entorhinal cortex. Nature 436: 801-806.

211. Green JD, Arduini AA (1954) Hippocampal electrical activity in arousal. J Neurophysiol 17: 533-557.

212. Schmidt R, Diba K, Leibold C, Schmitz D, Buzsaki G, et al. (2009) Single-trial phase precession in the hippocampus. J Neurosci 29: 13232-13241.

213. Huxter J, Burgess N, O'Keefe J (2003) Independent rate and temporal coding in hippocampal pyramidal cells. Nature 425: 828-832.

214. Geisler C, Robbe D, Zugaro M, Sirota A, Buzsaki G (2007) Hippocampal place cell assemblies are speed-controlled oscillators. Proc Natl Acad Sci U S A 104: 8149-8154.

215. Huxter JR, Senior TJ, Allen K, Csicsvari J (2008) Theta phase-specific codes for two-dimensional position, trajectory and heading in the hippocampus. Nat Neurosci 11: 587-594.

216. Skaggs WE, McNaughton BL, Wilson MA, Barnes CA (1996) Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. Hippocampus 6: 149-172.

217. Blum KI, Abbott LF (1996) A model of spatial map formation in the hippocampus of the rat. Neural Comput 8: 85-93.

218. Dragoi G, Buzsaki G (2006) Temporal encoding of place sequences by hippocampal cell assemblies. Neuron 50: 145-157.

219. Hafting T, Fyhn M, Bonnevie T, Moser MB, Moser EI (2008) Hippocampus-independent phase precession in entorhinal grid cells. Nature 453: 1248-1252.

220. Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL (1996) Population dynamics and theta rhythm phase precession of hippocampal place cell firing: a spiking neuron model. Hippocampus 6: 271-280.

221. Jensen O, Lisman JE (1996) Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. Learn Mem 3: 279-287.

222. Lengyel M, Szatmary Z, Erdi P (2003) Dynamically detuned oscillations account for the coupled rate and temporal code of place cell firing. Hippocampus 13: 700-714.

223. Magee JC (2001) Dendritic mechanisms of phase precession in hippocampal CA1 pyramidal neurons. J Neurophysiol 86: 528-532.

224. Harvey CD, Collman F, Dombeck DA, Tank DW (2009) Intracellular dynamics of hippocampal place cells during virtual navigation. Nature 461: 941-946.

225. Anderson J, Lampl I, Reichova I, Carandini M, Ferster D (2000) Stimulus dependence of two-state fluctuations of membrane potential in cat visual cortex. Nat Neurosci 3: 617-621.

226. Haider B, Duque A, Hasenstaub AR, Yu Y, McCormick DA (2007) Enhancement of visual responsiveness by spontaneous local network activity in vivo. J Neurophysiol 97: 4186-4202.

227. Curto C, Sakata S, Marguet S, Itskov V, Harris KD (2009) A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex. J Neurosci 29: 10600-10612.

228. Wolansky T, Clement E, Peters S, Palczak M, Dickson C (2006) Hippocampal slow oscillation: a novel EEG state and its coordination with ongoing neocortical activity. J Neurosci 26: 6213-6229.

229. Molle M, Marshall L, Gais S, Born J (2002) Grouping of spindle activity during slow oscillations in human non-rapid eye movement sleep. J Neurosci 22: 10941-10947.

230. Molle M, Yeshenko O, Marshall L, Sara SJ, Born J (2006) Hippocampal sharp wave-ripples linked to slow oscillations in rat slow-wave sleep. J Neurophysiol 96: 62-70.

231. Chen Z, Vijayan S, Barbieri R, Wilson MA, Brown EN (2009) Discrete- and continuous-time probabilistic models and algorithms for inferring neuronal UP and DOWN states. Neural Comput 21: 1797-1862.

232. Tokdar S, Xi P, Kelly RC, Kass RE (2010) Detection of bursts in extracellular spike trains using hidden semi-Markov point process models. J Comput Neurosci 29: 203-212.

233. Sanchez-Vives MV, Mattia M, Compte A, Perez-Zabalza M, Winograd M, et al. (2010) Inhibitory modulation of cortical up states. J Neurophysiol 104: 1314-1324.

234. Mukovski M, Chauvette S, Timofeev I, Volgushev M (2007) Detection of active and silent states in neocortical neurons from the field potential signal during slow-wave sleep. Cereb Cortex 17: 400-414.

235. Saleem AB, Chadderton P, Apergis-Schoute J, Harris KD, Schultz SR (2010) Methods for predicting cortical UP and DOWN states from the phase of deep layer local field potentials. J Comput Neurosci 29: 49-62.

236. Seamari Y, Narvaez JA, Vico FJ, Lobo D, Sanchez-Vives MV (2007) Robust off- and online separation of intracellularly recorded up and down cortical states. PLoS One 2: e888.

237. Gerkin RC, Clem RL, Shruti S, Kass RE, Barth AL (2010) Cortical up state activity is enhanced after seizures: a quantitative analysis. J Clin Neurophysiol 27: 425-432.

238. Margrie T, Brecht M, Sakmann B (2002) In vivo, low-resistance, whole-cell recordings from neurons in the anaesthetized and awake mammalian brain. Pflügers Archiv European Journal of Physiology 444: 491-498.

239. Baum L (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. Annals of Mathematical Statistics 41: 164-171.

240. Rabiner LR (1989) A Tutorial on Hidden Markov-Models and Selected Applications in Speech Recognition. Proceedings of the IEEE 77: 257-286.

241. Bishop CM (2006) Pattern Recognition and Machine Learning: Springer.

242. Ferguson JD. Variable Duration Models for Speech; 1980 October; Institute for Defense Analysis, Princeton, NJ. pp. 143-179.

243. Levinson S (1986) Continuosly variable duration hidden Markov models for automatic speech recognition. Computer Speech and Language 1: 29-45.

244. Yu S-Z (2010) Hidden semi-Markov models. Artificial Intelligence 174: 215-243.

245. Yu S-Z, Kobayashi H (2003) An efficient forward-backward algorithm for an explicit-duration hidden Markov model. Signal Processing Letters, IEEE 10: 11-14.

246. Yu S-Z, Kobayashi H (2006) Practical Implementation of an Efficient Forward-Backward Algorithm for an Explicit-Duration Hidden Markov Model. Ieee Transactions on Signal Processing 54: 1947-1951.

247. Mitchell C, Jamieson L. Modeling duration in a hidden Markov model with the exponential family; 1993 April; Minneapolis. pp. 11331-11334.

248. Mitchell C, Harper M, Jamieson L (1995) On the Complexity of Explicit Duration Hmms. Ieee Transactions on Speech and Audio Processing 3: 213-217.

249. Datta R, Hu J, Ray B. On efficient Viterbi decoding for hidden semi-Markov models; 2008 December. pp. 1-4.

250. Kasanetz F, Riquelme L, Murer M (2002) Disruption of the two-state membrane potential of striatal neurons during cortical desynchronisation in anaesthetized rats. J Physiology 532: 577-589.

251. Terrell G (1990) The maximal smoothing principle in density estimation. J Amer Statist Assoc 85: 470-477.

252. Choi S, Wette R (1969) Maximum Likelihood Estimation of the Parameters of the Gamma Distribution and Their Bias. Technometrics 11: 683-690.

253. Penny W (2001) KL-Divergences of Normal, Gamma, Direchlet and Wishart densities. Wellcome Department of Cognitive Neurology, University College London.

254. Bokil H, Andrews P, Kulkarni JE, Mehta S, Mitra PP (2010) Chronux: a platform for analyzing neural signals. J Neurosci Methods 192: 146-151.

255. Ephraim Y, Malah D, Juang B-H (1989) On the application of hidden Markov models for enhancing noisy speech. Ieee Transactions on Acoustics Speech and Signal Processing 37: 1846-1856.

256. Remondes M, Schuman E (2004) Role for a cortical input to hippocampal area CA1 in the consolidation of a long-term memory. Nature 431: 699-703.

257. 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.

258. Hargreaves E, Rao G, Lee I, Knierim J (2005) Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. Science 308.

259. Bokil H, Purpura K, Schoffelen JM, Thomson D, Mitra P (2007) Comparing spectra and coherences for groups of unequal size. J Neurosci Methods 159: 337-345.

260. Egorov A, Hamam B, Fransen E, Hasselmo M, Alonso A (2002) Graded persistent activity of entorhinal cortex neurons. Nature 420: 173-178.

261. Tahvildari B, Fransen E, Alonso A, Hasselmo M (2007) Switching between "on" and "off" states of persistent activity in lateral entorhinal layer III neurons. Hippocampus 17: 257-263.

262. Yoshida M, Fransen E, Hasselmo M (2008) mGluR-dependent persistent firing in entorhinal cortex layer III neurons. European Journal of Neuroscience 28: 1116-1126.

263. Fuster J (1973) Unit activity in prefrontal cortex during delayed-response performance: neuronal correlates of transient memory. J Neurophysiol 36: 61-78.

264. Young B, Otto T, Fox G, Eichenbaum H (1997) Memory representation within the parahippocampal region. J Neurosci 17: 5183-5195.

265. Suzuki W, Miller E, Desimone R (1997) Object and place memory in the macaque entorhinal cortex. J Neurophysiol 78: 1062-1081.

266. Mann E, Kohl M, Paulsen O (2009) Distinct roles of $GABA_A$ and $GABA_B$ receptors in balancing and terminating persistent cortical activity. J Neurosci 29: 7513-7518.

267. Major G, Tank D (2004) Persistent neural activity: prevalence and mechanisms. Curr Opin Neurobiol 14: 675-684.

268. Fransen E, Tahvildari B, Egorov A, Hasselmo M, Alonso A (2006) Mechanism of graded persistent cellular activity of entorhinal cortex layer V neurons. Neuron 49: 735-746.

269. Quilichini P, Sirota A, Buzsaki G (2010) Intrinsic circuit organization and theta-gamma oscillation dynamics in the entorhinal cortex of the rat. J Neurosci 30: 11128-11142.

270. Steward O, Scoville S (1976) Cells of origin of entorhinal cortical afferents to the hippocampus and fascia dentata of the rat. J Comp Neurol 169: 347-370.

271. Mehta M (2007) Cortico-Hippocampal interaction during up-down states and memory consolidation. Nature Neuroscience 10: 13-15.

272. Alonso A, Klink R (1993) Differential electroresponsiveness of stellate and pyramidal-like cells of medial entorhinal cortex layer II. J Neurophysiol 70: 128-143.

273. Dickson C, Mena A, Alonso A (1997) Electroresponsiveness of medial entorhinal cortex layer III neurons *in vitro*. Neuroscience 81: 937-950.

274. Dickson CT, Magistretti J, Shalinsky MH, Fransen E, Hasselmo ME, et al. (2000) Properties and role of I(h) in the pacing of subthreshold oscillations in entorhinal cortex layer II neurons. J Neurophysiol 83: 2562-2579.

275. Fransen E, Alonso AA, Dickson CT, Magistretti J, Hasselmo ME (2004) Ionic mechanisms in the generation of subthreshold oscillations and action potential clustering in entorhinal layer II stellate neurons. Hippocampus 14: 368-384.

276. Giocomo LM, Zilli EA, Fransen E, Hasselmo ME (2007) Temporal frequency of subthreshold oscillations scales with entorhinal grid cell field spacing. Science 315: 1719-1722.

277. Deshmukh SS, Yoganarasimha D, Voicu H, Knierim JJ (2010) Theta modulation in the medial and the lateral entorhinal cortices. J Neurophysiol 104: 994-1006.

278. Mitra PP, Pesaran B (1999) Analysis of dynamic brain imaging data. Biophys J 76: 691-708.

279. Kramis R, Vanderwolf CH, Bland BH (1975) Two types of hippocampal rhythmical slow activity in both the rabbit and the rat: relations to behavior and effects of atropine, diethyl ether, urethane, and pentobarbital. Exp Neurol 49: 58-85.

280. Tahvildari B, Alonso A (2005) Morphological and electrophysiological properties of lateral entorhinal cortex layers II and III principal neurons. J Comp Neurol 491: 123-140.

281. Duque A, Balatoni B, Detari L, Zaborszky L (2000) EEG correlation of the discharge properties of identified neurons in the basal forebrain. J Neurophysiol 84: 1627-1635.

282. Van der Werf YD, Witter MP, Groenewegen HJ (2002) The intralaminar and midline nuclei of the thalamus. Anatomical and functional evidence for participation in processes of arousal and awareness. Brain Res Brain Res Rev 39: 107-140.

283. Dickson CT, Alonso A (1997) Muscarinic induction of synchronous population activity in the entorhinal cortex. J Neurosci 17: 6729-6744.

284. Reisel D, Bannerman DM, Schmitt WB, Deacon RM, Flint J, et al. (2002) Spatial memory dissociations in mice lacking GluR1. Nat Neurosci 5: 868-873.

285. Zamanillo D, Sprengel R, Hvalby O, Jensen V, Burnashev N, et al. (1999) Importance of AMPA receptors for hippocampal synaptic plasticity but not for spatial learning. Science 284: 1805-1811.

286. Hoffman DA, Sprengel R, Sakmann B (2002) Molecular dissection of hippocampal theta-burst pairing potentiation. Proc Natl Acad Sci U S A 99: 7740-7745.

287. Frey MC, Sprengel R, Nevian T (2009) Activity pattern-dependent long-term potentiation in neocortex and hippocampus of GluA1 (GluR-A) subunit-deficient mice. J Neurosci 29: 5587-5596.

288. Andrasfalvy BK, Smith MA, Borchardt T, Sprengel R, Magee JC (2003) Impaired regulation of synaptic strength in hippocampal neurons from GluR1-deficient mice. J Physiol 552: 35-45.

289. Schmitt WB, Deacon RM, Seeburg PH, Rawlins JN, Bannerman DM (2003) A within-subjects, within-task demonstration of intact spatial reference memory and impaired spatial working memory in glutamate receptor-A-deficient mice. J Neurosci 23: 3953-3959.

290. McNaughton BL, O'Keefe J, Barnes CA (1983) The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. J Neurosci Methods 8: 391-397.

291. O'Keefe J (1976) Place units in the hippocampus of the freely moving rat. Exp Neurol 51: 78-109.

292. O'Keefe J, Conway DH (1978) Hippocampal place units in the freely moving rat: why they fire where they fire. Exp Brain Res 31: 573-590.

293. Buzsaki G, Horvath Z, Urioste R, Hetke J, Wise K (1992) High-frequency network oscillation in the hippocampus. Science 256: 1025-1027.

294. Buzsaki G, Czopf J, Kondakor I, Kellenyi L (1986) Laminar distribution of hippocampal rhythmic slow activity (RSA) in the behaving rat: current-source density analysis, effects of urethane and atropine. Brain Res 365: 125-137.

295. Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish AD (2005) Quantitative measures of cluster quality for use in extracellular recordings. Neuroscience 131: 1-11.

296. Fox SE, Ranck JB, Jr. (1981) Electrophysiological characteristics of hippocampal complex-spike cells and theta cells. Exp Brain Res 41: 399-410.

297. Buzsaki G, Leung LW, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. Brain Res 287: 139-171.

298. Lee I, Rao G, Knierim JJ (2004) A double dissociation between hippocampal subfields: differential time course of CA3 and CA1 place cells for processing changed environments. Neuron 42: 803-815.

299. Markus EJ, Barnes CA, McNaughton BL, Gladden VL, Skaggs WE (1994) Spatial information content and reliability of hippocampal CA1 neurons: effects of visual input. Hippocampus 4: 410-421.

300. Olypher AV, Lansky P, Muller RU, Fenton AA (2003) Quantifying location-specific information in the discharge of rat hippocampal place cells. J Neurosci Methods 127: 123-135.

301. Muller RU, Kubie JL (1989) The firing of hippocampal place cells predicts the future position of freely moving rats. J Neurosci 9: 4101-4110.

302. Battaglia FP, Sutherland GR, McNaughton BL (2004) Local sensory cues and place cell directionality: additional evidence of prospective coding in the hippocampus. J Neurosci 24: 4541-4550.

303. Hollup SA, Molden S, Donnett JG, Moser MB, Moser EI (2001) Accumulation of hippocampal place fields at the goal location in an annular watermaze task. J Neurosci 21: 1635-1644.

304. Rotenberg A, Mayford M, Hawkins RD, Kandel ER, Muller RU (1996) Mice expressing activated CaMKII lack low frequency LTP and do not form stable place cells in the CA1 region of the hippocampus. Cell 87: 1351-1361.

305. Kentros C, Hargreaves E, Hawkins RD, Kandel ER, Shapiro M, et al. (1998) Abolition of long-term stability of new hippocampal place cell maps by NMDA receptor blockade. Science 280: 2121-2126.

306. Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. Proc Natl Acad Sci U S A 94: 8918-8921.

307. Mehta M, McNaughton B (1997) Expansion and shift of hippocampal place fields: evidence for synaptic potentiation during behavior. In: Bower J, editor. Computational Neuroscience: Trends in Research. New York: Plenum Press. pp. 741-745.

308. Frank LM, Stanley GB, Brown EN (2004) Hippocampal plasticity across multiple days of exposure to novel environments. J Neurosci 24: 7681-7689.

309. Korotkova T, Fuchs EC, Ponomarenko A, von Engelhardt J, Monyer H (2010) NMDA receptor ablation on parvalbumin-positive interneurons impairs hippocampal synchrony, spatial representations, and working memory. Neuron 68: 557-569.

310. Malinow R, Malenka RC (2002) AMPA receptor trafficking and synaptic plasticity. Annu Rev Neurosci 25: 103-126.

311. Empson RM, Heinemann U (1995) The perforant path projection to hippocampal area CA1 in the rat hippocampal-entorhinal cortex combined slice. J Physiol 484 (Pt 3): 707-720.

312. Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. Neuron 25: 707-715.

313. Fuchs EC, Zivkovic AR, Cunningham MO, Middleton S, Lebeau FE, et al. (2007) Recruitment of parvalbumin-positive interneurons determines hippocampal function and associated behavior. Neuron 53: 591-604.

314. Pellegrini-Giampietro DE, Bennett MV, Zukin RS (1991) Differential expression of three glutamate receptor genes in developing rat brain: an in situ hybridization study. Proc Natl Acad Sci U S A 88: 4157-4161.

315. Jensen V, Kaiser KM, Borchardt T, Adelmann G, Rozov A, et al. (2003) A juvenile form of postsynaptic hippocampal long-term potentiation in mice deficient for the AMPA receptor subunit GluR-A. J Physiol 553: 843-856.

316. Mack V, Burnashev N, Kaiser KM, Rozov A, Jensen V, et al. (2001) Conditional restoration of hippocampal synaptic potentiation in Glur-A-deficient mice. Science 292: 2501-2504.

317. Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA (2001) NMDA receptor antagonism blocks experience-dependent expansion of hippocampal "place fields". Neuron 31: 631-638.

318. McHugh TJ, Blum KI, Tsien JZ, Tonegawa S, Wilson MA (1996) Impaired hippocampal representation of space in CA1-specific NMDAR1 knockout mice. Cell 87: 1339-1349.

319. Varvel SA, Lichtman AH (2002) Evaluation of CB1 receptor knockout mice in the Morris water maze. J Pharmacol Exp Ther 301: 915-924.

320. Robbe D, Buzsaki G (2009) Alteration of theta timescale dynamics of hippocampal place cells by a cannabinoid is associated with memory impairment. J Neurosci 29: 12597-12605.

321. Lenck-Santini PP, Holmes GL (2008) Altered phase precession and compression of temporal sequences by place cells in epileptic rats. J Neurosci 28: 5053-5062.

322. Berens P (2009) CircStat: A Matlab toolbox for circular statistics. Journal of Statistical Software 31.

323. Ego-Stengel V, Wilson MA (2007) Spatial selectivity and theta phase precession in CA1 interneurons. Hippocampus 17: 161-174.

324. Nishino H, Ono T, Muramoto K, Fukuda M, Sasaki K (1987) Neuronal activity in the ventral tegmental area (VTA) during motivated bar press feeding in the monkey. Brain Res 413: 302-313.

325. Schultz W, Romo R (1990) Dopamine neurons of the monkey midbrain: contingencies of responses to stimuli eliciting immediate behavioral reactions. J Neurophysiol 63: 607-624.

326. Ljungberg T, Apicella P, Schultz W (1992) Responses of monkey dopamine neurons during learning of behavioral reactions. J Neurophysiol 67: 145-163.

327. Richardson NR, Gratton A (1996) Behavior-relevant changes in nucleus accumbens dopamine transmission elicited by food reinforcement: an electrochemical study in rat. J Neurosci 16: 8160-8169.

328. Knutson B, Fong GW, Adams CM, Varner JL, Hommer D (2001) Dissociation of reward anticipation and outcome with event-related fMRI. Neuroreport 12: 3683-3687.

329. O'Doherty JP, Deichmann R, Critchley HD, Dolan RJ (2002) Neural responses during anticipation of a primary taste reward. Neuron 33: 815-826.

330. Schott BH, Minuzzi L, Krebs RM, Elmenhorst D, Lang M, et al. (2008) Mesolimbic functional magnetic resonance imaging activations during reward anticipation correlate with reward-related ventral striatal dopamine release. J Neurosci 28: 14311-14319.

331. Fiorillo CD, Tobler PN, Schultz W (2003) Discrete coding of reward probability and uncertainty by dopamine neurons. Science 299: 1898-1902.

332. Phillips PE, Stuber GD, Heien ML, Wightman RM, Carelli RM (2003) Subsecond dopamine release promotes cocaine seeking. Nature 422: 614-618.

333. Stuber GD, Wightman RM, Carelli RM (2005) Extinction of cocaine self-administration reveals functionally and temporally distinct dopaminergic signals in the nucleus accumbens. Neuron 46: 661-669.

334. Frey U, Matthies H, Reymann KG (1991) The effect of dopaminergic D1 receptor blockade during tetanization on the expression of long-term potentiation in the rat CA1 region in vitro. Neurosci Lett 129: 111-114.

335. Otmakhova NA, Lisman JE (1996) D1/D5 dopamine receptor activation increases the magnitude of early long-term potentiation at CA1 hippocampal synapses. J Neurosci 16: 7478-7486.

336. Lemon N, Manahan-Vaughan D (2006) Dopamine D1/D5 receptors gate the acquisition of novel information through hippocampal long-term potentiation and long-term depression. J Neurosci 26: 7723-7729.

337. Zhang JC, Lau PM, Bi GQ (2009) Gain in sensitivity and loss in temporal contrast of STDP by dopaminergic modulation at hippocampal synapses. Proc Natl Acad Sci U S A 106: 13028-13033.

338. Losonczy A, Zemelman BV, Vaziri A, Magee JC (2010) Network mechanisms of theta related neuronal activity in hippocampal CA1 pyramidal neurons. Nat Neurosci 13: 967-972.

339. Lichtman AH, Dimen KR, Martin BR (1995) Systemic or intrahippocampal cannabinoid administration impairs spatial memory in rats. Psychopharmacology (Berl) 119: 282-290.

340. Robbe D, Montgomery SM, Thome A, Rueda-Orozco PE, McNaughton BL, et al. (2006) Cannabinoids reveal importance of spike timing coordination in hippocampal function. Nat Neurosci 9: 1526-1533.

341. Benchenane K, Peyrache A, Khamassi M, Tierney PL, Gioanni Y, et al. (2010) Coherent theta oscillations and reorganization of spike timing in the hippocampal- prefrontal network upon learning. Neuron 66: 921-936.

342. Miller JD, Sanghera MK, German DC (1981) Mesencephalic dopaminergic unit activity in the behaviorally conditioned rat. Life Sci 29: 1255-1263.

343. Kiyatkin EA, Gratton A (1994) Electrochemical monitoring of extracellular dopamine in nucleus accumbens of rats lever-pressing for food. Brain Res 652: 225-234.

344. Breese CR, Hampson RE, Deadwyler SA (1989) Hippocampal place cells: stereotypy and plasticity. J Neurosci 9: 1097-1111.

345. Speakman A, O'Keefe J (1990) Hippocampal Complex Spike Cells do not Change Their Place Fields if the Goal is Moved Within a Cue Controlled Environment. Eur J Neurosci 2: 544-555.

346. Hok V, Lenck-Santini PP, Roux S, Save E, Muller RU, et al. (2007) Goal-related activity in hippocampal place cells. J Neurosci 27: 472-482.

347. Cagniard B, Balsam PD, Brunner D, Zhuang X (2006) Mice with chronically elevated dopamine exhibit enhanced motivation, but not learning, for a food reward. Neuropsychopharmacology 31: 1362-1370.

348. Puryear CB, Kim MJ, Mizumori SJ (2010) Conjunctive encoding of movement and reward by ventral tegmental area neurons in the freely navigating rodent. Behav Neurosci 124: 234-247.

349. Mehta MR (2001) Neuronal dynamics of predictive coding. Neuroscientist 7: 490-495.

350. Siapas AG, Lubenov EV, Wilson MA (2005) Prefrontal phase locking to hippocampal theta oscillations. Neuron 46: 141-151.

351. Raghavachari S, Kahana MJ, Rizzuto DS, Caplan JB, Kirschen MP, et al. (2001) Gating of human theta oscillations by a working memory task. J Neurosci 21: 3175-3183.

352. Buzsaki G (2002) Theta oscillations in the hippocampus. Neuron 33: 325-340.

353. Lee H, Simpson GV, Logothetis NK, Rainer G (2005) Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. Neuron 45: 147-156.

354. Steriade M (2003) The corticothalamic system in sleep. Front Biosci 8: d878-899.

355. Mokeichev A, Okun M, Barak O, Katz Y, Ben-Shahar O, et al. (2007) Stochastic emergence of repeating cortical motifs in spontaneous membrane potential fluctuations in vivo. Neuron 53: 413-425.

356. Gnatkovsky V, de Curtis M (2006) Hippocampus-mediated activation of superficial and deep layer neurons in the medial entorhinal cortex of the isolated guinea pig brain. J Neurosci 26: 873-881.

357. Brun VH, Otnass MK, Molden S, Steffenach HA, Witter MP, et al. (2002) Place cells and place recognition maintained by direct entorhinal-hippocampal circuitry. Science 296: 2243-2246.

358. Nakashiba T, Young JZ, McHugh TJ, Buhl DL, Tonegawa S (2008) Transgenic inhibition of synaptic transmission reveals role of CA3 output in hippocampal learning. Science 319: 1260-1264.

359. Brun VH, Leutgeb S, Wu HQ, Schwarcz R, Witter MP, et al. (2008) Impaired spatial representation in CA1 after lesion of direct input from entorhinal cortex. Neuron 57: 290-302.

360. Thompson LT, Best PJ (1989) Place cells and silent cells in the hippocampus of freely-behaving rats. J Neurosci 9: 2382-2390.

361. Sanderson DJ, Good MA, Seeburg PH, Sprengel R, Rawlins JN, et al. (2008) The role of the GluR-A (GluR1) AMPA receptor subunit in learning and memory. Prog Brain Res 169: 159-178.

362. Whittington MA, Traub RD, Jefferys JG (1995) Synchronized oscillations in interneuron networks driven by metabotropic glutamate receptor activation. Nature 373: 612-615.

363. Leung LS (1982) Nonlinear feedback model of neuronal populations in hippocampal CAl region. J Neurophysiol 47: 845-868.

364. Phillips PE, Robinson DL, Stuber GD, Carelli RM, Wightman RM (2003) Real-time measurements of phasic changes in extracellular dopamine concentration in freely moving rats by fast-scan cyclic voltammetry. Methods Mol Med 79: 443-464.