Published in final edited form as:

J Comput Neurosci. 2010 December ; 29(3): 581–597. doi:10.1007/s10827-010-0221-z.

Predicting stimulus-locked single unit spiking from cortical local field potentials

Edgar E. Galindo-Leon and

Department of Biology, Emory University, 1510 Clifton Road NE, Atlanta, GA 30322, USA

Robert C. Liu

Department of Biology, Emory University, 1510 Clifton Road NE, Atlanta, GA 30322, USA, Center for Behavioral Neuroscience, Georgia State University, PO Box 3966, Atlanta, GA 30302, USA

Edgar E. Galindo-Leon: egalind@emory.edu

Abstract

The rapidly increasing use of the local field potential (LFP) has motivated research to better understand its relation to the gold standard of neural activity, single unit (SU) spiking. We addressed this in an *in vivo*, awake, restrained mouse auditory cortical electrophysiology preparation by asking whether the LFP could actually be used to predict stimulus-evoked SU spiking. Implementing a Bayesian algorithm to predict the likelihood of spiking on a trial by trial basis from different representations of the despiked LFP signal, we were able to predict, with high quality and fine temporal resolution (2 ms), the time course of a SU's excitatory or inhibitory firing rate response to natural species-specific vocalizations. Our best predictions were achieved by representing the LFP by its wide-band Hilbert phase signal, and approximating the statistical structure of this signal at different time points as independent. Our results show that each SU's action potential has a unique relationship with the LFP that can be reliably used to predict the occurrence of spikes. This “signature” interaction can reflect both pre- and post-spike neural activity that is intrinsic to the local circuit rather than just dictated by the stimulus. Finally, the time course of this “signature” may be most faithful when the full bandwidth of the LFP, rather than specific narrow-band components, is used for representation.

Keywords

LFP; Spike prediction; Auditory cortex; Gamma band; Theta band; Beta band; Oscillation; Bayesian algorithm; A1; Evoked potentials; Electroencephalography; EEG; Hilbert transform; Single cortical cells; Phase; Despiking

1 Introduction

As this special volume makes clear, interest in the LFP as a measure of brain activity has recently exploded. Features of this low frequency extracellular potential have been shown in different brain regions to be correlated to sensory stimuli or motor outputs, including the auditory cortex (Kayser et al. 2007; Norena and Eggermont 2002), primary visual cortex (Belitski et al. 2008), visual area MT (Liu and Newsome 2006), somatosensory cortex (Ray et al. 2008), barrel cortex (Haslinger et al.), posterior parietal cortex (Asher et al. 2007), and

Correspondence to: Edgar E. Galindo-Leon, egalind@emory.edu.

Electronic supplementary material The online version of this article (doi:10.1007/s10827-010-0221-z) contains supplementary material, which is available to authorized users.

lateral intraparietal cortex (Pesaran et al. 2002). Activity within specific frequency bands of the LFP signal has been proposed to be important in controlling the timing of neural firing (Buzsaki and Draguhn 2004), and has been associated with higher-order processes like binding (Singer and Gray 1995), attention (Fries et al. 2001; Jensen and Colgin 2007), and memory (Lee et al. 2005; Mormann et al. 2005; Osipova et al. 2006; Pesaran et al. 2002). Finally, correlations between these frequency bands may also work as a coding scheme, such as during the retrieval of declarative memories (Canolty et al. 2006; Chrobak and Buzsaki 1998; Lakatos et al. 2005; Mormann et al. 2005).

Despite this rich phenomenology about the LFP and its more macroscopic cousin, the electroencephalogram (EEG), our understanding about its biophysical origin in the activity of individual neurons is still emerging. There is a general belief that the LFP reflects relatively large-scale, synchronized, slow currents associated with synaptic potentials, voltage-gated membrane potentials and spike afterpotentials (Logothetis 2003). As such, it is often compared to multiunit (MU) spiking, a different measure of local population activity that reflects the collective output of neurons within up to $\sim 300 \mu\text{m}$ of the electrode (Henze et al. 2000; Rasch et al. 2008). Although correlations have been found between the two (Belitski et al. 2008; Norena and Eggermont 2002; Rasch et al. 2008), such studies cannot reveal how membrane currents from an individual neuron might contribute to the LFP. MU activity is a weighted sum of contributions from many neurons, some of which may have overlapping spike waveforms. Despite this, correlating the LFP to the MU rather than SU activity has seemed reasonable because the relation of the LFP to any one neuron has been assumed to be weak, and the timing of any one spike to be irrelevant (Rasch et al. 2008). Yet recent studies suggest that the activity of an individual neuron may be more strongly reflected in the LFP than previously appreciated. For example, inducing bursting *in vivo* in a cortical neuron can change the state of the LFP recorded millimeters away (Li et al. 2009), presumably through its downstream connections. Moreover, the LFP around a SU inhibited by a stimulus can exhibit systematic differences compared to that around an excited SU (Galindo-Leon et al. 2009).

These conflicting perspectives led us to investigate the SU-LFP relation in more depth using a dataset in which well-isolated cortical SUs were recorded in an awake mouse's auditory cortex by means of a high-impedance electrode, and LFPs were low-pass filtered and recorded from the same electrode. Instead of examining correlations, we asked whether the relation between the two forms of neural activity could allow us to actually *predict* SU spiking from the LFP. To our knowledge, such spike prediction has only been addressed by (Rasch et al. 2008) for MU rather than SU spikes. Here, we employed a Bayesian algorithm to predict a SU's sound stimulus-evoked average firing rate with millisecond (ms) time resolution. We investigated this prediction using the wide-band LFP (4–100 Hz) as well as the filtered LFP in different frequency domains (θ , β and γ bands). Our results generally demonstrated surprisingly good predictions for a majority of neurons. We hypothesized that the wide-band representation would provide insights about the spike-LFP interaction that are not observable at narrow frequency bands, as suggested by other recent work (Manning et al. 2009). Indeed, our best predictions were achieved by representing the LFP as a wide-band phase signal. Hence, a given SU's action potentials were predictable based on its “signature” relationship with the LFP's time-varying fluctuations, which were presumably more temporally accurate when represented by the full spectral band. Interestingly though, combining these bands in an *independent* fashion to predict spiking did not do as well as the wide-band LFP, consistent with the additional notion that correlations between bands are important for spiking (Lakatos et al. 2005; Rasch et al. 2008).

2 Methods

2.1 Experimental procedures

The Emory University Institutional Animal Care and Use Committee approved all procedures. Electrophysiology experiments were carried out on 14 awake, head-restrained female CBA/CaJ mice, all between 14 and 24 weeks old at the time of initial head-implant surgery. A detailed description of the surgical procedure, experimental setup and acoustic stimulation can be found in (Galindo-Leon et al. 2009). Briefly, we performed recordings on each animal over a span of a week. Stimuli were generated using Tucker-Davis Technologies (TDT, Alachua, FL, USA) System 3 Gigabit hardware and software, and presented through the Brainware application. Thirty six pup and adult calls drawn from a large library of natural ultrasonic CBA/CaJ vocalizations (Liu et al. 2003) were played back along with a blank stimulus, up to 50 times each in random order. Since our purpose was to predict stimulus driven spikes, the blank trials were discarded for prediction, unless otherwise noted. Each trial was 600 ms long with the stimulus presented 200 ms after trial onset. Interrupted recordings with less than 25 presentations, or those with less than a total of 500 spikes were discarded due to insufficient data for estimating probability distributions (see below). Some of the data presented were part of a separate manuscript addressing plasticity in the coding of communication sounds (Galindo-Leon et al. 2009), a topic which is independent of the current work.

2.1.1 Extracellular recording—Electrophysiological recording locations were defined by a grid of dots stereotaxically arrayed in five columns and three rows over the left auditory cortex (Galindo-Leon et al. 2009). A single 4–6 M Ω tungsten electrode (FHC Inc, Bowdoin, ME) was advanced through a small (~ 150 μ m) hole drilled over each dot. The extracellular signal was sampled at a rate of 24414.0625 samples/s, divided into two channels and band-pass filtered for recordings of SUs ([0.3 to 6 kHz]) and LFPs ([2 to 300 or 1000 Hz] with a notch filter at 60 Hz). We focused on recording action potentials from well-isolated SU's at depths between 300 and 600 μ m. Several tests were applied to the recorded action potentials before they were classified as belonging to one SU. These included high signal-to-noise, positive and negative threshold-crossings, clustering of waveform features (Fig. 1(a)), and an absolute refractory period (1 ms) without spikes (Fig. 1(b)).

2.1.2 LFP pre-processing—Processing the LFP required special care to minimize distortion and attenuate action potential contamination. This and all subsequent data analysis was carried out in MATLAB (Mathworks, Natick, MA). To correct for phase delays introduced by data acquisition filters, the time-reversed LFP trace was filtered offline by the same type of Butterworth and notch filter. Each resulting (time-forward) LFP signal was then despiked as follow. First, a $[-0.5, 4]$ ms window around the time of each spike could be removed and replaced with a spline-interpolated signal. The LFP signal was then decimated (MATLAB function `decimate`, order 24), and low pass (`firpm`, Parks-McClellan optimal equiripple FIR filter, transition band between 90 and 100 Hz) forward and backward filtered (`filtfilt`) to further attenuate any high frequency contamination by action potentials without introducing more phase delays. This method consistently attenuated the residual power leaking into the low frequency region from large-amplitude spikes without introducing new contamination or artifacts (for further justification, see Supplemental Fig. S1).

Since some LFP trials at a given site could exhibit large amplitude transients most likely due to brief animal movements, we located such trials by finding when the LFP signal exceeded the range of values containing 99.5% of positive or negative LFP signal values, and automatically culled them. This constituted typically 3% of the trials at a given site. The resulting, pre-processed signal was termed the “wide-band original LFP.” Its power spectrum peaked at low frequencies (~ 5 – 10 Hz) and dropped off as $\sim 1/f^2$ (Fig. 1(c)), as is often observed

(Buzsaki and Draguhn 2004). We also considered narrowband versions of the LFP. Specifically, θ , β and γ -band original LFP signals were constructed with Parks-McClellan bandpass filters using the following transition bands: θ -band 4.0 to 4.5 and 10.0 to 10.5; β -band from 10.5 to 11.0 and 35 to 35.5 and γ -band from 35.5 to 36.0 Hz and 95.0 to 95.5 Hz. We did not consider frequencies lower than 4 Hz to avoid residual distortions near our 2 Hz high pass acquisition filter.

2.1.3 LFP representations—In addition to the original LFP, we applied a time-frequency analysis to decompose each trace into phase and amplitude components via a Hilbert transformation (hilbert), doing so for the wide-band as well as the θ , β and γ -band signals. The Hilbert transform is a linear operator that extends a real signal into the complex plane, uniquely breaking down the original signal into the product of an amplitude, $A(t)$ (Hilbert amplitude), and an oscillatory component, $e^{i\phi(t)}$ (with $\phi(t)$ as the Hilbert phase). It can be applied to both narrow bandwidth as well as wider band signals, although some care must be taken when interpreting the components in the latter situation (Boashash 1992; Pikovsky et al. 2001). This methodology is growing more common in the analysis of continuous neural signals in general (Bruns 2004), and LFP's in particular (Galindo-Leon et al. 2009; Haslinger et al. 2006). As discussed in (Galindo-Leon et al. 2009), it allows us to decompose the LFP into components that approximately describe the shape (Hilbert phase) and strength (Hilbert amplitude) of the local extracellular signal's movement between relative depolarization and hyperpolarization, and the wider bandwidth allows tracking of faster changes.

2.2 Prediction algorithm

2.2.1 Spiking likelihood—We predicted the probability of spiking from a specific LFP representation, X (any representation from any frequency band), on a trial-by-trial basis using a Bayesian algorithm. Bayes Theorem relates the conditional probabilities of two random variables (Cover and Thomas 1991). Our variables were the occurrence of a spike at absolute time t' (binary, $spk(t')$), and the trajectory of X over a time interval $t \in [t' + t_{ini}, t' + t_{fin}]$ relative to the spike, $spk(t')=1$. Here, t_{ini} and t_{fin} parameterized the start and end of a portion of the trajectory, relative to $t' - 50\text{ms} \leq t_{ini}$, $t_{fin} \leq 50\text{ms}$. This trajectory, which could be pre-, peri-, or post-spike, was denoted as a vector, $X(t' + t_{ini}, t' + t_{fin}) = (X(t' + t_{ini}), X(t' + t_{ini} + \Delta t), \dots, X(t' + t_{fin}))$, where Δt was the sampling period of the decimated signal, 0.9823 ms. The likelihood L of having a spike at time t' given a specific trajectory of the signal X was then defined as:

$$L(sp_k(t')) = \frac{P(\vec{X}(t' + t_{ini}, t' + t_{fin}) | sp_k(t') = 1)}{P(\vec{X}(t' + t_{ini}, t' + t_{fin}))} \times P(sp_k(t') = 1) \quad (1)$$

where $P(\vec{X}(t' + t_{ini}, t' + t_{fin}) | sp_k(t') = 1)$ was the spike-triggered probability distribution for trajectories (STPDT), $P(\vec{X}(t' + t_{ini}, t' + t_{fin}))$ was the *a priori* probability distribution for trajectories (PDT) regardless of when spikes occurred, and $P(sp_k(t')=1)$ was the overall spiking probability.

Assuming stationarity, the probability distributions in Eq. (1) are the same regardless of the absolute time of a spike t' , but could depend on the time relative to the spike, $(t-t')$. Estimating such distributions would require data sets growing exponentially with the number of dependent time points, making this problem experimentally and computationally intractable. Instead, here we explored two approximations. The first was the independent approximation, which assumed that the probability of signal X at time $t-t'$ was completely independent of probabilities at times $t-k\Delta t-t'$, where k is an integer. The second was the Markov approximation, which assumed time

dependence in the signal only between two consecutive time points; in other words, $P(X(t-t'))$ depended only on $X(t-k\Delta t-t')$, with $k=1$. We cannot claim that either of these approximations produces an optimal estimate of the spiking activity, or that the signals strictly obey these assumptions. Nevertheless, the results will demonstrate that our simple prediction algorithm is strikingly good, and will provide insight into the aspects of the LFP signal that are predictive of spiking.

The procedure to construct $P(\vec{X}(t' + t_{ini}, t' + t_{fin}) | spk(t') = 1)$ and $P(\vec{X}(t' + t_{ini}, t' + t_{fin}))$ depended on which approximation was used, as discussed below. In the following, probability distributions were constructed based on half of the trials (randomly selected), and spiking predictions were made for the other half of the trials. For computational purposes the values of X were discretized to fall into 24 equally-sized bins spanning the full range of values. The number of bins, and consequently the bin size, represents a tradeoff between the estimation error for each bin and the resolution of the STPDT. In other words, a large number of bins could better reflect small variations of the LFP trajectory, but would require more data to accurately estimate values for each bin. We chose 24 bins based on the phase representation since an initial survey of sites showed that this was sufficient to see fine structure in the STPDT relevant for predictions. We kept the same number of bins for the other representations so that the estimation error would be equivalent, thereby eliminating that as a possible confound. For the phase representation, the ordering of bins was circular.

2.2.2 Independent approximation—Assuming independence of the STPDT at different time points allowed us to rewrite the numerator in Eq. (1) as the product of the probabilities

of X at each individual time, or $\prod_{t=t'+t_{ini}}^{t'+t_{fin}} P(X(t) | spk(t'))$. With the same assumption for the PDT, Eq. (1) became

$$L(sp k(t')) = 1 | \vec{X}(t' + t_{ini}, t' + t_{fin}) = \frac{P(X(t' + t_{ini}) | spk(t') = 1)}{P(X(t' + t_{ini}))} \frac{P(X(t' + t_{ini} + \Delta t) | spk(t') = 1)}{P(X(t' + t_{ini} + \Delta t))} \dots \frac{P(X(t' + t_{fin}) | spk(t') = 1)}{P(X(t' + t_{fin}))} \times N. \quad (2)$$

The constant N was fit to equate the predicted and experimentally measured spontaneous firing rates, which we found to be more robust for scaling the absolute rate than simply using the measured value of $P(sp k(t') = 1)$. This point is addressed further in the Discussion.

The probabilities in (2) were constructed from the trial-by-trial LFP and SU recordings, as depicted in Fig. 1(d–f). A typical trace of the wide-band original LFP is shown along with its co-recorded spikes (vertical hashes). The STPDT was found by realigning the trajectories of this signal, X , from t_{ini} to t_{fin} relative to each individual spike time so that the positions of all the spikes coincided at time 0 (Fig. 1(d), inset). In order to avoid edges effects, spikes that occurred in the first and last 50 ms of any trial were discarded. This process was repeated across all trials that went into the STPDT construction, resulting in a matrix (Fig. 1(e)) with dark (light) colors denoting high (low) probabilities for whether X fell into a particular bin at a specific time relative to the spike. The PDT (Fig. 1(f)) was constructed by a similar procedure, except that “fake” spike times were used by choosing trajectories around the same number of randomly selected time points distributed uniformly across the (same) trials. Both the STPDT and PDT were smoothed (smooth, span=3). As expected, the PDT did not exhibit time-dependent structure, in contrast to the deflection in the STPDT around 0. Thus, each of the time-indexed probabilities in the numerator (denominator) of Eq. (2) was extracted from the corresponding column of the matrix in Fig. 1(e) (1F).

2.2.3 Markov approximation—Assuming a one-step Markov chain for the LFP signal X allowed us to rewrite Eq. (1) as

$$L(sp_k(t')=1|\vec{X}(t'+t_{ini}, t'+t_{fin})) = \frac{P(X(t'+t_{ini})|sp_k(t')=1)}{P(X(t'+t_{ini}))} \frac{P(X(t'+t_{ini}+\Delta t)|X(t'+t_{ini})|sp_k(t')=1)}{P(X(t'+t_{ini}+\Delta t)|X(t'+t_{ini}))} \dots \frac{P(X(t'+t_{fin})|X(t'+t_{fin}-\Delta t)|sp_k(t')=1)}{P(X(t'+t_{fin})|X(t'+t_{fin}-\Delta t))} \times N. \quad 3$$

The spiking probability in Eq. (1) was again replaced by the fitting parameter N .

The Markov STPDT and PDT were derived in a similar manner as in the independent approximation. In this case though, each term $P(X(t'+t_{ini}+(k+1)\Delta t)|X(t'+t_{ini}+k\Delta t); sp_k(t')=1)$ expressed the probability that X transitioned from a particular value at $t'+t_{ini}+k\Delta t$ to different possible values at time $t'+t_{ini}+(k+1)\Delta t$. Since these values were discretized in 24 bins, this created a square matrix (24×24) at each time point relative to the spike, so that the STPDT and PDT became three-dimensional arrays (Fig. 1(g) and (h), respectively). The strong diagonal bias indicated correlations across time in the signal, which the Markov approximation began to help us address.

Since results are reported here for different frequency bands, LFP representations and

prediction algorithms, the following notation is used: $\text{Band}_{\text{representation}}^{\text{algorithm}}$, where Band can be either W (wide-band), θ , β , γ , or C (independently combined θ , β and γ , see Section 3); representation can be orig (original), ampli (Hilbert amplitude) or phase (Hilbert phase); and algorithm can be ind (independent) or Mark (Markov).

2.2.4 Predicting the peri-stimulus time histogram (PSTH)—Although the likelihood function (Eqs. (1)–(3)) provided the time-dependent probability for spikes on individual trials, an arbitrary spiking mechanism would need to be introduced to generate “predicted spikes.” Instead, we chose to predict the average time-dependent spike rate, $R_p(t)$, rather than the trial-by-trial spikes. The predicted rate was simply the sum of the individual likelihoods over all predicted trials scaled by the normalization factor, N , fit to match the time- and trial-averaged spontaneous spike rate, R_{sp} , from 50 to 200 ms before stimulus onset. To quantify the prediction accuracy, we compared $R_p(t)$ to the experimental spike rate, as represented by the smoothed PSTH, $R_e(t)$. Smoothing was performed by upsampling the spike train time resolution by 20×, convolving spike trains with a Gaussian of standard deviation $\sigma=1, 2$ or 4 ms, down-sampling back to the original sample rate, and then averaging the trials together. Comparison was done by computing the linear correlation coefficient (CC, `corrcoef`) between $R_p(t)$ and $R_e(t)$ over a $T=200$ ms poststimulus period triggered by the stimulus. This provided a normalized measure of how well fluctuations away from the mean in the latter were reflected in the former. Note that it was insensitive to the absolute scaling of the firing rate. Finally, an alternative normalized mean square error measure was also evaluated, and gave substantially similar results (not reported here for simplicity).

3 Results

This study employed a Bayesian algorithm (see Section 2) to predict the stimulus-evoked spiking on a fine time resolution of well-isolated SUs from the co-recorded LFP activity in the left auditory cortex of the awake, head-restrained female mouse. Stimuli consisted of species-specific, infant and adult ultrasonic vocalizations varying in frequency (~60–80 kHz), duration (~12–65 ms) and frequency modulation (Galindo-Leon et al. 2009; Liu and Schreiner 2007). A total of 83 SUs were recorded, from which 23 SUs were culled due to insufficient data (see Section 2). An additional 11 SUs did not have a significantly stimulus-locked response to any

of the calls (greater than 2 standard deviation change in firing rate from spontaneous), and were excluded from population analyses. The remaining 49 SUs (59% of recordings) represented cases in which predictions of stimulus-evoked activity could be performed and meaningfully compared to experiments.

3.1 Predicting stimulus-evoked SU spiking

Predicting SU spiking from the LFP required careful consideration of how to represent the latter, since the relationship between the two is likely highly nonlinear. We generated predictions based on the wide-band original LFP voltage signal (Fig. 2(a)), as well as its Hilbert amplitude (Fig. 2(c)) and phase (Fig. 2(e)) representations. This allowed us to empirically judge whether certain representations of the LFP might systematically achieve better predictions. As noted in the Methods, the Hilbert transformation generates components approximately describing the local shape (phase) and strength (amplitude) of the fluctuations in the LFP. The local shape of these fluctuations are particularly relevant for spiking, since it has been found that spikes often preferentially fire at specific phases (e.g. in the θ band) of the LFP (Buzsaki and Draguhn 2004; Jacobs et al. 2007).

Using such a preferred representation, we then considered whether specific narrow frequency bands (θ , β , γ -bands) within the LFP signal were better for prediction. Furthermore, since approximations were necessary for computational tractability, we implemented two different methods (independent and Markov algorithms) to help ensure that our conclusions about the SU-LFP interrelation were not dependent on the details of the prediction algorithm.

The prediction process began by deriving the conditional probability distribution of trajectories X given spike at time t' (the spike-triggered probability distribution of trajectories, STPDT), as well as the a-priori probability distribution (PDT), directly from a random half of the trials (median total number of spikes per SU over all trials of 3919). The STPDT provided a picture of how a specific LFP representation varied around a spike. For the example SU in Figs. 1 and 2, the STPDT for the original LFP (Fig. 2(a)) and its Hilbert phase (Fig. 2(f), upper panel) under the independent approximation showed systematic structure consistent with a local valley in the LFP near the time of the spike. In particular, the Hilbert phase began concentrating around π (local valley) ~ 10 ms before the spike, evolved for ~ 30 ms up to $\sim 3\pi/2$, and then became progressively more random afterwards. On the other hand, the Hilbert amplitude STPDT was essentially unchanged around a spike (Fig. 2(d)), suggesting that this component carried little information about when a SU spike occurred.

The STPDT was then used on a trial-by-trial basis to compute the spiking likelihood (Eq. (1)) for each of the trials that were not used to derive the STPDT. The procedure is depicted in Figs. 2(e–h) for the Hilbert phase. The spiking likelihood at an absolute time t' was estimated based on the phase trajectory (Fig. 2(f), black lines) over the interval $[t' + t_{ini}, t' + t_{fin}]$ (Fig. 2(e), dashed box). Predictions were run separately for different values of t_{ini} between -50 ms to $+45$ ms in steps of 5 ms, and t_{fin} between $t_{ini} + 5$ ms to $+50$ ms in steps of 5 ms. For a given combination of (t_{ini}, t_{fin}) , the likelihood was determined at each time point t' during the trial, avoiding time points at the ends that would have required trajectories extending outside a trial. This produced a continuous spiking likelihood function (Fig 2(g)), whose high values ideally should have matched actual spiking activity in the SU (vertical hashes).

From the trial-by-trial likelihoods, we predicted the overall PSTH (Fig 2(h), black line) by summing the likelihood over all predicted trials and scaling the result so that the average pre-stimulus (50–200 ms) predicted rate equaled the experimental spontaneous rate during the same period. Hence, our PSTH prediction algorithm was entirely empirical, and required only the single fitting parameter N . The degree to which the predicted PSTH matched the experimental PSTH (Fig. 2(h), gray histogram, smoothed with 1 ms Gaussian) indicated the quality of our

prediction. To quantify this, we focused only on the stimulus-evoked period (first 200 ms after stimulus onset), and computed the CC. The CC is insensitive to the absolute scaling of the firing rate and provides a normalized, bounded measure of whether fluctuations in the experimental firing rate are consistently reflected in the prediction. We visualized this by plotting the predicted versus experimental firing rates on a time-point by time-point basis (Fig. 2(i)). Occasionally, the CC could be high even when the prediction itself did not capture the stimulus-evoked changes in *absolute* firing rate. In other words, the slope of the correlation could be far from 1 (perfect prediction). To refrain from judging these as good predictions, we required viable predictions to have slopes lying within a factor of 2 of unity when using the CC measure.

For each SU, the CC was evaluated for predictions run at all 210 possible combinations of t_{ini} and t_{fin} . Figure 3 shows how this STPDT time window affected the prediction quality for the example SU in the case of the independent approximation. The quality triangles plot the value of the CC as a gray scale, with darker shades indicating better predictions. Time windows for which the prediction did not satisfy the slope criterion are marked by “x.” The example demonstrates that the optimal STPDT window could be different for each of the LFP representations. In this case, the best prediction using the original LFP voltage signal was obtained for a window from [0, 5] ms around a spike, but was [0, 15] ms for Hilbert phase. The Hilbert amplitude generally performed poorly in predictions, regardless of the time window. From these plots, an SU's optimal STPDT window was determined for each of its LFP representations.

A final consideration in judging our predictions was the stochastic nature of the experimental PSTH itself. Since our algorithm predicted a firing rate function at a fine temporal resolution dictated by the downsampled period of our LFP ($\Delta t = 0.9823$ ms), prediction quality was always evaluated at this same resolution. However, lacking a spike generation mechanism, predicted firing rates were compared to *smoothed* experimental PSTHs. This raised the question of how the smoothing time constant affected the measures of prediction quality. To address this, we evaluated the CC for 3 different standard deviations for our Gaussian smoothing function: $\sigma = 1, 2$ and 4 ms. Figure 4 shows the predictions from the best STPDT windows for our example SU (Fig. 4(a)) as well as the population (Fig. 4(b)), for all three LFP representations, and both the independent and Markov algorithms. Increasing the time constant (Fig. 4(a), top 3 panels) smoothed the fluctuations in the experimental PSTH (gray histogram) and reduced the height of the transient firing rate peak. The CCs for 1 ms were generally lower than for 2 and 4 ms, but were only minimally changed between 2 and 4 ms. This was true regardless of the choice of LFP representation, or whether the independent or Markov approximations were used. Although it would provide the highest temporal resolution, the no-smoothing (or $\sigma = 0$ ms) case resulted in the poorest predictions because of estimation noise in the experimental firing rate (data not shown). Hence, we used a 2 ms smoothing time constant in subsequent analyses as a compromise between high temporal resolution and high CCs.

3.2 SU prediction by SU-LFP “signature”

Focusing on the independent approximation, the CC generally depended on the representation used for prediction. This was evident from comparing the time course for the predicted rate for the example SU (Fig. 4(a)) across the different representations. The phase-based prediction correctly tracked both the peak as well as the slight post-excitatory inhibition, whereas the original signal predicted a large rebound, and the amplitude predicted additional noisy peaks. Thus, in this particular example the Hilbert phase representation outperformed the original LFP, as was the case for most SUs in our population. Indeed, most of the points in the scatter plots of Fig. 4(b) were above the diagonals, as were the population-averaged CCs (black +s) for the different representations. The Hilbert amplitude representation generally performed

very poorly, while the original LFP occasionally did as well as the Hilbert phase. Note that the number of data points in each panel was less than the total number of SUs because the slope criterion eliminated predictions for some LFP representations for some SUs.

Although the wide-band Hilbert phase performed better than the others under the independent approximation, it should not be viewed as an intrinsically favored representation for relating LFP to SU activity. We based this on the predictions from the Markov algorithm (Fig. 4(a), bottom panel). Optimal STPDT windows were again found for each SU for each of the LFP representations. Hilbert amplitude was still the worst predictor both for our example SU (no viable STPDT window for amplitude) and across the population (Fig. 4(b), bottom panels). However, the Hilbert phase signal was not generally better for predictions than the original LFP (Fig. 4(b), bottom left panel, black+along diagonal) on average.

To understand this in more detail, we compared predictions made by the different LFP representations and algorithms by normalizing each CC by that SU's independent Hilbert phase CC (Fig. 5). The number of SUs included for different representations changed since the slope criterion differentially affected whether an optimal window could be found. On an average SU-by-SU basis, the independent phase (W_{phase}^{ind}) significantly outperformed both the independent original signal (median $W_{orig}^{ind} \sim 80\%$ of W_{phase}^{ind}), and the independent amplitude component ($W_{ampli}^{ind} \sim 50\%$ of W_{phase}^{ind}). The Markov amplitude component did even worse ($W_{phase}^{Mark} \sim 30\%$ of W_{phase}^{ind}), suggesting that the wide-band Hilbert amplitude is poorly related to spiking activity. On the other hand, despite variation from SU to SU (note wide interquartile ranges), the Markov original (W_{orig}^{Mark}) and the Markov phase (W_{phase}^{Mark}) were not systematically worse at predicting SU firing than W_{phase}^{ind} .

These results suggest that the relevant parameter for predicting SU spiking derives from how the local shape of the LFP changes instantaneously. The Hilbert phase is itself a measure of the instantaneous shape of the LFP, and it is sensitive to the relative difference in the LFP at nearby time points (Galindo-Leon et al. 2009). For example, when the Hilbert phase is just larger than π , the LFP is just past a local minimum and its value is increasing from one time point to another. Predictions based on the Hilbert phase (W_{phase}^{ind} or W_{phase}^{Mark}) were therefore inherently using information about the locally changing LFP. This was also the case for the W_{orig}^{Mark} predictions, which explicitly depended on the distribution of LFP changes at adjacent time points. Moreover, the fact that the Hilbert amplitude never performed well suggests that the absolute magnitude of the wide-band LFP is less important for spiking than the local, relative changes in the LFP.

Based on these conclusions, we were concerned that the predictions could have been affected by the strong LFP deflections evoked by stimuli. Thus, we re-ran the W_{phase}^{ind} predictions of the stimulus-evoked spiking using STPDTs derived only from spikes with ± 50 ms peri-spike windows falling only within the spontaneous period. In this case only, spontaneous spikes within 50–150 ms and 450–550 ms from *all* stimulus trials, and all spontaneous spikes from blank trials, were used to generate the STPDT, and prediction quality was again evaluated only for the stimulus-evoked portion (200–400 ms) of the trial. Even though there were fewer spikes (typically about 78% of that normally used) to estimate the STPDT, these predictions ($W_{Sp_{phase}^{ind}}$, last column in Fig. 5) were not significantly different from the W_{phase}^{ind} predictions. This suggests that the relationship between a SU's spikes and the LFP's instantaneous changes was intrinsic to the neural circuit, rather than dependent on the stimulus.

This idea was reinforced by recordings from several sites in which two SUs were simultaneously isolated for the same LFP signal. Such a case is illustrated in Fig. 6. The SUs had distinct spike waveforms (Fig. 6(d) and (f)), and no refractory spikes (Fig. 6(a) and (c)). SU1516 appeared to directly inhibit SU1517 with ~ 2 ms delay, as evidenced by the asymmetric suppression of interspike intervals between spikes from the two SUs (Fig. 6(b); compare histogram to shuffle-corrected expectation, gray line). The STPDTs for the two SUs (Fig. 6(e) and (g)) were noisy, but clearly different. At the time of the spike, the most likely phase (darkest shade) for SU1516 was around π (Fig. 6(e)), while it was around $\pi/2$ for SU1517 (Fig. 6(g)). Consequently, the two STPDTs predicted strikingly different firing rate time courses from the same LFP: SU1516 was excited by the stimulus (Fig. 6(h)), while SU1517 was inhibited (Fig. 6(i)), both in approximate agreement with their corresponding PSTHs. Hence, far from just indicating coarse-scale, stimulus-evoked spiking, the LFP has a “signature” relationship with individual SUs, dictated presumably by the local neural circuitry, which can then be used to predict that SU's spiking or suppression of spiking.

Further prediction examples are shown in Fig. 7(a) for W_{phase}^{ind} , which was more efficient ($\sim 10\times$ faster) to compute than the similarly performing W_{orig}^{Mark} and W_{phase}^{Mark} algorithms. As with the case discussed above, these predictions correctly captured both excitatory (Fig. 7(a), SU1005, SU1419, SU1361) as well as inhibitory (Fig. 7(a), SU1480, SU1361 and SU1362) components of the experimental PSTH. Both sustained (Fig. 7(a), SU1419) as well as transient (Fig. 2(h), SU1381) driven responses were well predicted. In addition, the optimal window did not necessarily include the time $t'=0$. Figure 7(a) shows cases where the optimal window was completely pre-spike (SU1005: window $[-25 -15]$) or post-spike (SU1362: window $[10 35]$).

Overall, of the 37 SUs with a valid Hilbert phase prediction, about half had a $CC(W_{phase}^{ind})$ greater than ~ 0.75 (Fig. 7(b)), indicating a remarkable success in predicting the PSTH on a fine time scale. Interestingly, when the prediction performed poorly, the corresponding SU's tended to be less well-driven by the stimuli (Fig. 7(a), SU1487), or completely non-responsive (Fig. 7(a), SU1329, not included in population data).

3.3 Optimal STPDT windows

An important step in the prediction procedure was the selection of the optimal STPDT window. If we had used the Hilbert phase only at the time of the spike, the CC would have on average been only 68% of the optimal Hilbert phase CC. Hence, our study extends the notion of a preferred LFP phase at which a SU fires spikes to an entire preferred time *trajectory* of phases. An analysis of the timing of the best window for this trajectory sheds further light on the SU-LFP relationship. If the optimal window encompassed times only before the spike, the LFP probably carried information about inputs that trigger spiking, such as synaptic currents. On the other hand, if the window was purely post-spike, then the LFP may have represented the membrane after-hyperpolarization induced by spiking. The data suggests a mix of these possibilities, since many of the optimal windows began before the spike and extended beyond the spike (Fig. 7(c)). A histogram of the initial times was weighted towards times before the spike (Fig. 7(d), black line), but included 11/37 (30%) of predictable SUs with post-spike only optimal windows (positive initial time). Similarly, the final time histogram was weighted mainly after the spike (Fig. 7(d), grey line), but 8/37 (22%) of predictable SUs had optimal windows that ended before the spike (negative final time). The remaining 48% of predictable SUs were best predicted when the STPDT window bracketed the spike. Since there was no systematic relation between the timing of the optimal window and the quality of the prediction as measured by the CC, our data suggests that the specific SU-LFP signature depended on the site and generally included both pre- and post-spike influences. Furthermore, the fact that most SUs did not have an optimal window that included the spike time suggests that our results were not simply due to a residual trace of the spike at time 0 in the LFP.

We checked whether this result was dependent on the LFP representation by looking at how the optimal window changed for the same cell across representation. We focused on the difference in the initial or final time between independent original and phase representations. Any systematic difference due to representation should be reflected in the difference $t_{ini_orig} - t_{ini_phase}$ or $t_{fin_orig} - t_{fin_phase}$. For example, if the optimal window for phase started systematically earlier than that for the original LFP, then the difference $t_{ini_orig} - t_{ini_phase}$ across the population would be mostly positive. We found that the distributions of $(t_{ini_orig} - t_{ini_phase})$ and $(t_{fin_orig} - t_{fin_phase})$ were both distributed around zero with a deviation of ~ 15 ms. Hence, the minor differences in optimal windows probably did not represent anything systematic about the types of representations. Instead, we suspect that the particular relationship that the LFP had with a SU reflected in part where we recorded the LFP relative to the physical location of the SU: sometimes it may have been more sensitive to synaptic input, and other times to after-hyperpolarizations, but these more or less affected the different representations similarly.

3.4 Predictions with θ , β and γ -bands

Although the wide-band LFP signal performed well in predicting SU activity, an important question was whether specific narrow-band regions of the LFP spectrum were responsible for this. We therefore repeated the prediction algorithm for θ , β or γ -band filtered versions of the LFP voltage signal. Here, the Hilbert phase algorithm is discussed in more detail, although all LFP representations were tested under the independent algorithm (Markov algorithms did not produce substantial improvements).

As expected, the STPDT exhibited preferred phase trajectories around the spike for each of the frequency bands (Supplemental Fig. S2). Optimal time windows were chosen for each SU's various narrow-band STPDTs, and these windows could be different for different frequency bands. The predictions for our example SU based on its θ , β and γ -band Hilbert phases are shown in Fig. 8(a), along with their corresponding CCs. The time scales of the predicted excitatory peaks were consistent with the time scales of oscillations within each frequency band. For example, the θ -band prediction produced a slowly varying firing rate (~ 200 ms period) that peaked near the onset of the stimulus, while the γ -band prediction featured faster oscillations (~ 23 ms period) with the largest peak aligned with the experimental peak. However, none of these predictions were as good as the wide-band phase (Fig. 2(i)).

This remained true across the population of SUs. None of the narrow-band Hilbert phase (Fig. 8(b), filled squares) nor amplitude (Fig. 8(b), open circles) predictions performed as well as their wide-band counterparts on a SU-by-SU basis. In fact, the phase in each higher band seemed less and less effective, consistent with reports that the θ , but not the γ phase is most correlated with MU spiking (Rasch et al. 2008). The amplitude was even worse for both θ and β -bands (Fig. 8(c)). Interestingly though, the γ -band amplitude noticeably outperformed the γ -band phase (Fig. 8(c), right)—the only case in which an amplitude signal produced better predictions than a phase signal. This indicates that the best representation for narrow-band LFPs may depend on the specific band. That is, spike-LFP interactions at low frequencies (4–35 Hz) may be better represented by the phase, whereas interactions at higher frequencies (36–100 Hz) may be better represented by the amplitude. This suggests that the spike-LFP interaction involves different mechanisms at different frequencies.

3.5 Independent combination of θ , β and γ -bands

In some respects, the fact that the prediction by the wide-band LFP was better than the narrow-band LFP might not be that surprising. If different frequencies in the LFP spectrum are statistically independent, then a larger range of frequencies carries more information, which should improve predictions. To test whether this could be the only source of the wide-band

predictor's success, we also made predictions for the 13 units in which all three narrow-band predictions were valid by combining their predictors under the assumption of frequency band independence (C-band). In other words, the overall likelihood was computed as simply the product of the likelihoods derived from each of the θ , β and γ -band phase components. This was done on a trial-by-trial basis, with the results summed together and scaled as before. Assuming independence allowed us to use optimal STPDT windows for each narrow-band contribution, chosen when each of the narrow-band frequency ranges were individually optimized. This meant there was no additional window optimization.

The C-band prediction for our example SU (Fig. 8(a), bottom right panel) correctly predicted the time of the transient peak in firing at the onset of the stimuli. However, the magnitude of this peak was more than an order of magnitude larger than the experimental firing rate. Such a prediction would normally have been deemed invalid by our slope criterion for the CC. In fact, across the SU population, none of the C-band predictions would have fallen within our slope criterion, indicating generally inaccurate absolute firing rate predictions. For argument sake though, we expanded the slope criterion to include C-band predictions with slopes falling within 0.2–5. In this case, the 4 SUs that satisfied this had an average C-band CC that was significantly lower than their corresponding wide-band CC. This was also the case if we completely ignored the slope criterion for the C-band predictions (Fig. 8(b), filled inverted triangle). These results suggest that the success of the wide-band LFP compared to the narrow-band LFP in predicting SU spiking was not simply due to a larger frequency range derived by the independent combination of constituent frequency bands.

4 Discussion

We demonstrated that it is possible to use the LFP signal to predict the time course of the stimulus-evoked spike rate of individual, well-isolated SUs with high temporal resolution. We did not argue that our algorithm provided an optimal predictor. Instead, we took an empirical approach—comparing predictions based on several different ways of representing the LFP signal—to arrive at insights into the relation between LFP and SU activity. We produced surprisingly good predictions of different types of firing responses, including cases where excitatory and inhibitory responses of different SUs were recorded simultaneously at the same site. This task would have been impossible for MU activity, where purely inhibitory responses are rarely observed.

Several algorithmic decisions contributed to this success. First, we limited our analyses to only the Hilbert amplitude and phase representations aside from the original LFP signal. This provides a unique decomposition of a signal into components that approximately correspond to its local strength and shape, respectively (see Section 2). Previous studies have found that both amplitude and phase can relate to spiking, particularly for MU activity (Haslinger et al. 2006; Whittingstall and Logothetis 2009). Moreover, stimuli consistently reset the LFP phase on a trial-by-trial basis (Yamagishi et al. 2008; Galindo-Leon et al. 2009; Kayser et al. 2008; Lakatos et al. 2009), making it a good choice as a reliable intermediary to predict stimulus-locked spiking. In fact, we found that how well spiking could be predicted by the LFP phase was correlated with the precision of that site's stimulus-evoked LFP phase ($r=0.47$; $p=0.007$). On the other hand, the original LFP was much more variable trial to trial, probably because of amplitude noise, so that its precision (measured by the inverse coefficient of variation at the peak LFP) did not correlate with prediction quality for the original LFP ($r=0.25$; $p=0.17$). Finally, other representations that describe the LFP shape, such as the first or second derivative of the LFP, would have required additional parameter optimization (e.g. a time step for the derivative), and were thus not examined. In any case, the known tendency for spikes to fire near the valley and not the peak of the LFP signal would also have made the first derivative alone a poor choice.

Second, we implemented a Bayesian algorithm based on the probability of trajectories of the LFP representation, rather than simply convolve the LFP with a spike-triggered average (STA) LFP. Predictions based on a STA generally assume that the underlying signal has a Gaussian white noise structure, which is clearly not true for the LFP. Although methods have been developed to address this through normalization by a signal's autocorrelation (Theunissen et al. 2001), this involves as much complexity as the approach used here. Furthermore, since phase is a circular variable, it does not have a meaningful average and would not have produced valid predictions by simple convolution, making the STA less generalizable to alternate representations.

Third, we optimized the timing of the window for the relationship between the LFP representation and the spike (STPDT), rather than assume that its value at the time of the spike was the only relevant parameter. This may have been necessary if the specific relationship between the LFP and the SU recorded at a particular site depended on the physical location of the electrode relative to the soma, dendrites and axon of the SU. Each of these may contribute differentially to pre- and post-spike activity in the LFP.

Fourth, we investigated predictions based on two complimentary assumptions concerning the nature of the correlations between LFP time points: independence, or a one-step Markov chain. Comparisons between these results revealed that a key ingredient for good SU predictions was the “signature” relation the SU spike had with the locally changing shape of the LFP signal, which is naturally reflected by the Hilbert phase component in the independent approach. We emphasize here that the two approaches taken together reveal that phase per se is not an intrinsically favored representation, but rather that the phase conveniently captures the shape of LFP fluctuations, which is the more relevant feature. The original LFP in the Markov approximation also captures an aspect of this shape since that algorithm depends on how the LFP changes from one point to another, and it did not perform significantly differently from independent phase algorithm. Additionally the fact that the SU-LFP relation during spontaneous activity can predict the stimulus locked spiking further supports the idea that the signature is intrinsic to the local circuit.

Finally, we employed a wide-band LFP in addition to looking within specific θ , β and γ bands. This was motivated by recent studies suggesting that features of different LFP frequency bands are coupled (e.g. phase-power, phase-phase correlations) (Canolty et al. 2006; Darvas et al. 2009; Lakatos et al. 2005), and that the wide-band LFP power can correlate with spiking (Manning et al. 2009). In our case, we found that the wide-band phase (under the independent assumption) gave the best predictions. This may have been due to the suggested coupling between bands, and/or the fact that “signatures” between spikes and LFPs could be tracked with better temporal fidelity.

A caveat to our findings is that in our prediction algorithm, we introduced the fitting parameter N , which scaled the predicted PSTH and absorbed the overall probability of spiking $p(\text{spike}(t)=1)$. In principle, the spiking probability could have been derived directly from the data, making our algorithm completely parameter free. This would have given us baseline firing rate estimates that was often of similar magnitude as our fitting procedure; the actual values of $p(\text{spike}(t)=1)$ were generally within a factor of 3 of N (for 65% of the SUs). Indeed, the two were highly significantly correlated with each other ($r=0.72$, $p=8\times 10^{-10}$, with a linear regression slope of 0.78). However, the two did not always match, with some values of $p(\text{spike}(t)=1)$ much larger than the corresponding fitted N . This suggests that our independent and Markovian approximations did not always correctly account for the absolute probability levels. Importantly though, our methods were sufficient to capture the relative probability changes driven by the stimulus, and our use of N did not affect our prediction quality measure since the CC is insensitive to the absolute scaling of the firing rate. Indeed, we obtained the same optimal

intervals and CC's regardless of whether TV or $p(\text{spike}(t)=1)$ is used, leaving our conclusions intact.

This work complements the recently published work by Rasch and collaborators (Rasch et al. 2008) in several aspects. First, in their work the authors predicted the coarse time-scale spiking of MU activity from the LFP in the primate visual system. They employed both support vector machine and linear regression algorithms to learn the LFP-MU relationship and used a binary classifier to predict the occurrence or absence of spiking in 5 ms time bins for continuous stimulus movies. They assessed prediction quality on a trial-by-trial basis by the rate of correctly classifying time bins as spiking or non-spiking. In contrast, our algorithm outputted spiking probability for each ~ 1 ms time bin, so that success was better judged by comparing against the PSTH instead of binary spiking. (Rasch et al. 2008) also quantified the Spearman rank order correlation between the Gaussian-smoothed (25 ms width) experimental and predicted spike trains, finding an average correlation of only ~ 0.35 for awake recordings. While this was considerably lower than our average 0.76 CC for $W_{\text{phase}}^{\text{ind}}$ predictions, our CC was based on trial-averaged rather than single trial spike trains, complicating a direct comparison. Nevertheless, one aspect of our results was notably different from (Rasch et al. 2008). They concluded that the precision of predicted spike times could not be better than ~ 25 ms. However, the quality of our PSTH predictions only began to degrade at a time resolution of 1 ms, with 2 and 4 ms smoothing time constants both producing similarly good results. One possible reason for the discrepancy may stem from the large difference in time scales for the structure of the LFP-spike relation. Our optimal wide-band phase STPDT windows were on average only 40 ms in duration, whereas the temporal width of their spike-triggered average LFP was ~ 500 ms—either because many of their MU spikes occurred in bursts, or because of intrinsic differences between visual and auditory coding.

Despite this obvious disparity in time scales, both the previous MU (Rasch et al. 2008) and the current SU study reached some similar conclusions concerning the LFP features relevant for predicting spikes. Both found that pre- and post-spike features in the LFP could be predictive to varying degrees for different sites, consistent with earlier interpretations of the biophysical origin of the LFP (Logothetis 2003). While the post-spike structure probably reflected membrane after-hyperpolarization, it may also have originated from synaptic or membrane currents associated with other neurons that were consistently active after the monitored SU. In either case, the LFP reflected the activity of the local neural circuit conditioned on the SU's spike, irrespective of how the ongoing external stimulus modulated neural activity. In fact, both (Rasch et al. 2008) and our study found that spontaneously occurring spikes (i.e. not stimulus-evoked) were just as effective at predicting other spikes. Hence, we speculate that the intrinsic local neural connectivity, which dominates the dynamics of cortical spiking under both spontaneous and stimulated conditions (Fiser et al. 2004; Kenet et al. 2003), underlies a “signature” of the SU-LFP interaction.

Both studies further found that the lower frequency phase was more useful than the phase in higher frequency bands. In (Rasch et al. 2008), the lowest frequencies fell in the δ -band (1–4 Hz), below the range we considered, but the θ -band phase also provided some information about spikes. Consistent with that, we found that the θ -band phase predicted spiking better than the other three frequency bands. On the other hand, (Rasch et al. 2008) determined that the γ -band power was also a key feature, whereas none of our narrow-band Hilbert amplitude signals predicted SU spiking very well. Curiously though, we found that in the γ -band only, the amplitude consistently outperformed the phase, so that LFP γ -band power was relatively more important than phase-locking for spiking. This, together with the observation that our γ -band Hilbert amplitude was modulated (by $\sim 10\%$) by the θ -band phase (data not shown), fits with reports that a hierarchical relationship exists between the phase in lower bands and the power in higher bands (Lakatos et al. 2005). This correlation between bands, although weak, may

help explain why the independent combination of the predictions based on the phase in the three narrow-band frequency ranges was insufficient to reproduce the same quality of prediction as the wide-band phase. Indeed, perhaps a combination of phase in lower bands and amplitude in higher bands might produce better predictions, but a comprehensive analysis of all such possible combinations is beyond the current scope.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Brian Kocher and Yongkui Zhang for electrophysiology assistance; Tamara Ivanova for mouse husbandry; anonymous reviewers for their constructive comments; and the NIDCD (008343) and the NSF CBN (IBN-9876754) for funding.

References

- Asher I, Stark E, Abeles M, Prut Y. Comparison of direction and object selectivity of local field potentials and single units in macaque posterior parietal cortex during prehension. *Journal of Neurophysiology* 2007;97:3684–3695. [PubMed: 17376847]
- Belitski A, Gretton A, Magri C, Murayama Y, Montemurro MA, Logothetis NK, et al. Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *Journal of Neuroscience* 2008;28:5696–5709. [PubMed: 18509031]
- Boashash B. Estimating and interpreting the instantaneous frequency of a signal—Part 1: Fundamentals. *Proceedings of the IEEE* 1992;80:520–538.
- Bruns A. Fourier-, Hilbert- and wavelet-based signal analysis: are they really different approaches? *Journal of Neuroscience Methods* 2004;137:321–332. [PubMed: 15262077]
- Buzsaki G, Draguhn A. Neuronal oscillations in cortical networks. *Science* 2004;304:1926–1929. [PubMed: 15218136]
- Canolty RT, Edwards E, Dalal SS, Soltani M, Nagarajan SS, Kirsch HE, et al. High gamma power is phase-locked to theta oscillations in human neocortex. *Science* 2006;313:1626–1628. [PubMed: 16973878]
- Chrobak JJ, Buzsaki G. Gamma oscillations in the entorhinal cortex of the freely behaving rat. *Journal of Neuroscience* 1998;18:388–398. [PubMed: 9412515]
- Cover, T.; Thomas, J. *Elements of information theory*. New York: Wiley and Sons; 1991.
- Darvas F, Miller KJ, Rao RP, Ojemann JG. Nonlinear phase-phase cross-frequency coupling mediates communication between distant sites in human neocortex. *Journal of Neuroscience* 2009;29:426–435. [PubMed: 19144842]
- Fiser J, Chiu C, Weliky M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 2004;431:573–578. [PubMed: 15457262]
- Fries P, Reynolds JH, Rorie AE, Desimone R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 2001;291:1560–1563. [PubMed: 11222864]
- Galindo-Leon EE, Lin FG, Liu RC. Inhibitory plasticity in a lateral band improves cortical detection of natural vocalizations. *Neuron* 2009;62:705–716. [PubMed: 19524529]
- Haslinger R, Ulbert I, Moore CI, Brown EN, Devor A. Analysis of LFP phase predicts sensory response of barrel cortex. *Journal of Neurophysiology* 2006;96:1658–1663. [PubMed: 16775200]
- Henze DA, Borhegyi Z, Csicsvari J, Mamiya A, Harris KD, Buzsaki G. Intracellular features predicted by extracellular recordings in the hippocampus in vivo. *Journal of Neurophysiology* 2000;84:390–400. [PubMed: 10899213]
- Jacobs J, Kahana MJ, Ekstrom AD, Fried I. Brain oscillations control timing of single-neuron activity in humans. *Journal of Neuroscience* 2007;27:3839–3844. [PubMed: 17409248]
- Jensen O, Colgin LL. Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Sciences* 2007;11:267–269. [PubMed: 17548233]

- Kayser C, Petkov CI, Logothetis NK. Tuning to sound frequency in auditory field potentials. *Journal of Neurophysiology* 2007;98:1806–1809. [PubMed: 17596418]
- Kayser C, Petkov CI, Logothetis NK. Visual modulation of neurons in auditory cortex. *Cerebral Cortex* 2008;18:1560–1574. [PubMed: 18180245]
- Kenet T, Bibitchkov D, Tsodyks M, Grinvald A, Arieli A. Spontaneously emerging cortical representations of visual attributes. *Nature* 2003;425:954–956. [PubMed: 14586468]
- Lakatos P, Shah AS, Knuth KH, Ulbert I, Karmos G, Schroeder CE. An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *Journal of Neurophysiology* 2005;94:1904–1911. [PubMed: 15901760]
- Lakatos P, O'Connell MN, Barczak A, Mills A, Javitt DC, Schroeder CE. The leading sense: supramodal control of neurophysiological context by attention. *Neuron* 2009;64:419–430. [PubMed: 19914189]
- Lee H, Simpson GV, Logothetis NK, Rainer G. Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* 2005;45:147–156. [PubMed: 15629709]
- Li CY, Poo MM, Dan Y. Burst spiking of a single cortical neuron modifies global brain state. *Science* 2009;324:643–646. [PubMed: 19407203]
- Liu J, Newsome WT. Local field potential in cortical area MT: stimulus tuning and behavioral correlations. *Journal of Neuroscience* 2006;26:7779–7790. [PubMed: 16870724]
- Liu RC, Schreiner CE. Auditory cortical detection and discrimination correlates with communicative significance. *PLoS Biol* 2007;5:e173. [PubMed: 17564499]
- Liu RC, Miller KD, Merzenich MM, Schreiner CE. Acoustic variability and distinguishability among mouse ultrasound vocalizations. *Journal of the Acoustical Society of America* 2003;114:3412–3422. [PubMed: 14714820]
- Logothetis NK. The underpinnings of the BOLD functional magnetic resonance imaging signal. *Journal of Neuroscience* 2003;23:3963–3971. [PubMed: 12764080]
- Manning JR, Jacobs J, Fried I, Kahana MJ. Broadband shifts in local field potential power spectra are correlated with single-neuron spiking in humans. *Journal of Neuroscience* 2009;29:13613–13620. [PubMed: 19864573]
- Mormann F, Fell J, Axmacher N, Weber B, Lehnertz K, Elger CE, et al. Phase/amplitude reset and theta-gamma interaction in the human medial temporal lobe during a continuous word recognition memory task. *Hippocampus* 2005;15:890–900. [PubMed: 16114010]
- Norena A, Eggermont JJ. Comparison between local field potentials and unit cluster activity in primary auditory cortex and anterior auditory field in the cat. *Hearing Research* 2002;166:202–213. [PubMed: 12062772]
- Osipova D, Takashima A, Oostenveld R, Fernandez G, Maris E, Jensen O. Theta and gamma oscillations predict encoding and retrieval of declarative memory. *Journal of Neuroscience* 2006;26:7523–7531. [PubMed: 16837600]
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA. Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nature Neuroscience* 2002;5:805–811.
- Pikovsky, A.; Rosenblum, M.; Kurths, J. *Synchronization: A universal concept in nonlinear sciences*. New York: Cambridge University Press; 2001.
- Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK. Inferring spike trains from local field potentials. *Journal of Neurophysiology* 2008;99:1461–1476. [PubMed: 18160425]
- Ray S, Hsiao SS, Crone NE, Franaszczuk PJ, Niebur E. Effect of stimulus intensity on the spike-local field potential relationship in the secondary somatosensory cortex. *Journal of Neuroscience* 2008;28:7334–7343. [PubMed: 18632937]
- Singer W, Gray CM. Visual feature integration and the temporal correlation hypothesis. *Annual Review of Neuroscience* 1995;18:555–586.
- Theunissen FE, David SV, Singh NC, Hsu A, Vinje WE, Gallant JL. Estimating spatio-temporal receptive fields of auditory and visual neurons from their responses to natural stimuli. *Network* 2001;12:289–316. [PubMed: 11563531]
- Whittingstall K, Logothetis NK. Frequency-band coupling in surface EEG reflects spiking activity in monkey visual cortex. *Neuron* 2009;64:281–289. [PubMed: 19874794]

Yamagishi N, Callan DE, Anderson SJ, Kawato M. Attentional changes in pre-stimulus oscillatory activity within early visual cortex are predictive of human visual performance. *Brain Research* 2008;1197:115–122. [PubMed: 18241846]

Fig. 1.

Relating the LFP to SU spikes. SUs were well isolated, as demonstrated by (a) SU1381's well-defined waveform (*all spikes shown*), and (b) lack of refractory spikes in the interspike interval (ISI) distribution. (c) The co-recorded LFP's power spectrum showed a peak around 5–10 Hz, decaying $\sim 1/f^2$. (d) Top panel: the temporal relation between SU spikes (*black dots*) and the co-recorded LFP in the wide-band original representation is shown for a set of trials. Acoustic stimulus is presented during the interval marked by the horizontal gray bar. Bottom panel: single trial describing the interaction SU-LFP; inset shows the LFP around each of the spikes in the trial. (e) The STPDT was obtained by realigning the LFP to the time of each spike, for all spikes in half the trials, and computing the probability that the LFP fell into a specific bin at a specific time relative to the spike. Here and in later figures, darker colors represent higher probabilities, while lighter colors denote lower probabilities. Under the independent algorithm, the distribution at each time bin relative to the spike was assumed independent from any other time bin. Spike-triggered average (STA) is shown as well (*dash-gray*). (f) The negative deflection in the STPDT contrasted with the temporally uniform *a priori* distribution based on random times in the trial (PDT). Under the Markov algorithm, the (g) STPDT and (h) PDT became 3-dimensional matrices, with the shading now representing the transition probability between the LFP at time t and $t+\Delta t$.

Fig. 2.

Predicting SU spiking from the LFP's Hilbert components. **(a)** LFP trace of a single trial and the co-recorded spikes (*vertical hashes*). **(b)** Original LFP STPDT (same trials as in Fig. 1(e)). **(c)** Hilbert amplitude representation of the same trial as in Fig. 1(d). **(d)** Hilbert amplitude STPDT. **(e)** correspondent Hilbert phase representation, and the corresponding **(f, upper panel)** STPDT and **(f, lower panel)** PDT. Dashed boxes outline the time window $[t_{ini} t_{final}]$ around $t'=0$ in the probability distributions that were convolved with the X trajectory (black line) for the Bayesian prediction (Eq. (1)). **(g)** spiking likelihood at each time point, t' , which ran from 50 ms to 550 ms, for each trial not used in estimating the STPDT. **(h)** The sum across trials of the likelihoods, normalized to agree with the pre-stimulus spontaneous firing rate, generated the predicted activity (*black line*). **(i)** The predicted rate was compared on a time-by-time basis with the experimental firing rate (gray histogram in (H)). Dashed lines represent slopes of 2 and 0.5, the limits for valid predictions

Fig. 3.

Quantifying the prediction quality and the selection of the optimal window. CC for original LFP representation (**a**), Hilbert phase representation (**b**) and Hilbert amplitude (**c**). Darker colors represent better prediction quality (high CC). From the triangle of quality values for different window combinations, we selected the one that outperformed the other intervals (*black circles*). Possible intervals that did not pass the slope condition are represented by “x”, indicating no assigned value

Fig. 4.

Comparison of predicted PSTHs across σ , representation (original, phase and amplitude) and algorithm (independent and Markovian). **(a)** Each panel shows the predicted PSTHs for SU1381 derived using the three representations with their respective optimal windows. The three top panels show the best prediction provided by the independent algorithm when the experimental PSTH (gray histogram) was smoothed with $\sigma=1, 2$ and 4 ms, respectively. The bottom panel shows the case with Markov algorithm and $\sigma=2$; note that the amplitude representation did not provide an acceptable slope in this case. **(b)** Each panel shows the comparison of the CC derived for different representations across the SU population. The plus symbol represents the median across SUs, and the star indicates the example, SU1381

Fig. 5.

Population summary of CC prediction quality, relative to the wide-band Hilbert phase representation using the independent algorithm, W_{phase}^{ind} . Upward triangles mark the population median, and error bars depict the interquartile range. For this and all subsequent figures, σ was set to 2 ms. W_{phase}^{ind} significantly outperformed other algorithms, except for W_{orig}^{Mark} and W_{phase}^{Mark} and WSp_{phase}^{ind} (Wilcoxon signrank test for significant difference from 1; $W_{orig}^{ind}/W_{phase}^{ind}$: signed rank statistic $S=22$, number of differences $nd=24$, $p=0.0003$; $W_{ampli}^{ind}/W_{phase}^{ind}$: $S=4$, $nd=24$, $p=0.00003$; $W_{orig}^{Mark}/W_{phase}^{ind}$: $S=73$, $nd=22$, $p=0.05$; $W_{phase}^{Mark}/W_{phase}^{ind}$: $S=77$, $nd=23$, $p=0.06$; $W_{ampli}^{Mark}/W_{phase}^{ind}$: $S=4$, $nd=16$, $p=0.0009$; and $WSp_{phase}^{ind}/W_{phase}^{ind}$: $S=103$, $nd=26$, $p=0.07$). The double asterisks indicate significance at the $p<0.01$ level

Fig. 6.

“Signature” SU-LFP interactions for two different co-recorded SUs, SU1516 (*dark gray*) and SU1517 (*light gray*). **(a)** and **(c)** The respective ISI distributions both lacked refractory spikes. **(b)** Spikes from SU1516 inhibited SU1517, as evidenced by the asymmetric cross correlation with no intervals at ~2 ms. The gray line shows the trial shuffled correlation. **(d)** and **(f)** The respective spike waveforms where clearly stereotyped. **(e)** and **(g)** The respective STPDs were noisy, but distinct from one another. **(h)** and **(i)**, The respective predicted PSTHs approximately matched their corresponding experimental PSTHs, which showed an excitatory response for SU1516 and inhibitory response for SU1517

Fig. 7.

Summary of wide-band phase predictions by the independent algorithm. **(a)** Six different examples of predictions spanning a wide range of CCs shows the algorithm can be successful for a wide range of SU response time courses. **(b)** Cumulative distribution of CC across 37 stimulus-driven cells with valid wide-band phase predictions. SUs from **A** are identified by their respective symbols. **(c)** The optimal window for prediction for each SU is depicted as a horizontal line at that SU's corresponding CC (lines for SUs in panel **A** are darker). There was no apparent relation between the timing of the optimal window and the CC. **(d)** Across the population, the initial time of the optimal window was mostly pre-spike, whereas the final times were mostly post-spike

Fig. 8.

Dependence of predictions on LFP frequency band. (a) PSTH predictions for SU1381 using the independent algorithm for the Hilbert phase within the θ -, β -, and γ - bands of the LFP. The corresponding STPDTs and PDTs were constructed for each frequency band in a similar way as in the wide-band case. To test whether success in the wide-band algorithm could be explained simply by the increased information carried by independent frequencies in the wider frequency range, a combined band (C-band) prediction was constructed. It was derived by the direct product of the θ -, β -, and γ - likelihoods. Note that in this example the predicted PSTH dramatically overestimates the magnitude of the excitatory transient, but does conserve the timing of the peak. (b) Population summary of the narrow-band and combined-band predictions, relative to the independent wide-band phase prediction. The figure shows median and interquartile ranges for the Hilbert amplitude (*open circles*) and Hilbert phase representations (*filled squares*) in each narrow-band case. For the C-band, only SUs with valid predictions in all three narrow-band ranges were included, and the slope criterion was not applied. In all cases, the ratios were significantly less than 1, indicating the independent wide-band phase again outperformed the other representations (Wilcoxon signrank test for significant difference from 1; θ -**phase**: $S=23$, $nd=28$, $p=0.00004$; θ -**ampli**: $S=3$, $nd=14$, $p=0.0006$; β -**phase**: $S=14$, $nd=24$, $p=0.0001$; β -**ampli**: $S=0$, $nd=16$, $p=0.0004$; γ -**phase**: $S=23$, $nd=23$, $p=0.0005$; γ -**ampli**: $S=15$, $nd=14$, $p=0.02$; **C-phase**: $S=2$, $nd=19$, $p=0.0002$). (c) In the γ -band only, the amplitude component performed significantly better than the phase on a SU-by-SU basis (θ -**band**: $S=13$, $nd=14$, $p=0.01$; β -**band**: $S=41$, $nd=19$, $p=0.03$; γ -**band**: $S=31$, $nd=17$, $p=0.03$). Asterisks indicate values significantly different from 1 (**: $p<0.01$; *: $p<0.05$).