

From Neurons to Circuits: Linear Estimation of Local Field Potentials

Malte Rasch,^{4,5} Nikos K. Logothetis,⁵ and Gabriel Kreiman^{1,2,3}

¹Department of Ophthalmology and Neuroscience, Children's Hospital Boston, Harvard Medical School, Boston, Massachusetts 02115, ²Center for Brain Science and ³Swartz Center for Theoretical Neuroscience, Harvard University, Cambridge, Massachusetts 02138, ⁴Graz University of Technology, A8010 Graz, Austria, and ⁵Max Planck Institute for Biological Cybernetics, D-72012 Tübingen, Germany

Extracellular physiological recordings are typically separated into two frequency bands: local field potentials (LFPs) (a circuit property) and spiking multiunit activity (MUA). Recently, there has been increased interest in LFPs because of their correlation with functional magnetic resonance imaging blood oxygenation level-dependent measurements and the possibility of studying local processing and neuronal synchrony. To further understand the biophysical origin of LFPs, we asked whether it is possible to estimate their time course based on the spiking activity from the same electrode or nearby electrodes. We used “signal estimation theory” to show that a linear filter operation on the activity of one or a few neurons can explain a significant fraction of the LFP time course in the macaque monkey primary visual cortex. The linear filter used to estimate the LFPs had a stereotypical shape characterized by a sharp downstroke at negative time lags and a slower positive upstroke for positive time lags. The filter was similar across different neocortical regions and behavioral conditions, including spontaneous activity and visual stimulation. The estimations had a spatial resolution of ~ 1 mm and a temporal resolution of ~ 200 ms. By considering a causal filter, we observed a temporal asymmetry such that the positive time lags in the filter contributed more to the LFP estimation than the negative time lags. Additionally, we showed that spikes occurring within ~ 10 ms of spikes from nearby neurons yielded better estimation accuracies than nonsynchronous spikes. In summary, our results suggest that at least some circuit-level local properties of the field potentials can be predicted from the activity of one or a few neurons.

Introduction

The brain is usually studied at multiple scales from molecules to systems. How signals at one scale relate to those at other scales is poorly understood. Of particular importance toward understanding how perception and action are orchestrated by neuronal signals is to characterize how the activity of neuronal circuits consisting of tens to thousands of neurons arises from their component units and their interactions. A major challenge toward achieving this goal is the experimental difficulty of recording simultaneously from multiple nearby neurons and the theoretical efforts required to characterize biophysically realistic neural networks.

The extracellular voltage recorded extracellularly through microwires is typically separated into two frequency bands (Logothetis, 2002): the multiunit spiking activity (MUA) and the local field potentials (LFPs) (see Fig. 1A). MUA represents a weighted sum of the action potentials of neurons within a radius of ~ 200 μm

around the electrode tip (Holt and Koch, 1999; Gold et al., 2006). Low-pass filtering the extracellular signal (with a corner frequency of ~ 100 – 300 Hz) yields the LFP. The biophysical nature of the MUA signal is much better understood than the origins of the LFP. The LFPs are thought to reflect the activity of large numbers of neurons in a sphere of one to several millimeters around the recording electrode (Mitzdorf, 1985; Juergens et al., 1999) (but see Katzner et al., 2009). Current source density analyses and simultaneous recordings of spikes and LFPs have suggested that LFPs are more strongly correlated with EPSPs, afterpotentials, and dendritic spikes than with the output action potentials of the surrounding neurons (Haberly and Shepherd, 1973; Mitzdorf, 1985; Logothetis, 2002). This has led to the notion that LFPs represent the input to and local processing within a given brain area.

Many approaches have been followed to compare spikes and LFPs (O'Keefe and Recce, 1993; Fries et al., 2001; Laurent, 2002; Pesaran et al., 2002; Bédard et al., 2004; Kreiman et al., 2006; Belitski et al., 2008; Montemurro et al., 2008; Rasch et al., 2008; Nauhaus et al., 2009, among others). Here we asked whether it is possible to estimate the time course of LFPs from spiking activity. At first glance, the large differences in terms of the spatial scales and biophysical origin between the two signals might suggest that this would be a challenging task. We use methods from “signal estimation theory” (Poor, 1994) to show that a linear filter operation on the activity of one or a few neurons can explain a significant fraction (but not all) of the LFP time course. The estimations have a spatial scale of ~ 1 mm and a temporal scale of ~ 200 ms. The linear

Received May 21, 2009; revised Aug. 17, 2009; accepted Sept. 11, 2009.

This research was sponsored by the Whitehall Foundation, the Klingenstein Fund, National Institutes of Health Grant 1R21EY019710-01, and the Children's Hospital Boston Faculty Development Award. We thank Fabrizio Gabbiani and Ulf Knoblich for comments on this manuscript and helpful discussions. We thank Stefano Panzeri, Rodrigo Quiroga, and Alberto Mazzoni for help with the spike sorting (supplemental Fig. S9, available at www.jneurosci.org as supplemental material). The V1 data used throughout the manuscript were recorded in the Logothetis laboratory. The inferior temporal cortex data shown in the last two columns of Figure 4A were recorded by Chou Hung and James DiCarlo.

Correspondence should be addressed to Gabriel Kreiman, Children's Hospital, Harvard Medical School, 3 Blackfan Circle, Boston, MA 02115. E-mail: gabriel.kreiman@tch.harvard.edu.

DOI:10.1523/JNEUROSCI.2390-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2913785-12\$15.00/0

filter estimation extrapolates across cortical regions and behavioral conditions. We also show that spikes that temporally coincide within ~ 10 ms with spikes from nearby neurons lead to better estimations. Together, these results suggest that the activity of individual neurons can be related to at least some of the circuit-level properties of local field potentials.

Materials and Methods

Electrophysiological recordings and data processing. Electrophysiological data recorded from seven anesthetized monkeys (*Macaca mulatta*) are included in the present study. The surgical methods and recording setup have been described previously (Rasch et al., 2008). Briefly, simultaneous recordings of neuronal activity were made from primary visual cortex (V1) using 8–16 electrodes configured in 4×4 or 8×2 matrices in a grid of 1–2 mm. We study data recorded during spontaneous activity (labeled “spont” throughout the text) and data recorded while the monkey was shown a commercial movie (labeled “stim” throughout the text). In the spontaneous condition, the V1 data studied here (7 monkeys, 109 electrodes) were recorded during periods when the input screen was blank for ~ 4 min. Multiple segments of spontaneous activity were recorded within each session. From each session, we considered five segments of spontaneous activity. Unless described otherwise in text, all electrodes recorded in each session were included in the analyses. This resulted in 545 (5×109) recorded time series of 4 min length, which we refer to as “trials” throughout the manuscript. In three sessions (two anesthetized monkeys), four electrodes were simultaneously placed in the lateral geniculate nucleus (LGN). In the visual stimulation condition (6 monkeys, 84 electrodes, 420 “trials”), a 4–6 min movie segment was shown to the monkey. The movie frames, synchronized to the monitor refresh rate (60 Hz), encompassed $7\text{--}12^\circ$ of the visual field (Rasch et al., 2008). The analyses for the spont and stim conditions are based on 4 min segments in which we discarded the initial and final 30 s to avoid potential nonstationarities. Examples of the power spectral density for the LFP recordings and the spike trains are shown in supplemental Figure S1 (available at www.jneurosci.org as supplemental material). The LFP power spectral densities show a power law decay similar to those reported in other studies (Henrie and Shapley, 2005; Nauhaus et al., 2009, among many others), and the log–log plots can be fitted with a line of slope close to 2 (Milstein et al., 2009) (supplemental Fig. S1, available at www.jneurosci.org as supplemental material). For comparison purposes, in Results (see Spontaneous activity versus visual stimulation) and Figure 4A, we show results obtained based on recordings from awake monkey inferior temporal (ITC) cortex during passive viewing conditions from the study reported by Kreiman et al. (2006).

The data preprocessing steps have been described in detail previously (Rasch et al., 2008). Briefly, electrode signals were decimated to 7 kHz. The 7 kHz signal was low-pass filtered with a cutoff frequency of 220 Hz and resampled at 500 Hz to obtain the LFPs. Spike times were detected by applying a threshold to the high-pass-filtered 7 kHz signal (fourth-order Butterworth, cutoff frequency of 500 Hz). Except when noted otherwise, the detection threshold was applied at 5 SDs of the noise component of the MUA signal. In supplemental Figure S9 (available at www.jneurosci.org as supplemental material), we examined the dependence of the estimation accuracies on the spike detection threshold. In supplemental Figure S9 (available at www.jneurosci.org as supplemental material), we also performed spike sorting using the algorithm described by Quiroga et al. (2004) to compare the performance of MUA against single-unit activity (SUA).

Linear estimation of time-varying signals from spike trains. We asked whether we could infer the time course of the local field potentials from the spike trains recorded from one cell or a small number of cells near the LFP electrode. One possible approach to this problem is to apply methods of signal estimation theory (Poor, 1994). Such methods have been used in neurobiology by several groups (Bialek et al., 1991; Gabbiani and Koch, 1996; Rieke et al., 1997; Kreiman et al., 2000) to study the transmission of information by peripheral sensory neurons. Here we analyzed

the data using the following algorithm (Bialek et al., 1991; Gabbiani and Koch, 1996) (see the scheme in Fig. 1A). Let

$$x(t) = \sum_i \delta(t - t_i) - x_0 \quad (1)$$

be the spike train obtained after subtracting the mean firing rate x_0 (where t_i are the spike occurrence times). Except in supplemental Figure S9B (available at www.jneurosci.org as supplemental material), $x(t)$ contains spikes from multiunit activity. A linear estimate, $L_{\text{est}}(t)$, of the local field potential, $L(t)$, given the spike train was obtained by convolving $x(t)$ with a filter, $h(t)$,

$$L_{\text{est}}(t) = \int_0^T d\tau h(t - \tau)x(\tau), \quad (2)$$

where T is the duration of the LFP recording, and τ is the integration variable spanning the convolution period. The Wiener–Kolmogorov (W–K) filter, $h(t)$, is chosen in such a way as to minimize the mean squared error, ε^2 , between the LFP and its estimate:

$$\varepsilon^2 = \frac{1}{T} \int_0^T dt [L(t) - L_{\text{est}}(t)]^2, \quad (3)$$

where the integration is over the LFP recording duration T . An explicit formula for this filter is given by Poor (1994):

$$h(t) = \int_{-f_c}^{f_c} \frac{P_{Lx}(-f)}{P_{xx}(f)} e^{-2i\pi ft} df, \quad (4)$$

where f_c is the cutoff frequency of the LFP, P_{Lx} is the Fourier transform of the cross-correlation between the LFP and the spike train, and P_{xx} is the Fourier transform of the spike train autocorrelation.

We considered three types of W–K filters: (1) “trial-specific” filters, (2) “electrode-specific” filters, and (3) “monkey-specific” filters (Fig. 1). For the trial-specific filter, each trial was divided into two segments of equal length. The first segment was used to estimate the W–K filter according to Equation 4. The second segment was used to estimate the LFP according to Equation 2. The performance of the filter in this test is labeled “estimation accuracy” throughout the text (see below for the computation of the estimation accuracy). We compared this estimation accuracy against the “reconstruction accuracy” obtained by using the same segment to compute the W–K filter and to reconstruct the LFP time course. The reconstruction accuracies are reported in the figures as black squares and likely contain some amount of overfitting. All the conclusions in this study are based on the estimation accuracies in which the data used to compute the filter are separate from the data used to estimate the LFP. For the electrode-specific filter, we used half of the trials recorded from each electrode (odd-numbered trials) to compute the W–K filter and used the remaining half of the trials (even-numbered trials) to estimate the LFPs (Fig. 1B). Unless stated otherwise, throughout the text, we report the estimation accuracies obtained by the electrode-specific filters. As described in Results, we observed that the W–K filters were very similar across electrodes and recording conditions. Therefore, we also evaluated the performance of a monkey-specific filter in which the data from half of the electrodes were used to compute the W–K filter, and the data from the remaining half of the electrodes were used to estimate the LFP (Fig. 1C).

The W–K filter in Equation 4 will not be causal in general [i.e., $h(t) \neq 0$ for $t > 0$]. Imposing the causality constraint on the filter h requires solving the causal Wiener–Hopf equation (Poor, 1994). Because h has a finite support in the time domain, causality can also be implemented by introducing a delay in the filter (Bialek et al., 1991). In the present study, we used this second method to impose causality (see Results, Causal filters) (see Fig. 8).

Data analysis. All the data were analyzed using MATLAB (MathWorks). The spike occurrence times were resampled at 500 Hz together

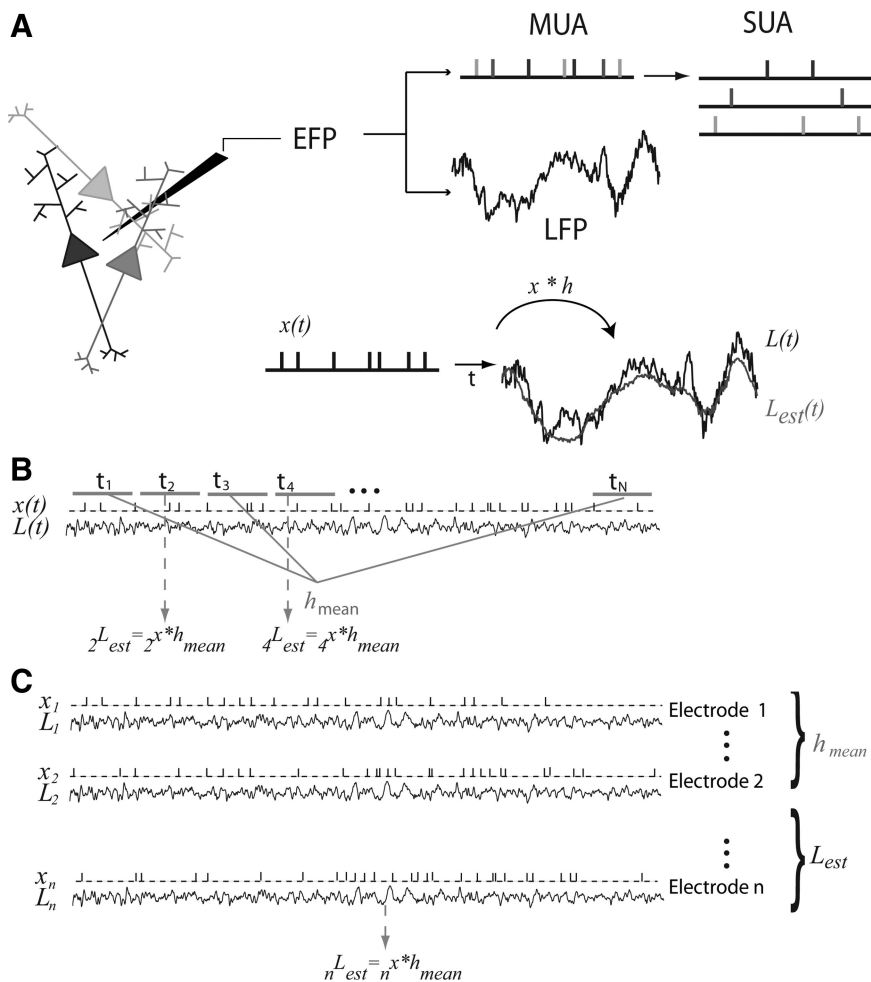


Figure 1. LFP estimation method. **A**, The extracellular field potential (EFP) is typically separated into two frequency bands: the LFPs and the MUA. The MUA signal can in turn be separated into multiple SUA clusters through spike sorting. Except for supplemental Figure S9 (available at www.jneurosci.org as supplemental material), all the figures in this manuscript use MUA. An estimate, $L_{est}(t)$, of the LFP, $L(t)$, was obtained from the spike train, $x(t)$, by convolving the spike train with a Wiener–Kolmogorov filter (for details, see Materials and Methods). The filter h was computed using half of the recorded LFP and spike train (labeled “trial-specific filter” in Fig. 3A). The estimation accuracy was quantified by computing the correlation coefficient (r) between the LFP and the estimate in the remaining half of the recordings. **B**, Construction of a general filter for the estimation of all LFPs collected with a given electrode in V1 (labeled “electrode-specific filter” in Fig. 3B). Several 4 min time series of spontaneous activity (trials) were simultaneously collected with 8–16 electrodes placed in V1 of seven monkeys. For each electrode, half of the trials (odd-numbered trials) were selected for the construction of a general filter h_{mean} . The LFPs of the remaining half of the trials collected with the same electrode were then estimated by convolving the general filter with the corresponding spike trains (for details, see Materials and Methods). **C**, In a similar manner to **B**, we considered the LFP estimation using a single general filter for all electrodes recorded from a given monkey. For this purpose, half of the electrodes were randomly selected to construct a general filter (labeled “monkey-specific filter” in Fig. 3C). The LFPs recorded from all the remaining electrodes were then estimated by convolving the general filter with each spike train.

with the LFPs. Estimates of the LFP and spike train power spectra were obtained with a fast Fourier transform algorithm and Bartlett windowing using $nfft = 2048$. As expected, the estimation accuracies, and in particular the reconstruction accuracies, depend on $nfft$ (supplemental Fig. S2A, available at www.jneurosci.org as supplemental material). Small values of $nfft$ do not capture enough information of the LFP time course, and very large values lead to overfitting. Throughout the text, we use $nfft = 2048$.

The same analysis was performed to obtain estimates of the cross-correlation between spike trains and LFPs. The cutoff frequency of the LFP was estimated by fitting the squared gain of the low-pass (Butterworth) filter transfer function to the power spectrum of the LFP, $L(t)$. The optimal W–K filter, $h(t)$, was obtained by deconvolving the cross-correlation of the spike train and the LFP with the power spectrum of the spike train according to Equation 4. The LFP estimations [$L_{est}(t)$] were obtained by computing the convolution of the W–K filter and the spike

train, $x(t)$ (see Eq. 2), in the frequency domain with the use of a fast Fourier transform.

Performance measures. Because the filter h is derived by minimizing the mean squared error between the LFP and the estimated LFP (see Eq. 3), a natural measure for the quality of the estimations is the mean squared error, ϵ^2 . Another commonly used performance measure is the Pearson’s correlation coefficient (denoted by r throughout the manuscript). Given that, in this case, both $L(t)$ and $L_{est}(t)$ are normalized so that they have 0 mean and SD 1, a linear relationship exists between the mean squared error and the correlation coefficient: $\epsilon^2 = 2(1 - r)$ (for the derivation of this expression, see supplemental data and Fig. S1F, available at www.jneurosci.org as supplemental material). Throughout the text, we report the r values and refer to them as the estimation accuracy.

To assess the statistical significance of the estimations, we compared the results against a null hypothesis in which no correlation existed between the temporal structure of the LFPs and the spike trains. Under the null hypothesis, we generated random spike trains with the same mean firing rate as the experimental spike trains but with spike times governed by a Poisson process [$x_{rand}(t)$]. We then followed the same procedure in Equation 4 to compute the optimum W–K filter to estimate the LFP time course from the random spike train, h_{rand} :

$$L_{est}^{rand}(t) = \int_0^T d\tau h_{rand}(t - \tau) x_{rand}(\tau), \quad (5)$$

with $x_0^{rand} = x_0$. If the temporal structure of the spike trains conveys no information about the LFP time course, we would expect the correlation between the estimated LFP and the actual LFP obtained from the experimental spike trains to be close to the correlations obtained from the random spike trains. We repeated this procedure 50 times; the average estimation accuracies obtained under the null hypothesis are shown as black triangles throughout the figures. To assess the statistical significance of the LFP estimations, we performed two-tailed t tests, comparing the estimation accuracies obtained under the null hypothesis against those obtained with the actual spike trains.

Results

Linear LFP estimations based on single trials

We considered 545 spontaneous activity “trials” recorded from V1 (109 electrodes, 7 monkeys), of ~ 4 min length (we discarded the first and last 30 s to avoid any potential border effects). To estimate the LFP from the spike train, we computed a linear filter that minimizes the difference between the LFP and its estimate (Eqs. 2–4) (Fig. 1A). To avoid overfitting the data, the estimation filter h was constructed (as described in Materials and Methods, Linear estimation of time-varying signals from spike trains and Data analysis) with the first half of the LFP and spike train of each trial. The estimation, $L_{est}(t)$, was computed by convolving h with the second half of each spike train. Therefore, there was no over-

lap between the data used to compute the filter and the data used to test the filter by estimating the LFP time course and measuring the correlation between $L(t)$ and $L_{\text{est}}(t)$ (see Materials and Methods, Performance measures).

Three 1 s segments of typical LFP traces and their estimations are shown in Figure 2. These three examples illustrate the range of estimation accuracies from good (Fig. 2A, where the estimation captures, to a large degree, the overall structure of the LFP time course) to poor (Fig. 2C, where the estimate is only weakly correlated with the actual LFP). We show another example electrode in Figure S3A (available at www.jneurosci.org as supplemental material), and we show more data segments for the electrode illustrated in Figure 2A in Figure S3B (available at www.jneurosci.org as supplemental material). We quantified the accuracy of the LFP estimation by computing the Pearson's correlation coefficient between the estimation and the actual LFP [referred to as the estimation accuracy (r) throughout the manuscript; see above, Performance measures]. In the examples shown in Figure 2, r has a value of 0.61 (A), 0.29 (B), and 0.01 (C). On the right in Figure 2, we show the corresponding W–K filters for time lags between -800 and $+800$ ms [the actual number of points in the W–K filter was $n\text{fft} + 1$ (2049), but there were only small deviations from 0 outside the ± 800 ms range]. As illustrated by the examples in Figure 2, wider W–K filters tended to yield higher estimation accuracies. Typical features of the W–K filter include a sharp downstroke for tens of milliseconds for negative time lags followed by a slower upstroke lasting >100 ms for positive time lags (see also the example W–K filters in Figs. 3C, 4B, 8 and supplemental Figs. S3, S7, S9, available at www.jneurosci.org as supplemental material).

Figure 3A shows the estimation accuracy for each one of the 109 electrodes recorded from V1 during spontaneous activity after averaging across all trials (five trials). We found that r varies significantly across trials and electrodes. Whereas the estimation accuracy for some trials is higher than 0.6, it is almost 0 for others (see the overall distribution in supplemental Fig. S2B, available at www.jneurosci.org as supplemental material). Overall, most of the electrodes (95 of the 109 electrodes) yielded a statistically significant estimation accuracy compared with the null hypothesis obtained by creating a Poisson spike train with the same mean firing rates ($p < 0.01$, two-tailed t test) (see above, Performance measures). The average estimation accuracy across V1 was 0.36 ± 0.15 (mean \pm SD, $n = 109$ electrodes) (Fig. 3A). On average, the estimations obtained under the null hypothesis led to a correlation of $\bar{r}_{\text{rand}} = 0.001 \pm 0.035$ (mean \pm SD, $n = 109$ electrodes) (Fig. 3A, triangles and bottom dotted line). Despite the variability across trials and electrodes, we found that the estimations were highly significant [two-sample Kolmogorov–Smirnov

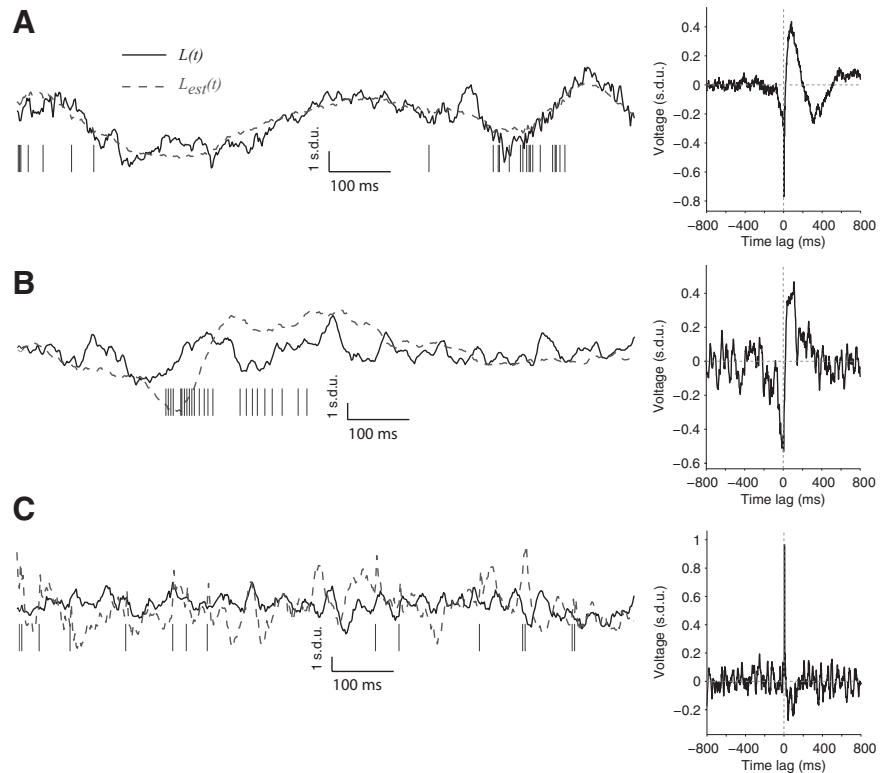


Figure 2. Examples of linear estimation of LFPs. **A–C**, Three examples of LFP estimations in three different electrodes recording V1 spontaneous activity. A random 1 s sample is shown for each electrode; the firing rates and estimation accuracies are computed using a 240 s recording. The solid line shows the actual LFP, and the dashed line shows the estimated LFP. The vertical ticks show the action potentials (MUA). On the right, we show the Wiener–Kolmogorov filter for each electrode (h_{mean}). We only show the W–K filter for time lags between -800 and $+800$ ms; the number of points in the W–K filter was $n\text{fft} + 1$ ($n\text{fft} = 2048$). These examples illustrate the range of estimation levels that we observed (**A**, $r = 0.61$, firing rate of 46 Hz; **B**, $r = 0.29$, firing rate of 10 Hz; **C**, $r = 0.01$, firing rate of 13 Hz). More data segments for the electrode in **A** are shown in supplemental Figure S3B (available at www.jneurosci.org as supplemental material).

test comparing the distribution in supplemental Figure S2B (available at www.jneurosci.org as supplemental material) against the corresponding distribution obtained under the null hypothesis, $p < 10^{-10}$].

To gain additional insight into what contributes to the variability in the estimation accuracies across trials and across electrodes, we plotted the estimation accuracy as a function of basic properties of the spike trains and LFPs (supplemental Fig. S4, available at www.jneurosci.org as supplemental material). The estimation accuracy increased with higher firing rates (correlation coefficient of 0.49) (Fig. S4A,D, available at www.jneurosci.org as supplemental material), with the total power of the LFP signal (correlation coefficient of 0.44) (supplemental Fig. S4B,E, available at www.jneurosci.org as supplemental material) and also with the coefficient of variation (CV) of the interspike interval distribution (correlation coefficient of 0.59) (supplemental Fig. S4C,F, available at www.jneurosci.org as supplemental material). It should be noted that these relatively large CV values are based on MUA and not SUA. Based on supplemental Figure S4 (available at www.jneurosci.org as supplemental material), subsequent analyses in the manuscript focused on those electrodes that had a firing rate of at least 5 spikes/s and a CV of at least 1 (88 of the 109 electrodes).

Given the intrinsic noise in the spike trains and LFPs, we were interested in estimating an upper bound for the estimation accuracy r . For this purpose, we asked how well this linear algorithm

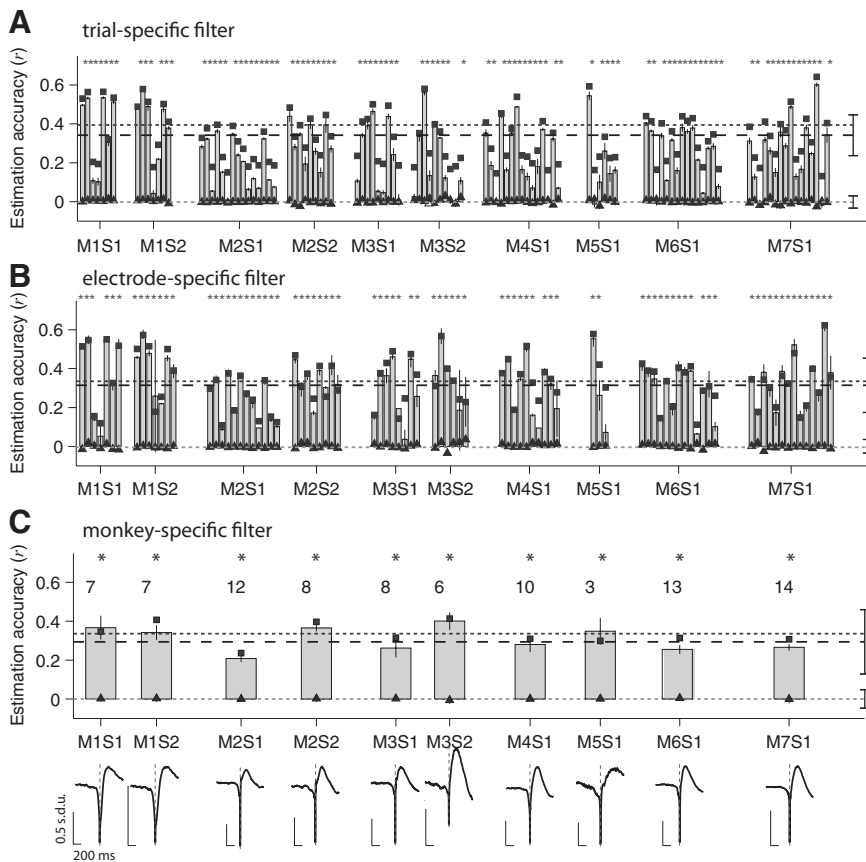


Figure 3. Linear estimation of LFPs for all electrodes. **A**, LFP time course estimations based on trial-specific W–K filters (see Results and Fig. 1A). Each 4 min trial was split into two segments: one of these was used to build the W–K filter, and the other segment was used to compute the LFP estimate. Each vertical bar reports the mean estimation accuracy (labeled *r* throughout text) for all test trials in one electrode, and error bars represent 1 SEM (5 trials). The data are blocked according to recording session (S) and monkey (M). The black squares represent the mean reconstruction accuracy obtained when the same data were used to compute the filter and reconstruct the LFP (see Materials and Methods). The triangles indicate the estimation accuracy obtained under the null hypothesis using random spike trains with the same firing rates (see Materials and Methods, Performance measures). **p* < 0.01 based on a two-tailed *t* test comparing the estimation accuracies against the null hypothesis. The dashed line indicates the average estimation accuracy (\bar{r}) across all electrodes and all monkeys ($\bar{r} = 0.36 \pm 0.15, n = 109$). The dark gray dotted line indicates the average reconstruction accuracy [average of the black squares ($\bar{r} = 0.39 \pm 0.12, n = 109$)]. The light gray dotted line indicates the average estimation accuracy under the null hypothesis (average of the triangles, $\bar{r}_{\text{rand}} = 0.001 \pm 0.035, n = 109$). **B**, LFP time course estimations based on electrode-specific W–K filters (see Results and Fig. 1B). For each electrode, half of the trials were used to compute the W–K filter, and the spike trains in the remaining trials were used to estimate the LFP (Fig. 1B). Each bar shows the estimation accuracy for a different electrode averaged over the trials used to estimate the LFP (error bars denote 1 SEM). The conventions are the same as in **A**. Only those electrodes with firing rate > 5 spikes/s and CV > 1 are shown here and in subsequent figures (*n* = 88 electrodes of the total of 109 electrodes). **C**, LFP time course estimations based on monkey-specific filters (see Results and Fig. 1C). For each monkey, the data from half of the electrodes were used to compute the W–K filter, and the spike trains from the remaining electrodes were used to estimate the LFP. For each recording session, here we show the average estimation accuracy across all the electrodes used to estimate the LFP (error bars denote 1 SEM). The conventions are the same as in **A**. The number of electrodes recorded in each session is indicated above each bar. The monkey-specific W–K filters are shown for each recording session.

performs when estimating the LFP time course by estimating the LFP measured in one electrode, LFP₁, from the LFP recorded from a nearby electrode, LFP₂ (supplemental Fig. S6A, available at www.jneurosci.org as supplemental material). Because the recordings were performed simultaneously, two electrodes that are nearby are likely to be measuring a similar signal from slightly different places. As expected, the LFP–LFP estimation decreased with the distance between the two electrodes (supplemental Fig. S6C, available at www.jneurosci.org as supplemental material). For a distance of 1 mm (the smallest

distance available in our dataset), the mean LFP–LFP estimation was 0.55 ± 0.05 (mean \pm SD, *n* = 109 electrodes; compare with 0.36 for the spike–LFP estimation accuracy with a distance of 0 mm).

Toward a general function to map spikes into LFPs

Is the map between spikes and LFPs specific to each trial/electrode or is there a general filter that can extrapolate across different trials or even different recording electrodes? We asked whether we could find a general filter that, when applied to a spike train in V1 recorded from a given electrode, would give us an estimate of the corresponding LFP recorded from the same electrode. We addressed this question in three steps, with increasing degrees of extrapolation. First, we constructed an electrode-specific filter to predict the LFPs of all the trials recorded from a given electrode after computing the W–K filter using different trials recorded from the same electrode (Fig. 1B, one filter per electrode as opposed to one filter per trial as used in the previous section). Second we constructed a monkey-specific filter to predict the LFPs of all the trials and electrodes of each monkey after computing the W–K filter using different electrodes recorded from the same monkey and same brain region (Fig. 1C, one W–K filter per monkey). Finally, we assessed whether a W–K filter computed from all the electrodes in a given monkey could be used to estimate the LFPs recorded in a different monkey. Note that, in all cases, the LFP estimate for a given electrode was computed from the spike trains recorded from the same electrode (using Eq. 2). The extrapolation refers to the way in which the W–K filter is computed (using Eq. 4 and data from different trials, different electrodes, or different monkeys).

We constructed the electrode-specific filter for a given electrode by minimizing the sum of the mean squared errors in the estimations of individual trials recorded from that electrode (Fig. 1B,C). We find the filter (h_{mean}) that minimizes

$$\epsilon^2 = \frac{1}{N} \sum_j \int_0^T dt \left[L_j(t) - \int_0^T h_{\text{mean}}(t - \tau) x_j(\tau) d\tau \right]^2, \quad (6)$$

where *N* is half of the total number of trials for the electrode, *L_j* is the LFP corresponding to trial *j*, and *x_j* is the spike train of trial *j* with the mean subtracted. The minimization of the error leads to an expression for the Fourier transform of the filter

(for a derivation of this expression, see supplemental data, available at www.jneurosci.org as supplemental material),

$$\hat{h}_{\text{mean}}(f) = \frac{\sum_{j=1}^N j P_{Lx}(-f)}{\sum_{j=1}^N j P_{xx}(f)}, \quad (7)$$

where jP_{Lx} and jP_{xx} are the cross power spectrum and power spectrum, respectively, of the LFP and spike train in trial j . In a similar manner, we constructed a single monkey-specific filter for each monkey by using Equations 6 and 7 but summing over all the trials and electrodes used to construct the filter (half of the electrodes for each monkey; see below).

In both cases (electrode-specific W–K filter and monkey-specific W–K filter), we estimated the LFP in each trial by convolving the spike train in the same trial with h_{mean} (instead of using the filter h computed for each trial as in the previous section). When building an electrode-specific W–K filter, we used different trials to compute the filter and to compute the estimation: we picked half of the trials for a given electrode to build the filter according to Equation 7 and used the remaining trials from the same electrode to predict the LFP time course (Fig. 1B). When building a monkey-specific W–K filter, we used different electrodes to compute the filter and to compute the estimation (Fig. 1C): we randomly picked half of the electrodes to build the filter according to Equation 7 and used the remaining electrodes to predict the LFP time course.

Using electrode-specific filters, the mean estimation accuracy was $\bar{r} = 0.31 \pm 0.14$ (Fig. 3B, dashed line). Using monkey-specific filters, the mean estimation accuracy was $r = 0.30 \pm 0.16$ (Fig. 3C, dashed line). As expected, the mean estimation accuracy in this case is lower than the values reported in Figure 3A, yet it is remarkable that a general linear filter can achieve a rather accurate estimation of the LFP time course. The distribution of LFP estimations were highly significant compared with those obtained using Poisson spike trains with the same mean rate [$p < 10^{-6}$ (electrode-specific filters) and $p < 10^{-4}$ (monkey-specific filters), two-sample Kolmogorov–Smirnov test]. The shape of each of the monkey-specific W–K filters is shown in Figure 3C.

The properties of the Wiener–Kolmogorov filter in the context of encoding time-varying signals by spike trains have been studied using idealized neuron models (Gabbiani, 1996; Gabbiani and Koch, 1996). Using an integrate-and-fire model, assuming a linear filter, neglecting the refractory period and assuming an exponentially distributed spike threshold, it has been shown that the spike-triggered average (STA) of the signal yields an estimation filter that converges onto the W–K optimal filter in the limit of low firing rates (Gabbiani and Koch, 1996). This result also holds for experimental recordings (Wessel et al., 1996). We therefore quantified how well we could estimate the LFP time course using Equation 2 from a filter built by computing the spike-triggered average of the LFP instead of the optimal Wiener–Kolmogorov filter. Supplemental Figure S7, A and B (available at www.jneurosci.org as supplemental material), compares the shape of the STA with the shape of the average electrode-specific optimal W–K filter. The STA shows a sharp negativity for negative time lags and an upswing for positive time lags. This stereotypical structure of the STA is common across V1 electrodes and resembles the shape of the optimal filter h_{mean} . The STA yielded a good estimate of the LFP time course (supplemental Fig. S7C, available at www.jneurosci.org as supplemental material). As expected, this estimate was slightly worse than the one obtained using the optimal W–K filter. Thus, to a coarse approximation

and in the limit of low-firing rates, we can think of the W–K filter as the average LFP surrounding a spike.

Spontaneous activity versus visual stimulation

The results presented thus far correspond to recordings during spontaneous activity. We asked whether the linear estimation of LFPs from spike trains would extend to conditions in which V1 neurons are activated by visual input (as opposed to spontaneous firing). To examine the influence of visual stimulation on the spike–LFP relationship, we repeated the analyses in recordings from 84 electrodes in V1 while the anesthetized monkeys were shown commercial movies (see Materials and Methods). Consistent with the previous findings, we could also linearly estimate the LFP time course during visual stimulation (Fig. 4A). The shape of the W–K filter for the visual stimulation condition was similar to the corresponding shape during spontaneous activity. The mean estimation accuracy during the stimulation condition was 0.32 ± 0.13 using electrode-specific filters. Furthermore, there was a strong correlation between the estimation accuracies during spontaneous activity and those during stimulus-driven activity (correlation coefficient of 0.64) (supplemental Fig. S5, available at www.jneurosci.org as supplemental material).

Given the similarity across conditions (spontaneous activity vs visual stimulation) for a given monkey, we asked whether the LFP estimations could extrapolate across monkeys. Using the same approach described by Equations 6 and 7, we computed the W–K filter using all the recordings from a given monkey (Eqs. 6–7) and used this filter to estimate the LFPs recorded in a different monkey (Eq. 2). Figure 4B shows that the W–K filter from one monkey can be used to estimate the LFPs recorded from the same area (V1) in a different monkey (in all cases, the LFP estimate for a given electrode is based on the spike trains from the same electrode). This observation is attributable to the stereotypical shape of the W–K filters across monkeys (Fig. 4B, bottom).

Is the relationship between spikes and LFPs specific to a given brain area or is there a universal mapping between these two scales of neural analysis? We addressed this question by considering physiological recordings from the macaque monkey ITC described by Kreiman et al. (2006). During the ITC recordings, awake monkeys passively viewed a stream of grayscale objects. Using the same approach as in Figure 3, the mean correlation between the estimated LFP and the actual LFP for 125 ITC electrodes was $\bar{r} = 0.27 \pm 0.17$. Given the similarity between the V1 W–K filters and the ITC W–K filters, we asked whether we could use the V1 W–K filters to estimate the ITC LFPs based on the ITC spike trains. Using the general filter computed with the V1 data to estimate the ITC LFP from the ITC spike trains, we obtained a mean estimation accuracy of $\bar{r} = 0.25 \pm 0.12$. In other words, using a general filter computed with data from V1, we obtained an estimation accuracy of 0.30 in the remaining half of the V1 electrodes and an estimation accuracy of 0.25 in the ITC data. This shows a remarkable degree of extrapolation and suggests a general relationship between spikes and LFPs across different neocortical areas.

In contrast to the extrapolation across conditions (spontaneous activity and stimulus-driven activity in V1) (Fig. 4A), monkeys (Fig. 4B), and neocortical areas (V1 and ITC) (Fig. 4A), the W–K filter showed a poor performance in the LGN recordings. In three recording sessions in two monkeys, eight additional electrodes were simultaneously positioned in the LGN. Using 95 trials recorded from the LGN, we estimated the LGN LFPs from the LGN spiking activity as described in Materials and Methods (supplemental Fig. S8, available at www.jneurosci.org as supplemental material) shows the estimation accuracy over all LGN trials

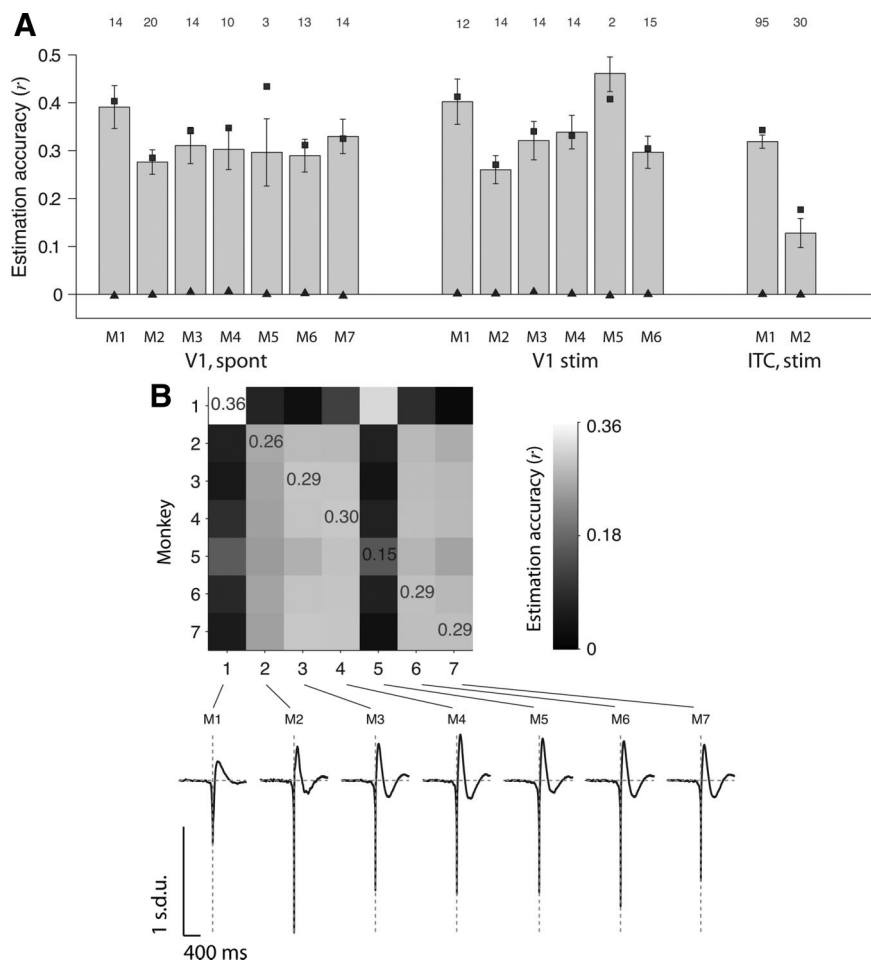


Figure 4. Linear estimation of LFPs across monkeys, stimuli and locations. **A**, Average estimation accuracies for all electrodes in V1 under spontaneous activity in seven monkeys (left), all electrodes in V1 during visual stimulation in six monkeys (middle), and all electrodes in ITC during visual stimulation in two monkeys (right). The V1 visual stimulation consisted of natural movies (see Materials and Methods). The ITC visual stimulation consisted of static grayscale images presented for 100 ms during passive viewing (Kreiman et al., 2006). Each bar indicates the mean estimation accuracy averaged over all the electrodes in each monkey (using electrode-specific filters; see Fig. 1B). The error bars denote 1 SEM. The number of electrodes averaged in each case is indicated above each column. The conventions are the same as in Figure 3; the squares indicate the reconstruction accuracy, and the triangles show the estimation accuracies expected by chance. **B**, Estimation accuracies computed across monkeys (V1 data, spontaneous activity). Entry in row *i*, column *j* indicates the estimation accuracy obtained when the W–K filter was computed using all the electrodes recorded in monkey *i* and the estimation accuracy was computed using the spike trains and LFPs from monkey *j*. Note that the LFP estimate for a given electrode is always computed here using the spike train from the same electrode (using Eq. 2). The W–K filter is computed using data from a different monkey (using Eq. 7). The values reported correspond to the average across all electrodes in monkey *j*. We indicate the estimation accuracies for the diagonal entries for reference. The bottom part shows the W–K filters for each monkey.

using electrode-specific filters (compare with Fig. 3B). The algorithm performed substantially worse when trying to estimate LGN LFPs than when trying to estimate V1 LFPs: the average estimation accuracy for the LGN was $r = 0.03$. We observed that the STA of the LFP was essentially flat and noisy in the LGN, which further indicates the lack of a clear relationship between spikes and LFPs in the LGN.

Frequency analysis

The results presented above quantify the estimation of the entire LFP time course. We asked whether there were differences in the estimation accuracy across different LFP frequency bands. For this purpose, we filtered the LFP into different frequency bands between 0.1 and 100 Hz and computed a separate W–K filter to estimate the different frequency components (Fig. 5). The estimation accuracies were statistically significant compared with the

null hypothesis for all frequency bands ($p < 0.01$, based on a two-sample Kolmogorov–Smirnov test). The lower frequency band (i.e., 0.1–20 Hz) yielded the best estimation accuracy ($p < 10^{-5}$ for the spontaneous condition and $p < 10^{-5}$ for the stimulation condition, based on a two-sample Kolmogorov–Smirnov test comparison against the null hypothesis).

These observations suggest that the spiking activity contains more information about low frequencies of the LFP than about high frequencies. This result is consistent with the results reported previously (Rasch et al., 2008). Rasch et al. showed that spikes seem to be locked at the onset of low frequency oscillations in the LFP and that the spike–LFP coherence level is higher for low frequencies than for high frequencies. Indeed, the spike–LFP coherence (Pesaran et al., 2002) can be used as a measure of the error in the estimations. High coherence in a given frequency band translates into small LFP estimation errors (for additional details about the relationship between W–K filtering and spike–LFP coherence, see supplemental data, available at www.jneurosci.org as supplemental material).

Robustness of the estimations to spike time jitter

Intuitively, the possibility of estimating the LFP time course from the spike train depends on the temporal structure of the spike train. This is further emphasized by the poor LFP estimations when using the null model consisting of a spike train with the same mean rate but random spike times. To further investigate the robustness of the LFP estimates to spike time distortions, we recomputed the W–K filter after adding temporal jitter to the spike trains. We created synthetic spike trains from the experimental ones by randomly jittering the spike times. The LFP was then estimated from these synthetic spike trains. For each spike train, the spike times were moved from their actual occurrence times by a random amount taken from a 0 mean Gaussian distribution with an SD σ_{jitter} .

The robustness of the estimations to time jittering was evaluated by plotting the mean estimation accuracy as a function of σ_{jitter} and computing the amount of temporal distortion required to cause a 50% drop in the estimation accuracy ($_{50}\sigma_{\text{jitter}}$) (Fig. 6). The value of $_{50}\sigma_{\text{jitter}}$ was obtained by fitting the following equation: $r = r_0 - \frac{r_0 \sigma_{\text{jitter}}^n}{\sigma_{\text{jitter}}^n + 50 \sigma_{\text{jitter}}^n}$, where r_0 is the mean estimation accuracy when $\sigma_{\text{jitter}} = 0$, and $_{50}\sigma_{\text{jitter}}$ and n are free parameters. Figure 6A shows the mean estimation accuracy (\bar{r}) as a function of σ_{jitter} for an example electrode in V1. To compare the decrease in r with time jitter across electrodes, we normalized the esti-

mation accuracy by r_0 (Fig. 6B). The distribution of ${}_{50}\sigma_{\text{jitter}}$ for all V1 electrodes is shown in Figure 7C; the mean value was 172 ± 89 ms.

The robustness of the LFP estimate to time jitter was strongly correlated with the width of the optimal W–K filter. We found a linear relationship between the width of the optimal filter and ${}_{50}\sigma_{\text{jitter}}$ (slope of 0.35, correlation coefficient of 0.64). The relatively long values of ${}_{50}\sigma_{\text{jitter}}$ are consistent with the higher values of \bar{r} at low frequencies described in the previous section (Fig. 5).

Dependence on distance between spike trains and LFP electrode

Spikes constitute a measure of activity within a small vicinity of the electrode, on the order of 200 μm (Holt and Koch, 1999; Logothetis, 2002). In contrast, LFPs measure neural activity within a larger area (Mitzdorf, 1985; Kruse and Eckhorn, 1996; Juergens et al., 1999; Logothetis, 2002; Kreiman et al., 2006; Belitski et al., 2008; Liu et al., 2008; Nauhaus et al., 2009) (but see Katzner et al., 2009). Because recordings were performed simultaneously with multiple electrodes, we were able to study the spatial resolution of the relationship between spikes and LFPs by considering spiking activity and LFPs recorded from separate electrodes. For this purpose, we considered two electrodes separated by a distance D , and we used the spike train in one electrode to estimate the LFP in the other electrode according to Equations 2–4. In Figure 7, the estimation accuracy is plotted as a function of the distance between the spike electrode tip and the LFP electrode tip ($D = 0$ corresponds to the case in which LFPs and spikes were recorded from the same electrode as discussed in the previous sections). Figure 7A shows the results of this analysis for an example recording session where electrodes spanned a distance of up to 4 mm. We fitted the function $r = r_0 - \frac{r_0 D}{D_r + D_{50}}$, where D_{50} is a free parameter (Fig. 7A, dotted line). We show the distribution of D_{50} values in Figure 7B. To compare the decrease in r with distance across electrodes, we normalized the estimation accuracy by r_0 (Fig. 7C). The estimation accuracy decreases by $\sim 50\%$ when the two electrodes are 1 mm apart (this is the minimum electrode distance in these data). Beyond 1 mm, we found a weak dependence of the estimation accuracy with distance. This significant drop with distance was also observed when we considered only the 0.1–20 Hz frequency band of the LFP.

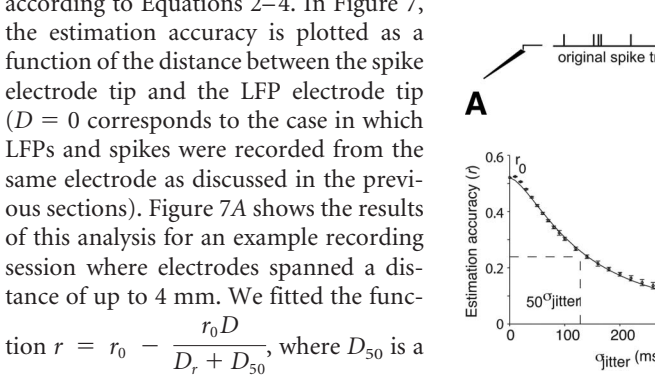


Figure 6. Robustness to spike time jitter. **A**, Mean estimation accuracy as a function of the spike time jittering for an example electrode in V1 (spontaneous activity). For each value of σ_{jitter} , we randomly shifted each spike train by an amount taken from a Gaussian distribution with mean of 0 and standard deviation σ_{jitter} (see scheme at the top). We show the mean estimation accuracy averaged across all test trials. Error bars indicate 1 SEM. r_0 denotes the estimation accuracy obtained for $\sigma_{\text{jitter}} = 0$. The continuous line indicates a fit with the following function: $r = r_0 - \frac{r_0 \sigma_{\text{jitter}}^n}{\sigma_{\text{jitter}}^n + 50 \sigma_{\text{jitter}}^n}$, where ${}_{50}\sigma_{\text{jitter}}$ and n are free parameters. **B**, For each electrode, we repeated the procedure described in **A** and normalized r by dividing it by r_0 (“Normalized estimation accuracy”). We show the mean normalized estimation accuracy across all electrodes ($n = 88$) using the electrode-specific filters. **C**, Distribution of ${}_{50}\sigma_{\text{jitter}}$ for all V1 electrodes recorded during spontaneous activity ($n = 88$). The black arrow shows the position of the example from **A**. The dashed vertical line indicates the mean value: the mean value of ${}_{50}\sigma_{\text{jitter}}$ was 172 ± 89 ms (mean \pm SD). Bin size of 25 ms.

Causal filters

In the analyses presented thus far, at any given time point t , both the spikes before t and the spikes after t contribute to the estimate of the LFP (see Eq. 2 and the shape of the W–K filters in Figs. 2, 3). We conjectured that there may be a temporal asymmetry such that the LFP estimate at time t using those spikes before t may yield a different estimation accuracy than those occurring after t . To evaluate this possibility, we constructed two different causal filters and compared their performance to the estimation accu-

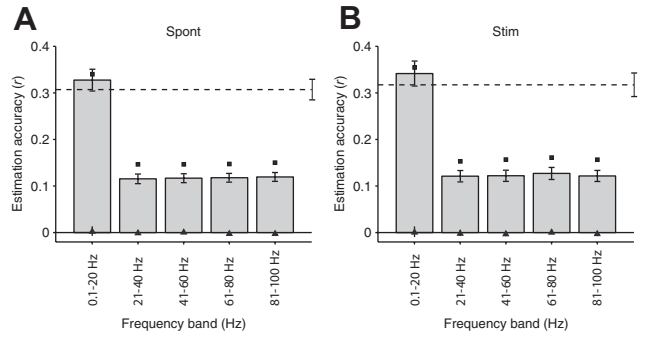


Figure 5. Frequency dependence. Mean estimation accuracies for five different frequency bands between 0.1 and 100 Hz for all V1 trials during spontaneous activity (**A**, $n = 88$) and visual stimulation (**B**, $n = 71$). We filtered the LFPs in the corresponding frequency band, we computed the W–K filters, and we estimated the LFPs by convolving the electrode-specific filters with each spike train (see Results, Toward a general function to map spikes into LFPs). The horizontal dashed lines represent the mean estimation accuracy for the unfiltered LFPs (see Fig. 3B). The conventions are the same as in Figure 3. Error bars denote 1 SEM. The black squares correspond to the reconstruction accuracies, and the triangles indicate the chance performance levels. The LFP estimations are significantly better than the ones obtained under the null hypothesis (triangles) at all frequency bands ($p < 0.01$, based on a two-sample Kolmogorov–Smirnov test), and the 0.1–20 Hz frequency band yielded significantly higher estimation accuracies than all other frequency bands ($p < 0.01$, based on a two-sample Kolmogorov–Smirnov test).

racies from the non-causal (NC) filter used in the previous sections.

In the first case (C^+), we constructed a filter that was set to 0 for negative time lags [i.e., $h_{C^+}(\alpha) = 0$ for $\alpha < 0$]. In the second case (C^-), we constructed a filter that was set to 0 for positive time lags [i.e., $h_{C^-}(\alpha) = 0$ for $\alpha > 0$]. In both cases, after imposing the causality constraint, the estimates were computed as described above in Materials and Methods. Figure 8 shows r for the three different filters, NC, C^+ , and C^- , for 88 V1 electrodes during spontaneous activity. We found that the predictions computed with NC and C^+ were statistically indistinguishable ($p > 0.1$, two-sample Kolmogorov–Smirnov test). In contrast, the LFP estimations using

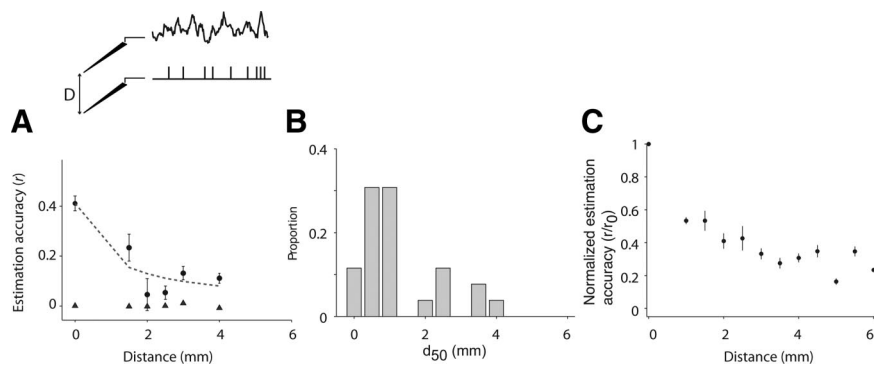


Figure 7. Dependence on LFP-spike distance. **A**, Example electrode showing the estimation accuracy as a function of the distance D between the spike train and the LFP (see scheme at the top). Here the spike train and the LFP were simultaneously recorded using different electrodes, separated by the distance given in the x -axis (distance = 0 corresponds to the estimations reported so far in which the spike train and the LFP were recorded from the same electrode). For each spike–LFP pair, we repeated the procedure described in Figure 1 *B* to compute the W–K filter and estimate its performance. Error bars denote 1 SEM. The triangles show the estimation accuracy obtained under the null hypothesis (generating a Poisson spike train with the same number of spikes). The dashed line corresponds to a fit using the following function: $r = r_0 - \frac{r_0 D}{D + D_{50}}$. **B**, Distribution of D_{50} values for the 65 electrodes in which we obtained a reliable fit ($p < 0.01$ in a χ^2 goodness-of-fit test). Bin size of 0.5 mm. **C**, We followed the procedure in **A** and normalized r by r_0 for all the V1 electrodes recorded during spontaneous activity ($n = 88$). The estimation accuracies were binned according to the spike–LFP distance using a bin size of 0.5 mm. Error bars correspond to 1 SEM. The mean estimation accuracy decreased by $\sim 50\%$ when the distance between the LFP and spike train location is 1 mm. We did not have any electrode pair with distance < 1 mm in these data.




Figure 8. Causal filters. Mean estimation accuracies (electrode-specific filters, V1, spontaneous activity) for three different filters: NC (non-causal filter, as used in Figs. 2–7), C^+ [only has information for positive time lags; $h_{C^+}(\alpha) = 0$ for $\alpha < 0$], and C^- [only has information about negative time lags; $h_{C^-}(\alpha) = 0$ for $\alpha > 0$]. The conventions are the same as in Figure 3. Error bars denote 1 SEM ($n = 88$ electrodes). The squares indicate the mean reconstruction accuracy. The triangles indicate the mean estimation accuracies under the null hypothesis (generating a Poisson spike train with the same number of spikes). All the LFP estimation accuracies shown here are statistically significant compared with the estimations obtained under the null hypothesis ($p < 0.01$, two-sample Kolmogorov–Smirnov test). The estimations obtained with NC and C^+ are statistically indistinguishable ($p > 0.1$), and there is a significant difference between the estimations obtained with either NC or C^+ , and C^- ($p < 0.01$).

the C^- filter were significantly worse than those based on either NC or C^+ (Fig. 8) ($p < 0.01$). The estimations based on C^- were still significantly better than chance ($p < 0.01$, two-sample Kolmogorov–Smirnov test).

Estimations based on multiple temporally clustered spike trains

It has long been assumed that local field potentials reflect the synchronous activity of ensembles of neurons (Mitzdorf, 1985; Fries et al., 2001; Pesaran et al., 2002; Berens et al., 2008). Based on this assumption, we hypothesized that spikes clustered in the time dimension (i.e., co-occurring within short time windows) across multiple electrodes could show an enhanced contribution to the LFP time course than isolated spikes (i.e., those spikes that do not co-occur with other spikes within short time windows). To test this hypothesis, we considered an electrode recording LFPs and all simultaneously recorded spike trains from nearby electrodes within a sphere of radius R . For a given time window τ , we separated those spike times that coincided within τ ms with spikes in any other electrode within the sphere (“time-clustered” spikes) and those spikes that occurred more than τ ms away from spikes in any other electrode within this sphere (“time-isolated” spikes). We used

the same procedure in Equations 2–4 to estimate the W–K filter and the LFP based on these two types of spike models. In Figure 9A–C, we show that, overall, the time-clustered spikes (red circles) yielded higher estimation accuracies than the time-isolated spikes for all values of R and most values of τ . In those cases in which the time-isolated spikes yielded better estimation accuracies, we noted that there was a wide discrepancy in the number of time-isolated and time-clustered spikes (e.g., $\tau = 2$ ms in Fig. 9A). Given the correlation between estimation accuracy and firing rate (supplemental Fig. S4A, available at www.jneurosci.org as supplemental material), we repeated the analysis restricting to those values of R , τ and those trials in which the number of spikes for the time-isolated spikes was within 20% of the number of spikes in the time-clustered spikes (Fig. 9D–F). We observed that time-clustered spikes yielded a significantly higher estimation accuracy for $\tau = 8$ and $\tau = 16$ ms (we could not test this hypotheses for smaller values of τ because in those cases there were always many more spikes in the time-isolated condition) (Fig. 9A–C). Overall, spikes that are temporally clustered with nearby spikes within ~ 10 ms contain more information about the LFP time course than isolated spikes.

Discussion

There has been increased interest recently in local field potentials (Fries et al., 2001; Logothetis, 2002; Pesaran et al., 2002; Mehring et al., 2003; Bédard et al., 2004; Kreiman et al., 2006; Kraskov et al., 2007; Nir et al., 2007; Belitski et al., 2008; Nauhaus et al., 2009, among others) attributable to the correlation between LFPs and functional imaging measurements (Logothetis, 2002; Kayser et al., 2004; Mukamel et al., 2005), the use of LFPs as a measure of local synchrony (Fries et al., 2001; Womelsdorf et al., 2007), and the potential use of LFPs for prosthetic applications (Pesaran et al., 2002; Mehring et al., 2003). Additionally, understanding the biophysical origins of LFPs could provide insights into the computations performed in local cortical circuits.

The biophysics underlying the generation of LFPs is not clearly understood. Although several pieces of evidence suggest

that the spatial extent of the LFPs is larger than the spatial extent of the spiking signals (Mitzdorf, 1985; Kruse and Eckhorn, 1996; Juergens et al., 1999; Logothetis, 2002; Bédard et al., 2004; Kreiman et al., 2006; Logothetis et al., 2007; Katzner et al., 2009; Nauhaus et al., 2009), the exact quantitative limits remain unclear. Additionally, current source density analyses and simultaneous recordings of LFPs and spikes have suggested that the LFPs are better understood in terms of synaptic potentials, afterpotentials, and other membrane potentials that show a slower temporal resolution than spikes themselves (Mitzdorf, 1985; Kamondi et al., 1998). Despite the large differences in spatial and temporal scales, our observations show that a linear filter operation on the spiking activity of one or a few neurons can explain a significant fraction of the variations in the LFP time course.

The shape of the W–K filter was highly stereotypical (Figs. 2, 3, 4B, 8 and supplemental Figs. S3, S7, S9, available at www.jneurosci.org as supplemental material; see also the spike-triggered average LFP in Fig. S7A, available at www.jneurosci.org as supplemental material): a sharp downstroke for negative time lags was followed by a slower upstroke for positive time lags. This shape is remarkably similar to the waveform shape in extracellular action potentials (see the experimental recordings as well as the computational simulations by Gold et al., 2006, their Fig. 1). The biophysics of the signals that give rise to the extracellular action potential have been characterized in multiple studies (Koch, 1999). For example, the simulations in the study by Gold et al. (2006) show that this prototypical shape can be generated by combining well known channels, most prominently fast inactivating Na⁺ channels and a variety of slower K⁺ channels. Although it is dangerous to extrapolate from the waveform of an extracellular action potential to the W–K filter in our study, it is tempting to speculate that part of the linear estimation that we report here might be accounted for by (weighted) linear and temporally smoothed averaging over many transmembrane potentials.

For an individual neuron, there can be strong nonlinearities in the generation of action potentials from the incoming synaptic currents and the integration of subthreshold dendritic events (Koch, 1999). The linear relationship that we report here concerns the estimation of an average over large numbers of dendritic events across many neurons. Furthermore, there is a topographical organization in V1 whereby neurons in the vicinity of the electrode registering the spiking activity may share similar properties. It is possible that this topography plays an important role in the structure of the LFP signal. If this conjecture holds, then we might expect that the relationship between spikes and LFPs may be more complex in areas that do not show such a strong topography as V1.

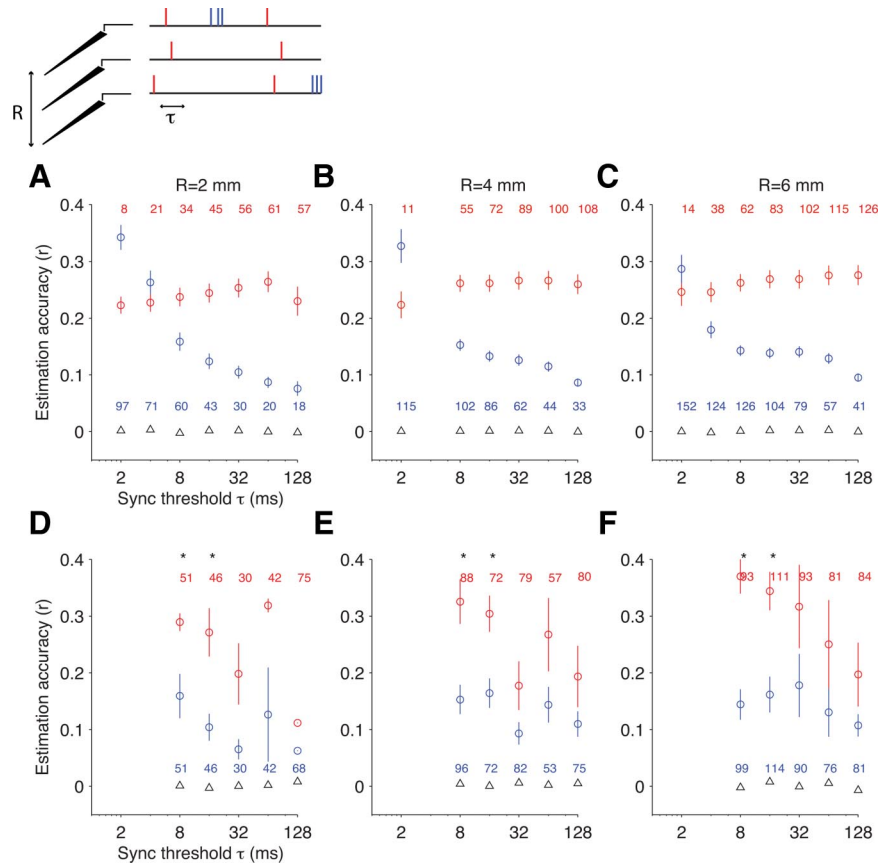


Figure 9. LFP estimation from time-clustered spikes. For each LFP recording site, we considered all the spike trains recorded within a sphere of radius R ($r = 2$ mm in **A**, **D**; $r = 4$ mm in **B**, **E**; $r = 6$ mm in **C**, **F**). For each spike, we asked whether there were spikes recorded from other electrodes within a distance R of the LFP and within τ ms (x -axis). We built a model spike train that contains all the spikes within R mm of the LFP and within τ ms of any nearby spike (temporally clustered spikes; red spikes in the scheme on top and red circles in **A–F**) and a separate model spike train that contains all the spikes within R mm of the LFP but without other spikes within τ ms (isolated spikes; blue spikes in the scheme on top and blue circles in **A–F**). These model spike trains were used to estimate the LFP using Equations 2–4. Here we report the estimation accuracies as a function of τ for different values of R for the two types of spike train models. The numbers indicate the average firing rates for the corresponding circles. The triangles indicate the estimation accuracy obtained under the null hypothesis. **D–F**, Part of the differences between temporally clustered spikes and isolated spikes in **A–C** are attributable to the different firing rates between the two models. To control for the number of spikes, in **D–F**, we only considered those values of R , τ and those trials in which the number of spikes for the red circles was within 20% of the number of spikes for the blue circles.

The local spiking activity analyzed here in the form of MUA consists of action potentials derived from multiple neurons. The exact number of neurons in the MUA is not known. However, in contrast with the LFPs, the spatial extent of the MUA decays rapidly with distance (Kreiman et al., 2006), suggesting that the MUA has a spatial extent of ~ 200 μ m and that the number of components is unlikely to be larger than hundreds of neurons. Additionally, we performed spike sorting on the MUA (supplemental Fig. S9, available at www.jneurosci.org as supplemental material). Even with the best available methods for spike sorting, we cannot guarantee that the SUA consists only of a single neuron, but the SUA signal is unlikely to consist of more than a few neurons (Lewicki, 1998). The small difference in the estimation accuracies using different spike cutoff thresholds (supplemental Fig. S9A, available at www.jneurosci.org as supplemental material) as well as the small difference in the estimation accuracies between SUA and MUA (supplemental Fig. S9B, available at www.jneurosci.org as supplemental material) further supports the idea that part of the LFP time course can be understood from highly local spiking signals. Therefore, it seems unlikely that the current

estimation accuracies based on linear filtering are a consequence of the averaging of a noisy signal over tens to hundreds of neurons in the MUA. Carrying the argument to the extreme, even a single neuron might carry significant amounts of information about the LFP, a macroscopic property of the local circuit (Katzner et al., 2009).

We chose to use a linear filter because this transformation provides the simplest possible insights into the nature of the relationship between spikes and LFPs. It is conceivable that more complex nonlinear operations can account for an even higher fraction of the LFP time course. We emphasize that there remains a significant fraction of the LFP time course that is not accounted for by our linear estimates. The values reported throughout the text correspond to the mean estimation accuracies. There are cases in which the correlation between the linear estimate and the LFP was as high as ~ 0.6 (see the distribution in Fig. 3 and supplemental Fig. S2B, available at www.jneurosci.org as supplemental material). This value is close to the best estimation accuracy of the LFP time course in one electrode from the LFP signal in a nearby electrode.

How generic is the map between spikes and LFPs? Several pieces of evidence argue that at least part of the LFP time course can be explained by a rather general linear function of the local spike trains. First, the estimation accuracies were quite stationary over time: for a given electrode, a filter computed in one trial performed quite well on a separate trial. Second, the estimations extrapolated across electrodes for a given monkey and even across monkeys (Fig. 4). Third, we also showed that a W–K filter computed using spikes and LFPs recorded from primary visual cortex could estimate the LFP time course in inferior temporal cortex using the inferior temporal cortex spikes (see Results, Spontaneous activity versus visual stimulation) (we further note that these are recordings performed in different labs, with different behavioral conditions, and using different tasks).

The generality of the linear W–K filter used here seems to break down outside of neocortex. First, the linear W–K filter did not work well when we attempted to estimate the LFP time course in LGN (supplemental Fig. S8, available at www.jneurosci.org as supplemental material) or in the human hippocampus (data not shown). It seems that, outside of neocortex, the relationship between these two signals is more complicated and might require nonlinear operations, extensive recordings from more nearby units, or building more complex models that incorporate other aspects of the circuit architecture. The LFP in the LGN during spontaneous activity may be generated or governed by modulatory inputs that are less related to the spiking of cells. Another nonexclusive possibility is that the special geometry of neocortex, with six layers and a columnar architecture, is an important component of the linear relationship between spikes and LFPs. As noted above, it is possible that topography also plays a role in the relationship between spikes and LFPs. These observations constrain future development of biophysical models of the origin of local field potential signals.

When we used causal filters (Fig. 8), we noted a significant asymmetry: the LFP estimations based on a filter that was constrained to be 0 for negative time lags (C^+) were significantly higher than those estimations based on a filter that was constrained to be 0 for positive time lags (C^-). Furthermore, the estimations based on the C^+ filter were almost as high as the ones obtained when using the non-causal filter (Fig. 8).

The possibility of linearly estimating the time course of the local field potentials from spike trains during spontaneous activity extrapolates to conditions of visual stimulation (Fig. 4A). We

also observed that the linear estimation holds in another brain area (inferior temporal cortex) under a passive viewing task in awake monkeys (see Results, Spontaneous activity versus visual stimulation) (Kreiman et al., 2006). Therefore, the basic map between spikes and LFPs holds for different brain areas (VI, ITC), stimulation conditions (spontaneous activity, visual stimulation), and behavioral states (anesthesia, awake). The estimation accuracies and, in particular, the estimation accuracies for different frequency bands may show a strong dependence on the experimental conditions. Several investigators have noted that the stimulus, task, and behavioral state of the animal (e.g., attention, short-term memory) can influence the relationship between spikes and LFPs (Fries et al., 2001; Pesaran et al., 2002; Liu and Newsome, 2005; Buschman and Miller, 2007; Womelsdorf et al., 2007; Belitski et al., 2008). In particular, there are many cases in which spikes and LFPs (or particular frequency bands of the LFPs) can be decorrelated (Kreiman et al., 2006; Belitski et al., 2008; Berens et al., 2008).

It has been assumed that LFPs may reflect synchronous activity across ensembles of local neurons (Fries et al., 2001; Laurent, 2002; Logothetis, 2002; Womelsdorf et al., 2006; Montemurro et al., 2008). Our observation that temporally clustered spikes yield better estimation accuracies than temporally isolated spikes (Fig. 9) seems to be compatible with the idea that LFPs may reflect computations that are local both in space as well as in the time domain. However, the estimation accuracies were robust to significant amounts of spike time jittering (Fig. 6). More research is necessary to further our understanding of the biophysical signals that give rise to the LFPs, how many neurons are involved, and how neuronal events across local ensembles give rise to the field potential.

As in other systems, a multiscale analysis of neuronal circuits provides rich insights that cannot be achieved by focusing on a single level only. Many efforts in neuroscience aim to understand perception and behavior using only macroscopic signals such as blood oxygenation level-dependent measurements. As emphasized previously (Logothetis, 2008), it is not always trivial to interpret these measurements without a detailed understanding of the underlying architecture and cellular function. Other efforts in neuroscience correlate perception and behavior with the activity of single neurons. Circuits of neurons may show emergent properties that are not always easy to visualize by looking at individual neurons without studying their interactions (Koch and Segev, 1989). The combination of spikes and LFPs provides an ideal resolution to understand the relationship between individual neurons and local neuronal circuits.

References

- Bédard C, Kröger H, Destexhe A (2004) Modeling extracellular field potentials and the frequency-filtering properties of extracellular space. *Biophys J* 86:1829–1842.
- Belitski A, Gretton A, Magri C, Murayama Y, Montemurro MA, Logothetis NK, Panzeri S (2008) Low-frequency local field potentials and spikes in primary visual cortex convey independent visual information. *J Neurosci* 28:5696–5709.
- Berens P, Keliris GA, Ecker AS, Logothetis NK, Tolias AS (2008) Comparing the feature selectivity of the gamma-band of the local field potential and the underlying spiking activity in primate visual cortex. *Front Syst Neurosci* 2:2.
- Bialek W, Rieke F, de Ruyter van Steveninck RR, Warland D (1991) Reading a neural code. *Science* 252:1854–1857.
- Buschman TJ, Miller EK (2007) Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science* 315:1860–1862.
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillation

- tory neuronal synchronization by selective visual attention. *Science* 291:1560–1563.
- Gabbiani F (1996) Coding of time varying signals in spike trains of linear and half-wave rectifying signals. *Network Comput Neural Syst* 7:61–85.
- Gabbiani F, Koch C (1996) Coding of time varying signals in spike trains of integrate and fire neurons with random threshold. *Neural Comput* 8:44–66.
- Gold C, Henze DA, Koch C, Buzsáki G (2006) On the origin of the extracellular action potential waveform: a modeling study. *J Neurophysiol* 95:3113–3128.
- Haberly LB, Shepherd GM (1973) Current-density analysis of summed evoked potentials in opossum prepyriform cortex. *J Neurophysiol* 36:789–802.
- Henrie JA, Shapley R (2005) LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *J Neurophysiol* 94:479–490.
- Holt GR, Koch C (1999) Electrical interactions via the extracellular potential near cell bodies. *J Comput Neurosci* 6:169–184.
- Juergens E, Guettler A, Eckhorn R (1999) Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical and EEG-potentials but not in human EEG. *Exp Brain Res* 129:247–259.
- Kamondi A, Acsády L, Buzsáki G (1998) Dendritic spikes are enhanced by cooperative network activity in the intact hippocampus. *J Neurosci* 18:3919–3928.
- Katzner S, Nauhaus I, Benucci A, Bonin V, Ringach DL, Carandini M (2009) Local origin of field potentials in visual cortex. *Neuron* 61:35–41.
- Kayser C, Kim M, Ugurbil K, Kim DS, König P (2004) A comparison of hemodynamic and neural responses in cat visual cortex using complex stimuli. *Cereb Cortex* 14:881–891.
- Koch C (1999) *Biophysics of computation*. New York: Oxford UP.
- Koch C, Segev I (1989) *Methods in neuronal modeling: from synapses to networks*. Cambridge, MA: MIT.
- Kraskov A, Quiroga RQ, Reddy L, Fried I, Koch C (2007) Local field potentials and spikes in the human medial temporal lobe are selective to image category. *J Cogn Neurosci* 19:479–492.
- Kreiman G, Krahe R, Metzner W, Koch C, Gabbiani F (2000) Robustness and variability of neuronal coding by amplitude sensitive afferents in the weakly electric fish *Eigenmannia*. *J Neurophysiol* 84:189–204.
- Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ (2006) Object selectivity of local field potentials and spikes in the inferior temporal cortex of macaque monkeys. *Neuron* 49:433–445.
- Kruse W, Eckhorn R (1996) Inhibition of sustained gamma oscillations (35–80 Hz) by fast transient responses in cat visual cortex. *Proc Natl Acad Sci U S A* 93:6112–6117.
- Laurent G (2002) Olfactory network dynamics and the coding of multidimensional signals. *Nat Rev Neurosci* 3:884–895.
- Lewicki MS (1998) A review of methods of spike sorting: the detection and classification of neural action potentials. *Network Comput Neural Syst* 9:R53–R78.
- Liu H, Agam Y, Madsen JR, Kreiman G (2008) Stationarity of the neural code for vision. *Soc Neurosci Abstr* 34:568.15.
- Liu J, Newsome WT (2005) Correlation between speed perception and neural activity in the middle temporal visual area. *J Neurosci* 25:711–722.
- Logothetis NK (2002) The neural basis of the blood-oxygen-level-dependent functional magnetic resonance imaging signal. *Philos Trans R Soc Lond B Biol Sci* 357:1003–1037.
- Logothetis NK (2008) What we can do and what we cannot do with fMRI. *Nature* 453:869–878.
- Logothetis NK, Kayser C, Oeltermann A (2007) In vivo measurements of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron* 55:809–823.
- Mehring C, Rickert J, Vaadia E, Cardoso de Oliveira S, Aertsen A, Rotter S (2003) Inference of hand movements from local field potentials in monkey motor cortex. *Nat Neurosci* 6:1253–1254.
- Milstein J, Mormann F, Fried I, Koch C (2009) Neuronal shot noise and Brownian 1/f² behavior in the local field potential. *PLoS One* 4:e4338.
- 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.
- 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.
- Mukamel R, Gelbard H, Arieli A, Hasson U, Fried I, Malach R (2005) Coupling between neuronal firing, field potentials, and fMRI in human auditory cortex. *Science* 309:951–954.
- Nauhaus I, Busse L, Carandini M, Ringach DL (2009) Stimulus contrast modulates functional connectivity in visual cortex. *Nat Neurosci* 12:70–76.
- Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A, Fried I, Malach R (2007) Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to inter-neuronal correlations. *Curr Biol* 17:1275–1285.
- O’Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- 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.
- Poor HV (1994) *An introduction to signal detection and estimation*. New York: Springer.
- Quiroga RQ, Nadasdy N, Ben-Shaul Y (2004) Unsupervised spike sorting with wavelets and superparamagnetic clustering. *Neural Comput* 16:1661–1687.
- Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK (2008) Inferring spike trains from local field potentials. *J Neurophysiol* 99:1461–1476.
- Rieke F, Warland D, van Steveninck R, Bialek W (1997) *Spikes: exploring the neural code*. Cambridge, MA: MIT.
- Wessel R, Koch C, Gabbiani F (1996) Coding of time-varying electric field amplitude modulations in a wave-type electric fish. *J Neurophysiol* 75:2280–2293.
- Womelsdorf T, Fries P, Mitra PP, Desimone R (2006) Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* 439:733–736.
- Womelsdorf T, Schoffelen JM, Oostenveld R, Singer W, Desimone R, Engel AK, Fries P (2007) Modulation of neuronal interactions through neuronal synchronization. *Science* 316:1609–1612.