J Comput Neurosci
DOI 10.1007/s10827-009-0208-9

Local field potentials indicate network state and account for neuronal response variability

Ryan C. Kelly · Matthew A. Smith · Robert E. Kass · Tai Sing Lee

Received: 28 July 2009 / Revised: 25 November 2009 / Accepted: 22 December 2009

Abstract Multineuronal recordings have revealed that neurons in primary visual cortex (V1) exhibit coordinated fluctuations of spiking activity in the absence and in the presence of visual stimulation. From the perspective of understanding a single cell's spiking activity relative to a behavior or stimulus, these network fluctuations are typically considered to be noise. We show that these events are highly correlated with another commonly recorded signal, the local field potential (LFP), and are also likely related to global network state phenomena which have been observed in a number of neural systems. Moreover, we show that attributing a component of cell firing to these network fluctuations via explicit modeling of the LFP improves the recovery of cell properties. This suggests that the impact of network fluctuations may be estimated using the LFP, and that a portion of this network activity is unrelated to the stimulus and instead reflects ongoing cortical activity. Thus, the LFP acts as an easily accessible bridge between the network state and the spiking activity.

Keywords Local field potential · Correlation · Network state · Spontaneous activity · Multielectrode array · Decoding · Population coding

Action Editor: Alain Destexhe

R. C. Kelly (✉) · M. A. Smith · R. E. Kass · T. S. Lee
Center for the Neural Basis of Cognition,
Pittsburgh, PA, USA
e-mail: rkelly@cs.cmu.edu

R. C. Kelly · T. S. Lee
Computer Science Department,
Carnegie Mellon University,
Pittsburgh, PA, USA

M. A. Smith
Department of Neuroscience,
University of Pittsburgh,
Pittsburgh, PA, USA

R. E. Kass
Department of Statistics,
Carnegie Mellon University,
Pittsburgh, PA, USA

R. E. Kass
Machine Learning Department,
Carnegie Mellon University,
Pittsburgh, PA, USA

1 Introduction

One of the most striking features of spike trains is their variability—that is, the same visual stimulus does not elicit the same spike pattern for repeated trials. This variability is often considered "noise," the connotation being that it is due to unknown factors. Identifying these factors should improve understanding and enable better characterization of neural response. One aspect of trial-to-trial fluctuation that has been examined is correlation between neurons. From neuroanatomy, we know that the visual cortex contains a rich pattern of connections. Because of these connections, the spiking activity of groups of cortical neurons is not independent—instead it is correlated on a range of time scales (Smith and Kohn 2008). The correlation among neurons has most often been measured by simultaneous recording of nearby pairs. In the retina, it has recently become possible to record from a nearly complete population of certain types of ganglion cells in a region

Published online: 22 January 2010

and identify the correlation structure of this population (Shlens et al. 2009). However, in cerebral cortex recording a full population of individual neurons in a region has been impossible, and large scale recordings *in vivo* have been rare. Correlated variability has a strong influence on population coding—it can enhance or diminish the information encoded within that population depending on a number of factors (Zohary et al. 1994; Shadlen and Newsome 1998; Abbott and Dayan 1999; Averbeck et al. 2006). Although it is possible for a signal to be encoded in the correlation strength itself (Kohn and Smith 2005; Samonds and Bonds 2005; Smith and Kohn 2008; Huang and Lisberger 2009), correlated variability is more often thought of as noise which is to be removed in order to better reveal the signal. Experimenters typically deal with this problem through multiple trial repetitions and simply averaging the neuronal responses, assuming that any noise is thereby diminished or removed.

An understanding of neuronal population activity is incomplete when measured only with spiking activity from a small subset of the total neurons in a region. Intracellular measurements of subthreshold activity have revealed that membrane potentials are continuously correlated between nearby cells in primary visual cortex (V1) (Lampl et al. 1999). In addition, the firing of neurons has been strongly linked to ongoing population activity measured with optical imaging (Tsodyks et al. 1999) and LFP (Nauhaus et al. 2009). This link to the state of the local population is an influential force affecting the variability in a cell's spiking behavior. Indeed, groups of neurons transition between "Up" (depolarized) and "Down" (hyperpolarized) states, which leads to cycles of higher and lower than normal firing rates (for review, see Destexhe and Contreras 2006). Though they are often referred to as discrete states, this is somewhat deceptive—neuronal populations transition smoothly between epochs of high and low responsivity. These state transitions occur in sleeping and anesthetized animals, as well as in cortical slices (Johnson and Buonomano 2007). Similar fluctuations in responsivity have also been reported in awake animal subjects (Leopold et al. 2003; Luczak et al. 2007) as well as awake human patients (He et al. 2008; Nir et al. 2008). Thus, a signal which reflects these states, and in turn helps account for some variability, would be an important tool for decoding a neuron's response.

The local field potential (LFP) is one such candidate signal. It is thought to reflect the average synaptic input to a region near the electrode (Mitzdorf 1987; Buzsaki 2004), and is known to be correlated with the subthreshold membrane potential fluctuations in nearby neurons (Eggermont and Smith 1995; Petersen et al. 2003). Furthermore, slow dynamics of spiking activity can be inferred from the LFP (Rasch et al. 2008), and the phase of the LFP is predictive of multi-unit activity (Haslinger et al. 2006). However, there are a number of reasons why the LFP might not be a suitable signal for this purpose. Estimates of the spatial spread of LFP range from a few hundred microns to a few millimeters (Mitzdorf 1987; Kruse and Eckhorn 1996; Kreiman et al. 2006; Liu and Newsome 2006; Berens et al. 2008; Katzner et al. 2009; Xing et al. 2009), and depend on laminae (Xing et al. 2009) and frequency (Frien et al. 2000; Henrie and Shapley 2005; Liu and Newsome 2006; Siegel and Koenig 2003). Furthermore, the feature selectivity of the LFP is broader than that of spiking activity (Kreiman et al. 2006; Liu and Newsome 2006; Berens et al. 2008; Katzner et al. 2009). Though it is clear there is a relationship between LFP, neuronal states, and spiking activity, it is not known to what extent these factors can be disentangled to better decode spike–stimulus relationships. We therefore recorded spiking activity and LFP simultaneously from a population of cells with a multielectrode array, and we used the LFP to decompose neuronal firing into stimulus-dependent and network-state dependent components. We used a generalized linear model to examine the correlation of network fluctuations with LFP and to correct the neural response to the stimulus for the effects of LFP. Generalized linear models have become a standard framework for spike train analysis (Paninski et al. 2007, 2009; Pillow 2007; Pillow et al. 2008). We found that the stimulus-dependent firing rate function reveals greater signal to noise in orientation tuning than the same measure computed via the usual mean firing rate function. An improvement in signal to noise was also present for stimulation with natural movies, demonstrating the generality of the effect. These results show that each spike's relationship to the stimulus can be estimated using merely the LFP from the same electrode—a new connection between the LFP, the ongoing network state, and spiking activity.

2 Recording methods

This experiment involved extracellular recording from microelectrode arrays implanted in V1 of three adult male cynomolgus (*Macaca fascicularis*) macaque monkeys. The techniques we use to record from the visual cortex of anesthetized, paralyzed macaques have been described in detail elsewhere (Cavanaugh et al. 2002; Smith and Kohn 2008). Briefly, we maintained anesthesia throughout the experiment by a continuous intravenous infusion of sufentanil citrate (typically

6–18 μg/kg, adjusted as needed for each animal). To minimize eye movements, the animal was paralyzed with a continuous intravenous infusion of vecuronium bromide (0.1 mg/kg/h). We monitor vital signs continuously (EEG, ECG, blood pressure, end-tidal P_{CO_2}, temperature and lung pressure). The pupils were dilated with topical atropine and the corneas protected with gas-permeable hard contact lenses. We used supplementary lenses to bring the retinal image into focus by direct ophthalmoscopy, and then adjusted the refraction further to optimize the response of recorded units. Experiments typically lasted 4–5 days. All experimental procedures complied with guidelines approved by the Albert Einstein College of Medicine of Yeshiva University Animal Welfare Committee.

We recorded neural activity with the "Utah" Array (Blackrock Microsystems, Salt Lake City, Utah) using techniques reported previously (Kelly et al. 2007). This device is a 10 × 10 grid of silicon microelectrodes (1 mm in length) spaced 400 μm apart, covering 12.96 mm^2. The details of the array insertion have been described in detail elsewhere (Smith and Kohn 2008). Briefly, we inserted the array 0.6 mm into cortex using a pneumatic insertion device (Rousche and Normann 1992), which led to recordings confined mostly to layers 2–3 of parafoveal V1 (receptive fields within 5° of the fovea). Signals from each microelectrode were amplified and bandpass filtered (250 Hz to 7.5 kHz) to acquire spiking data. Waveform segments that exceeded a threshold (set as a multiple of the rms noise on each channel) were digitized (30 kHz) and sorted off-line. We first performed a principal components analysis by waveform shape (Shoham et al. 2003) and then refined the output by hand with custom time-amplitude window discrimination software (written in MATLAB; MathWorks). This yielded 105–129 candidate neural units from the three array implants, with roughly one half estimated to be well isolated single units similar in quality to those recorded with single microelectrodes (Kelly et al. 2007). LFPs were acquired from the same electrodes simultaneously by low pass filtering the raw signal (0.3 to 250 Hz) and sampling at 1 kHz. The LFP signal was then smoothed with a 100 ms boxcar filter for all subsequent analysis (50 and 200 ms filters produced the same results).

All stimuli were generated with custom software on a Silicon Graphics Octane2 Workstation and displayed at a resolution of 1,024 × 768 pixels and frame rate of 100 Hz on a CRT monitor (stimulus intensities were linearized in luminance). Three kinds of stimuli were presented: gratings, natural movies, and a blank screen. The full-contrast sinusoidal gratings (presented in two of the three array implants) consisted of a pseudo-randomly chosen sequence of directions of drift, each lasting 300 ms and then proceeding to the next direction with no blank frames between. We used a total of 98 directions, evenly spaced, resulting in a sequence lasting roughly 30 s. Because we never repeated the same direction in a block, this 30 s stimulus comprised a unique "movie" of gratings in which no frame was repeated. This movie, with the same pseudo-random order of drift direction, was repeated 120 times. The spatial frequency (1.3 cpd) and temporal frequency (6.25 Hz) were chosen to fit the preference of parafoveal V1 neurons (DeValois et al. 1982; Foster et al. 1985; Smith et al. 2002), while the position and size (8° diameter circular aperture) of the grating were sufficient to cover the receptive fields of all the neurons. The natural movie stimulus (presented in one of the three implants) was a square 5° 30-s movie of a consumer film, repeated 120 times. The grating and movie stimuli were surrounded by a gray field of average luminance. We also recorded spontaneous activity in each of the three array implants (mean luminance approximately 40 cd/m^2), consisting of 30 min of responses measured in the presence of a blank gray screen of average luminance.

3 Correlated fluctuations in spiking and the local field potential

We simultaneously recorded single- and multiple-unit neuronal activity (sorted offline) and LFPs from 100-electrode Utah arrays implanted in V1. We measured responses to sequences of sinusoidal gratings and natural movies as well as spontaneous activity with a mean gray screen. By pairing the spiking response on each electrode with the LFPs recorded from each of 96 active electrodes, we were able to measure and model the dependence of spike trains on the LFP.

In the absence of a visual stimulus, neurons in V1 respond at a baseline rate which varies from nearly no activity to tens of spikes per second. This rate is of course not constant—neuronal activity changes in a seemingly random pattern over time. However, some of that apparent randomness can be explained when the simultaneous responses of other neurons are observed. In Fig. 1(a), the responses from 129 single and multi-unit recordings from a single array are shown in a raster plot for a 15 s period. If the neurons were responding randomly and independently, there would be no structure apparent in this plot. However, there are a number of prominent vertical bands present (for instance, one at roughly 8 s). These bands are produced by the temporally coordinated responses of a large

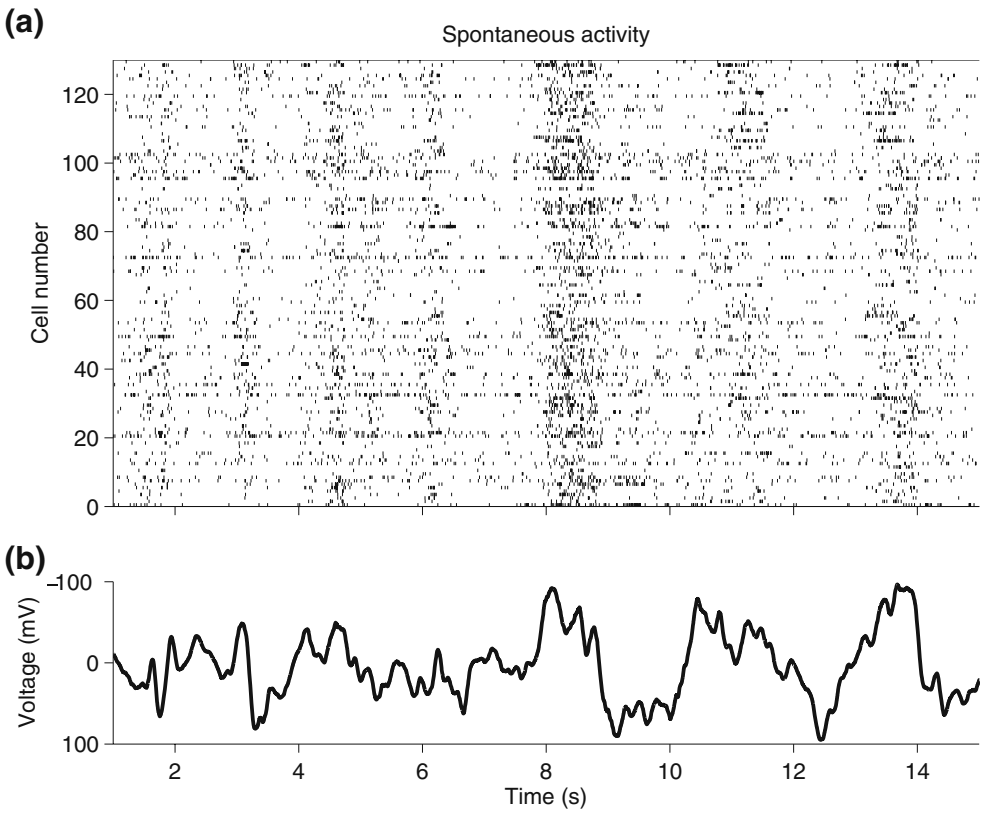

Fig. 1 The activity of a neuronal population recorded with the array. In the *upper plot*, each row depicts the firing of one cell during spontaneous activity (the room is dark). Most cells in the population exhibit correlated activity on a slow time scale. The local field potential, shown *below*, typically contains peaks during the epochs of high correlated firing. The local field potential is averaged across all electrodes in this figure, and the y-axis is flipped for illustrative purposes

portion of the neurons in the population. This kind of correlation in spiking has been reported previously in spontaneous and evoked activity in both anesthetized and awake animals (Zohary et al. 1994; Shadlen and Newsome 1998; Bair et al. 2001; Kohn and Smith 2005; Smith and Kohn 2008; Huang and Lisberger 2009).

Patterns of network activity such as these, or network "states", are well characterized in a number of systems and recording preparations (for review, see Kohn et al. 2009). One means of measuring the state of a system is through the field potentials, which are thought to represent the sum of all synaptic inputs to a region of cortex (Legatt et al. 1980; Gray et al. 1995). We compared the LFP (Fig. 1(b) recorded simultaneously) with the spontaneous spiking activity in a population of neurons. We found that the states which contained a large number of spikes ("Up", or depolarized states) tended to correspond to large negative fluctuations in the LFP. Similarly, the absence of spikes ("Down", or depolarized states) correspond to positive trends in the LFP.

Although Fig. 1 only displayed correlated fluctuations in spontaneous activity, they were also present during the presentation of a visual stimulus. Figure 2 shows two sets of 5 raster plots, each created in the same way as Fig. 1(a), for five repeated presentations of a sequence of gratings (Fig. 2(a)) or a natural movie (Fig. 2(b)). In panel (a), the sequence of gratings is identical in trials 1–5. Thus, the presence of a period of high activity on one of the trials that does not occur at the same time on all of the trials represents an "Up" state. Because the neurons prefer different orientations, there is no period in the grating sequence during which all of the neurons tended to fire in response to the visual stimulus. For this reason, the "Up" states were easy to identify. For the five identical natural movie repeats (Fig. 2(b)), stimulus-dependent and independent correlated activity were more difficult to separate. Although correlated "Up" states are present, there were also periods of high activity in the entire population corresponding to stimulus events, presumably abrupt scene transitions, that altered the response of a large group of the recorded neurons. This mingling of stimulus-driven and state-driven activity makes it difficult to decode the signals sent by the spikes from a single neuron. However, with information about the population activity, it is possible to separate these two factors. Such a separation is certain to be imperfect—it is difficult or impossible to assess the "true" meaning of an individual spike. Nonetheless, we use a quantitative measurement (orientation tuning, Section 6) which allows us to infer that the signal due to the stimulus has

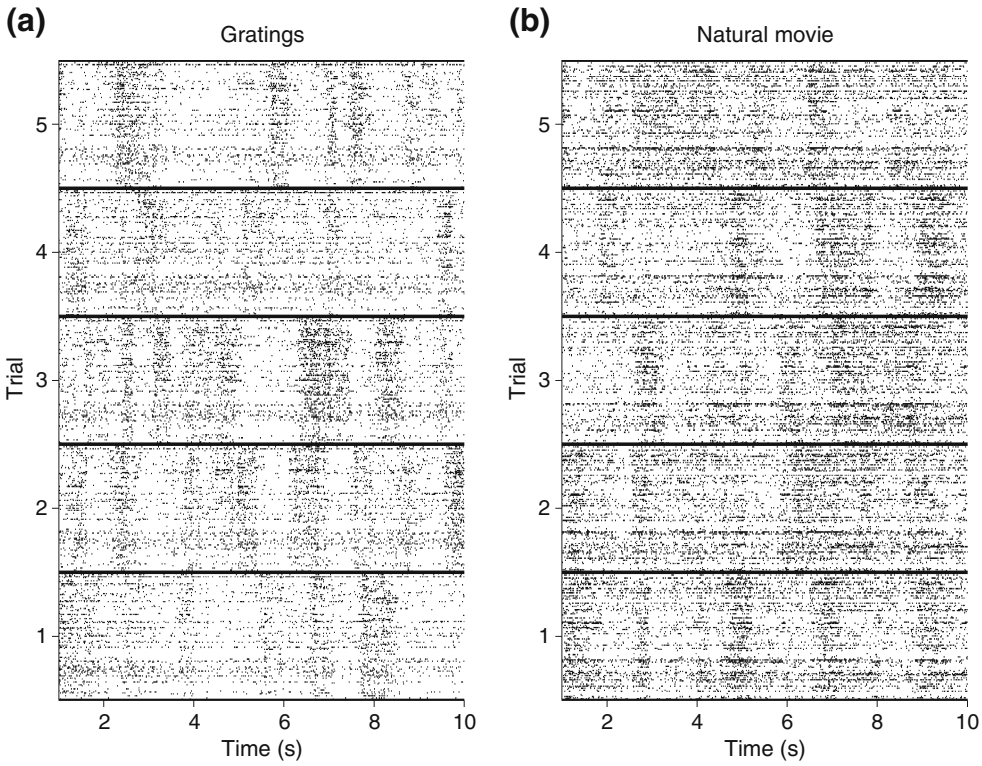

Fig. 2 Network fluctuations differ across repeated identical visual stimulation. (**a**) The stimulus is a movie of drifting gratings of different orientations, 300 ms per orientation. For each trial, the stimulus is identical. Stimulus independent correlated activity occurs differently on each trial. (**b**) The stimulus is 10 s of a single natural movie, for a different population than the one shown in (**a**). Stimulus independent correlated activity is apparent and also some stimulus-dependent activity is visible. For example, at 9 s into the stimulus a feature of the movie that drives most of the population elevates activity on every trial

been increased relative to the network "noise", utilizing the assumption that the stimulus and network components are independent components which together comprise the neural response.

4 A generalized linear model for the effects of network state

Each spike train is a sequence of spike times and the natural statistical framework for spike train analysis is that of point processes, which may take a rather general form (Paninski et al. 2009). It is simplest to assume the spike trains follow inhomogeneous Poisson processes, and for our analyses it is unlikely that more complicated models would change our conclusions. We used time bins of duration 1 ms, ensuring that each bin has at most one spike. A spike train may thereby be considered a binary sequence $Y = Y_1 \dots Y_T$, each value Y_t indicating whether or not the cell fired at time t. According to the Poisson assumption the probability of observing y_t spikes at time t is given by

$$P(Y_t = y_t) = (r_t\Delta)^{y_t} \exp(-r_t\Delta) \tag{1}$$

where r_t is a rate parameter for the process at time t and $\Delta = 1$ ms. The likelihood of observing the entire spike train $Y = y_1 \dots y_T$ is the product of the likelihoods for each time bin, where these observations are independent:

$$P(Y = y_1 \dots y_T) = \prod_t (r_t\Delta)^{y_t} \exp(-r_t\Delta) \tag{2}$$

In our context, the rate parameter r is an unobservable property of the neuron, and we proceed by fitting r so that this likelihood value is maximized. Equivalently, and for computational convenience, we can maximize the log-likelihood

$$L = \sum_t^T \left(y_t \log r_t - r_t\right) \tag{3}$$

We make the assumption that the relationship between the rate r and the observables for any given cell is a linear relationship, followed by a nonlinearity, the exponential function. Paninski (2004) gives the rationale for using this kind of model as a trade-off between neural system faithfulness and computational tractability. The observed components are a network state component and a stimulus component.

$$\log r_t^{(i)} = s_t^{(i)} + \beta^{(i)} x_t^{(i)} \tag{4}$$

where $r_t^{(i)}$ is the expected number of spikes at time t from cell i, $s_t^{(i)}$ is the stimulus-dependent component of cell i's firing rate, $x_t^{(i)}$ is an independent variable corresponding to the LFP value (taken to be an indication of the local network activity) at time t, and $\beta^{(i)}$ is a constant which controls the extent to which the network state affects the firing of cell i.

Thus, the values $s_t^{(i)}$ and $\beta^{(i)}$ which maximize the log-likelihood function

$$L^{(i)} = \sum_t^T \left[y_t^{(i)} \left(s_t^{(i)} + \beta^{(i)} x_t^{(i)} \right) - \exp\left(s_t^{(i)} + \beta^{(i)} x_t^{(i)} \right) - \log y_t^{(i)}! \right] \quad (5)$$

make up the best model of $r_t^{(i)}$ given the data. This is a convex function of $\beta^{(i)}$ and $s^{(i)}$, and we can easily fit it with iteratively reweighted least squares (IRLS), which is equivalent to the Newton–Raphson method (Hardin and Hilbe 2007) in this context.

We further made the assumption that $\beta^{(i)}$ is not dependent on the stimulus. Since $s_t^{(i)}$ was just a constant $c^{(i)}$ when the stimulus is constant (a mean gray screen), we fit $\beta^{(i)}$ with spontaneous activity data, which is the simpler formula

$$L^{(i)} = \sum_t^T \left[y_t^{(i)} \left(c^{(i)} + \beta^{(i)} x_t^{(i)} \right) - \exp\left(c^{(i)} + \beta^{(i)} x_t^{(i)} \right) - \log y_t^{(i)}! \right] \quad (6)$$

involving finding the values of only two parameters, $c^{(i)}$ and $\beta^{(i)}$ for each cell i. Larger $|\beta^{(i)}|$ indicates a greater dependence of firing on the value of the LFP.

In Fig. 3(a), the model fits are shown schematically for two example neurons. On the left, spikes are shown along with the corresponding LFP from the same electrode. The predicted firing rate function after the fitting procedure is shown on the right along with the actual spike train. The value of β varied across neurons in the population—some cells tended to be more influenced by network state than others. To give context for the values of β, a value $\beta = -0.017$ means that an LFP change of -100 mV predicts a 5.5 fold increase in firing rate. A value $\beta = 0.014$ predicts for the same LFP change, a 4 fold increase in firing rate. The raster plot in Fig. 3(b), sorted by β, demonstrates this trend.

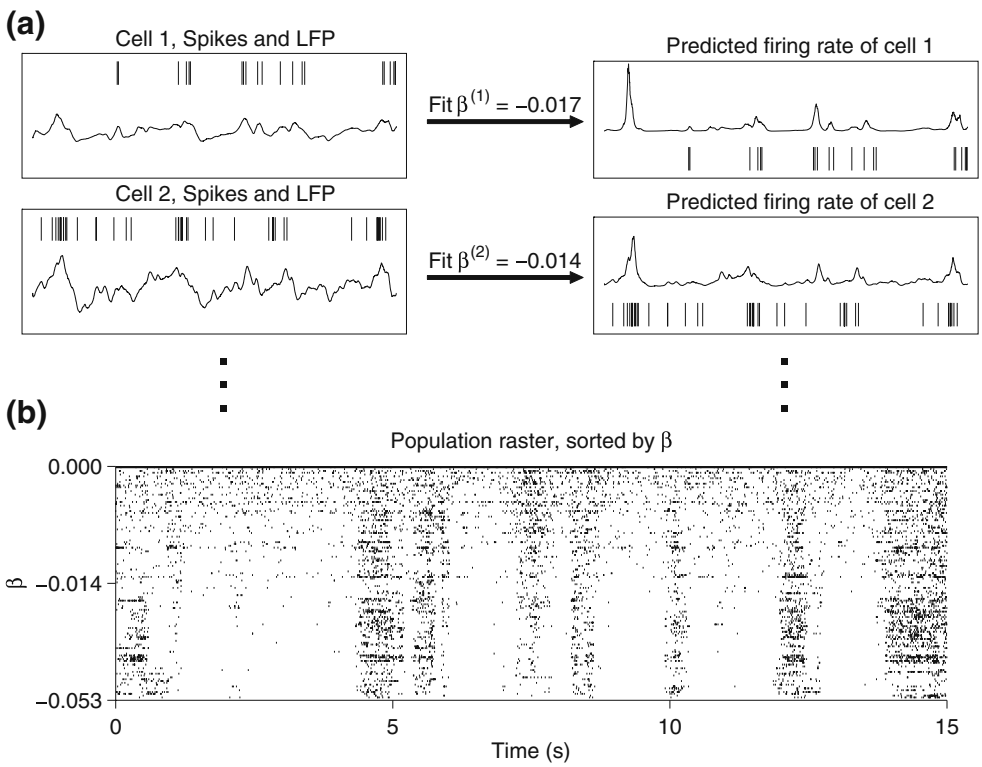

Fig. 3 Description of the model. (**a**) For each cell, $\beta^{(i)}$ is fit using the generalized linear model and spontaneous spiking data, with simultaneous LFP data. (**b**) Cells with low $|\beta^{(i)}|$ are at the *top* and generally are not driven by the large correlated events. Cells with high $|\beta^{(i)}|$ are at the *bottom* and are highly susceptible to the network state

Neurons in the upper rows of this plot tended to not be very influenced by network state, and thus had low β values. They also tended to have higher tonic firing rates ($r = 0.48$, $p < 0.0001$). Neurons in the lower rows of the plot had a strong tendency to fire during negative inflections of the LFP, and thus had highly negative β values. In the extreme, those neurons tended to fire only during "Up" states.

5 Spatial distribution of state dependence

Correlation is known to decrease as the distance between recorded neurons grows (Smith and Kohn 2008). While LFPs reflect signal from a larger region of tissue, they also tend to cohere more for nearby electrodes than distant ones (Leopold et al. 2003). If the network-driven activity we describe here is one source of correlation between spiking neurons, it is likely to also show a distance dependence. In Fig. 3, β was computed using the LFP recorded from the same electrode as the spiking activity. We tested the spatial dependence of spiking on the LFP by using LFPs recorded from each electrode in the array to compute β.

Figure 4(a) and (b) show the spatial distribution of β for two example neurons. The asterisks indicate the electrode from which spiking activity was recorded for each example. The largest negative values of β, indicated by deep blue colors, were achieved when nearby electrodes supplied the LFP. Stated another way, the spiking activity on a given electrode tends to be most influenced by the LFP of nearby electrodes. The precise spatial distribution of β values varied between cells—in some cases there was a smaller region of high beta values than others. In addition, the overwhelming trend was for β values to be negative around the recording

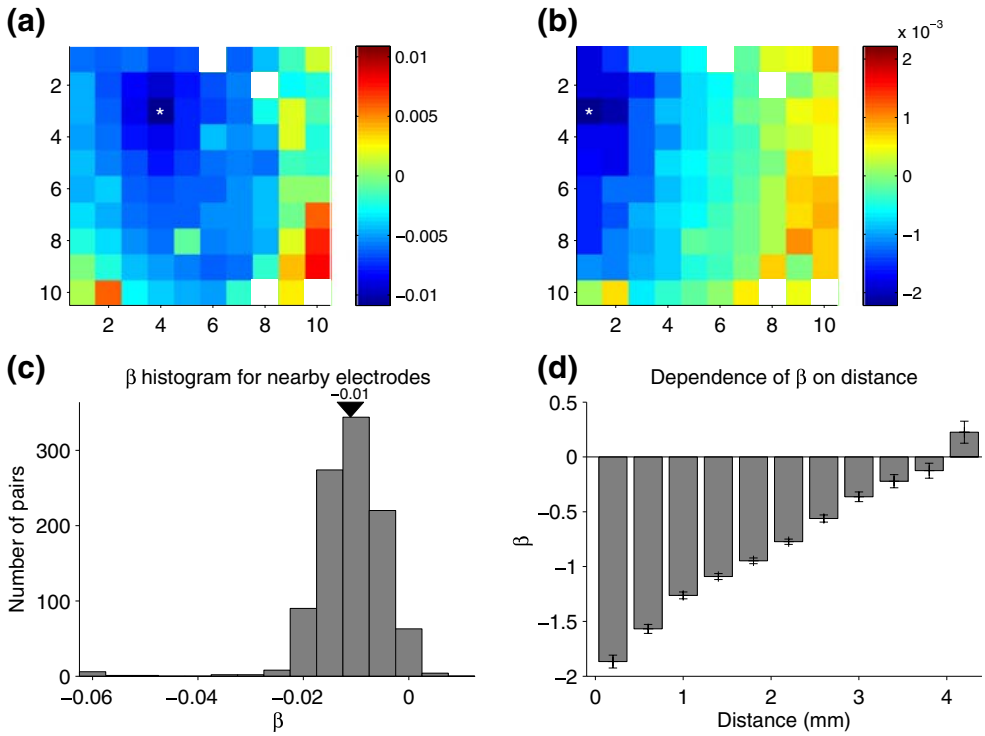

Fig. 4 Distribution of β. (**a**), (**b**) Examples of β_i for two examples cells. β_i was fit using spontaneous data. These plots depict the 10×10 grid of electrodes on the array. The two cells are recorded from the electrodes labeled with white asterisks. β_i was fit using the LFP from each electrode by itself, and the values are plotted here in the color maps. When β_i is highly negative, negative LFP values from that electrode indicate a higher firing rate of the cell. When β is positive, positive LFP values indicate a higher firing rate. The electrodes closer to the electrode recording the cell's activity tend to carry more information about the cell's firing rates. (**c**) The distribution of β values across the entire population of cells. The data shown in this plot are all possible pairings of cells and LFPs from the same or adjacent electrode. So each cell may be represented up to nine times in this histogram, depending on its location in the array. For almost all cells, β is negative for nearby electrodes. (**d**) For each cell, β was computed independently using data from each LFP electrode, and grouped into bins based on the distance between the two electrodes. The bar values are the mean values across the cells. There is a smooth falloff of β with increasing distance between the electrodes

electrode (Fig. 4(c)). This reflects the fact that spiking activity tended to occur when the LFP voltages were negative (as seen in the figures above).

To summarize its spatial dependence, we computed the average β for all pairs of spiking electrode and LFP electrode across the entire array. The results of this analysis are shown in Fig. 4(d), which plots average β as a function of the distance between the spiking and LFP electrodes. The largest negative values of β were for the same electrode. β was reduced in magnitude with increasing distance throughout the range in which we could make reliable measurements (4 mm).

6 Improvement in orientation selectivity after accounting for state dependence

The observed spiking activity of a neuron in response to a stimulus such as a grating, as we demonstrated above in Fig. 2(a), includes effects of both the stimulus and the network state. Our model factor βx_t is an attempt to assess the influence of the network state on each neuron. Once we determined the value of $\beta^{(i)}$ for each neuron i during spontaneous activity, we assumed that the relationship between spiking and the network state (LFP value $x_t^{(i)}$) remained the same during visual stimulation. We thus fit $s_t^{(i)}$ using the LFP and spiking data from repeated trials of drifting grating recordings, now holding $\beta^{(i)}$ fixed. The expression is

$$\log r_{t,n}^{(i)} = s_t^{(i)} + \beta^{(i)} x_{t,n}^{(i)} \tag{7}$$

where $r_{t,n}$ and $x_{t,n}$ are with respect to time and the trial number n, while the stimulus-dependent component $s^{(i)}$ only varies with time, since it is the same on each trial.

One caveat is that during visual stimulation across trials of the same stimulus, the LFP has a significant component of activity which is related to the stimulus. We therefore subtracted the mean LFP across repeated trials from each trial's data before computing the fixed offset term $\beta^{(i)} x_t^{(i)}$.

$$\bar{x}_t^{(i)} = \frac{1}{N} \sum_{n}^{N} x_{t,n}^{(i)} \tag{8}$$

$$\log r_{t,n}^{(i)} = s_t^{(i)} + \beta^{(i)} \left[x_{t,n}^{(i)} - \bar{x}_t^{(i)} \right] \tag{9}$$

Figure 5(a) shows some trial data for an example cell.

We solved for $s_t^{(i)}$, and $\exp(s_t^{(i)})$ is the firing rate function with the LFP effect taken into account, in spikes/ms. Figure 5(b) shows the raw peri-stimulus time histogram (PSTH), or $r_t^{(i)}$ in our model, and the LFP-adjusted firing rate function (aPSTH). The aPSTH can be thought of as the stimulus-dependent component of firing, since the component of firing that can be attributed to the LFP, or network effect, has been removed. For points in time which happened to have more network state influence, the aPSTH estimate was very different than the PSTH value, and conversely some time points had a small amount of influence. Even with a large number of trials (N = 120), network fluctuations were not evenly distributed in time, a fact that can bias any computed estimate of orientation tuning. Figure 5(c) shows the orientation tuning for this cell, computed both from the PSTH and from the aPSTH. The aPSTH estimate of orientation tuning has a more pronounced difference between preferred and non-preferred orientations, possibly reflecting the removal of action potential contributions that are unrelated to the stimulus. The orientation tuning curves are normalized by their mean firing rate for illustrative purposes.

We characterized the variability of the raw and LFP-adjusted firing by computing Fano factors for each. Specifically, for each neuron we computed the Fano factor (ratio of variance to mean) for each orientation with respect to the spike counts obtained from each repetition of that orientation. Because we found no dependence of the Fano factor on orientation, we averaged across these conditions to produce a single Fano factor for each cell. We then similarly computed the Fano factor among the LFP-adjusted firing rates. Examining the Fano factors across neurons in the two populations shows a substantial and highly significant reduction in variability ($p < 10^{-6}$, paired t-test) in the LFP-adjusted data. Figure 5(d) and (e) are histograms of the differences between the two variability measures for the two neural populations. For both populations, the LFP-adjusted component tended to have a smaller Fano factor, or more reliability, than the raw spike train. This reduction was not specific to responses evoked with grating stimuli. We performed the same analysis on the natural movie data shown in Fig. 2 and determined there is also a significant reduction in Fano factor for spikes evoked by this stimulus (mean = -0.32, $p < 10^{-6}$, paired t-test).

Figure 6(a) and (b) show the orientation tuning curves measured from two other neurons. In both of these examples, again it is apparent that the orientation selectivity improved after we accounted for network state—the peaks are higher and the troughs are lower. Note that the values on the y axis are relative to the mean firing rate across all conditions, where a value of

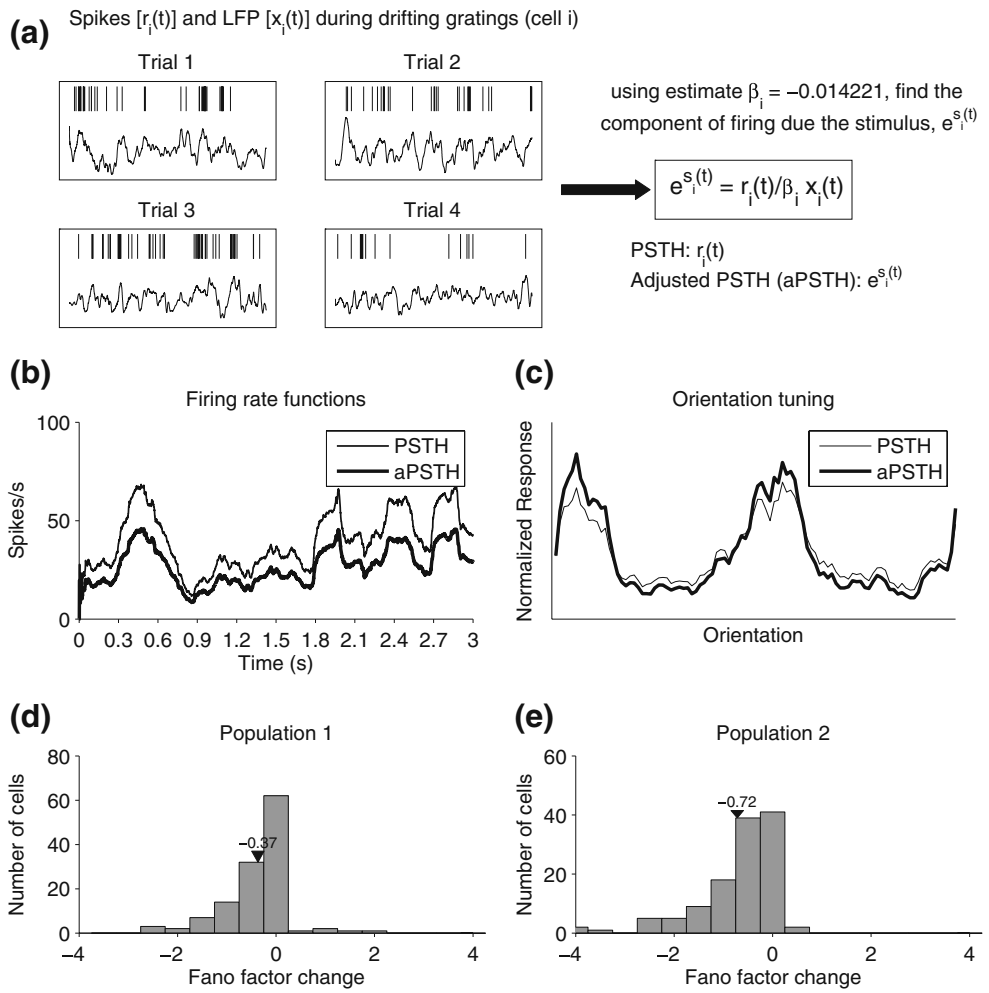

Fig. 5 Using the model to find orientation tuning with the network state removed. (**a**) Data recorded during four trials of a single grating movie. The stimulus consisted of drifting gratings of different orientations, ordered randomly. Each orientation was presented for 300 ms. Shown in the plots are the spike trains for each trial and the LFP from the same electrode. The data is used to compute $\exp(s_i(t))$, the component of firing due to the stimulus. It is essentially a LFP-weighted version of the PSTH. (**b**) The firing rate function was calculated by fitting the model for the component of firing due to the stimulus, given a fixed β_i, which were calculated from the spontaneous data as in Fig. 3. There were 120 trials in total. (**c**) The orientation tuning was calculated by averaging the firing rate function over bins corresponding to the 300 ms grating orientation stimuli epochs. (**d**), (**e**) In the two cell populations for which the grating stimulus was shown, the Fano factors were calculated for the raw spike train and also the model adjusted spike train. Most cells show a Fano factor decrease after the network state is accounted for, indicating a reduction in variability

1 is exactly the mean firing rate and a value of v is v times the mean firing rate. We evaluated this selectivity change quantitatively using vector averaging:

$$\text{Selectivity} = \frac{|\sum_j \rho_j \exp(2i\theta_j)|}{\sum_j \rho_j} \tag{10}$$

θ_j is a specific orientation presented, and ρ_j is the firing rate corresponding to that orientation. $i = \sqrt{-1}$. Selectivity is 1 when the neural firing is concentrated at one or opposite orientations (e.g. 90° and 270°). We found an average improvement of 0.04 (Fig. 6(c)), which was statistically significant ($p < 10^{-6}$, paired t-test). This improvement in selectivity was strongest when we used the LFP from the same or nearby electrodes (Fig. 6(d)).

We have used the LFP as a signal to reflect network state and account for neuronal variability, and through our model aid in decoding neuronal responses. It is also the case that the spikes themselves could serve as such a signal. Because the LFP has significant differences

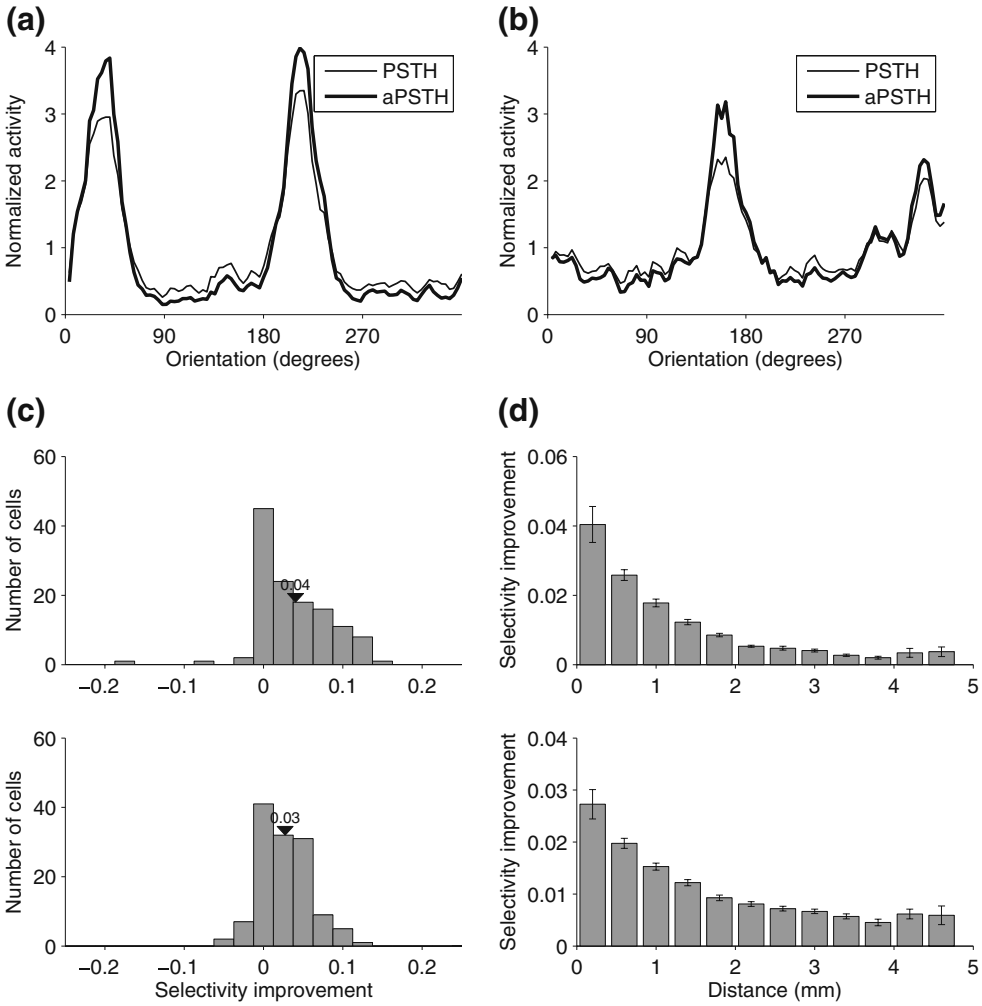

Fig. 6 Increased orientation selectivity. (**a**), (**b**) Examples for two cells of orientation tuning, computed from the raw firing rate and from the estimate c_i (the firing due to the stimulus). The estimate c_i is computed here using the LFP from the same electrode that recorded the cell's activity. (**c**) Improvement in two array populations. For most cells in the populations, selectivity was improved when accounting for the ongoing network activity with the LFP factor. The LFP was from the same electrode that recorded the cell's activity. (**d**) Falloff of improvement with distance in the same two populations. Improvement was calculated using the LFP from each electrode, and this histogram shows the average improvement for electrodes at given distances

from spiking activity in terms of spatial extent and tuning selectivity, it was not known whether it would produce better results in our model. We therefore made a direct comparison between the orientation selectivity improvement using the LFP and spike trains in the GLM. We treated each spike train as a time series, and performed smoothing in exactly the same manner as for the LFP time series data. The rest of the analysis was identical, with the smoothed spike train used instead of the LFP as the measure of network state (x_t). In order to be sure that the distances were equal and that sorting errors did not contribute to the shared variability, we used the spike trains from only the neighboring electrodes. We found that at this distance (equivalent to the second bin in Fig. 6(d)), the average improvement in selectivity when using spike trains was half the magnitude), significantly lower than what we observed using the LFP (0.013 vs. 0.026, $p < 10^{-4}$, paired t-test).

7 Discussion

We recorded spiking activity and the LFP simultaneously from a group of neurons with a microelectrode array implanted in primary visual cortex. We found that large negative fluctuations in the LFP tended to coincide with times of heightened spiking activity, or "Up" states, in neuronal firing. By using the value of the LFP as the indicator of network state, we attempted to separate neuronal responses into stimulus-dependent and state-dependent components. After adjustment for LFP effects the variability of individual neuronal responses was reduced, leading to increased SNR in orientation tuning. This finding demonstrates the value of adjusting firing rates for LFP effects, which are likely due to ongoing network activity rather than a visual stimulus.

Slow, correlated fluctuations in the firing rate of neurons were present in both spontaneous and evoked

activity. This type of correlated variability, often measured by the Pearson's correlation between the spike counts of pairs of neurons over many repeats of the same stimulus (r_{sc}), has been previously reported in a number of experiments in both awake and anesthetized animals (Zohary et al. 1994; Shadlen and Newsome 1998; Bair et al. 2001; Kohn and Smith 2005; Smith and Kohn 2008; Huang and Lisberger 2009). Nonetheless, the visual effect of such correlation on the pattern of spikes in a large neuronal population is quite striking (Fig 1(a)) and has been rarely visualized before.

The epochs of correlated activity, lasting typically 200 to 800 ms, were highly correlated with large negative peaks in the LFP (Fig 1(b)). This pattern is similar to "Up" and "Down" states, which occur with roughly the same frequency (0.3–1 Hz) and have been reported in a number of different species, cortical areas, and experimental conditions (for review, see Destexhe and Contreras 2006). While it is not clear that a common mechanism underlies all of these fluctuations in cortical responsivity, it is known that global network states are intimately related to anatomical connectivity, such as among orientation columns (Areili et al. 1996; Tsodyks et al. 1999). However, while spontaneous activity might reveal the underlying cortical network in which the neurons are embedded, it poses a problem for characterizing cell tuning properties, and for decoding stimulus information based on the spiking activity of the neurons.

The activity of neurons is affected by both the dynamical changes in these cortical states and the tuning-based response to input stimuli. Here, we proposed a simple generalized linear model technique as a first attempt to estimate the contribution of the global cortical state during responses to a visual stimulus. The resulting weighting of the spikes allows a more efficient estimate of the response properties of the neurons, and it demonstrates the relative shortage of stimulus information in the spikes which occur during "Up" states. We used the LFP as an indicator of the network state, a signal which can be obtained simultaneously with spiking activity from the same electrode. We found that variability decreased after accounting for the LFP in both stimulus conditions, gratings and natural movies, demonstrating that this result generalizes. In addition, the LFP is a better indicator of the network state than the spike trains of other neurons, and produces larger decreases in variability using our model. Fluctuations in field potentials have long been known to be related to spiking activity (Eggermont and Smith 1995; Petersen et al. 2003; Rasch et al. 2008). However, the full relationship between the LFP and spiking activity has been the subject of much recent interest and some controversy (Liu and Newsome 2006; Berens et al. 2008; Katzner et al. 2009; Xing et al. 2009), and it was not certain that a simple model such as the one presented here could successfully improve the reliability of a stimulus dependent activity measure.

The model makes a number of simplifying assumptions about the relationship between spiking behavior and the LFP, and is limited in its ability to model precise spike timing due to its having only a single parameter β derived from spontaneous activity. It is also likely that this parameter β is somewhat different during stimulation, and fitting β with data acquired during visual stimulation would improve the model. Moreover, it does not allow for modeling the way a visual stimulus could be interacting with the network state—the components of the model are independent. It is known that neuronal response properties, such as receptive field structure (Ringach et al. 2002; David et al. 2004; Körding et al. 2004), can change depending on the visual stimulus with which they are measured. In our case, we can't be sure that the state-dependence measured from spontaneous activity is the best predictor of state-dependence to gratings, or to other visual stimuli.

Nevertheless, this simple model shows that the LFP can indicate the state of the network, and it takes a significant step toward a true factorization of a spike train into stimulus-dependent and independent components. Our results demonstrate that it is possible to harness the information in the LFP to explicitly decouple these components to a significant extent. The model may be enhanced by incorporating additional effects, including supplementary indicators of the network state such as the full grid of LFP electrode responses, variables identifying high-frequency components of LFP, history effects, effects of other cells, or nonlinear interactions. All such effects could be incorporated as in other applications of generalized linear models (Kass and Ventura 2001; Pillow 2007; Pillow and Latham 2008; Pillow et al. 2008; Paninski et al. 2009).

In this paper, we have laid out a procedure that is generally applicable to neural data which is subject to network state effects on action potential generation. While it was necessary for us to record many neurons and field potentials simultaneously in order to directly observe the population activity fluctuations, our basic finding is a general one that is relevant to single unit recording data. It can be implemented with spikes and LFP from even a single electrode, and does not require a large array. Despite the simplicity of the model, we have shown that using this procedure essentially boosts the signal to noise ratio of the orientation tuning provided by individual neurons. Thus, our findings are a powerful demonstration of the way in which field

potential measurements relate to spiking activity. This kind of modeling can be used to compute a stimulus-dependent firing rate estimate in preparation for other computations as well, from receptive field estimation to latency calculations. It is especially appropriate for data with a limited number of trials, which is subject to large trial by trial variations in the effect of the network state. Providing a better model of these spontaneous states and their transitions, as well as further elucidating their origin and the connection with the underlying network, is an important goal of future experimental and theoretical research.

Acknowledgements This work was supported by a National Science Foundation (NSF) Integrative Graduate Education and Research Traineeship to RCK (DGE-0549352), National Eye Institute (NEI) grants EY015958 and EY018894 to MAS, National Institute of Mental Health (NIMH) Grant MH64445 and NSF CISE IIS 0713206 to TSL, and NIMH grant MH064537 to REK. Data was collected by RCK, MAS and Adam Kohn in his laboratory as a part of a collaborative effort between the Kohn laboratory at Albert Einstein College of Medicine and the Lee laboratory at Carnegie Mellon University. We thank Adam Kohn for collaboration, and we are also grateful to Amin Zandvakili, Xiaoxuan Jia and Stephanie Wissig for assistance in data collection.

References

Abbott, L. F., & Dayan, P. (1999). The effect of correlated variability on the accuracy of a population code. *Neural Computation, 11*, 91–101.

Areili, A., Sterkin, A., Grinvald, A., & Aertson, A. (1996). Dynamics of ongoing activity: Explanation of the large variability in evoked cortical responses. *Science, 273*(5283), 1868–1871.

Averbeck, B. B., Latham, P. E., & Pouget, A. P. (2006). Neural correlations, population coding and computation. *Nature Reviews. Neuroscience, 7*, 358–366.

Bair, W., Zohary, E., & Newsome, W. T. (2001). Correlated firing in macaque visual area MT: Time scales and relationship to behavior. *Journal of Neuroscience, 21*, 1676–1697.

Berens, P., Keliris, G., Ecker, A., Logothetis, N., & Tolias, A. (2008). Feature selectivity of the gamma-band of the local field potential in primate primary visual cortex. *Frontiers in Neuroscience, 2*, 199–207.

Buzsaki, G. (2004). Large-scale recording of neuronal ensembles. *Nature Neuroscience, 7*, 446–451.

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002). Nature and interaction of signals from the receptive field center and surround in macaque V1 neurons. *Journal of Neurophysiology, 88*, 2530–2546.

David, S., Vinje, W., & Gallant, J. (2004). Natural stimulus statistics alter the receptive field structure of V1 neurons. *Journal of Neuroscience, 24*, 6991–7006.

Destexhe, A., & Contreras, D. (2006). Neuronal computations with stochastic network states. *Science, 314*, 85–90.

DeValois, R. L., Albrecht, D. G., & Thorell, L. G. (1982). Spatial frequency selectivity of cells in macaque visual cortex. *Vision Research, 22*, 545–559.

Eggermont, J., & Smith, G. (1995). Synchrony between single-unit activity and local field potentials in relation to periodicity coding in primary auditory cortex. *Journal of Neurophysiology, 73*, 227–245.

Foster, K. H., Gaska, J. P., Nagler, M., & Pollen, D. A. (1985). Spatial and temporal frequency selectivity of neurones in visual cortical areas V1 and V2 of the Macaque monkey. *Journal of Physiology, 365*, 331–363.

Frien, A., Eckhorn, R., Bauer, R., Woelbern, T., & Gabriel, A. (2000). Fast oscillations display sharper orientation tuning than slower components of the same recordings in striate cortex of the awake monkey. *European Journal of Neuroscience, 12*, 1453–1465.

Gray, C. M., Maldonado, P. E., Wilson, M., & McNaughton, B. (1995). Tetrodes markedly improve the reliability and yield of multiple single-unit isolation from multi-unit recordings in cat striate cortex. *Journal of Neuroscience Methods, 63*, 43–54.

Hardin, J. W., & Hilbe, J. (2007). *Generalized linear models and extensions*. College Station: Stata.

Haslinger, R., Ulbert, I., Moore, C., Brown, E., & Devor, A. (2006). Analysis of LFP phase predicts sensort response of barrel cortex. *Journal of Neurophysiology, 96*, 1658–1663.

He, B., Snyder, A., Zempel, J., Smyth, M., & Raichle, M. (2008). Electrophysiological correlates of the brains intrinsic large-scale functional architecture. *Proceedings of the National Academy of Sciences of the United States of America, 105*, 16039–16044.

Henrie, J., & Shapley, R. (2005). LFP power spectra in V1 cortex: The graded effect of stimulus contrast. *Journal of Neurophysiology, 94*, 479–490.

Huang, X., & Lisberger, S. (2009). Noise correlations in cortical area MT and their potential impact on trial-by-trial variation in the direction and speed of smooth pursuit eye movements. *Journal of Neurophysiology, 101*, 3012–3030.

Johnson, H., & Buonomano, D. (2007). Development and plasticity of spontaneous activity and up states in cortical organotypic slices. *Journal of Neuroscience, 27*(22), 5915–5925.

Kass, R., & Ventura, V. (2001). A spike-train probability model. *Neural Computation, 13*, 1713–1720.

Katzner, S., Nauhaus, I., Benucci, A., Bonin, V., Ringach, D., & Carandini, M. (2009). Local origin of field potentials in visual cortex. *Neuron, 61*, 35–41.

Kelly, R. C., Smith, M. A., Samonds, J. M., Kohn, A., Bonds, A. B., Movshon, J. A., et al. (2007). Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. *Journal of Neuroscience, 27*, 261–264.

Kohn, A., & Smith, M. A. (2005). Stimulus dependence of neuronal correlation in primary visual cortex of the Macaque. *Journal of Neuroscience, 25*, 3661–3673.

Kohn, A., Zandvakili, A., & Smith, M. A. (2009). Correlations and brain states: From electrophysiology to functional imaging. *Current Opinion in Neurobiology, 19*, 434–438.

Körding, K., Kayser, C., Einhäuser, W., & König, P. (2004). How are complex cell properties adapted to the statistics of natural stimuli? *Journal of Neurophysiology, 91*, 206–212.

Kreiman, G., Hung, C., Kraskov, A., Quiroga, R., Poggio, T., & DiCarlo, J. (2006). Object selectivity of local field potentials and spikes in the Macaque inferior temporal cortex. *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. *Proceedings of the National Academy of Sciences, 93*, 6112–6117.

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

Legatt, A. D., Arezzo, J., & Vaughan, H. G. (1980). Averaged multiple unit activity as an estimate of phasic changes in local neuronal activity: Effects of volume-conducted potentials. *Journal of Neuroscience Methods, 2*, 203–217.

Leopold, D. A., Murayama, Y., & Logothetis, N. K. (2003). Very slow activity fluctuations in monkey visual cortex: Implications for functional brain imaging. *Cerebral Cortex, 13*, 422–433.

Liu, J., & Newsome, W. (2006). Local field potential in cortical area MT: Stimulus tuning and behavioral correlations. *Journal of Neuroscience, 26*, 7779–7790.

Luczak, A., Bartho, P., Marguet, S., Buzsaki, G., & Harris, K. (2007). Sequential structure of neocortical spontaneous activity *in vivo*. *Proceedings of the National Academy of Sciences of the United States of America, 104*, 347–352.

Mitzdorf, U. (1987). Properties of the evoked potential generators: Current source-density analysis of visually evoked potentials in the cat cortex. *International Journal of Neuroscience, 33*, 33–59.

Nauhaus, I., Busse, L., Carandini, M., & D.L., R. (2009). Stimulus contrast modulates functional connectivity in visual cortex. *Nature Neuroscience, 12*, 70–76.

Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., Fisch, L., et al. (2008). Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature Neuroscience, 11*(9), 1100–1108.

Paninski, L. (2004). Maximum likelihood estimation of cascade point-process encoding models. *Network: Computation in Neural Systems, 15*, 243–262.

Paninski, L., Brown, E., Iyengar, S., & Kass, R. (2009). Statistical models of spike trains. In C. Liang, & G. Lord (Eds.), *Stochastic methods in neuroscience* (pp. 278–303). Oxford: Clarendon.

Paninski, L., Pillow, J., & Lewi, J. (2007). Statistical models for neural encoding, decoding, and optimal stimulus design. *Progress in Brain Research, 165*, 493.

Petersen, C., Grinvald, A., & Sakmann, B. (2003). Spatiotemporal dynamics of sensory responses in layer 2/3 of rat barrel cortex measured *in vivo* by voltage-sensitive dye imaging combined with whole-cell recordings and neuron reconstructions. *Journal of Neuroscience, 23*, 1298–1309.

Pillow, J. (2007). Likelihood-based approaches to modeling the neural code. In K. Doya, S. Ishii, A. Pouget, & R. Rao, (Eds.), *Bayesian brain: Probabilistic approaches to neural coding* (pp. 53–70). Cambridge: MIT.

Pillow, J., & Latham, P. (2008). Neural characterization in partially observed populations of spiking neurons. *Advances in Neural Information Processing Systems, 20*, 1161–1168.

Pillow, J., Shlens, J., Paninski, L., Sher, A., Litke, A., Chichilnisky, E., et al. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature, 454*, 995–999.

Rasch, M., Gretton, A., Murayama, Y., Maass, W., & Logothetis, N. (2008). Inferring spike trains from local field potentials. *Journal of Neurophysiology, 99*, 1461–1476.

Ringach, D., Hawken, M., & Shapley, R. (2002). Receptive field structure of neurons in monkey primary visual cortex revealed by stimulation with natural image sequences. *Journal of Visualization, 2*, 12–24.

Rousche, P. J., & Normann, R. A. (1992). A method for pneumatically inserting an array of penetrating electrodes into cortical tissue. *Annals of Biomedical Engineering, 20*, 413–422.

Samonds, J. M., & Bonds, A. B. (2005). Gamma oscillation maintains stimulus structure-dependent synchronization in cat visual cortex. *Journal of Neurophysiology, 93*, 223–236.

Shadlen, M. N., & Newsome, W. T. (1998). The variable discharge of cortical neurons: Implications for connectivity, computation, and information coding. *Journal of Neuroscience, 18*, 3870–3896.

Shlens, J., Field, G., Gauthier, J., Greschner, M., Sher, A., Litke, A., & Chichilnisky, E. (2009). The structure of large-scale synchronized firing in primate retina. *Journal of Neuroscience, 29*, 5022–5031.

Shoham, S., Fellows, M., & Normann, R. (2003). Robust, automatic spike sorting using mixtures of multivariate t-distributions. *Journal of Neuroscience Methods, 127*, 111–122.

Siegel, M., & Koenig, P. (2003). A functional gamma-band defined by stimulus-dependent synchronization in area 18 of awake behaving cats. *Journal of Neuroscience, 23*, 4251–4260.

Smith, M. A., Bair, W., & Movshon, J. A. (2002). Signals in macaque V1 neurons that support the perception of Glass patterns. *Journal of Neuroscience, 22*, 8334–8345.

Smith, M. A., & Kohn, A. (2008). Spatial and temporal scales of neuronal correlation in primary visual cortex. *Journal of Neuroscience, 28*, 12591–12603.

Tsodyks, M., Kenet, T., Grinvald, A., & Arieli, A. (1999). Linking spontaneous activity of single cortical neurons and the underlying functional architecture. *Science, 286*(5446), 1943–1946.

Xing, D., Yeh, C., & Shapley, R. (2009). Spatial spread of the local field potential and its laminar variation in visual cortex. *Journal of Neuroscience, 29*, 11540–11549.

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature, 370*, 140–143.