

Inferring Cortical Variability from Local Field Potentials

Yuwei Cui,¹  Liu D. Liu,²  James M. McFarland,¹ Christopher C. Pack,² and Daniel A. Butts¹

¹Department of Biology and Program in Neuroscience and Cognitive Science, University of Maryland, College Park, Maryland 20815, and ²Montréal Neurological Institute, McGill University, Montréal, Quebec H3A 2B4, Canada

The responses of sensory neurons can be quite different to repeated presentations of the same stimulus. Here, we demonstrate a direct link between the trial-to-trial variability of cortical neuron responses and network activity that is reflected in local field potentials (LFPs). Spikes and LFPs were recorded with a multielectrode array from the middle temporal (MT) area of the visual cortex of macaques during the presentation of continuous optic flow stimuli. A maximum likelihood-based modeling framework was used to predict single-neuron spiking responses using the stimulus, the LFPs, and the activity of other recorded neurons. MT neuron responses were strongly linked to gamma oscillations (maximum at 40 Hz) as well as to lower-frequency delta oscillations (1–4 Hz), with consistent phase preferences across neurons. The predicted modulation associated with the LFP was largely complementary to that driven by visual stimulation, as well as the activity of other neurons, and accounted for nearly half of the trial-to-trial variability in the spiking responses. Moreover, the LFP model predictions accurately captured the temporal structure of noise correlations between pairs of simultaneously recorded neurons, and explained the variation in correlation magnitudes observed across the population. These results therefore identify signatures of network activity related to the variability of cortical neuron responses, and suggest their central role in sensory cortical function.

Key words: correlations; LFP; model; population; prediction; variability

Significance Statement

The function of sensory neurons is nearly always cast in terms of representing sensory stimuli. However, recordings from visual cortex in awake animals show that a large fraction of neural activity is not predictable from the stimulus. We show that this variability is predictable given the simultaneously recorded measures of network activity, local field potentials. A model that combines elements of these signals with the stimulus processing of the neuron can predict neural responses dramatically better than current models, and can predict the structure of correlations across the cortical population. In identifying ways to understand stimulus processing in the context of ongoing network activity, this work thus provides a foundation to understand the role of sensory cortex in combining sensory and cognitive variables.

Introduction

Sensory neuron responses to multiple presentations of identical stimuli can be highly variable (Tolhurst et al., 1983; Softky and Koch, 1993; Faisal et al., 2008; Masquelier, 2013). While some of this variability may simply reflect noise that originates at early stages of processing (Zohary et al., 1994; Faisal et al., 2008), an-

other potential source of variability is coordinated patterns of ongoing activity in the cortical network (Arieli et al., 1996). Such patterns are often referred to as the “cortical state,” which has in turn been linked to a variety of cognitive and behavioral processes (Engel et al., 2001; Cohen and Newsome, 2008; Morishima et al., 2009), including working memory (Mendoza-Halliday et al., 2014), the allocation of attention (Saalmann et al., 2007), and the general state of arousal (Constantinople and Bruno, 2011). Thus, rather than reflecting neuronal noise, trial-to-trial variability may represent the influence of the cortical state, which is of clear importance for sensory cortical function.

A key challenge in uncovering the role of the cortical state in sensory function is being able to relate cortical state to concrete, experimentally observable quantities. Cortical state might be directly ascribed to changes in sensory processing, and thus inferred as “hidden variables” (Niell and Stryker, 2010; Goris et al., 2014; Pachitariu et al., 2015; Rabinowitz et al., 2015). However, the dynamics of cortical activity often can be too complex to reduce to a few variables—particularly in the awake animal (Tan

Received June 30, 2015; revised Jan. 15, 2016; accepted Feb. 22, 2016.

Author contributions: Y.C., C.C.P., and D.A.B. designed research; Y.C. and L.D.L. performed research; J.M.M. and D.A.B. contributed unpublished reagents/analytic tools; Y.C. analyzed data; Y.C. and D.A.B. wrote the paper.

This work was supported by National Science Foundation Grants IIS-0904430/CNS-103331 (Y.C., J.M.M., C.C.P., D.A.B.) and IIS-1350990 (D.A.B.); Canadian Institutes of Health Research Grant MOP-115178 (C.C.P.); National Institutes of Health Grant F32-EY-023921 (J.M.M.); a Canadian Graduate Scholarship 121719 (L.D.L.); and a William Hodos Dissertation Research Assistantship (Y.C.). We thank Drs. Theodoros Zanos and Curtis Baker for helpful comments on the manuscript.

The authors declare no competing financial interests.

Correspondence should be addressed to Daniel A. Butts, Department of Biology, 1210 Biology–Psychology Building #144, University of Maryland, College Park, MD 20815. E-mail: dab@umd.edu.

DOI:10.1523/JNEUROSCI.2502-15.2016

Copyright © 2016 the authors 0270-6474/16/364121-15\$15.00/0

et al., 2014; Pachitariu et al., 2015). Alternatively, multielectrode recordings provide signals that could be used to infer the more complex network variables predictive of trial-to-trial variability. For example, simultaneously recorded neurons often have shared variability, or “noise correlations” (Zohary et al., 1994; Cohen and Kohn, 2011), which can be used in predicting the activity of single neurons (Pillow et al., 2008; Ohiorhenuan et al., 2010; Okun et al., 2015; Schölvinck et al., 2015). Relevant elements of network activity might also be reflected in lower-frequency fluctuations in the extracellular voltage known as local field potentials (LFPs). LFPs are thought to largely reflect synchronous synaptic inputs to a given area, thus potentially containing information about the specific aspects of network activity related to inputs into a given area (Buzsáki and Draguhn, 2004; Khawaja et al., 2009; Buzsáki et al., 2012; Einevoll et al., 2013). Likewise, LFPs recorded in multiple brain areas are often related to behavioral or cognitive variables (Fries et al., 2001; Lee et al., 2005; Zanos et al., 2015).

Multielectrode recordings thus provide a large number of possible signals that might be correlated with cortical state and its influence on sensory processing. However, many of these different signals will exhibit intertwined correlations that have typically been studied separately, such as between neural activity and stimulus, between LFP and stimulus (Belitski et al., 2008; Montemurro et al., 2008), between different bands of the LFP (Jensen and Colgin, 2007; Canolty and Knight, 2010), between LFPs across depth (Xing et al., 2009), and between LFP and neural activity (Rasch et al., 2008). Thus, determining aspects of recorded activity that are related to trial-to-trial variability, and by extension the modulation of stimulus processing in the cortex, requires adopting an integrated analysis approach, which goes beyond looking at these different relationships in isolation.

Here, we describe a structured statistical model that predicts neural activity recorded in the middle temporal (MT) area using the visual stimulus, LFPs, and multiunit activity (MUA) recorded simultaneously on other electrodes. This modeling approach naturally accounts for the different correlations among stimulus, MUA, and LFP signals, and thus is able to find the combination of factors that yield the best predictions of the observed MT neuron spikes. We found that information in the LFP can offer dramatic improvements (on average, fivefold) in predicting MT neuron responses over using the stimulus alone. MT neurons have very consistent selectivity to two frequency bands in the LFP, gamma (30–70 Hz) and delta (1–4 Hz), with consistent phase preferences across the population. By comparison, the use of simultaneously recorded neurons to infer cortical variability offered only twofold improvements over a model that used the stimulus alone. The inferred relationships of neural responses to ongoing network activity can also predict the form of noise correlations among MT neurons. These results thus provide a means to characterize cortical neuron activity in terms of a combination of stimulus processing and network modulation evident in LFPs, and more generally suggest a central role of ongoing network activity in modulating sensory neuron function.

Materials and Methods

Electrophysiology. Data were recorded from two female adult rhesus macaque monkeys, prepared using standard surgical techniques that have been described previously (Mineault et al., 2012). Eye movements were monitored at 500 Hz by an infrared eye tracker (EyeLink II; SR Research). Extracellular recordings were performed on 108 well isolated single units during a passive fixation task, of which 93 units were recorded using a multisite linear electrode array (U-probe; Plexon) with 16 or 24 record-

ing sites separated by 100 μm . An additional 15 units recorded with single electrodes contributed to the measurements of stimulus-locked response power (Fig. 1). Signals were amplified, bandpass filtered, sorted on-line, and resorted off-line, using spike-sorting software (Plexon) to isolate single units. We recorded LFPs using dedicated custom hardware that had wideband analog filters (two-pole high-frequency cutoff at 2.5 kHz and low-frequency cutoff of 0.7 Hz) and a sampling rate of 10 kHz. Recording broadband signals allowed us to remove spurious correlations between spikes and LFPs using a Bayesian spike removal algorithm (Zanos et al., 2011). All aspects of the experiments were approved by the Animal Care Committee of the Montreal Neurological Institute and were conducted in compliance with regulations established by the Canadian Council on Animal Care.

Stimulus presentation. Animals were trained to fixate on a small fixation point during the presentation of a continuously evolving visual stimulus. The animals were given a liquid reward periodically for maintaining fixation within 2° of the fixation point. Stimuli consisted of continuously varying optic flow (Fig. 1A) (Mineault et al., 2012; Cui et al., 2013). Briefly, the stimulus was composed of moving dots whose velocity field was generated as a random combination of the following six optic flow components: horizontal/vertical translation, expansion, rotation, and horizontal/vertical shears. The stimulus was displayed in a slowly moving aperture with a diameter ranging from 8° to 20° , the position of which slowly evolved, following a Gaussian noise distribution that was low-pass filtered in time with a cutoff of 0.05–0.10 Hz, centered on the receptive fields of the recorded neurons. All stimuli were presented on a CRT monitor with a display resolution of 1600×1000 pixels ($49^\circ \times 36^\circ$ of visual field at a distance of 50 cm) and a 60 Hz frame rate.

All neurons were presented with a long continuous stimulus lasting between 10 and 36 min (mean = 18 min), as well as a short segment of the stimulus (5 s) that was centered on the receptive field of the cell and repeated between 60 and 240 times (mean, 96 trials). The repeats were presented continuously (with no gap in between successive trials) to minimize transients due to onset responses. The long continuous stimulus presentations were used to estimate model parameters, and model performance was evaluated on the repeat trials (i.e., for cross-validation). In addition, we used these repeated presentations to measure the response reliability and to distinguish between stimulus-locked and trial-variable elements of the response.

For both continuous trials and repeated trials, we excluded all data associated with lapses in fixation (when the gaze location of the animal deviated by $>1^\circ$ from the fixation point), from 100 ms before the lapse to 500 ms after the recovery of fixation. Only periods with fixations that were longer than 1 s were used for model fitting and evaluation, resulting in an average of 10.1 ± 3.2 min of usable data for each unit.

Estimating the amount of firing rate variance. While it is clear that MT neuron spike trains are not consistent across repeated presentations of the stimulus (Fig. 1C), there is no established method to estimate exactly how much neural activity is predictable from the stimulus alone versus “trial-variable” elements, and thus attributable to other network influences. Specifically, if the firing rate on the i th trial is represented by $\lambda^{(i)}(t)$, where t is the time relative to the beginning of each trial, the law of total variance gives the following:

$$\text{Var}_i[\lambda^{(i)}(t)] = \text{Var}_i[E_i[\lambda^{(i)}(t)]] + E_i[\text{Var}_i[\lambda^{(i)}(t)]], \quad (1)$$

where E represents the “expectation” (mean with infinite data) over the indexed variable, and thus $E_i[\lambda^{(i)}(t)] = \lambda_{\text{psth}}(t)$ is the trial-averaged firing rate [i.e., peristimulus time histogram (PSTH)]. Thus, the total firing rate variance (left side of Eq. 1) is the sum of the stimulus-locked firing rate variance and trial-variable variance.

The challenge is that the spike count $n^{(i)}(t)$, rather than the firing rate $\lambda^{(i)}(t)$, is observed on individual trials. Therefore, assumptions are necessary to estimate how much of the observed variance in the spike count is attributable to firing rate variance (Amarasingham et al., 2015); i.e., estimating $\text{Var}_i[\lambda^{(i)}(t)]$ (rightmost term in Eq. 1) from the data. Because the variance of the stimulus-locked firing rate $\text{Var}_i[\lambda_{\text{psth}}(t)]$ (the other term in Eq. 1) is the variance of the PSTH (which we adjust for finite-sample bias using Sahani and Linden, 2003), summed together this gives

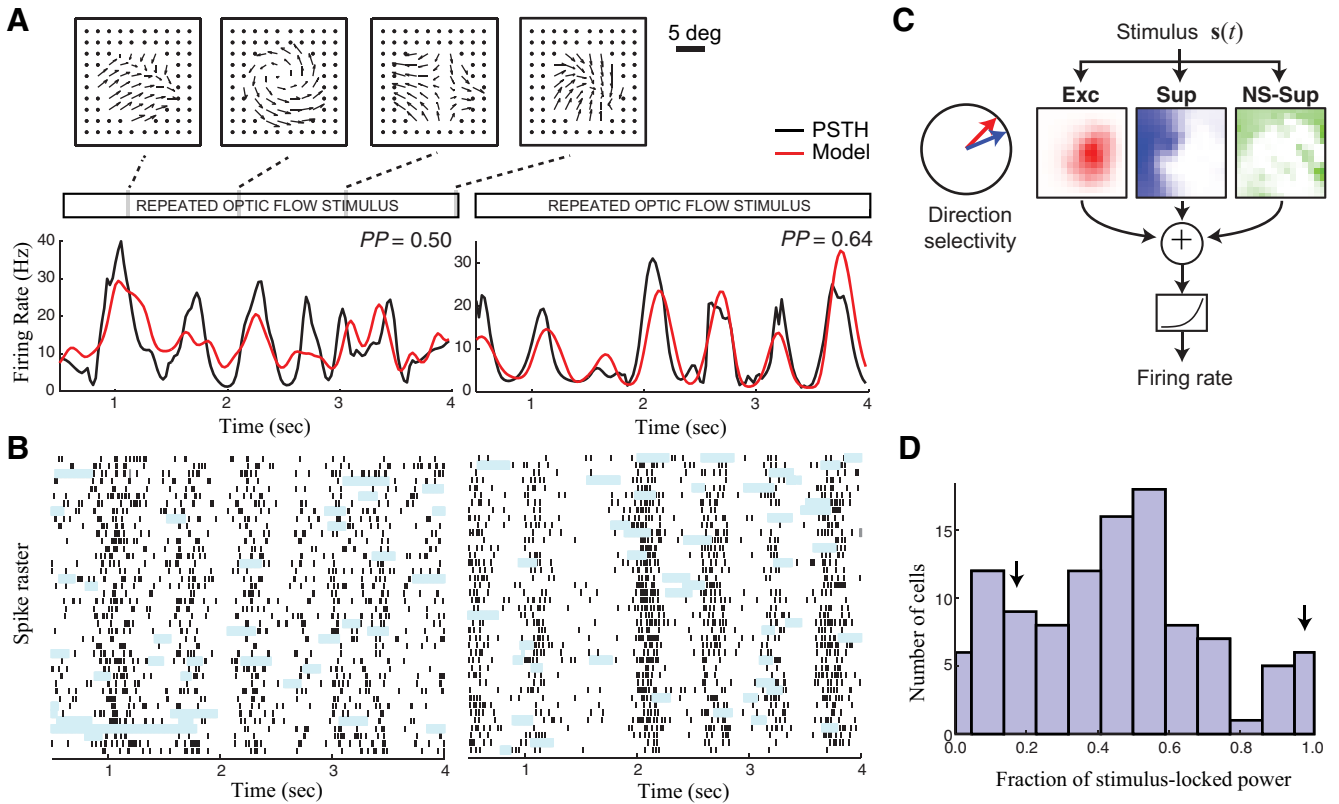


Figure 1. Response variability of MT neurons. **A**, The responses of two MT neurons to repeated presentations of a continuously varying optic flow stimulus, with four example frames from a section of the repeated stimulus segment (top). The peristimulus time histograms (PSTHs, bottom, black) are compared with the firing rate predictions of the stimulus-processing model (red), labeled with the fraction of stimulus-locked response explained by the model [predictive power (PP)]. **B**, Spike rasters of the same neurons as in **A**, from which the PSTHs are generated. The shaded areas mark epochs where data are excluded from analysis due to periods when eye position was outside the fixation window. **C**, Model schematic illustrating the main components of the stimulus-processing model: the stimulus composed is first processed locally by direction- and speed-selective local subunits, and then gets pooled across space separately by Exc (left), Sup (middle), and NS-Sup (right) components. Red and blue arrows (left) indicate direction selectivity for the Exc and Sup terms. Finally, a spiking nonlinearity is applied to this signal to transform it into a firing rate prediction (bottom). **D**, The distribution of the stimulus-locked rate variation across MT neurons ($n = 108$, mean = $45.9 \pm 25.9\%$; see Materials and Methods). The black arrows indicate the location of the two example neurons in **A** (stimulus-locked variance = 16.5% and 96.4%).

the total firing rate variance (left side of Eq. 1), from which we can estimate the fraction of MT neuron activity that is stimulus locked (Fig. 1D; see Fig. 3D), as well as how much of this is explained by the models we consider (see Fig. 6E).

To estimate the trial-variable firing rate variance $\text{Var}_i[\lambda^{(i)}(t)]$ of a given MT neuron with observed spike counts $n^{(i)}(t)$, we constructed a surrogate model that is constrained to have the same stimulus-locked firing rate as the observed data $\lambda_{\text{psth}}(t)$, as well as to reproduce the observed interspike interval distribution, which is an essential component for constraining the spike count variability (Czanner et al., 2008). Specifically, we fit a rate-modulated Bernoulli process with a conditional spike probability given by the following:

$$\lambda_{\text{mod}}^{(n)}(t) = F[g_{\text{psth}}(t) + g_{\text{spk}}(t) - \theta], \quad (2)$$

where θ is the spiking threshold, and $F[g] = 1/[1 + \exp(-g)]$ is a sigmoid spiking nonlinearity function used to describe the Bernoulli process (Haslinger et al., 2012). The two other terms in the model correspond to the outputs of a stimulus-driven term $g_{\text{psth}}(t)$ and a spike history term $g_{\text{spk}}(t)$. They are generated using linear combinations of B-spline basis functions, such that:

$$g_{\text{psth}}(t) = \sum_i k_i^{\text{psth}} B_{i,2}(t), \quad (3)$$

where k_i^{psth} are individual amplitudes of each piecewise linear (i.e., second-order) spline basis functions $B_{i,2}(t)$, which have a 25 ms spacing; and

$$g_{\text{spk}}(t) = \sum_{i_1} \sum_{i_2} h_{i_1} B_{i_1}(t - t_{i_1}), \quad (4)$$

where t_{i_1} is time of each spike, h_{i_1} is a parameter for the spike history term (Paninski, 2004), and $B_{i,1}(t)$ is a unit-pulse (i.e., first-order) spline basis function, with knots spaced at 1, 2, 3, 4, 5, 6, 7, 8, 10, 12, 16, 20, 28, and 36 ms. The h_{i_1} values were all constrained to be negative. The parameters k_i^{psth} and h_{i_1} were estimated as a generalized linear model using maximum-likelihood estimation (i.e., with logistic regression; Haslinger et al., 2012). We simulated the model at 1 ms resolution (where spike refractoriness is important) to generate a spike train for each trial, and pooled the resulting data into 25 ms bins to compare with the spike counts of the recorded data.

The fitted model accurately captured the PSTH, as well as the measured refractoriness for MT neurons (i.e., interspike interval distribution). Note that, even though the 25 ms bin size is long relative to the spike refractory period of the neurons, the presence of spike history effects at shorter timescales contributes significantly to the reliability of the spike count in these longer time windows. By construction, the surrogate model had no trial-variable firing rate, but shared the same stimulus-locked firing rate variance and equivalent noise in spike generation. As a result, we estimated the trial-variable firing rate variance by taking the difference of the variances of the observed and simulated spiking, as follows:

$$\text{Var}_{i1}[\lambda_{\text{TV}}^{(i)}(t)] = \text{Var}_{i1}[n^{(i)}(t)] - \text{Var}_{i1}[n_{\text{mod}}^{(i)}(t)], \quad (5)$$

where $n_{\text{mod}}^{(i)}(t)$ is the simulated spikes from the surrogate model (Eq. 2). The fraction of stimulus-locked rate variation was then given by $\text{Var}_r[\lambda_{\text{psth}}(t)]/(\text{Var}_r[\lambda_{\text{psth}}(t)] + \text{Var}_r[\lambda_{\text{V}}^{(i)}(t)])$ (Fig. 1D).

We validated these methods on simulated data, where spike counts were generated given a known trial-variable firing rate, $\lambda^{(i)}(t)$. The estimated fraction of stimulus-locked rate variation based on the method agreed very well with the ground truth. Such estimates of the relationship between these quantities represent an improvement in describing the data over heuristics using constant variance-to-mean ratios typical of a Poisson neuron (Churchland et al., 2011; Goris et al., 2014). For example, we find that, for all neurons, different time points in the experiment (i.e., in response to different stimuli) will have different variance-to-mean ratios in their spike count, which can be predicted by our surrogate model. Our method shares many of the same assumptions with a previously proposed method (Czanner et al., 2008), but that method could not be directly used due to the rapidly varying stimulus-driven firing rate variations in our data on the order of 25 ms (Cui et al., 2013).

Statistical modeling framework for describing MT neuron responses. We used a maximum likelihood-based framework that included both stimulus-dependent terms and network activity-dependent terms to model MT neuron responses. We assumed that neuron responses are generated by an inhomogeneous Poisson process with an instantaneous rate $\lambda_{\text{mod}}(t)$. Measured spike trains were binned in time at 5 ms resolution to obtain the observed response rate $r_{\text{obs}}(t)$. The log-likelihood (LL) of the model is then given (up to an additive constant) by:

$$LL[r_{\text{obs}}(t), r(t)] = \sum_t [r_{\text{obs}}(t) \log \lambda_{\text{mod}}(t) - \lambda_{\text{mod}}(t)], \quad (6)$$

where $r_{\text{obs}}(t)$ is the measured neuronal response, and $\lambda_{\text{mod}}(t)$ is the model predicted firing rate (Paninski, 2004). Note that to maximize this quantity, the model predicted rate should be high when the observed rate is high (maximizing the first term), and otherwise is as low as possible (minimizing the second term).

All models used a fixed spiking nonlinearity, $F[\cdot]$, that acts on the sum of the separate outputs of the stimulus-processing component $g_{\text{stim}}(t)$, the LFP component $g_{\text{LFP}}(t)$, and the multiunit activity component $g_{\text{MUA}}(t)$:

$$\lambda_{\text{mod}}(t) = F[g_{\text{stim}}(t) + g_{\text{LFP}}(t) + g_{\text{MUA}}(t) - \theta], \quad (7)$$

where θ is the spiking threshold and the spiking nonlinearity function $F[\cdot]$ was chosen to be of the form $\log[1 + \exp(\cdot)]$. Such a model can be optimized within the maximum-likelihood framework efficiently (Cui et al., 2013; McFarland et al., 2013).

The stimulus-processing component of MT neuron model. For neurons recorded during the continuous optic flow stimuli, we based the stimulus-processing model component on a previously developed model (Cui et al., 2013). Briefly, the stimulus selectivity of MT neurons is based on nonlinear combinations of localized processing by V1-like inputs. There are three stimulus-processing terms (Fig. 1B): direction-selective excitation (Exc), direction selective suppression (Sup), and non-direction-selective suppression (NS-Sup), each with unique direction selectivity (except NS-Sup), speed sensitivity, and spatial and temporal integration. All components of the stimulus model were fit to the data at 25 ms resolution, as previously described (Cui et al., 2013). However, for the models that include network terms, described below, the output of the stimulus component was upsampled to 5 ms resolution using linear interpolation.

The network components of the MT neuron model. Our model inferred the relevant network inputs from both the LFPs and MUA recorded from the electrode array. To process the LFPs, a continuous wavelet transformation was applied either to the LFP signals themselves or, in the case of multielectrode array recordings, to their second spatial derivative, using a complex Morlet mother wavelet at 16 logarithmically spaced scales, corresponding to center frequencies from 0.5 to 70 Hz. All data were processed at 5 ms resolution. Calculation of the spatial derivatives of the LFPs is closely related to current source density analysis (Nicholson and Freeman, 1975), and yielded a more localized distribution of model

weights compared with those without this spatial transform, although this addition had a negligible effect on overall model performance.

The instantaneous phase $\phi_{ab}(t)$ and amplitude $\psi_{ab}(t)$ of the (transformed) LFPs at a given depth (d) and a frequency band (b) were derived from each component of the complex-valued wavelet transform. To model the relationship between the response of a neuron and the LFP signals, we used amplitude-modulated sine and cosine terms as linear predictors, with model coefficients α_{db} and β_{db} :

$$g_{\text{LFP}}(t) = \sum_{d,b} \psi_{ab}(t) [\alpha_{db} \cos[\phi_{ab}(t)] + \beta_{db} \sin[\phi_{ab}(t)]] = \mathbf{k}_{\text{LFP}} \cdot \mathbf{X}_{\text{LFP}}(t), \quad (8)$$

where $\mathbf{X}_{\text{LFP}}(t)$ is a design matrix with components $\psi_{ab}(t) \cos[\phi_{ab}(t)]$ and $\psi_{ab}(t) \sin[\phi_{ab}(t)]$, and \mathbf{k}_{LFP} is the LFP-receptive field vector that contains the corresponding model parameters α_{db} and β_{db} . Such a linear model can be viewed as fitting an (amplitude-modulated) sinusoidal dependence on the LFP at each frequency and depth, with model weight $\sqrt{\alpha_{db}^2 + \beta_{db}^2}$ and preferred phase $\arctan(\beta_{db}/\alpha_{db})$.

We also used simultaneously recorded MUA to predict neuronal responses with a causal coupling filter. The generating signal based on MUA is:

$$g_{\text{spk}}(t) = \sum_{d,\tau} k_{d\tau} r_d(t - \tau) = \mathbf{k}_{\text{spk}} \cdot \mathbf{X}_{\text{spk}}(t), \quad (9)$$

where $r_d(t)$ is the MUA from depth d and $k_{d\tau}$ reflects the dependence of the spike response on MUA at d and time τ in the past. The MUA from the same channel as the recorded neuron considered was always excluded to avoid potential contamination, and for other channels with well isolated single units, the isolated single unit replaced the MUA signal.

Regularization of the LFP model. Because we sampled the LFP signal at 16 frequencies, the LFP model has 32 parameters for each depth and 512 parameters across 16 depths. The MUA filters include 15 depths and 15 lags for each unit. To avoid overfitting, we applied L2 regularization to the model coefficients (McFarland et al., 2013), using a penalty term added to the LL that penalizes nonsmoothness of the filter. For the LFP model, this term can be expressed as:

$$LL_{\text{smooth}}^{\text{LFP}} = \eta_d^{\text{LFP}} (\|\mathbf{L}_d \alpha_{db}\|_2 + \|\mathbf{L}_d \beta_{db}\|_2) + \eta_b^{\text{LFP}} (\|\mathbf{L}_b \alpha_{db}\|_2 + \|\mathbf{L}_b \beta_{db}\|_2), \quad (10)$$

where \mathbf{L}_d and \mathbf{L}_b are discrete Laplacian operators with respect to the depth and frequency dimensions. The hyperparameters η_d^{LFP} and η_b^{LFP} control smoothness across frequency bands and across cortical depth, respectively. For the MUA filter, the penalty term is similarly expressed:

$$LL_{\text{smooth}}^{\text{spk}} = \eta_d^{\text{spk}} \|\mathbf{L}_d k_{d\tau}\|_2 + \eta_\tau^{\text{spk}} \|\mathbf{L}_\tau k_{d\tau}\|_2. \quad (11)$$

The hyperparameters η_d^{LFP} , η_b^{LFP} , η_d^{spk} , and η_τ^{spk} were adjusted using a nested cross-validation scheme, where 20% of the fitting data were randomly selected and reserved during optimization of the LFP filter. The regularization parameters that gave the best performance on the reserved data were used to evaluate performance in the cross-validation data (which were not used to fit the model coefficients nor the hyperparameters).

Evaluation of model performance. In this article, we measured model performance using two different metrics, based on likelihood and variance. The standard metric for statistical models of this sort is the cross-validated log-likelihood (LL_x), which is the log-likelihood of the model (Eq. 5) evaluated on set-aside data not used to fit the model. We report a log-likelihood relative to a null model, $LL_{\text{model}} - LL_{\text{null}}$, where the null model simply predicts the overall average firing rate. Although the cross-validated log-likelihood is a more sensitive measure of model performance relative to variance-based measures such as explained variance (R^2), it lacks the intuition that variance-based measures provide, and thus we report both.

Traditional variance-based measures can be dominated by “unexplainable” spike count noise, and standard methods to correct for this depend on the assumption of no trial-to-trial rate fluctuations (Sahani and Linden, 2003). Thus, to gauge the performance of the model in

capturing trial-variable response power, we used the variance of the model output as a proxy for explained variance, as the maximum-likelihood estimation optimizes the predicted firing rate to match the elements of the observed response it can explain. As validation of this metric, we find that likelihood improvements closely match this variance-based metric (see Fig. 3C), and thus we use it in this article to translate the likelihood-based measures of improvement into more familiar terms. It also allows for the decomposition of model-predicted stimulus-locked and total power to compare with our estimates of observed response power described above (see Fig. 6E).

In all cases, cross-validated model performance was based on the data recorded in response to the stimulus repeats. Because the overall gain and offset of the spiking nonlinearity often drifts over long periods, we fit these two parameters from the overall model (fit using long continuous stimuli) using the odd repeat trials. We thus excluded the odd repeats from cross-validation and used only the even repeated trials to measure model performance.

Measurement and prediction of noise correlations. To estimate the noise correlation between pairs of simultaneously recorded neurons, we first calculate the cross-covariance of the neural responses on repeat trials over a range of latencies τ , as follows:

$$\text{Cov}_{ij}(\tau) = \frac{1}{N(T - \tau)} \sum_{k=1}^N \sum_{t=1}^{T-\tau} \left[r_i^{(k)}(t) - \langle r_i^{(k)} \rangle \right] \left[r_j^{(k)}(t + \tau) - \langle r_j^{(k)} \rangle \right], \quad (12)$$

where, $r_i^{(k)}(t)$ is the response of the i th neuron on the k th trial, $\langle r_i^{(k)} \rangle$ is the average response rate of neuron i on the k th trial, N is the number of trials, and T is the duration of each trial. To isolate the “noise covariance,” we subtract estimates of the stimulus-driven firing rate comodulation $\text{Cov}_{ij}^{\text{stim}}(\tau)$ using a shift predictor that considers neural responses on adjacent pairs of trials, as follows:

$$\text{Cov}_{ij}^{\text{stim}}(\tau) = \frac{1}{N(T - \tau)} \sum_{k=1}^N \sum_{t=1}^{T-\tau} \left[r_i^{(k)}(t) - \langle r_i \rangle \right] \left[r_j^{(k+1)}(t + \tau) - \langle r_j \rangle \right]. \quad (13)$$

The noise correlation between neurons i and j is then given as follows:

$$\rho_{ij}(\tau) = \frac{\text{Cov}_{ij}(\tau) - \text{Cov}_{ij}^{\text{stim}}(\tau)}{\sigma_i \sigma_j}. \quad (14)$$

Noise correlation estimates are normalized by the product of SDs of the spike count response for each neuron in the pair, which for the i th neuron is as follows:

$$\sigma_i = \frac{1}{N} \sum_{k=1}^N \sqrt{\frac{1}{T} \sum_{t=1}^T [r_i^{(k)}(t) - \langle r_i \rangle]^2}. \quad (15)$$

The noise correlation predicted by the model was calculated by replacing the measured responses in Equations 12 and 13 with the model-predicted neural responses. In all cases, we excluded fixation break periods from the calculation. For example, if fixation was lost at any time on the k th trial between t and $t + \tau$, this data point was excluded from the summation in Equation 12, and the normalization factor was adjusted accordingly.

Statistical analyses. We used robust statistics when comparing performances of different models, as the log-likelihoods were often not normally distributed. Paired comparisons of group medians were performed with two-sided Wilcoxon signed rank tests. The variability of medians was estimated using bootstrapping techniques (random sampling with replacement, using 1000 repetitions). Correlations between measured and predicted noise correlations were computed using Pearson’s product-moment correlation coefficient.

Measurement of phase-locking strength. We measure the phase-locking strength using standard approaches (Lachaux et al., 1999). It is defined as follows:

$$\text{PLS}_{db} = \frac{1}{N_{\text{spk}}} \left| \sum_{j=1}^{N_{\text{spk}}} \exp(j\phi_{db}(t_{\text{spk}}^j)) \right|, \quad (16)$$

where t_{spk}^j is the time of the j th spike, N_{spk} is the total number of spikes, j is the imaginary number $\sqrt{-1}$, and $|\cdot|$ is the magnitude of the resulting complex number. The maximum value of this quantity is unity, when all spikes occur at the same phase, and each complex vector adds constructively. The minimum value is zero, when spikes are uniformly distributed across phase, and the complex vectors cancel. This quantity was calculated across both frequency bands and depths independently.

Results

Response variability of MT neurons

We presented a naturalistic optic flow stimulus (Fig. 1A, top) to macaques during a passive fixation task while recording from neurons in the MT area with a laminar multielectrode array. This stimulus was designed for the purpose of performing detailed characterization of MT neuron selectivity to visual motion (Cui et al., 2013). Using a recently developed model of MT stimulus processing, we could predict neural responses to these stimuli accurately, as is evident from comparing the model-predicted firing rates to the across-trial average response (or PSTH; Fig. 1A, bottom). However, the success in explaining the trial-averaged response did not translate into the ability of the model to predict spiking patterns on individual trials, which could differ significantly between repeated presentations of the same stimulus (Fig. 1B).

Such trial-to-trial variability is fundamentally unpredictable, or unexplainable, given the stimulus alone. Thus, such variability is usually ignored in measures of model performance, which focus on reproducing the trial-averaged firing rate. For example, the MT neuron stimulus-processing model derived in our previous study (Fig. 1C) had a predictive power, or a fraction of explainable variance correctly predicted (Sahani and Linden, 2003), of 0.33 ± 0.26 (mean \pm SD; $n = 108$), which is comparable to the performance of other successful models of stimulus processing in visual cortex (David and Gallant, 2005; Willmore et al., 2010; Nishimoto and Gallant, 2011). However, if other experimentally uncontrolled variables such as “cortical network state” were important for the function of that MT neuron, the predictive power would overestimate model performance, because some modulation in the firing rate of the neuron would be, in principle, explainable if these variables were known. Thus, an important question is how much of the trial-to-trial response variability is due to fluctuations in the underlying firing rates of neurons? In other words, we would like to know how much across-trial response variability might be predicted given appropriate knowledge of the relevant ongoing cortical processes.

We thus estimate how much trial-to-trial variability in the observed spiking data is due to underlying firing rate modulations. Although there is generally no unique way to measure the firing rate of a neuron from a single trial of recorded spiking data (Amarasingham et al., 2015), here we adopt a reasonable set of assumptions to estimate this quantity. Namely, we simulate spiking data using a model that generates the observed trial-averaged firing rate (PSTH) and has the observed interspike interval distribution (see Materials and Methods). Such a model will have different patterns of spikes on each trial, but the variance in its spike train will be consistent with a single underlying rate that is common across all trials. Due to the additive properties of these different sources of variance (i.e., rate modulation vs spike generation, given the rate), the difference between the response vari-

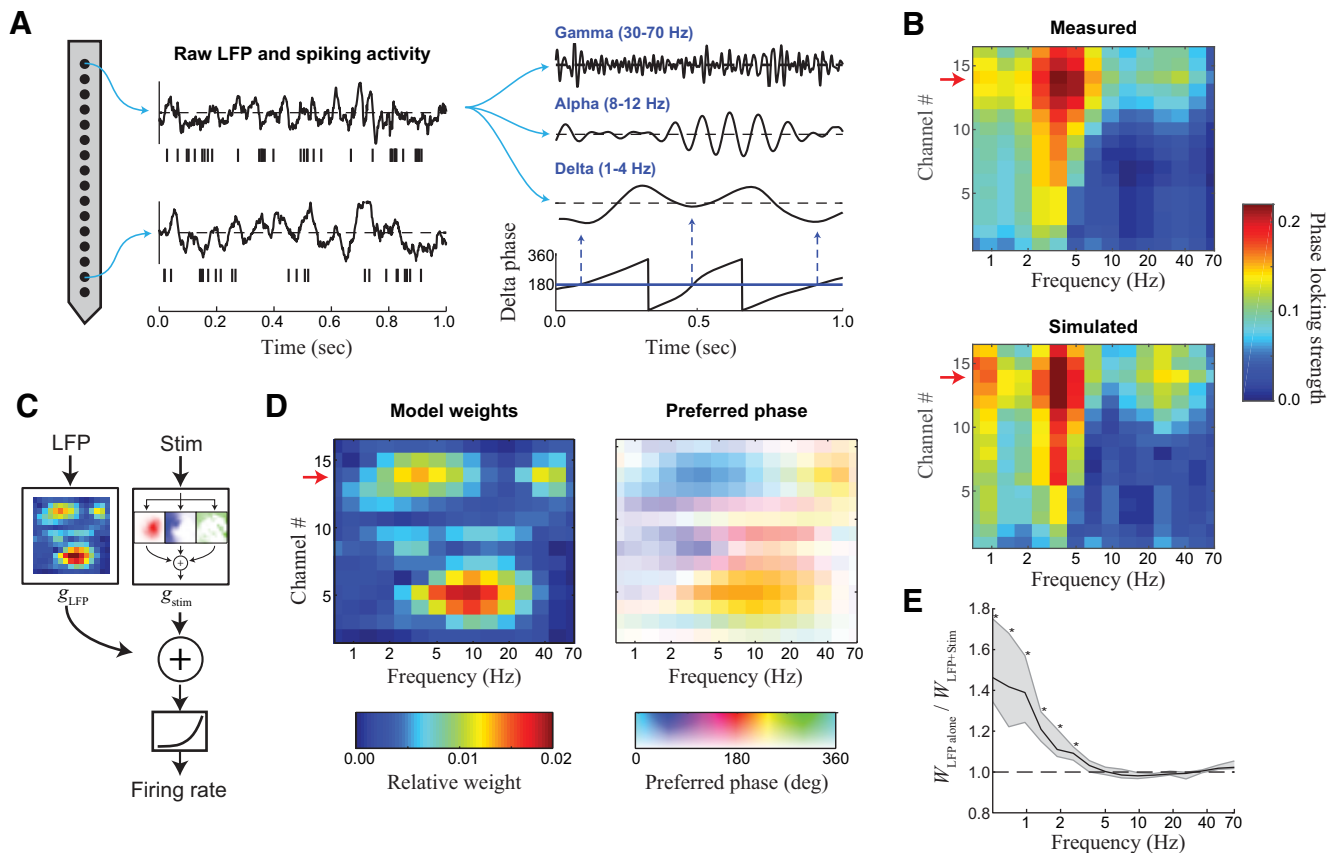


Figure 2. Predicting MT responses from the stimulus and inferred network activity. **A**, Time–frequency analysis of the LFPs. Left, LFP signals and spiking activity were recorded with a 16-channel multielectrode array and used as inputs to a model that predicts a given single unit on the array (described below). Middle, To illustrate how each LFP was processed for use in the model, we show an LFP segment bandpassed in the gamma (30–70 Hz), alpha (8–12 Hz), and delta (1–4 Hz) bands. Right, To process each LFP for input to our model, we use a continuous wavelet transform to decompose the LFP signal into its amplitude and phase in each frequency band. To demonstrate how phase relates to a given bandpassed signal, the delta phase is shown below the delta band-filtered LFP (bottom middle); with times that phase is 180° (blue arrows) corresponding to local minima in the bandpassed signal. **B**, The phase-locking strength measured across frequency band and depth for an example neuron (top), demonstrating sensitivity in the delta and gamma bands. The phase-locking strength of simulated spikes generated by the LFP model (bottom) is nearly identical, demonstrating how the localized model components in **D** combine with correlations between LFP band and depth to result in empirical phase-locking measurements. **C**, Schematic of the LFP model, which includes a stimulus-processing component (from Fig. 1B), and the LFP-based component. **D**, Example LFP model component. Model weights (left) and preferred phases (right) are displayed as a function of both frequency bands (horizontal axis) and cortical depth (vertical axis). Phase values are depicted by hue, with the color brightness representing the corresponding weights to emphasize frequency bands and depths that are most important. **E**, To determine the effect of simultaneously fitting multiple model components, we measured the ratio of model weights for single-electrode LFP models fit with and without the stimulus-processing model. This shows that LFP models fit without the stimulus-processing component typically have much larger weights in low-frequency bands (<4 Hz, * $p < 0.001$, $n = 108$).

ance of the observed data and that of this model neuron provides an estimate of the total variance associated with these unobserved trial-to-trial firing rate fluctuations. This variance then adds to the variance of the trial average (“stimulus-locked”) response to arrive at an estimate of the total firing rate variance of the neuron. We validated such an estimate with simulated data, as well as on other datasets (see Materials and Methods).

We used this measure of total firing rate variance to ask how much of each MT neuron response was, in principle, predictable given the stimulus (i.e., the fraction of the firing rate of each neuron that is consistent across trials). Across the population of MT neurons, there was a wide range for this quantity (Fig. 1D). While some neurons exhibited little trial-variable differences in rates (Fig. 1B, right), most neurons exhibit significant trial-to-trial variability in their rates (Fig. 1B, left). On average, less than half of the firing rate fluctuations were stimulus locked (0.46 ± 0.26 , $n = 108$), meaning that even a “perfect” model of MT stimulus processing could explain less than half of the firing rate of the neuron.

Dramatic improvement of model predictions using the LFP

We hypothesized that this trial-to-trial firing rate variability might reflect the influence of ongoing patterns of activity in the cortical network (Arieli et al., 1996; Masquelier, 2013). Such coordinated activity patterns can be detected in LFPs recorded simultaneously with spiking activity (Banerjee et al., 2012), suggesting that LFP signals might be used to predict trial-to-trial variability (Kelly et al., 2010; Haslinger et al., 2012). Because only some components of the LFP signal are predictive of cortical neuron activity, we first separated the LFP into different frequency bands, which are thought to represent distinct physiological sources as well as unique aspects of network dynamics (Buzsáki and Draguhn, 2004). We also used LFP signals across cortical depths to further resolve the different sources contributing to the LFP (Schroeder and Lakatos, 2009; de Cheveigné et al., 2013; Fig. 2A).

We first used standard methods to quantify the relationship between MT neuron spiking and these different components of the LFP signals. In particular, we measured the “phase-locking

strength” (Lachaux et al., 1999; Fig. 2B, top, example neuron), which quantifies the extent to which the spikes of a neuron occur at a consistent phase of a given LFP component and is small when spikes are unrelated to LFP phase. Most MT neurons demonstrated significant phase locking to a wide range of LFP bands, concentrated in the delta (1–4 Hz) and gamma (30–70 Hz) bands. Note that because LFP signals at different frequencies and cortical depths are correlated (Buzsáki and Draguhn, 2004), measures such as phase-locking strength that treat each LFP component independently will reflect these correlations rather than distinct sources of information about MT spikes, and thus extend across depth and have broad frequency tuning (Fig. 2B, top). Furthermore, measuring the spike–LFP relationship without taking stimulus processing into account cannot isolate the components of the LFP that are most useful for capturing the trial-to-trial firing rate fluctuations in particular.

An alternative approach, which addresses the above limitations, is to build an explicit model that predicts the neural response given both the stimulus and the observed LFP signals (Fig. 2C). The LFP component of this model assigns a weight and preferred phase to each LFP frequency band and depth, and is combined with the previously derived stimulus-processing model to obtain an optimal prediction given both stimulus and LFP. These parameters can be selected using maximum likelihood methods such that these different elements are combined appropriately to best explain the MT neuron response, automatically accounting for the strong correlations among the various explanatory factors.

In particular, the weights assigned by the model to different LFP components are much more localized to particular depths and frequency bands than standard measures such as phase-locking strength (Fig. 2D), which reflect the strong correlations among LFP components. Despite the large apparent differences between model-based and standard measures of spike–LFP coupling, we verified that our models correctly predicted the observed pattern of phase-locking strength (Fig. 2B, bottom). Furthermore, because the components of the model accounting for the LFP and stimulus are optimized simultaneously, the LFP component weights are selected to specifically predict trial-to-trial variability and ignore elements of the LFP that are “redundant” with the stimulus. In the absence of the stimulus-processing component, the parameters of the LFP component had larger weights for the low-frequency bands (Fig. 2E), reflecting a relationship between low-frequency LFP components and the stimulus itself (Belitski et al., 2008; Montemurro et al., 2008).

The addition of the LFP component dramatically improved the ability to predict the observed MT spike trains. To quantify the improvement in the ability of the model to predict MT spike trains using the LFP in addition to the stimulus, compared with using the stimulus alone, we first compare the “likelihoods” that each model assigns to the data, using appropriate methods for cross-validation (see Materials and Methods). We measured model performance first at 25 ms resolution (Fig. 3A, left) in order to make a fair comparison to models that only use the stimulus (and thus only predict the firing rate at the slower timescales of temporal filtering in MT) (Cui et al., 2013). In this case, the addition of the LFP component resulted in a nearly threefold increase (2.91 ± 0.43 ; $p < 10^{-12}$, $n = 93$) in the LL_x . Considering the model performance at a time resolution of 5 ms revealed further improvement (Fig. 3A, right; 4.5 ± 0.8 -fold increase in LL_x ; $p < 10^{-12}$), because the gamma component of the LFP helps to predict spiking activity at these finer timescales, which cannot be predicted by the more slowly changing stimulus component.

While in all cases the models used LFP signals from all available electrodes, a large fraction of this model improvement was obtained using only the LFP signal recorded on the same channel as the recorded unit (3.02 ± 0.27 -fold increase over the stimulus-processing model; $p < 10^{-12}$, $n = 93$).

While LL_x is the most accurate way to compare model predictions using single-trial spiking responses, its specific values are less interpretable compared with trial-averaged measures such as R^2 and predictive power, which provide a measure of absolute model performance relative to the perfect model. To provide an additional, more intuitive, measure of model performance, we next compared the variance of the model predicted firing rates with the total firing rate variance estimated above. Because the magnitude of model-predicted firing rates is implicitly adjusted to best match the data, the variance of model-predicted rates provides a useful proxy for the explained variance, given that the “true” trial-to-trial firing rates are unobserved. Indeed, the model-predicted variance (Fig. 3B) was highly correlated with LL_x improvements across neurons (Pearson’s correlation coefficient = 0.984, $p < 10^{-67}$; Fig. 3C), suggesting that the dramatic improvements in LL_x directly translate into improvements in the proportion of response variance explained by the models.

To ensure that the ability of the model to predict spiking from the LFPs was not due to contamination of the LFPs by the spike waveforms themselves, we used a robust spike-removal algorithm to subtract spikes from the LFP signal (see Materials and Methods; Zanos et al., 2011). Furthermore, we performed an additional control by comparing models that were fit without using the LFP signal from the electrode where the unit was recorded, using only signals from neighboring electrodes that contain negligible spike artifacts from the recorded neuron (Zanos et al., 2012). We indeed found that the LFP model parameters, including weights and preferred phases, were comparable when using LFPs from the same electrode versus a neighboring electrode 100 μm away (data not shown).

Overall, these results suggest that much of the trial-variable activity of MT neurons—which is typically considered “noise” in the context of stimulus processing—is the result of network activity that is reflected in the LFPs. Indeed, the importance of the LFPs for the model predictions was larger for cells that had more trial-to-trial variability (i.e., smaller percentages of stimulus-locked power; Fig. 3D). This demonstrates that neurons with large trial-to-trial fluctuations in activity, rather than necessarily being unreliable, might simply be more strongly coupled to ongoing network activity, and obtaining an accurate picture of their function requires taking such influences into account.

Components of the LFP predictive of spiking activity

These models also identify the particular components of the LFP that are useful in predicting MT neuron responses, providing an alternative to standard spike–LFP measurements (Fig. 2). We found that LFP weights in a given frequency band were often concentrated at a particular depth (relative to the channel from which the spikes were obtained), but there was significant variation in the LFP frequencies and depths used to predict neural responses (Fig. 4A). Note that we could only extract depth relative to the recorded unit (vs absolute depth in the cortex), given that our electrode penetrations were generally not orthogonal to the cortical layers. To look for consistent features of the LFP model weights across the population, we aligned the weight profiles relative to the channel where the spikes were recorded (Fig. 4B). This procedure revealed consistently stronger weights in the

gamma band (30–70 Hz) at the same depth as the recorded unit, and more diffuse weights across depth in the delta band (1–4 Hz) and beta band (12–30 Hz).

To quantify the relative importance of different LFP frequency bands in predicting spiking responses, we measured the impact of removing individual frequency bands on model performance (Fig. 4C). This analysis showed that the delta (1–4 Hz) and gamma (30–70 Hz range) bands provided the largest contributions, consistent with simpler measures of spike–LFP relationships such as phase-locking strength (Fig. 2B). Furthermore, across the population, MT neurons had very consistent preferred phases of the local LFP signal (Fig. 4D), despite wide variation in their stimulus tuning and degree of coupling to the LFP.

Using MUA to predict MT neuron responses

Given that the response variability of different cortical neurons is correlated (Bair et al., 2001; Cohen and Kohn, 2011), one should also be able to predict, to some extent, the trial-to-trial variability of a neuron from the simultaneously recorded activity of neighboring neurons (Fig. 5A, top). Indeed, recent work (Okun et al., 2015; Schölvinc et al., 2015) has shown that simply taking the average firing rate among a group of simultaneously recorded neurons [“population rate” (PR)] can be highly predictive of the trial-to-trial variability of a neuron (Fig. 5A, bottom). We thus tested whether the recorded MUA might offer improvements to the prediction of single-neuron responses similar to those of the LFP-based model.

First, we considered the coupling between a given MT neuron and the population rate (Fig. 5B), which we modeled with a linear temporal filter acting on the PR. Consistent with previous work (Okun et al., 2015), the coupling between single-neuron responses and PR typically decayed as a function of latency (Fig. 5C), with a median decay time constant of 15.5 ± 2.4 ms ($n = 93$). Inclusion of the PR-coupling term significantly improved model performance relative to the stimulus-processing model alone (1.51 ± 0.14 -fold increase in LL_x , $p < 10^{-11}$, $n = 93$, Fig. 5D), although this was less than half of the increase gained from the LFP component (Fig. 3A).

Another possibility is that only the activity of a subset of recorded neurons would be useful for predicting the activity of a given neuron. To test this possibility, we fit models with separate temporal filters applied to each recorded MUA signal (Fig. 5E),

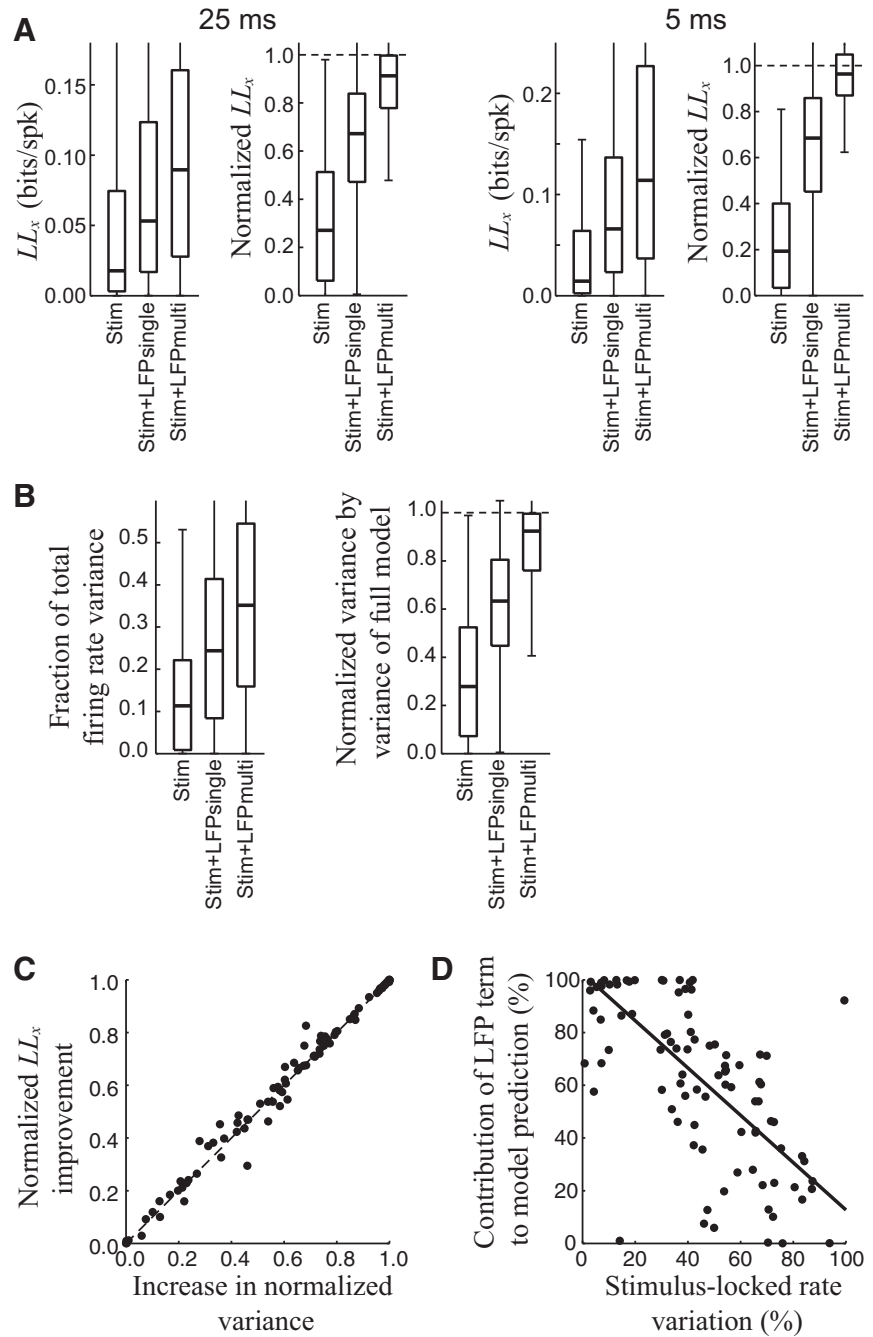


Figure 3. Dramatic improvement of model performance with the inclusion of the LFP component. **A**, LL_x values calculated for stimulus-processing models (left), and models that included either the LFP from the same electrode (middle) or all electrodes (right), for all neurons in the study. LL_x values were calculated at both 25 ms (left) and 5 ms (right) time resolution. To demonstrate the relative performance of each model across all neurons, we also normalized each LL_x value by that of the best model for that neuron (introduced below), and thus no model median reaches the (dashed) unity line here. The inclusion of the LFP component yields a more than threefold improvement in performance over the stimulus-processing model at 25 ms resolution (approximately the time course of the stimulus-locked rate). When considered at a higher time resolution (5 ms), the LFP model yields a median fivefold performance improvement, due to information about finer spike timing gained from the gamma-band LFP. **B**, The total variance of predicted firing rates from different models, normalized by the measured variance (left) or the variance of the best model for that neuron (right). **C**, Model performance improvement is tightly correlated with the increase in the variance of model outputs ($r = 0.984$, $p < 10^{-67}$). **D**, Improvements in model performance were greater for neurons that were less stimulus locked ($r = -0.64$, $p < 10^{-10}$).

rather than a single filter applied to the PR. This analysis revealed that neurons recorded from nearby electrodes contribute most to predicting the single-neuron response (Fig. 5F), suggesting that the most useful interneuronal correlations are locally clustered

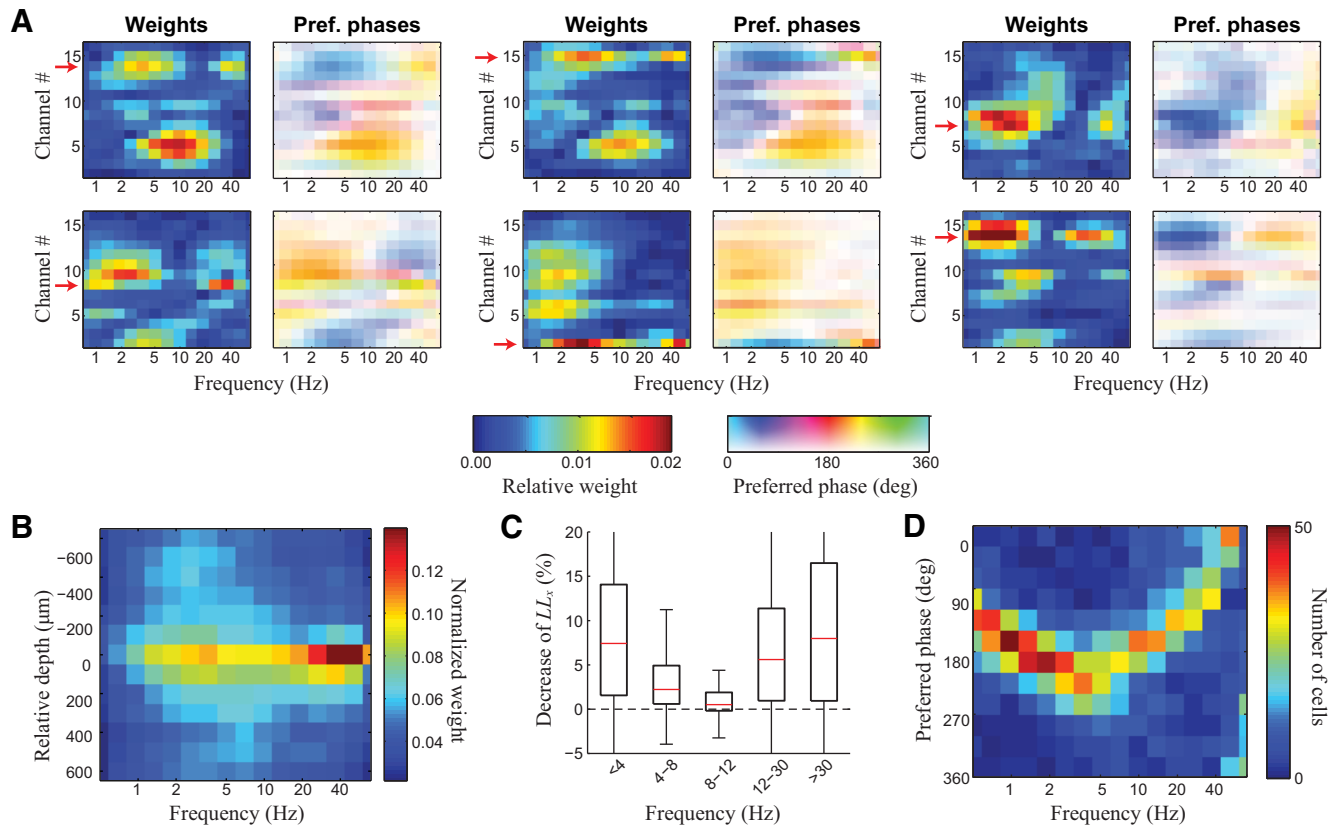


Figure 4. Frequency and phase preferences of MT neurons. **A**, There is a diverse array of tuning captured by the LFP components of different neurons, shown for six examples, with the red arrow designating the depth of the unit being modeled. Model weights (left) and preferred phases (right) are displayed as a function of both frequency bands (horizontal axis) and cortical depth (vertical axis) in the same way as in Figure 2C. **B**, While the absolute depth of the unit could not be determined in our experiments, we aligned model weights relative to the depth of the unit being modeled, and averaged these weights across neurons. This shows that model weights from the gamma band (30–70 Hz) are mostly localized in depth, whereas weights from the delta band (1–4 Hz) are distributed across multiple channels ($n = 93$). **C**, The relative impact of each LFP frequency band on model performance was computed by fitting each LFP model while omitting a given band. The effect on normalized LL_x is shown across all neurons ($n = 93$), showing that the removal of gamma (30–70 Hz) and delta (1–4 Hz) bands has the biggest impact on model performance. **D**, The consistency across neurons of phase preferences in each LFP band is demonstrated by a two-dimensional histograms of the preferred phase of each neuron in each frequency band from the single-electrode LFP model ($n = 108$).

(Smith and Kohn, 2008). However, although the full MUA model term had many more parameters than the PR term, it provided only moderate improvements in model performance (Fig. 5D; 1.09 ± 0.2 -fold increase over the PR model, $p < 10^{-5}$, $n = 93$). This suggests that a single coupling term in the PR can largely capture the variability in single-unit firing that might be extracted from MUA (Okun et al., 2015).

The relative contributions of MUA and LFP signals to predicting variability

To compare the contributions of the two signals reflecting cortical “network activity” (LFPs and MUA) to predicting MT neuron response variability, we combined the stimulus-processing, LFP, and MUA components into a single model (Fig. 6A). By simultaneously fitting the parameters of all three components, the combined model automatically accounts for the correlations among these various signals, assigning weights only to those elements that best explain the MT neuron response. Thus, sequentially adding model components can reveal the degree to which each component contributes unique information that is not present in other components. Direct comparisons of model improvement (Fig. 6B) revealed that the MUA signals contained largely redundant information with the LFP, because adding MUA components to the stimulus–LFP model yielded very little improvement in model performance (1.04 ± 0.02 -fold increase in LL_x , $p =$

0.003, $n = 93$). This suggests that LFPs contain richer information about the elements of ongoing network activity relevant to trial-to-trial variability relative to MUA.

To gain further insight into this, we separately analyzed the contributions of LFP and MUA model components to predicting both the stimulus-locked and trial-variable components of neural activity, again using trials with repeated presentations of frozen noise stimuli (Fig. 6C). Model components producing the same output on every trial can only contribute to predicting the stimulus-locked component of the response, while model components whose output is variable from trial to trial will contribute to capturing the trial-variable response. By construction, the stimulus-dependent component of the model was fixed across trials (Fig. 6D, second column). In contrast, the output of the LFP component had very few stimulus-locked features (Fig. 6D, third column), suggesting that it mainly contributed to trial-variable components of the spiking response. Finally, the MUA model component was smaller in magnitude, but typically had clear stimulus-locked and trial-variable components, with similarities to both the stimulus-dependent and LFP component outputs (Fig. 6D, right).

To quantify these different elements of the model predictions, we decomposed the output of each model component into its across-trial average (i.e., its contribution to predicting the PSTH of the neuron; Fig. 6D, bottom, black curves) and its trial-to-trial

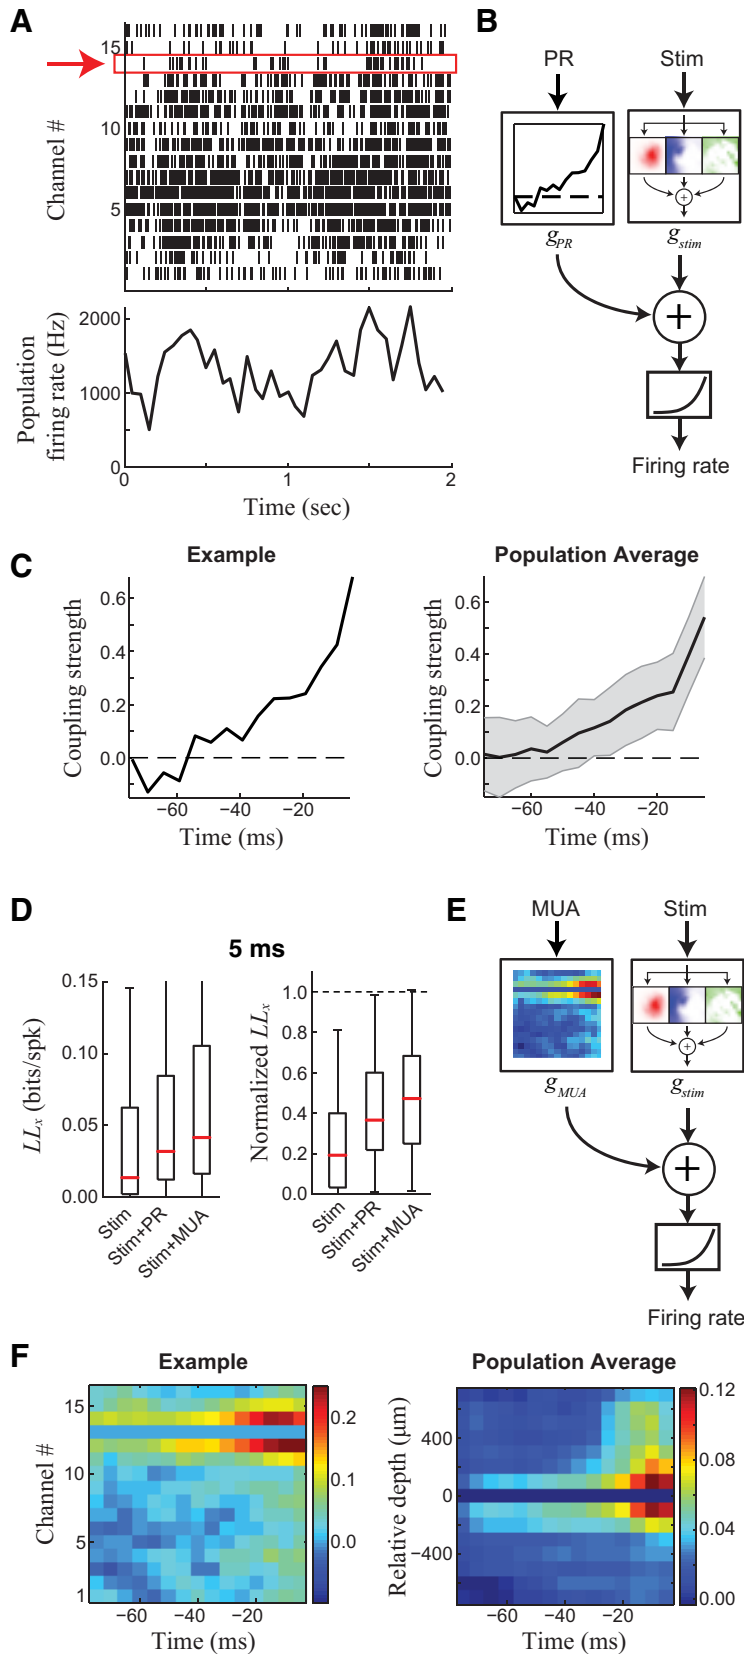


Figure 5. Modeling MT neuron response with multiunit activity. **A**, Example multiunit spiking activity (MUA, top) and the resulting population firing rate (PR; bottom) recorded from the multi-electrode array, with the red box highlighting the single unit being modeled (which was excluded from the model). **B**, Schematic of the population rate model, which sum the output of PR-based and stimulus-processing components, passing the result through a spiking nonlinearity. The contribution of the PR was modeled using a linear temporal filter applied to the recent history of PR. **C**, Left, A typical PR filter, showing coupling between the example neuron and PR as a function of time lag (horizontal axis). Right, Average of all PR model components across the population

variability (e.g., red curves minus black curves), and compared the variance of each to the total firing rate variances estimated from the measured data (Fig. 6E). To begin with, the stimulus-processing model alone accounted for $29 \pm 23\%$ of the stimulus-locked rate variance, and none of the trial-variable variance. Both the LFP and MUA components contributed to the stimulus-locked variance, producing a combined model that captured $42 \pm 23\%$ of the observed stimulus-locked variance.

However, the bulk of the improvement gained by the full model came from the LFP component capturing trial-to-trial rate fluctuations. Such fluctuations were not predicted at all by the stimulus-processing model (0%), but jumped to $44 \pm 35\%$ of the estimated trial-variable variance (see Materials and Methods) for the stimulus-MUA-LFP model. The LFP component was responsible for the bulk of this contribution, and the stimulus-LFP model (without the MUA component) predicted $39 \pm 34\%$ of the trial-variable variance. The predictions of the full stimulus-MUA-LFP model captured $43 \pm 28\%$ of the total variance of MT neuron firing rates.

LFP signals predict correlated response of pairs of neurons

Trial-to-trial response variability is often shared among simultaneously recorded neurons, and exhibits complex structure across a range of spatial and temporal scales (Zohary et al., 1994; Cohen and Kohn, 2011). Because LFPs are thought to reflect synchronous synaptic inputs to an area (Buzsáki and Draguhn, 2004; Enevoll et al., 2013), the LFP models might also explain the correlated trial-to-trial variability among neurons (known as noise correlations). To test this possibility, we first calculated the cross-

of recorded neurons ($n = 93$). **D**, Left, The LL_x of the PR and MUA models across the population of recorded neurons. Right, The same data, showing the normalized LL_x across all neurons (as in Fig. 3), where the LL_x of each model was normalized by the LL_x of the best model for that neuron (see below). The normalized LL_x is thus on the same scale as the LFP models considered earlier (Fig. 3), demonstrating that the PR-based models have less than half the performance of the LFP-based models. **E**, Schematic of the MUA model, where the MUA component replaces the PR component considered above (**B**). **F**, Left, A typical MUA model component, showing weights on MUA as a function of time lag (horizontal axis) and electrode depth (vertical axis) for the same neuron as the top left panel in **B**. Right, Average of all MUA model components across the population of recorded neurons.

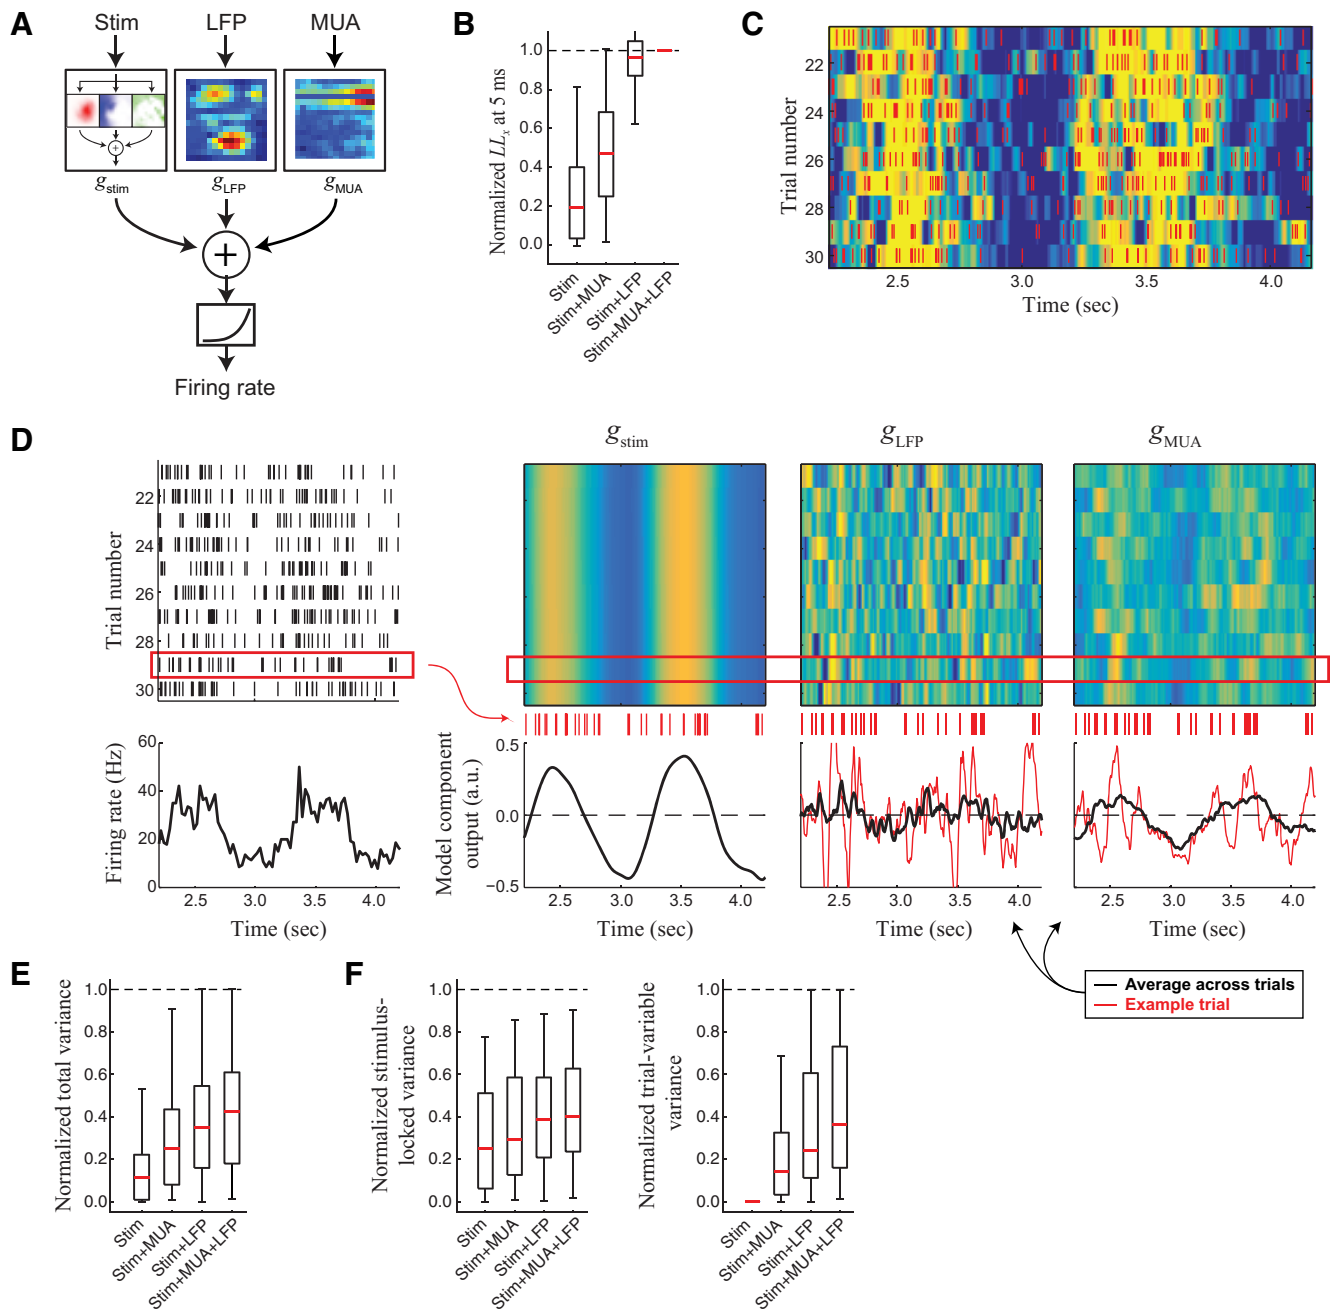


Figure 6. Network inputs predict trial-to-trial variability. **A**, Schematic of the full model, which includes a stimulus-processing component, and two “network-based” components: the LFP and MUA components. **B**, The model performance of the range of models considered, using the cross-validated log-likelihood relative to the “full” Stim–LFP–MUA model (normalized LL_x). This demonstrates that, while the addition of the MUA component can yield a nearly twofold improvement over the stimulus-processing (Stim) model, it has only a very small impact when added to the Stim–LFP model. **C**, The summed output of the three model components of the full model explains the spike response (red rasters) of the MT neuron on a trial-by-trial basis, as shown by the correspondence between observed spikes and model output. Colors are scaled such that blue is the most negative model output, and yellow is the most positive (arbitrary units). **D**, Spike raster across repeated trials (left), and the output of the three model components (right) across repeated trials of the same motion stimulus. The stimulus component output is identical across trials; the output of the LFP-dependent term is highly variable across trials, and the MUA component output has both stimulus-locked and trial-variable elements. The trial-averaged spike response and output of the three model components are shown on the bottom (black curve), together with an example trial in red. **E**, Variance of the model-predicted firing rate, normalized by the total amount of estimated firing-rate variance. **F**, The total variance of the model outputs are further separated into the contributions to the model firing rate of the stimulus-locked (left) and trial-variable (right) components, and are normalized by the corresponding measured stimulus-locked and estimated trial-variable variance. For example, the Stim model cannot predict any trial-variable variance (first column on right), while the full model (Stim–LFP–MUA, fourth column) predicts a median of 38% of the trial-variable variance of each neuron.

correlation between the responses of simultaneously recorded pairs of neurons on repeat trials, correcting for stimulus-driven effects using a “shift predictor” (see Materials and Methods). This revealed a variety of noise correlation structures across the population (Fig. 7A, left column). We then calculated the noise correlations predicted from the network-based models for each pair

of neurons. The LFP-predicted noise correlation functions consistently captured the temporal structure of the observed noise correlations, such that the correlation between observed and predicted noise correlation functions had a median value of 0.63 ± 0.04 (Fig. 7B). By comparison, the PR model (Okun et al., 2015; Schölvinck et al., 2015) was largely limited to capturing correla-

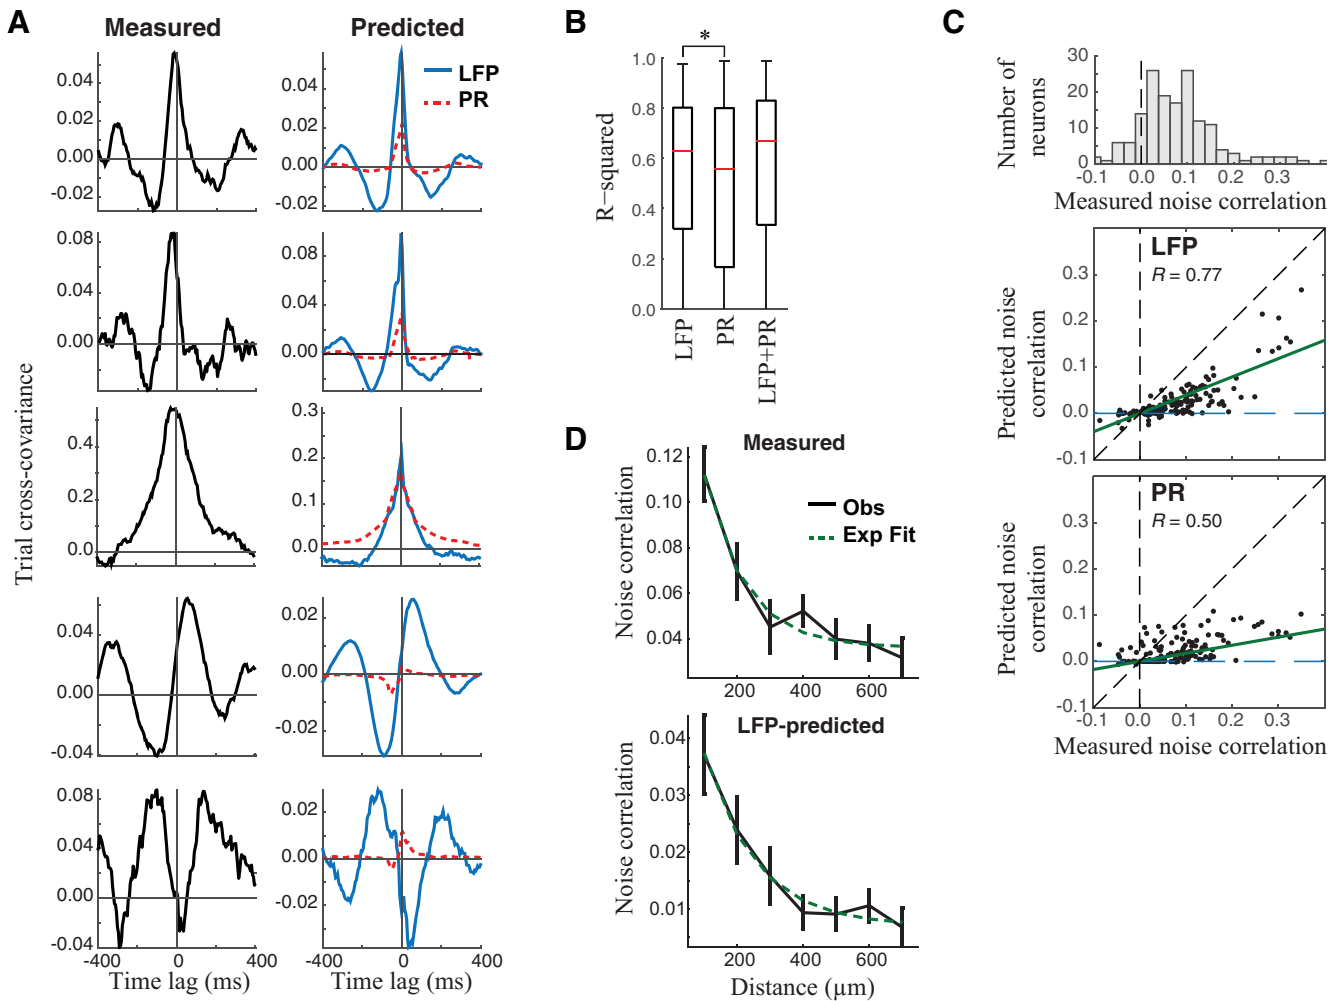


Figure 7. LFP signals predict noise correlation. **A**, Examples demonstrating a range of measured noise correlations between pairs of MT neurons (left), which largely match the shape of the noise correlations predicted by the LFP models fit separately to each neuron in the pair (right, blue), compared with the predictions of the PR model (dashed red). **B**, The correlation coefficients between measured and predicted noise correlation functions demonstrate that the LFP models are capturing the shape of the noise correlation functions (median, $r = 0.63 \pm 0.04$). The PR model has significantly worse performance (median, $r = 0.56 \pm 0.06$, $p < 0.05$), whereas the combined PR plus LFP model does not perform significantly better than the LFP model (median, $r = 0.67 \pm 0.04$, $p = 0.60$). **C**, Top, Simultaneously recorded pairs of neurons had a range of noise correlations (histogram shown across all pairs). Middle, bottom, The LFP model could predict these correlations on a neuron-by-neuron basis (middle, $r = 0.77$), again with the PR model not achieving as high performance (bottom, $r = 0.50$). **D**, The LFP model also could predict the dependence of noise correlation on relative depth: measured noise correlation (top) declines as a function of distance between MT neuron pairs, which is captured by the LFP model (bottom). Dashed lines show exponential fits to this relationship.

tion functions with a single peak at or around zero latency (Fig. 7A, red), which was less accurate over all pairs than the LFP model at capturing the shape of the correlation function (Fig. 7B; median correlation coefficient = 0.56 ± 0.06 , $p < 0.05$).

The LFP-based prediction of noise correlations could also account for the magnitude of noise correlations on a pair-by-pair basis (Fig. 7C; Pearson correlation coefficient = 0.766, slope = 0.40, $p < 10^{-14}$), which was much better than the PR model (Pearson correlation coefficient = 0.500, slope = 0.17, $p < 10^{-14}$). This ability to predict the relative coupling strength between pairs of neurons was intact even though the magnitudes of the predicted noise correlation were a factor of three smaller than the observed correlations, likely because each model explained a fraction of the response variance of individual neurons. As observed in previous studies (Smith and Kohn, 2008), pairwise noise correlation strength decreased systematically with the distance between the neurons, and the LFP models were able to capture this relationship (Fig. 7D). This suggests that components of the LFP can predict much of the observed noise correlation structure across pairs of neurons.

Discussion

While the function of sensory neurons is typically evaluated based on their responses to sensory stimuli, a large fraction of the activity in sensory cortex is not predictable from the stimulus (Ermentrout et al., 2008; Masquelier, 2013; Goris et al., 2014). Here we show that in area MT of the awake macaque cortex, on average more than half of the firing rates of neurons are not explainable from the stimulus alone during a passive viewing of a naturalistic motion stimulus (Cui et al., 2013). As neuronal variability is often correlated between cortical neurons (Averbeck et al., 2006; Cohen and Kohn, 2011), we hypothesized that such variability might be due to stimulus-independent inputs that can be captured through particular features of the measured LFPs (Kelly et al., 2010; Haslinger et al., 2012; Ecker et al., 2014). Indeed, a model that used LFP signals in addition to the stimulus produced, on average, a nearly fivefold improvement in predicting MT neuron responses. This dramatic performance improvement was largely due to the ability of these models to predict close to half of the trial-to-trial firing rate variability, leveraging the

trial-variable elements of the LFP signals. This suggests that the identified elements of the LFP represent coordinated inputs that were largely distinct from stimulus-driven inputs. Likewise, the aspects of the neural responses inferred from the LFP, as a result of being almost entirely trial variable, were clearly distinct from the activity of simultaneously recorded neurons and, thus, likely reflected unique information about inputs onto the recorded neurons.

Inference of relevant network activity using LFPs

A fundamental barrier to the investigation of how ongoing network activity influences the stimulus processing of a cortical neuron is the difficulty in recording from the many different inputs of a neuron, which will generally originate from both nearby cortical neurons as well as from projections from other brain areas. LFPs are potentially able to provide this information, because they are thought to be strongly influenced by synchronous synaptic inputs and, as a result, will often reflect inputs from other areas that are not necessarily apparent in local spiking activity (Buzsáki and Draguhn, 2004; Khawaja et al., 2009; Einevoll et al., 2013). Different sources contributing to the LFP will often be overlapping in time but will exhibit dynamics on different timescales (Buzsáki and Draguhn, 2004), and particular oscillation frequencies have been attributed to the activation of different cortical microcircuits, which can have distinctive signatures as a function of depth within the cortical column (Schroeder and Lakatos, 2009; Xing et al., 2009; de Cheveigné et al., 2013). Our statistical framework allowed for all depths and frequencies to be analyzed simultaneously to best infer the components of network activity relevant to a given cortical neuron.

Furthermore, by fitting a joint stimulus–LFP model, this approach specifically identified frequency bands of the LFP relevant for predicting the components of neural activity that cannot be explained by the stimulus alone (i.e., trial-to-trial variability). In this way, our model-based analysis has an important advantage relative to simple analyses of spike–LFP coupling, such as coherence and phase locking (Lachaux et al., 1999; Fries et al., 2001), in that it is not confounded by the strong, and generally complex, relationships that exist among the stimulus, the activity of other neurons, and the various components of the LFP signals (i.e., different frequency bands and depths). As a result, our models typically focused on a small subset of these LFP components (relative to measures such as phase-locking strength) and specifically identified elements of the LFPs that are predictive of the trial-variable component of the responses of neurons. Note that the functional form of the LFP component that we considered here resembles previous treatments (Haslinger et al., 2012) and also is mathematically equivalent to processing in the time domain (Kelly et al., 2010).

Our analysis identified the following two bands in the LFP that were most important in predicting MT neuron responses: the delta band (1–4 Hz) and the gamma band (30–70 Hz). While the influence of the gamma band was primarily confined to the same or neighboring electrodes relative to the recorded unit, the influence of low-frequency bands was often strongest at other depths. Although the geometry of our multielectrode recordings prevented determining which cortical lamina the different signals arose from, it is likely that such influences will have a consistent laminar structure (Schroeder and Lakatos, 2009; Xing et al., 2009; de Cheveigné et al., 2013). While gamma-band activity has been associated with cortical neuron firing rates in numerous contexts (Nir et al., 2007; Burns et al., 2010; Buzsáki and Wang, 2012; Jia et

al., 2013), the delta band is only of recent interest (Lakatos et al., 2008; Nácher et al., 2013).

Inference of relevant network activity using simultaneously recorded neurons

We found that incorporating the effects of simultaneously recorded MUA could nearly double the performance of model predictions over the stimulus-processing model alone, consistent with a number of recent studies using neighboring neural activity to improve model predictions in the retina (Pillow et al., 2008) and cortex (Kelly et al., 2010; Ohiorhenuan et al., 2010; Ecker et al., 2014; Okun et al., 2015). However, such improvements were significantly smaller than the fivefold improvements provided by the LFP in predicting MT neuron responses. Furthermore, the addition of the MUA term into a stimulus–LFP model offered only negligible improvements, suggesting that the MUA contributions were already predicted by the LFP. We do not believe that having more simultaneously recorded units would further improve the stimulus–MUA model performance, because individually weighting the different recorded units to maximize model performance revealed that the largest influences on MT neuron firing were from neighboring neurons. Individually weighting each recorded unit only offered small improvements over pooling over all recorded neurons to use a “population firing rate,” as suggested in other recent work modeling variability (Lin et al., 2015; Okun et al., 2015).

Why did the MUA-based model offer a much smaller improvement than the LFP-based model? One clear advantage of the LFP model is that it has extra degrees of flexibility in selecting the magnitude of coupling to different bands (as well as depths), which allows for each neuron to couple to a variety of components and dynamics, whereas the MUA only provides a single temporally varying element. Such flexibility is likely necessary, for example, to explain the nonsymmetric noise correlations observed for some simultaneously recorded neuron pairs (Fig. 7). In contrast, the MUA term can only represent shared firing rate fluctuations across the population, which necessarily must average together shared stimulus processing and shared trial-variable inputs (Fig. 6).

Intrinsic to the increased flexibility of the LFP model is the fact that the LFP itself can represent multiple elements of cortical dynamics (Buzsáki and Draguhn, 2004) and, in particular, is thought to be strongly influenced by synchronous synaptic inputs (Buzsáki and Draguhn, 2004; Khawaja et al., 2009; Einevoll et al., 2013), which are likely sources of MT neuron modulation. In contrast, MUA activity will be limited to representing only those inputs that result in synchronous spiking across the population.

Function of cortical variability

The visual cortex has been long recognized to be “noisier” than earlier processing stages (Kara et al., 2000). While such neuronal variability can be indicative of noise that is disruptive to cortical function (Faisal et al., 2008), it also might be a signature of unexplored aspects of cortical function that are related to ongoing processes not controlled by the stimulus (Reynolds and Chelazzi, 2004; Stein et al., 2005; Masquelier, 2013). The success of the LFP model in identifying the trial-to-trial variability of single neurons, as well as the form of noise correlations between pairs of neurons (Fig. 7), indeed suggests that such variability is the result of modulation by shared cortical inputs that are independent of the stimulus. In the context of the passive fixation task considered here, it is not clear whether there is a functional “role” of such

modulation. However, changes in the properties of LFPs, as well as changes in the relationship between spikes and LFPs, have been associated with cognitive and behavioral variables in a large array of previous experiments (Fries et al., 2008; Schroeder and Lakatos, 2009; Banerjee et al., 2012; Nácher et al., 2013), and likewise the LFP is an effective signal for the control of brain–machine interfaces (Hwang and Andersen, 2013). The trial-to-trial variability of the activity of MT neurons during perceptual discrimination tasks is also correlated with the subject's decisions (Britten et al., 1996; Purushothaman and Bradley, 2005; Nienborg et al., 2012). The ability of this approach to identify signatures of neural modulation in the LFP related to trial-to-trial variability in these passive conditions might thus provide a window into the influences of higher-level cognitive processes when applied in task-specific conditions.

References

- Amarasingham A, Geman S, Harrison MT (2015) Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proc Natl Acad Sci U S A* 112:6455–6460. [CrossRef Medline](#)
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871. [CrossRef Medline](#)
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366. [CrossRef Medline](#)
- Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *J Neurosci* 21:1676–1697. [Medline](#)
- Banerjee A, Dean HL, Pesaran B (2012) Parametric models to relate spike train and LFP dynamics with neural information processing. *Front Comput Neurosci* 6:51. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765. [Medline](#)
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA (1996) A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Vis Neurosci* 13:87–100. [CrossRef Medline](#)
- Burns SP, Xing D, Shapley RM (2010) Comparisons of the dynamics of local field potential and multiunit activity signals in macaque visual cortex. *J Neurosci* 30:13739–13749. [CrossRef Medline](#)
- Buzsáki G, Draguhn A (2004) Neuronal oscillations in cortical networks. *Science* 304:1926–1929. [CrossRef Medline](#)
- Buzsáki G, Wang XJ (2012) Mechanisms of gamma oscillations. *Annu Rev Neurosci* 35:203–225. [CrossRef Medline](#)
- Buzsáki G, Anastassiou CA, Koch C (2012) The origin of extracellular fields and currents—EEG, ECoG, LFP and spikes. *Nat Rev Neurosci* 13:407–420. [CrossRef Medline](#)
- Canolty RT, Knight RT (2010) The functional role of cross-frequency coupling. *Trends Cogn Sci* 14:506–515. [CrossRef Medline](#)
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819. [CrossRef Medline](#)
- Cohen MR, Newsome WT (2008) Context-dependent changes in functional circuitry in visual area MT. *Neuron* 60:162–173. [CrossRef Medline](#)
- Constantinople CM, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical networks. *Neuron* 69:1061–1068. [CrossRef Medline](#)
- Cui Y, Liu LD, Khawaja FA, Pack CC, Butts DA (2013) Diverse suppressive influences in area MT and selectivity to complex motion features. *J Neurosci* 33:16715–16728. [CrossRef Medline](#)
- Czanner G, Eden UT, Wirth S, Yanike M, Suzuki WA, Brown EN (2008) Analysis of between-trial and within-trial neural spiking dynamics. *J Neurophysiol* 99:2672–2693. [CrossRef Medline](#)
- David SV, Gallant JL (2005) Predicting neuronal responses during natural vision. *Network* 16:239–260. [CrossRef Medline](#)
- de Cheveigné A, Edeline JM, Gaucher Q, Gourévitch B (2013) Component analysis reveals sharp tuning of the local field potential in the guinea pig auditory cortex. *J Neurophysiol* 109:261–272. [CrossRef Medline](#)
- Ecker AS, Berens P, Cotton RJ, Subramanian M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82:235–248. [CrossRef Medline](#)
- Einevoll GT, Kayser C, Logothetis NK, Panzeri S (2013) Modelling and analysis of local field potentials for studying the function of cortical circuits. *Nat Rev Neurosci* 14:770–785. [CrossRef Medline](#)
- Engel AK, Fries P, Singer W (2001) Dynamic predictions: oscillations and synchrony in top-down processing. *Nat Rev Neurosci* 2:704–716. [CrossRef Medline](#)
- Ermentrout GB, Galán RF, Urban NN (2008) Reliability, synchrony and noise. *Trends Neurosci* 31:428–434. [CrossRef Medline](#)
- Faisal AA, Selen LP, Wolpert DM (2008) Noise in the nervous system. *Nat Rev Neurosci* 9:292–303. [CrossRef Medline](#)
- Fries P, Reynolds JH, Rorie AE, Desimone R (2001) Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* 291:1560–1563. [CrossRef Medline](#)
- Fries P, Womelsdorf T, Oostenveld R, Desimone R (2008) The effects of visual stimulation and selective visual attention on rhythmic neuronal synchronization in macaque area V4. *J Neurosci* 28:4823–4835. [CrossRef Medline](#)
- Goris RL, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17:858–865. [CrossRef Medline](#)
- Haslinger R, Pipa G, Lima B, Singer W, Brown EN, Neuenschwander S (2012) Context matters: the illusive simplicity of macaque V1 receptive fields. *PLoS One* 7:e39699. [CrossRef Medline](#)
- Hwang EJ, Andersen RA (2013) The utility of multichannel local field potentials for brain-machine interfaces. *J Neural Eng* 10:046005. [CrossRef Medline](#)
- Jensen O, Colgin LL (2007) Cross-frequency coupling between neuronal oscillations. *Trends Cogn Sci* 11:267–269. [CrossRef Medline](#)
- Jia X, Tanabe S, Kohn A (2013) gamma and the coordination of spiking activity in early visual cortex. *Neuron* 77:762–774. [CrossRef Medline](#)
- Kara P, Reinagel P, Reid RC (2000) Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27:635–646. [CrossRef Medline](#)
- Kelly RC, Smith MA, Kass RE, Lee TS (2010) Local field potentials indicate network state and account for neuronal response variability. *J Comput Neurosci* 29:567–579. [CrossRef Medline](#)
- Khawaja FA, Tsui JM, Pack CC (2009) Pattern motion selectivity of spiking outputs and local field potentials in macaque visual cortex. *J Neurosci* 29:13702–13709. [CrossRef Medline](#)
- Lachaux JP, Rodriguez E, Martinerie J, Varela FJ (1999) Measuring phase synchrony in brain signals. *Hum Brain Mapp* 8:194–208. [CrossRef Medline](#)
- Lakatos P, Karmos G, Mehta AD, Ulbert I, Schroeder CE (2008) Entrainment of neuronal oscillations as a mechanism of attentional selection. *Science* 320:110–113. [CrossRef Medline](#)
- Lee H, Simpson GV, Logothetis NK, Rainer G (2005) Phase locking of single neuron activity to theta oscillations during working memory in monkey extrastriate visual cortex. *Neuron* 45:147–156. [CrossRef Medline](#)
- Lin IC, Okun M, Carandini M, Harris KD (2015) The nature of shared cortical variability. *Neuron* 87:644–656. [CrossRef Medline](#)
- Masquelier T (2013) Neural variability, or lack thereof. *Front Comput Neurosci* 7:7. [CrossRef Medline](#)
- McFarland JM, Cui Y, Butts DA (2013) Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS Comput Biol* 9:e1003143. [CrossRef Medline](#)
- Mendoza-Halliday D, Torres S, Martinez-Trujillo JC (2014) Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nat Neurosci* 17:1255–1262. [CrossRef Medline](#)
- Mineault PJ, Khawaja FA, Butts DA, Pack CC (2012) Hierarchical processing of complex motion along the primate dorsal visual pathway. *Proc Natl Acad Sci U S A* 109:E972–E980. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Morishima Y, Akaishi R, Yamada Y, Okuda J, Toma K, Sakai K (2009) Task-specific signal transmission from prefrontal cortex in visual selective attention. *Nat Neurosci* 12:85–91. [CrossRef Medline](#)

- Nácher V, Ledberg A, Deco G, Romo R (2013) Coherent delta-band oscillations between cortical areas correlate with decision making. *Proc Natl Acad Sci U S A* 110:15085–15090. [CrossRef Medline](#)
- Niell CM, Stryker MP (2010) Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* 65:472–479. [CrossRef Medline](#)
- Nienborg H, Cohen MR, Cumming BG (2012) Decision-related activity in sensory neurons: correlations among neurons and with behavior. *Annu Rev Neurosci* 35:463–483. [CrossRef Medline](#)
- Nir Y, Fisch L, Mukamel R, Gelbard-Sagiv H, Arieli A, Fried I, Malach R (2007) Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. *Curr Biol* 17:1275–1285. [CrossRef Medline](#)
- Nishimoto S, Gallant JL (2011) A three-dimensional spatiotemporal receptive field model explains responses of area MT neurons to naturalistic movies. *J Neurosci* 31:14551–14564. [CrossRef Medline](#)
- Ohiorhenuan IE, Mechler F, Purpura KP, Schmid AM, Hu Q, Victor JD (2010) Sparse coding and high-order correlations in fine-scale cortical networks. *Nature* 466:617–621. [CrossRef Medline](#)
- Okun M, Steinmetz NA, Cossell L, Iacaruso MF, Ko H, Barthó P, Moore T, Hofer SB, Mrcic-Flogel TD, Carandini M, Harris KD (2015) Diverse coupling of neurons to populations in sensory cortex. *Nature* 521:511–515. [CrossRef Medline](#)
- Pachitariu M, Lyamzin DR, Sahani M, Lesica NA (2015) State-dependent population coding in primary auditory cortex. *J Neurosci* 35:2058–2073. [CrossRef Medline](#)
- Paninski L (2004) Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15:243–262. [CrossRef Medline](#)
- Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP (2008) Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature* 454:995–999. [CrossRef Medline](#)
- Purushothaman G, Bradley DC (2005) Neural population code for fine perceptual decisions in area MT. *Nat Neurosci* 8:99–106. [CrossRef Medline](#)
- Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP (2015) Attention stabilizes the shared gain of V4 populations. *Elife* 4:e08998. [CrossRef Medline](#)
- Rasch MJ, Gretton A, Murayama Y, Maass W, Logothetis NK (2008) Inferring spike trains from local field potentials. *J Neurophysiol* 99:1461–1476. [CrossRef Medline](#)
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647. [CrossRef Medline](#)
- Rust NC, Mante V, Simoncelli EP, Movshon JA (2006) How MT cells analyze the motion of visual patterns. *Nat Neurosci* 9:1421–1431. [CrossRef Medline](#)
- Saalmann YB, Pigarev IN, Vidyasagar TR (2007) Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316:1612–1615. [CrossRef Medline](#)
- Sahani M, Linden JF (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems* (Becker S, Thrun S, Obermayer K, eds), pp 125–132. Cambridge, MA: MIT.
- Schölvinc ML, Saleem AB, Benucci A, Harris KD, Carandini M (2015) Cortical state determines global variability and correlations in visual cortex. *J Neurosci* 35:170–178. [CrossRef Medline](#)
- Schroeder CE, Lakatos P (2009) Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci* 32:9–18. [CrossRef Medline](#)
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28:12591–12603. [CrossRef Medline](#)
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350. [Medline](#)
- Stein RB, Gossen ER, Jones KE (2005) Neuronal variability: noise or part of the signal? *Nat Rev Neurosci* 6:389–397. [CrossRef Medline](#)
- Tan AY, Chen Y, Scholl B, Seidemann E, Priebe NJ (2014) Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* 509:226–229. [CrossRef Medline](#)
- Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23:775–785. [CrossRef Medline](#)
- Treue S, Martínez Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579. [CrossRef Medline](#)
- Willmore BD, Prenger RJ, Gallant JL (2010) Neural representation of natural images in visual area V2. *J Neurosci* 30:2102–2114. [CrossRef Medline](#)
- Xing D, Yeh CI, Shapley RM (2009) Spatial spread of the local field potential and its laminar variation in visual cortex. *J Neurosci* 29:11540–11549. [CrossRef Medline](#)
- Zanos S, Zanos TP, Marmarelis VZ, Ojemann GA, Fetz EE (2012) Relationships between spike-free local field potentials and spike timing in human temporal cortex. *J Neurophysiol* 107:1808–1821. [CrossRef Medline](#)
- Zanos TP, Mineault PJ, Pack CC (2011) Removal of spurious correlations between spikes and local field potentials. *J Neurophysiol* 105:474–486. [CrossRef Medline](#)
- Zanos TP, Mineault PJ, Nasiotis KT, Guitton D, Pack CC (2015) A sensorimotor role for traveling waves in primate visual cortex. *Neuron* 85:615–627. [CrossRef Medline](#)
- Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* 370:140–143. [CrossRef Medline](#)