

Low-Frequency Local Field Potentials and Spikes in Primary Visual Cortex Convey Independent Visual Information

Andrei Belitski,^{1*} Arthur Gretton,^{1*} Cesare Magri,^{2,3*} Yusuke Murayama,¹ Marcelo A. Montemurro,⁴ Nikos K. Logothetis,^{1,5} and Stefano Panzeri^{2,4}

¹Max Planck Institute for Biological Cybernetics, D-72076 Tübingen, Germany, ²Italian Institute of Technology, Robotics, Brain and Cognitive Sciences Department, I-16163 Genoa, Italy, ³Dipartimento di Informatica e Comunicazione, Università degli Studi di Milano, I-20135 Milan, Italy, ⁴Faculty of Life Sciences, University of Manchester, Manchester M60 1QD, United Kingdom, and ⁵Imaging Science and Biomedical Engineering, University of Manchester, Manchester M13 9PT, United Kingdom

Local field potentials (LFPs) reflect subthreshold integrative processes that complement spike train measures. However, little is yet known about the differences between how LFPs and spikes encode rich naturalistic sensory stimuli. We addressed this question by recording LFPs and spikes from the primary visual cortex of anesthetized macaques while presenting a color movie. We then determined how the power of LFPs and spikes at different frequencies represents the visual features in the movie. We found that the most informative LFP frequency ranges were 1–8 and 60–100 Hz. LFPs in the range of 12–40 Hz carried little information about the stimulus, and may primarily reflect neuromodulatory inputs. Spike power was informative only at frequencies <12 Hz. We further quantified “signal correlations” (correlations in the trial-averaged power response to different stimuli) and “noise correlations” (trial-by-trial correlations in the fluctuations around the average) of LFPs and spikes recorded from the same electrode. We found positive signal correlation between high-gamma LFPs (60–100 Hz) and spikes, as well as strong positive signal correlation within high-gamma LFPs, suggesting that high-gamma LFPs and spikes are generated within the same network. LFPs <24 Hz shared strong positive noise correlations, indicating that they are influenced by a common source, such as a diffuse neuromodulatory input. LFPs <40 Hz showed very little signal and noise correlations with LFPs >40 Hz and with spikes, suggesting that low-frequency LFPs reflect neural processes that in natural conditions are fully decoupled from those giving rise to spikes and to high-gamma LFPs.

Key words: local field potential; macaque; neural coding; V1; movie; multiunit activity

Introduction

The neural signals commonly measured with extracellular micro-electrodes consist of time-varying spatial distributions of action potentials (“spikes”) superimposed on relatively slow varying field potentials, which relate well to subthreshold integrative processes in areas such as dendrites that are otherwise inaccessible. Population spiking and subthreshold activity can be to some extent studied distinctly by using band-separation techniques. Usually, the former type of activity, known as multiple-unit activity (MUA), is estimated by examining the power variation of the signal in the high-frequency range (typically 400–3000 Hz), whereas the so-called local field potential (LFP) is assessed by the power variation of the low-frequency range (e.g., 1–250 Hz). In addition, the spikes of a small population of neurons can be de-

tected and classified, providing information on single-unit activity.

A large number of experiments have suggested that, although MUA and single-unit spikes are primarily attributable to spiking activity of large pyramidal neurons, and thus measure the cortical output [for details and references, see Logothetis (2003) and Logothetis and Wandell (2004)], LFPs reflect the input and intracortical processing in a cortical area. In particular, LFPs are thought to reflect population synaptic potentials (Mitzdorf, 1987; Juergens et al., 1999), and other types of slow activity unrelated to synaptic events, including voltage-dependent membrane oscillations (Kamondi et al., 1998), and spike afterpotentials (Granit et al., 1963; Harada and Takahashi, 1983; Gustafsson, 1984; Buzsaki and Kandel, 1998; Buzsaki, 2002).

Both LFPs and spike rates encompass a range of different frequencies. The former are traditionally divided into bands, initially introduced in the human EEG literature, which correlate with distinct behavioral states (Lindsley and Wicke, 1974; Steriade and Hobson, 1976; Basar, 1980; Steriade, 1991), and are thought to originate from different types of neural events triggered by different processing pathways, such as sensory pathways or neuromodulation. Later physiological studies suggested other frequency parcellations thought to reflect the modulation of thalamocortical loops by the ascending network system and basal forebrain (Steriade et al., 1993).

Received Aug. 22, 2007; revised March 13, 2008; accepted April 16, 2008.

This work was supported by The Max Planck Society, The Royal Society, and Engineering and Physical Sciences Research Council EP/C010841, EP/E002331, and EP/E057101. M.A.M. was supported by a Medical Research Council Fellowship in Neuroinformatics. We thank N. Brunel, A. Mazzoni, and G. Notaro for useful discussions.

*A.B., A.G., and C.M. contributed equally to this work.

Correspondence should be addressed to either of the following: Stefano Panzeri, Italian Institute of Technology, Robotics, Brain and Cognitive Sciences Department, via Morego, 30, I-16163 Genoa, Italy, E-mail: stefano.panzeri@iit.it; or Nikos K. Logothetis, Max Planck Institute for Biological Cybernetics, Spemannstrasse 38, D-72076 Tübingen, Germany, E-mail: nikos.logothetis@tuebingen.mpg.de.

DOI:10.1523/JNEUROSCI.0009-08.2008

Copyright © 2008 Society for Neuroscience 0270-6474/08/285696-14\$15.00/0

The knowledge about how cortical LFPs encode sensory stimuli is still very limited compared with that regarding spikes. Recent studies (Pesaran et al., 2002; Kayser and König, 2004; Henrie and Shapley, 2005; Liu and Newsome, 2006) have begun to report how LFPs and spike responses in several frequency bands are tuned to one or few simple stimulus features. However, not much is yet known about which LFPs and spike frequencies are most important for the encoding of complex naturalistic sensory stimuli, and which LFP and spike frequencies encode either similar or independent features of naturalistic stimuli. To investigate these questions, we recorded LFPs, MUA, and spikes from the primary visual cortex of anesthetized macaques while repeatedly presenting a color movie lasting several minutes. We computed how much information about all existing visual features in the movie is conveyed by changes in power of LFP, MUA, and spikes, and we investigated which frequencies are tuned to similar stimulus features or share common sources of variability. To avoid any assumption regarding optimal band separation, we examined the characteristics and the stimulus selectivity of each individual frequency rather than the traditional EEG-partitioned fluctuations.

Materials and Methods

Four adult rhesus monkeys (*Macaca mulatta*) participated in these experiments. All procedures were approved by the local authorities (Regierungspräsidium) and were in full compliance with the guidelines of the European Community (EUVD 86/609/EEC) for the care and use of laboratory animals. Before the experiments, form-fitting head posts and recording chambers were implanted during an aseptic and sterile surgical procedure.

Implantation of recording chamber. The surgical procedures have been described in detail previously (Logothetis et al., 2002). Briefly, a skull form-specific head-holder and recording chamber were made out of magnetic resonance-compatible PEEK (polyether etherketone) (TecaPEEK) or medical-grade titanium, and implanted stereotaxically on the cranium of each animal under general anesthesia [balanced anesthesia consisting of isoflurane (1.3%) and fentanyl (3 µg/kg, i.v.) injections, with 1.8 L/min N₂O and 1.0 L/min O₂].

Animal preparation, induction and maintenance of anesthesia. After the animals were premedicated with glycopyrolate (0.01 mg/kg, i.m.) and ketamine (15 mg/kg, i.m.), an intravenous catheter was inserted and vital monitors (HP OmniCare/CMS; Hewlett Packard; electrocardiogram, noninvasive blood pressure, CO₂, SpO₂, temperature) connected. The monkeys were preoxygenated and short anesthesia was induced with fentanyl (3 µg/kg), thiopental (5 mg/kg), and succinylcholine chloride (3 mg/kg) for the intubation of the trachea. The animals were ventilated using an Ohmeda anesthesia machine (Ohmeda), maintaining an end-tidal CO₂ of 33 mmHg and oxygen saturation >95%. Balanced anesthesia was maintained with remifentanyl (typical, 1 µg · kg⁻¹ · min⁻¹). Muscle relaxation was achieved with mivacurium (5 mg · kg⁻¹ · h⁻¹). Body temperature was kept constant, and lactated Ringer's solution was given at a rate of 10 ml · kg⁻¹ · h⁻¹. During the entire experiment, the vital signs of the monkey and the depth of anesthesia were continuously monitored so that we could respond accordingly.

Drops of 1% ophthalmic solution of anticholinergic cyclopentolate hydrochloride were instilled into each eye to achieve cycloplegia and mydriasis. Refractive errors were measured and contact lenses (hard PMMA lenses; Wöhlk) were put on the monkey's eyes with continuous drops of saline throughout the experiment to prevent the eyes from drying. The lenses with the appropriate dioptric power were used to bring the animal's eyes into focus on the stimulus plane.

Electrophysiological recordings. The electrodes were arranged in a 4 × 4 square matrix (interelectrode spacing varied from 1 to 2.5 mm) and introduced each experimental session into the cortex through the overlying dura mater by a microdrive array system (Thomas Recording). We refer to the study by Eckhorn and Thomas (1993) for more details. Electrode tips were typically (but not always) positioned in the upper or middle cortical layers. The impedance of the electrode varied from 300 to

800 kΩ. Both spontaneous and stimulus-induced neural activity was collected and recorded for periods up to 6 min. Signals were amplified using an Alpha Omega amplifier system (Alpha Omega). Recordings were performed in a darkened booth (Illtec; Illbruck Acoustic). The site receptive fields were plotted manually and the position and size of each field were stored together with the stimulus parameters and the acquisition data.

Visual stimulation. A windows machine equipped with an OpenGL graphics card (Wildcat series; 3DLABS) was used to produce visual stimulation of 640 × 480 resolutions with 24 bit true color at 60 Hz for each eye. Stimulus images were generated by our own software written in C/Tcl and used OpenGL implementation. Frame rate of played movies was 30 Hz using hardware double buffering to provide smooth animation. The VGA outputs drove the video interface of a fiber-optic stimulus presentation system (Silent Vision; Avotec), and also drove the experimenter's monitor. The field of view of the system was 30 (horizontal) × 23 (vertical) degrees of visual angle, and the focus was fixed at 2 diopters. Binocular presentations were accomplished through two independently positioned plastic, fiber-optic glasses. Positioning was aided by a modified fundus camera (RC250; Carl Zeiss) that has an attachment to hold the projector on the same axis of the camera lens. After observing the foveal region, the projector was fixed relative to the animal. This process was done for each eye to bring the two projectors into the same visual alignment centered at the fovea. The contact lenses for the eyes had matched diopter with the Avotec projector, to focus images on the retina. To ensure accurate control of stimulus presentation, a photodiode was attached to the experimenter's monitor permitting the recording of the exact presentation time of every single frame.

The timing of stimulus presentation and of the acquisition of images was controlled by an industrial PC with one Pentium CPU (Advantec) running the QNX real-time operating system (QNX Software Systems) and our own software for experiment control and data acquisition. The control program presented the stimuli and acquired a number of physiological signals, such as respiration flow, inhaled and exhaled airway pressure, and plethysmogram.

Data filtering and extraction of field potentials and spiking activity. During recording, the signals were filtered in real time by the Alpha Omega Engineering processing system. The signals were high-pass filtered by a preamplifier and an on-line digital two-pole Butterworth filter with a cutoff frequency of 1 Hz, to remove fluctuations at DC level and increase the signal resolution to a level suitable for spike recordings.

The LFPs were extracted from the raw recordings by bandpass filtering the neural signal between the frequencies of 1 and 250 Hz. The filtering procedure was as follows. First, the neural signal was digitized, and its sampling rate was reduced from the original sampling rate by a factor of 3 from 20,835 to 6945 Hz. It was then bandpass filtered and down-sampled in two steps: first to a sampling rate of 1.5 kHz with a fourth-order Butterworth filter (500 Hz cutoff edge), and then to a rate of 500 Hz using a Kaiser window between frequencies of 1 and 250 Hz, sharp transition bandwidth (1 Hz), very small passband ripple (0.01 dB), and high stopband attenuation (60 dB). The two-step procedure was computationally more efficient than a single filtering operation to the final sampling rate. The sharpness of the second filter was used to avoid aliasing, without requiring a higher sampling rate attributable to a broad filter transition band, which would increase the computational cost of all subsequent operations. Forward and backward filtering was used to eliminate phase shifts introduced by the filters.

The band-limited LFP signals plotted in Figure 1 were generated using the same Kaiser filters as above with the appropriate bandwidth settings (1–5, 28–32 Hz, etc.). From these bandpassed signals, we first computed the Hilbert transform, and then we computed the instantaneous power (used for the plot in Fig. 1 only) as the square of the modulus of the Hilbert transform.

Power-line noise (50 Hz) was reduced during data acquisition by high grade of grounding in the electrical infrastructure of the recording room, by grounding every device to a single point, and by keeping the grounding short. The data recorded and analyzed here appeared only marginally affected by line noise contamination, as described next. The power spectrum of LFPs, for all 78 channels analyzed here, never presented a visible

sharp peak of power at ~ 50 Hz, suggesting that the amount of line noise contamination was small (for a representative example, channel 7 animal D04, of the LFP power spectrum, see Fig. 2A). The level of line noise contamination, measured (independently for each channel) as the difference between the value of the local peak of the spectrum ~ 50 Hz and the average of the spectrum in a 5 Hz region surrounding the peak but not including the peak value itself, was never more than 1% of the mean LFP power around the 50 Hz frequency region. Furthermore, the information carried by LFPs (see Fig. 3A) showed only very small local variations at 50 Hz and its harmonics 100 Hz (typically $<2\text{--}3\%$ of the mean LFP information around the considered frequency). Thus, line noise did not affect our conclusions in an appreciable way.

To extract the MUA, the 6945 Hz signal was first high-pass filtered at 100 Hz using Butterworth fourth-order filter, and then bandpass filtered in the range of 400–3000 Hz using a Kaiser window filter with a transition band of 50 Hz, stopband attenuation of 60 dB, and passband ripple of 0.01 dB. The absolute value of the signal was then taken, and it was decimated by a factor of 8 to reduce computation. Finally, it was low-pass filtered at 250 Hz and resampled at 500 Hz to match the sampling of the LFP signal. The MUA obtained in this way is thought to represent a weighted average of the extracellular spikes of all neurons within a sphere of $\sim 140\text{--}300$ μm around the tip of the electrode (Logothetis, 2003).

To extract spike times, the 6945 Hz signal was filtered in the range of 500–3500 Hz. The threshold for spike detection was set at 3.5 SDs. A spike was recognized as such only if the last spike occurred >1 ms earlier. This simple method does not always permit the isolation of single units. For the purpose of this study, however, recording the activity of a small group of neurons provided a sufficient (if not better) estimate of spike rate changes.

Because we found that spikes and MUA were very strongly correlated, in this study we will report almost entirely results obtained using spikes.

Spectral estimation. To assess how the power of LFP and spike rate oscillations changed over different segments of the movies, we divided each movie into nonoverlapping time windows of length T (unless otherwise stated, we used windows of length $T = 2.048$ s comprising $L = 1024$ datapoints sampled at 500 Hz). The neural signals measured in each window were evoked by the visual features presented up to that point. Thus, each window could be considered as containing the response to a different “stimulus,” labeled by an index s . The resulting neural oscillation in each stimulus time window s was quantified by computing the power spectrum r_f at each frequency f , independently for each trial.

The power spectrum in each window was obtained using the multitaper technique (Percival and Walden, 1993), which provides an efficient way to simultaneously control the bias and variance of spectral estimation by using multiple Slepian data tapers. The use of Slepian functions minimizes the bias, whereas the use of multiple orthogonal tapers on the same data minimizes the variance. The Slepian functions are defined in terms of their length L in time and their bandwidth W in frequency. For each choice of L and W , up to $K = 2LW - 1$ tapers are highly concentrated in frequency, having $>90\%$ of their power within the interval $[-W, W]$, and can be averaged for spectral estimation. To reduce the spectral bias, the average over tapers was computed using the adaptive procedure described by Percival and Walden (1993).

A simplified way of conceptualizing the multitaper method is that it provides an average over the local frequency ensemble with a range $2W$ (Percival and Walden, 1993). The value of W should be chosen on the basis of empirical considerations. Increasing W has the advantage of decreasing the variance of a single-trial spectral analysis, and thus increasing the information that can be extracted knowing the single-trial spectrum. Too large a W , however, may reduce the resolution of the spectral estimate, and may also introduce correlations between spectra at different frequencies because of excessive smoothing. The choice of $LW = 2$ appeared to achieve a good tradeoff, by providing low variance and relatively high peak information values across frequencies, while at the same time always correctly providing a result of zero trial-by-trial correlation between powers at different frequencies when we tested the method on simulated broadband signals that did not contain such correlations across frequencies. Thus, in the following, we will present results obtained using $LW = 2$. That said, our information analysis results

were robust to the choice of LW in the tested range $LW = 1$ to $LW = 4$ (results not shown).

The average LFP and Spike spectra plotted in Figure 2, A and C, respectively, were generated by averaging the power spectra, computed using the procedure described above, obtained in each stimulus time window and trial.

Information theoretic analysis. To determine how well the power r_f (of either LFP or spikes) at a certain frequency f encodes the visual features in the movie, we computed the mutual information $I(S; R_f)$ (Shannon, 1948) between the stimulus window in the movie and the power at frequency f as follows:

$$I(S; R_f) = \sum_s P(s) \sum_{r_f} P(r_f|s) \log_2 \frac{P(r_f|s)}{P(r_f)}, \quad (1)$$

where $P(s)$ is the probability of presentation of the stimulus window s (here equal to the inverse of the total number of stimulus windows), $P(r_f|s)$ is the probability of observing a power r_f at frequency f in response to a single trial to stimulus s , and $P(r_f)$ is probability of observing power r_f across all trials in response to any stimulus. $I(S; R_f)$ quantifies the reduction of the uncertainty about the stimulus that can be gained from observing, in a single-trial, the power at frequency f . Because we use base 2 logarithms, $I(S; R_f)$ is expressed in units of bits: 1 bit of information means that, on average, observation of the neuronal response in one trial reduces the observer’s stimulus uncertainty by a factor of 2.

The above single-frequency information analysis can be extended to compute how much information about the stimulus we can obtain when combining together the power r_{f_1} and r_{f_2} at two different frequencies [note that r_{f_1} and r_{f_2} may come from the same neural signal (e.g., two different LFP bands) or from two different neural signals (e.g., r_{f_1} from LFP and r_{f_2} from spikes)]. The mutual information that the joint knowledge of the powers r_{f_1} and r_{f_2} conveys about the stimulus is as follows:

$$I(S; R_{f_1} R_{f_2}) = \sum_s P(s) \sum_{r_{f_1} r_{f_2}} P(r_{f_1} r_{f_2}|s) \log_2 \frac{P(r_{f_1} r_{f_2}|s)}{P(r_{f_1} r_{f_2})}, \quad (2)$$

where $P(r_{f_1} r_{f_2} | s)$ is the probability of observing powers r_{f_1} and r_{f_2} at frequencies f_1 and f_2 , respectively, in response to a single trial with stimulus s ; and $P(r_{f_1} r_{f_2})$ is the probability of observing powers r_{f_1} and r_{f_2} across all stimuli.

The computation of the above information quantities requires the estimation of the stimulus–conditional response probabilities. These probabilities are not known a priori and must be measured experimentally from a finite number of trials. In our case, the number N_s of trials available to compute the stimulus–conditional response probabilities was equal to the number of repetitions of the movie (ranging from 12 to 44), whereas the number N of trials available to compute the unconditional response probabilities was N_s times the total number of stimulus windows S (because we used ~ 2 s stimulus windows and the movie was ~ 5 min long, S was ~ 150). The estimated probabilities suffer from finite sampling errors, which induce a systematic error (bias) in estimates of the information (Panzeri et al., 2007). We used the following procedure to correct for the bias. First, to facilitate the sampling of its probability, we discretized the response space by binning the power at each frequency into $R = 6$ equipopulated bins. We first computed the information by simply “plugging” the empirical binned probabilities into the information equations. The binned information measures still suffer from limited sampling bias. We corrected for this bias as follows. First, we used a quadratic extrapolation procedure (Strong et al., 1998) to estimate and subtract out the bias of each information quantity. When computing the information $I(S; R_{f_1} R_{f_2})$ conveyed by the joint observation of two powers, we also applied the “shuffling procedure” described by Montemurro et al. (2007) and Panzeri et al. (2007), which greatly reduces the bias of multidimensional information estimates. We then checked for residual bias by a “bootstrap procedure”: stimuli and responses were paired at random, and the information for these random pairings was computed. Because in this random case the information should be zero, the resulting value is an indication of a residual error: the bootstrap estimate of this error was therefore removed. To avoid information losses introduced by

too coarse a binning, it is useful to increase the number R of bins as much as possible. However, increasing R makes it progressively more difficult to control and correct for sampling bias. We chose $R = 6$ because it was the biggest number that consistently led to robust unbiased results under the sampling conditions of our experiment [we investigated this by extensive computer simulations on synthetic data with statistics close to that of the recorded power (results not shown); the performance of this bias correction procedure on realistically simulated neural spike trains has been reported in detail previously (Montemurro et al., 2007; Panzeri et al., 2007)]. We observed, however, that varying R over the range 3–10 did not change the structure of information as function of frequency; it just mildly affected the overall scale.

For all the above calculations of the mutual information, to speed up computation we considered only 65 positive Fourier frequency coefficients, spaced by ~ 3.8 Hz in the 1–250 Hz range.

Coefficients of variations of signal and noise. The information analysis tells which frequency ranges allow better discrimination among stimuli on a single trial, but it does not tell whether the information increase at certain frequencies is attributable to greater reliability across trials of the responses at these frequencies, or to an increased modulation by the stimulus. To separate out the contribution of stimulus modulation and of response variability, it is useful to characterize the response r_f in each stimulus window as “signal plus noise” (Gawne and Richmond, 1993; Averbeck et al., 2006) as follows:

$$r_f = \bar{r}_f + n_f \quad (3)$$

where we took as the “signal” \bar{r}_f the trial-averaged power (the bar denotes the average across trials at fixed stimulus) and took as the “noise” the trial-by-trial fluctuations n_f of the response around their averaged across trials. We stress that such “noise” does not necessarily reflect only noise in the real sense, but reflects all types of variations at fixed stimulus, which may well include the effect of various types of potentially important neural contributions such as modulation from a common ascending pathway.

To quantify the degree to which the signal changes with the stimulus, and thus can potentially encode the stimulus, we computed, independently for each frequency and channel, the signal CV, defined as the coefficient of variation (CV) of the trial-averaged power across the stimulus windows in the movie, as follows:

$$\text{SignalCV} = \frac{\text{std}_{\text{stim}}(\bar{r}_f)}{\langle \bar{r}_f \rangle_{\text{stim}}} \quad (4)$$

where $\langle \dots \rangle_{\text{stim}}$ and std_{stim} denote the mean and SD over stimulus windows, respectively. To quantify the averaged unreliability of the power across trials at fixed stimulus, we then computed, again independently for each frequency and channel, the “noise CV”: this is the CV across trials of the fluctuations of the power at a given frequency about its mean for each stimulus as follows:

$$\text{NoiseCV} = \left\langle \frac{\text{std}_{\text{tr}}(n_f)}{\bar{r}_f} \right\rangle_{\text{stim}} \quad (5)$$

where std_{tr} denotes the SD across trials at fixed stimulus window.

These signal and noise correlation quantities can be generalized to the two-dimensional case as follows:

$$2\text{DSignalCV} = \frac{\sqrt{\det(\text{cov}_{\text{stim}}\{\bar{r}_{f1}; \bar{r}_{f2}\})}}{\langle \bar{r}_{f1} \rangle_{\text{stim}} \langle \bar{r}_{f2} \rangle_{\text{stim}}} \quad (6)$$

$$2\text{DNoiseCV} = \left\langle \frac{\sqrt{\det(\text{cov}_{\text{tr}}\{n_{f1}; n_{f2}\})}}{\bar{r}_{f1} \bar{r}_{f2}} \right\rangle_{\text{stim}} \quad (7)$$

where cov_{stim} and cov_{tr} denote the covariance matrix across stimuli and trials, respectively. The two-dimensional expressions are analogous to the one-dimensional case, with the SD in the numerator replaced by the square root of the determinant of the covariance matrix, and the normalization factor in the denominator consisting of the product of the mean of two frequencies. The use of the determinant of the covariance matrix

ensures that these measures take into account that the covariation may be affected by correlations across the power at different frequencies of the type described next.

Signal, noise, and overall correlation coefficients across different frequencies. To determine which frequencies have related stimulus selectivity and which have shared sources of variability, we performed a linear analysis of correlations across frequencies of both the signal and the noise, as follows.

The correlations of the mean responses across different stimuli of two frequencies are called “signal correlations” (Gawne and Richmond, 1993; Panzeri et al., 1999; Averbeck et al., 2006) because they are entirely attributable to stimulus selectivity. The signal correlation coefficient was computed, for each frequency pair and channel, as the Pearson correlation across stimuli of the trial-averaged responses. Positive values indicate that the two frequencies have similar stimulus preference, whereas a zero value indicates that the two frequencies prefer totally uncorrelated stimuli.

Correlations manifested as covariations of the trial-by-trial fluctuation around the mean response to the stimulus are traditionally called “noise correlations” (Gawne and Richmond, 1993; Averbeck et al., 2006). Because these noise covariations are measured at fixed stimulus, they ignore all effects attributable to shared stimulation. To quantify the strength of noise correlations, we computed the Pearson correlation coefficient (across trials at fixed stimulus) of the trial-average-subtracted powers n_{f1} and n_{f2} , and then we averaged it over all stimulus windows. This quantifies the correlations of the variations around the mean at each trial and stimulus window. Positive values of noise correlation mean that when the power of one frequency fluctuates over its mean values, the power in the other frequency is also more likely to do so.

The overall amount of correlation across all trials and stimulus windows between the responses at frequencies f_1 and f_2 , is the result of both noise and signal correlations. The strength of correlations was quantified as the Pearson correlation coefficient (across all trials and stimulus windows) of the powers r_{f1} and r_{f2} . The overall correlation is typically higher than the noise correlation if both noise and signal are positively correlated, whereas it can be smaller than the noise correlation if the noise correlation is positive and the signal correlation is negative or null.

Results

LFP and spike responses during movie stimulation

We recorded from 44 sites in the primary visual cortex (V1) from four anesthetized monkeys, using an array of electrodes. From each site, we recorded LFPs and spikes during the presentation of a 3.5- to 6-min-long sequence from commercially available movies. Each movie extract was repeated 12–40 times to sample the probability distribution over the neural responses to each part of the film. In some cases, we were able to measure the responses of a recording site to repeated presentations of a second or third different movie. In such cases, the responses of the same site to different movies were analyzed separately. This gave us a total of 78 cases. Because the overall spectral response properties were highly repeatable across sites and movies, we collected all 78 cases together in our population analysis.

Before analyzing the neural data, it is interesting to consider the timescales of changes in the sensory stimulus by computing the average spectrum of frame-by-frame contrast variations in each of the movies used for stimulus presentation. Results are reported in supplemental Figure S1 (available at www.jneurosci.org as supplemental material). The power in the stimulus dropped with increasing frequency in all the movies.

We next analyzed how the neural signals responded to the movie. Figure 1, *A*, *C*, and *E*, shows the bandpassed LFP responses (in the 1–5, 28–32, and 72–76 Hz frequency ranges, respectively) for several trials, from a representative example recording site (electrode 7 in monkey D04) during the presentation of a 16-s-long movie sequence. The presentation of the movie elicited LFP

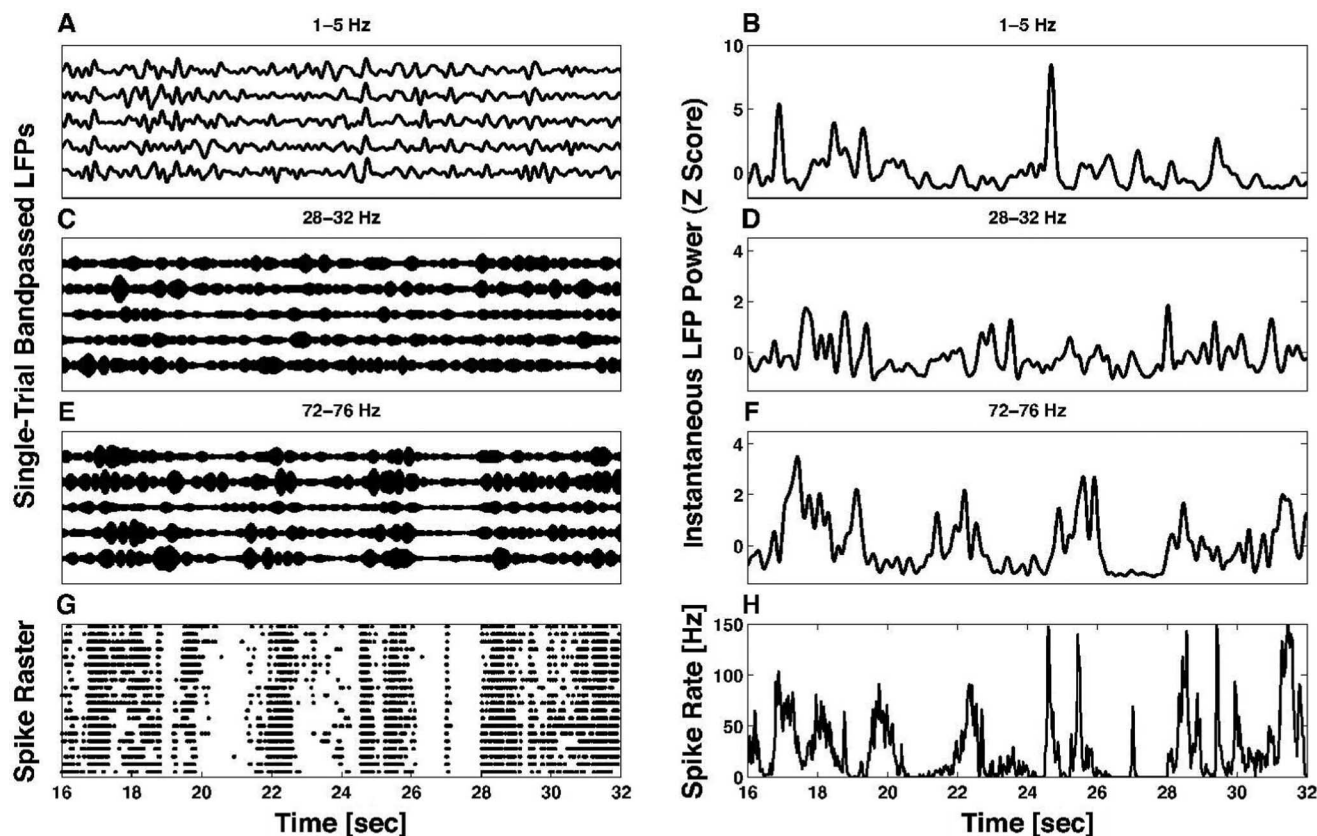


Figure 1. Illustration of the time course of the bandpassed LFP and of spiking activity during movie presentation, observed on a single electrode (animal D04, electrode 7). *A*, LFP traces (bandpassed in the 1–5 Hz frequency range) from five presentations of a 16-s-long movie extract. The traces were displaced on the vertical axis to make them distinguishable. *B*, The average over all trials of the time course of the instantaneous power of the 1–5 Hz LFP. The instantaneous power is normalized as a Z score (i.e., we subtracted the mean over the movie and divided by the SD over the movie of the instantaneous power). *C*, Time courses of the 28–32 Hz bandpassed LFP to the same five presentations of a 16-s-long movie extract as in *A*. *D*, The average over all trials of the time course of the instantaneous power of the 28–32 Hz LFP, with conventions as in *B*. *E*, Time courses of the 72–76 Hz bandpassed LFP to the same five presentations of a 16-s-long movie extract as in *A*. *F*, The average over all trials of the time course of the instantaneous power of the 72–76 Hz LFP, with conventions as in *B*. *G*, Raster plot of spike times (indicated by dots) resulting from repeated presentation of the selected 16 s movie extract. The five trials in *A*, *C*, and *E* correspond to the first five trials in *G*. *H*, Spike rate, averaged over all movie trials and computed in 4-ms-long sliding time bins, during the 16 s movie extract.

patterns that, both in the low-frequency (1–5 Hz) range (Fig. 1*A*) and in the high-gamma (72–76 Hz) range (Fig. 1*E*), were clearly modulated by the movie and repeatable across trials: episodes of high instantaneous power were elicited reliably in correspondence of certain scenes in the movie. In contrast, the LFP waveforms in the intermediate frequency range of 28–32 Hz (Fig. 1*C*) showed patterns that were not reliably associated to certain times during the movie. The average across all the 40 movie presentations of the instantaneous power (see Materials and Methods) is summarized in Figure 1, *B*, *D*, and *F*: the 1–5 and 72–76 Hz LFPs show high and significant modulations of power at several points of the movie, whereas the 28–32 Hz LFP does not. As shown in Figure 1, *G* (single-trial raster plot) and *H* (trial-averaged time-dependent spike rate), the spike rates clearly encoded the movie. The high spike rate episodes appeared to be associated more closely with episodes of high LFP power in the high-gamma LFP frequency range than at lower LFP frequencies. This suggests that gamma LFPs may be more closely related the stimulus-modulated spiking activity than low LFP frequencies.

To quantify more systematically the characteristics of LFP fluctuations at different frequencies, we computed the average LFP spectrum (see Materials and Methods). This is useful to document the relative change in magnitude of these fluctuations under different stimulation conditions and for different frequency bands that are differentially sensitive to stimulus–neural

or neural–neural interactions. Figure 2*A* shows the averaged spectrum of the LFP recorded from electrode 7 in monkey D04 (which was representative of the whole dataset). The highest LFP power was at low frequencies (≤ 8 Hz), and the power decreased steeply at increasing frequencies. The very low frequencies of the spectrum are likely to capture at least in part a succession of transient evoked responses (Fig. 1*A*). We compared the averaged LFP spectrum evoked during the movie to the LFP spectrum of the same electrode during spontaneous activity (measured in the absence of visual stimulation). There was an increase of power of ~ 4 dB with movie stimulation at very low frequencies (< 12 Hz). The evoked and spontaneous LFP spectrograms were similar at frequencies between 12 and 24 Hz, and at frequencies > 24 Hz there was more power with movie presentation. Consistent with previous studies (Frien et al., 2000; Henrie and Shapley, 2005), we found the most substantial power increase over spontaneous activity of the movie-evoked LFP in the gamma frequency region 40–120 Hz.

To quantify better how different parts of the movie modulate the fluctuations of LFPs, we divided the movie time into non-overlapping 2.048-s-long windows (each considered as a visual “stimulus”) and computed, independently for each trial, the spectrogram in each window. The temporal evolution of the trial-averaged LFP spectrogram in each stimulus time window during movie presentation is shown in Figure 2*B*. The trial-averaged LFP

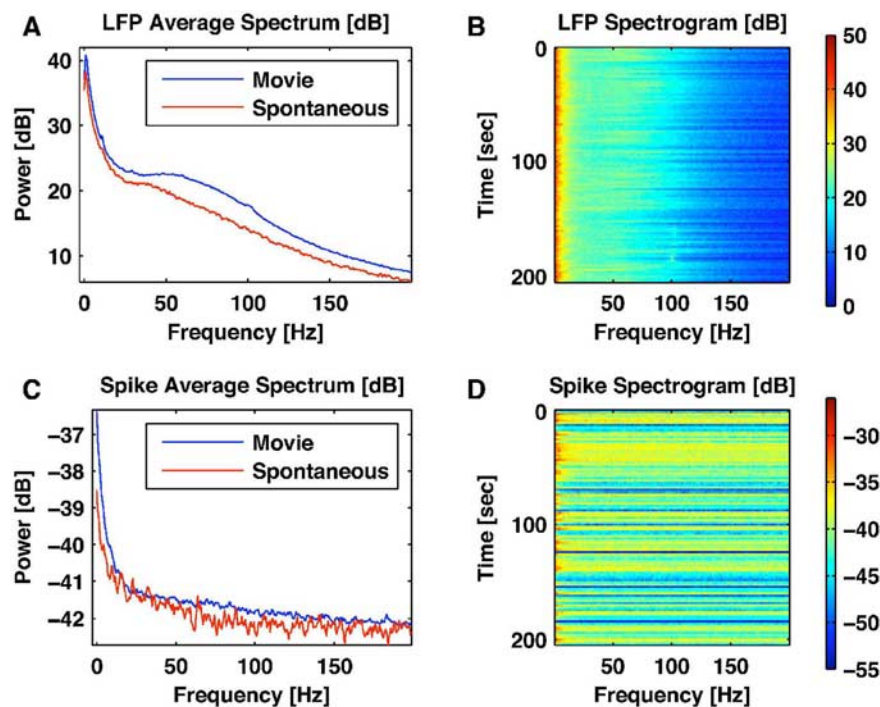


Figure 2. Spectral properties of LFPs and spikes during movie stimulation and spontaneous activity. All data from this figure were taken from the example channel (electrode 7, animal D04). **A**, The blue line plots the trial-averaged LFP spectrum over the entire movie presentation. The red line plots the average LFP spectrum during spontaneous activity. **B**, Trial-averaged spectrogram of the LFP (units of decibels) along the movie presentation. **C**, The blue line plots the trial-averaged spike spectrum over the entire movie presentation. The red line plots the average spike spectrum during spontaneous activity. **D**, Trial-averaged spectrogram of the spiking activity as a function of time during movie presentation (units of decibels).

power clearly changes over the course of the movie presentation, thus potentially conveying information about the movie, over a wide range of frequencies. Figure 2*B* indicates that frequencies in certain ranges (e.g., in the 70–120 Hz range) change more with the stimulus than other frequencies (such as those in the 15–30 Hz range). A detailed quantitative study of the stimulus modulation of the LFP power at each frequency will be presented in the next subsection.

A similar qualitative analysis was also performed for the spikes. The trial-averaged spectra of spikes evoked during movie presentation and during spontaneous activity were similar to each other in shape, with an overall spike power increase during movie stimulation (Fig. 2*C*). As with LFPs, the spike spectrum showed a high power at low frequencies, but the power decreased at higher frequencies to a nonzero saturating value. This saturation is commonly observed in spectral analysis of spike trains and is related to the near-Poissonian statistics of cortical spike trains, which causes the power spectrum of a spike train at high frequencies to be approximately proportional to the total number of spikes in the train (Bair et al., 1994; Pesaran et al., 2002). Figure 2*D* shows the trial-averaged time-dependent spectrogram of the spike train recorded in the same example channel. The power at lower frequencies (≤ 12 Hz) varied considerably during the movie, and was modulated by the movie to a greater degree than the power > 25 Hz. At any fixed time point in the movie, the frequencies > 25 Hz all appeared to have very similar mean power (Fig. 2*D*, horizontal stripes).

The information conveyed by the power at different frequencies

We next quantified the mutual information $I(S; R_f)$ (Eq. 1) that the power of either LFP or spikes at frequency f conveys about

which section of the movie was being presented. This characterization of information about the movie is independent of any assumption about which visual features in the movie are encoded by the neural signals: it thus reflects information about all the possible visual attributes occurring in the movie (de Ruyter van Steveninck et al., 1997). From here on, unless otherwise stated, we will report all results as mean (\pm SEM) across the entire population because individual sessions from different animals gave highly consistent results.

Figure 3*A* shows the mutual information $I(S; R_f)$ between the movie and the LFP power at frequency f . The first informative LFP region was the low-LFP-frequency range (1–8 Hz). The peak information value in this region was 0.22 ± 0.01 bits, and was reached at 4 Hz (Fig. 3*A*). The second highly informative LFP frequency range was in the high-gamma range of 60–100 Hz (Fig. 3*A*). The peak information value at high frequencies was 0.22 ± 0.01 bits. The position of the high-frequency peak was relatively stable across sessions, and occurred at $74 \text{ Hz} \pm 15 \text{ Hz}$ (mean \pm SD across different sessions). The information peak in the high-gamma range was broad in all experiments: the width (computed as the distance between points at 90% of the maximum information) was $31 \pm 11 \text{ Hz}$

(mean \pm SD across sessions). Because the movie frame rate was 30 Hz, and the movie had most of the power at low frequencies (supplemental Fig. S1, available at www.jneurosci.org as supplemental material), the highly informative LFP frequencies in the 60–100 Hz range were much more rapid than any variations in the stimulus time course. This demonstrates that cortical LFPs use a form of temporal encoding for representing naturalistic stimuli. It is interesting to note that the dependence of LFP information on frequency correlated well with the difference between movie-driven and spontaneous power (supplemental Fig. S2*A*, available at www.jneurosci.org as supplemental material), but did not correlate well with the ratio between movie-driven and spontaneous power. A potential explanation of this finding is that what conveys information is the stimulus modulation of the additional amount of power evoked during movie presentation (for an illustration, see supplemental Fig. S3, available at www.jneurosci.org as supplemental material), rather than the stimulus modulation of the proportion of power increase.

Despite their relatively high power, LFP frequencies in the range of 12–50 Hz conveyed a small amount of information (Fig. 3*A*), on average in the range of 0.06–0.07 bits (significantly non-null; $p < 0.05$; bootstrap test). The only noticeable feature in this middle frequency range was the tendency of LFP information to increase at ~ 30 Hz. Because the movie frame rate was 30 Hz (see Materials and Methods), this small feature (reflected only in LFPs and not in spikes) (see below) will not be considered in later discussion because it may reflect LGN discharge locked to the frame rate.

Given these observations, it is of interest to determine whether the increased information in the 1–8 Hz and 60–100 Hz ranges is attributable to greater reliability of the responses at these fre-

quencies across different trials, or to these signals being more readily modulated by the stimulus. To address this question, we computed, independently for each frequency and channel, the signal CV (Eq. 4), which quantifies the degree to which the trial-averaged response changes with the stimulus, and noise CV (Eq. 5), which quantifies the average unreliability of the power across trials at fixed stimulus. Figure 3B shows that (on average across all channels) the signal CV of LFPs, however, is maximal for low frequencies (1–8 Hz) and high-gamma frequencies (60–100 Hz). Figure 3B also shows that LFP frequencies in the intermediate range of 8–20 Hz have very little information because (despite the presence of some signal modulation) they have a particularly high noise CV. Thus, the maximally informative LFP frequency ranges correspond to the frequency intervals whose trial-averaged powers are most modulated by the movie, and the minima of LFP information correspond to the most noisy frequencies.

We next computed the mutual information $I(S; R_f)$ between the movie and the power of spikes at each frequency f (Fig. 3C). Unlike for LFPs, the information carried by spike power decreased monotonically with frequency, leveling off after ~ 20 Hz. The peak value of information (0.35 ± 0.03 bits) was reached at 1 Hz (the lowest frequency analyzed) for all experiments. Information in spike power saturated at a constant value of ~ 0.24 for frequencies > 20 Hz (Fig. 3C). As explained above, this is because at high frequencies the spike spectrum mainly reflects the global spike count, not high-frequency oscillations of the time-dependent spike rate. To correct for this effect, we normalized the spike spectrogram to one in each trial (Salinas et al., 2000). This procedure makes the spike power mostly independent of the total spike count. We recomputed the information carried by this normalized power as a function of frequency. After normalization, the information carried by spike power at frequencies > 20 Hz was very small, suggesting that only the spike power at low frequencies conveys information independent of that conveyed by the total spike count in the window.

Figure 3D reports the signal and noise CV of the spike power as function of the frequency. At low frequencies, there was a marked increase in signal CV, which was accompanied by a proportionally smaller increase in noise CV. Thus, the spike information peak at low frequencies corresponds to an increase in signal modulation.

The joint information carried by two different LFP frequencies and their signal and noise correlations

After our determination of which frequencies in the LFP and spike signals convey the most information, the next question is to investigate which LFP frequencies are tuned to similar features of the natural stimulus; which share sources of variability unrelated to the stimulus; and which are instead decoupled under natural conditions. Because different LFP frequencies are associated with different neural phenomena, addressing this question is crucial to

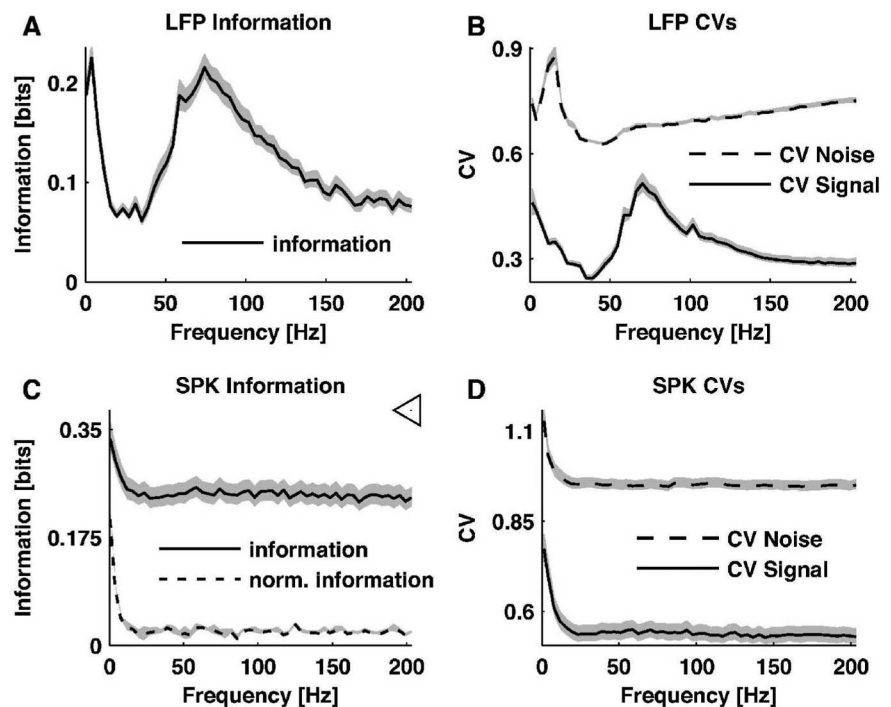


Figure 3. Information about the movie stimulus carried by the power at different frequencies of LFPs and spikes. **A**, The information $I(S; R_f)$ that the power of LFPs at frequency f conveys about the movie. **B**, The signal CV (solid line) and the noise CV (dashed line) of LFPs as function of frequency. **C**, The black solid line shows the information $I(S; R_f)$ that the power of spikes at frequency f conveys about the movie. The dashed line is the spike power information corrected for the effect of the overall spike count on the spectrum (see text). The arrowhead points to the value of the information about the movie carried by spike counts. **D**, The signal CV (solid line) and the noise CV (dashed line) of spikes as function of frequency. In all panels, the lines report the mean over the entire dataset, and the gray area around it shows its SEM. SPK, Spike.

understanding how these neural mechanisms complement each other and participate in stimulus coding.

We start by considering the power of LFPs at two different frequencies f_1 and f_2 , and compute how much information about the stimulus we can obtain from the joint observation of the powers r_{f_1} and r_{f_2} . This information is denoted as $I(S; R_{f_1}R_{f_2})$. Figure 4A reports the average over the dataset of $I(S; R_{f_1}R_{f_2})$ as a function of f_1 and f_2 . We recall that, on average, the single-frequency information (Fig. 3A) reached an identical value of 0.22 bits for low ($f \sim 1$ –8 Hz) or high-gamma ($f \sim 60$ –100 Hz) frequencies, respectively. Thus, it is of interest to understand how best to extract information about the stimulus by combining the observation at a pair of frequencies. We see from Figure 4A that the maximum of $I(S; R_{f_1}R_{f_2})$ is 0.43 ± 0.018 bits, and is attained by combining one low ($f_1 = 4$ Hz) frequency and a second high ($f_2 = 70$ Hz) frequency. This value of $I(S; R_{f_1}R_{f_2})$ is almost exactly equal to $I(S; R_{f_1}) + I(S; R_{f_2})$, the sum of the information carried by each the frequencies f_1 and f_2 individually. A slightly smaller information value (0.41 ± 0.019 bits) is achieved by taking both frequencies in the low range (f_1 and f_2 , 1–8 Hz). However the joint information about the stimulus in two high frequencies (f_1 and f_2 , ~ 60 –100 Hz) is much lower (0.33 bits) than the sum of the information carried by the two frequencies individually. Figure 4, B and C, reports the values of the two-dimensional signal and noise CVs of all frequency pairs. As in the single frequency case considered above, the frequency pairs with the highest information ($f_1 \approx 4$ Hz and $f_2 \approx 70$ Hz) were the ones with the highest signal CV. Figure 4B shows that the reduction of the joint information when considering two gamma frequencies (when compared with the maximal information case $f_1 \approx 4$ Hz




Figure 4. The information and the coefficient of variations carried by a pair of different LFP frequencies. *A*, The information about the movie $I(S; R_{f_1}, R_{f_2})$ that can be extracted from two joint LFP frequencies f_1 and f_2 . *B*, The two-dimensional signal CV quantifying the stimulus modulation of the trial-averaged power of two joint LFP frequencies f_1 and f_2 . *C*, The two-dimensional noise CV quantifying the reliability across trials of the power of two joint LFP frequencies f_1 and f_2 . Results in all panels reported the average over the entire dataset.

and $f_2 \approx 70$ Hz) is in part attributable to a reduction of the corresponding signal CV. This reduction in signal CV may be explained by the presence of the strong signal correlations in the gamma range that we will report below. The frequency pairs with a very high noise CV information (f_1 and f_2 in the 16–40 Hz range) also carry low information about the movie.

Thus, the best way to extract information about the stimulus is to combine one low frequency and one high frequency rather than two high frequencies. Why is it the case, and why does the joint information $I(S; R_{f_1} R_{f_2})$ conveyed by one low and one high frequency sum up, whereas the information conveyed by two high-gamma frequencies does not? If two frequencies were tuned to completely different stimulus features, and they did not share any source of noise, then we would expect $I(S; R_{f_1} R_{f_2})$ to be equal to $I(S; R_{f_1}) + I(S; R_{f_2})$, the sum of the information that each frequency conveys separately. The difference $I(S; R_{f_1}) + I(S; R_{f_2}) - I(S; R_{f_1} R_{f_2})$ is called information “redundancy” (Panzeri et al., 1999; Pola et al., 2003; Schneidman et al., 2003). When redundancy is positive, the two frequencies are said to convey redundant information about the stimulus; when redundancy is zero, the two frequencies are said to convey independent information.

The behavior of the joint information carried by pairs of LFP frequencies can be understood by quantifying the amount of signal and noise correlations shared by each pair of LFP frequencies. Signal correlations are known to contribute to decrease the joint information, because they mean that the two frequencies are tuned in similar way to stimuli and thus convey partly redundant messages. Noise correlations (i.e., correlations in the trial-to-trial fluctuations around the mean) can instead either decrease or increase the joint information, depending on their relationship to signal correlations. When signal and noise correlations have a similar sign, the joint information decreases, whereas when signal and noise correlations have opposite signs, the joint information increases. When signal correlation is zero or small, noise correlations have no or little impact on the joint information, no matter what their strength or sign is (Panzeri et al., 1999; Pola et al., 2003; Schneidman et al., 2003; Averbach et al., 2006).

The signal and noise correlations between frequencies of LFP activity in V1 over the entire 1–170 Hz frequency range are reported in Figure 5. In the following, we concentrate the discussion mostly on the correlation between frequency belonging to the low-frequency range and the high-gamma range, because there are the most informative ranges.

We first consider the signal and noise correlations shared by two frequencies f_1 and f_2 in the high-gamma (60–100 Hz) region. This is the region in which the signal correlation is maximal: on average over all data, signal correlation was ~ 0.7 – 0.8 when both frequencies are in the high-gamma 60–100 Hz (Fig. 5A). Thus, high-gamma frequencies are all similarly tuned to visual features. However, the amount of noise correlation (Fig. 5B) was small in this frequency band. The fact that there is small noise correlation but there is a very high degree of signal correlation explains why the information about the movie carried by pairs of high-gamma frequencies is much less than the sum provided by each individual frequency (i.e., why high-gamma frequencies are redundant).

This pattern of correlation is illustrated more intuitively in Figure 6, second row, by means of scatterplots of the power recorded from the example of electrode 7 from monkey d04nm1. To illustrate signal correlation, we produced a scatterplot of the joint trial-averaged LFP powers in each stimulus time window at two example frequencies (72 vs 76 Hz), both in the high-gamma region. A large positive linear signal correlation of the mean re-

sponses to each stimulus was clearly visible (Fig. 6D). To illustrate noise correlation, we produced a scatterplot of the variations around the mean at each trial and stimulus window (Fig. 6E). This scatterplot showed no correlation (r^2 value, not significant; $p > 0.3$; bootstrap test) between the “noise” (i.e., variations about the mean) at fixed stimulus. The overall effect of stimulus and noise correlations is a positive linear overall correlation between both frequencies (Fig. 6F).

Let us consider the LFP bands at low frequencies (<24 Hz). These frequencies were the only ones to present a high amount of noise correlation, which reached a peak value of 0.37 ± 0.02 in this range. In this low-frequency region, there was positive signal correlation. Thus, frequencies in the range of 1–24 Hz share both sources of noise and of signal. However, signal correlation was much less strong than in the high-gamma range. Figure 5C shows that signal correlation adds to the noise correlation and contributes to increase the overall correlation. The fact that the amount of noise correlation is almost equal to the overall amount of correlation means common sources of noise are the most important contribution to covariations of low frequencies, and shared signal information has a weaker quantitative impact. These results are confirmed and explicitly illustrated in the scatterplots of the powers taken from the example recording channel (Fig. 6, first row). It is well known that positive noise correlations reduce the joint information when they are coupled with positive signal correlation. However, the effect of noise correlation on information is very small if signal correlation is small (Averbeck et al., 2006). Thus, the relatively low amount of signal correlation between two low LFP frequencies explains why the information carried by two low LFP frequencies is close to the sum of the information they carry individually, even in the presence of strong noise correlations.

Finally, we consider the relationship between low LFP frequencies and high LFP frequencies. We found there was a cutoff at ~ 40 Hz that separates unrelated regions: any power modulation at a frequency <40 Hz does not share any signal or noise correlation with the power modulation at higher frequencies (Fig. 5A,B). This strongly suggests that frequencies at different sides of the 40 Hz cutoff do not originate from any common neural mechanisms. As a consequence, there is negligible information redundancy between frequencies on either side of 40 Hz. This explains why the best way to extract information about the stimulus is to combine one low LFP frequency and one high LFP frequency.

To investigate the difference between sensory-evoked and spontaneous activity, we recorded spontaneous neural activity from the same sites during 5–10 repeated 5-min-long periods without visual stimulation (blank screen). We then computed the correlation across all frequency pairs, using the previously described techniques. This spontaneous correlation should be compared with the noise correlation obtained during visual stimulation because the overall correlation also reflects signal correlations, which cannot be present during spontaneous activ-




Figure 5. Correlations among pairs of different LFP frequencies. **A**, The signal correlation between the trial-averaged powers observed at two different LFP frequencies f_1 and f_2 during movie presentation. **B**, The noise correlation between the LFP powers at frequencies f_1 and f_2 during movie presentation. **C**, The overall correlation between the LFP power at frequencies f_1 and f_2 during movie presentation. **D**, The correlation between the LFP power at frequencies f_1 and f_2 during spontaneous activity. Results in all panels reported the average over the entire dataset.

ity (note that for spontaneous activity, noise and overall correlations are equal because there is no stimulus). We found (Fig. 5D) that the correlation was still highest for low frequencies (1–24 Hz). The extent of the region with high correlation within the 1–24 Hz region was slightly reduced compared with the noise correlation observed in the same frequency range during movie stimulation. This suggests that most, but not all, components of the noise correlations among low LFP frequencies are robust to the removal of sensory stimulation.

Strong signal and noise correlations across all spiking frequencies

The single frequency analysis suggested that the power of the spike train at high frequency primarily reflected spike counts. We strengthen this conclusion by analyzing the joint information in spike power fluctuations at pairs of different frequencies.

Results (averaged over the population) are reported in Figure 7. The joint information $I(S; R_{f_1}R_{f_2})$ (Fig. 7A) that can be gained from the knowledge of any two spike frequencies is maximal when one of the two frequencies is very low (1 Hz). The CV analysis (Fig. 7B,C) reveals that, once again, the maxima of $I(S; R_{f_1}R_{f_2})$ corresponded to the maxima of the joint signal CV. Figure 7, D and E, shows that there was a large positive amount of both signal correlation (~ 0.85) and noise correlation (~ 0.4), which was again approximately constant across all spike frequencies >20 Hz. This shows that all high-frequency spike powers are all highly redundant to each other, most likely because they are all redundant to the spike count. To demonstrate the relationship between spike count and high-frequency spike power fluctuation, we computed the signal and noise correlation between the spike count and spike power at high frequencies: signal correlation was 0.9 and noise correlation was 0.6. Together, these results

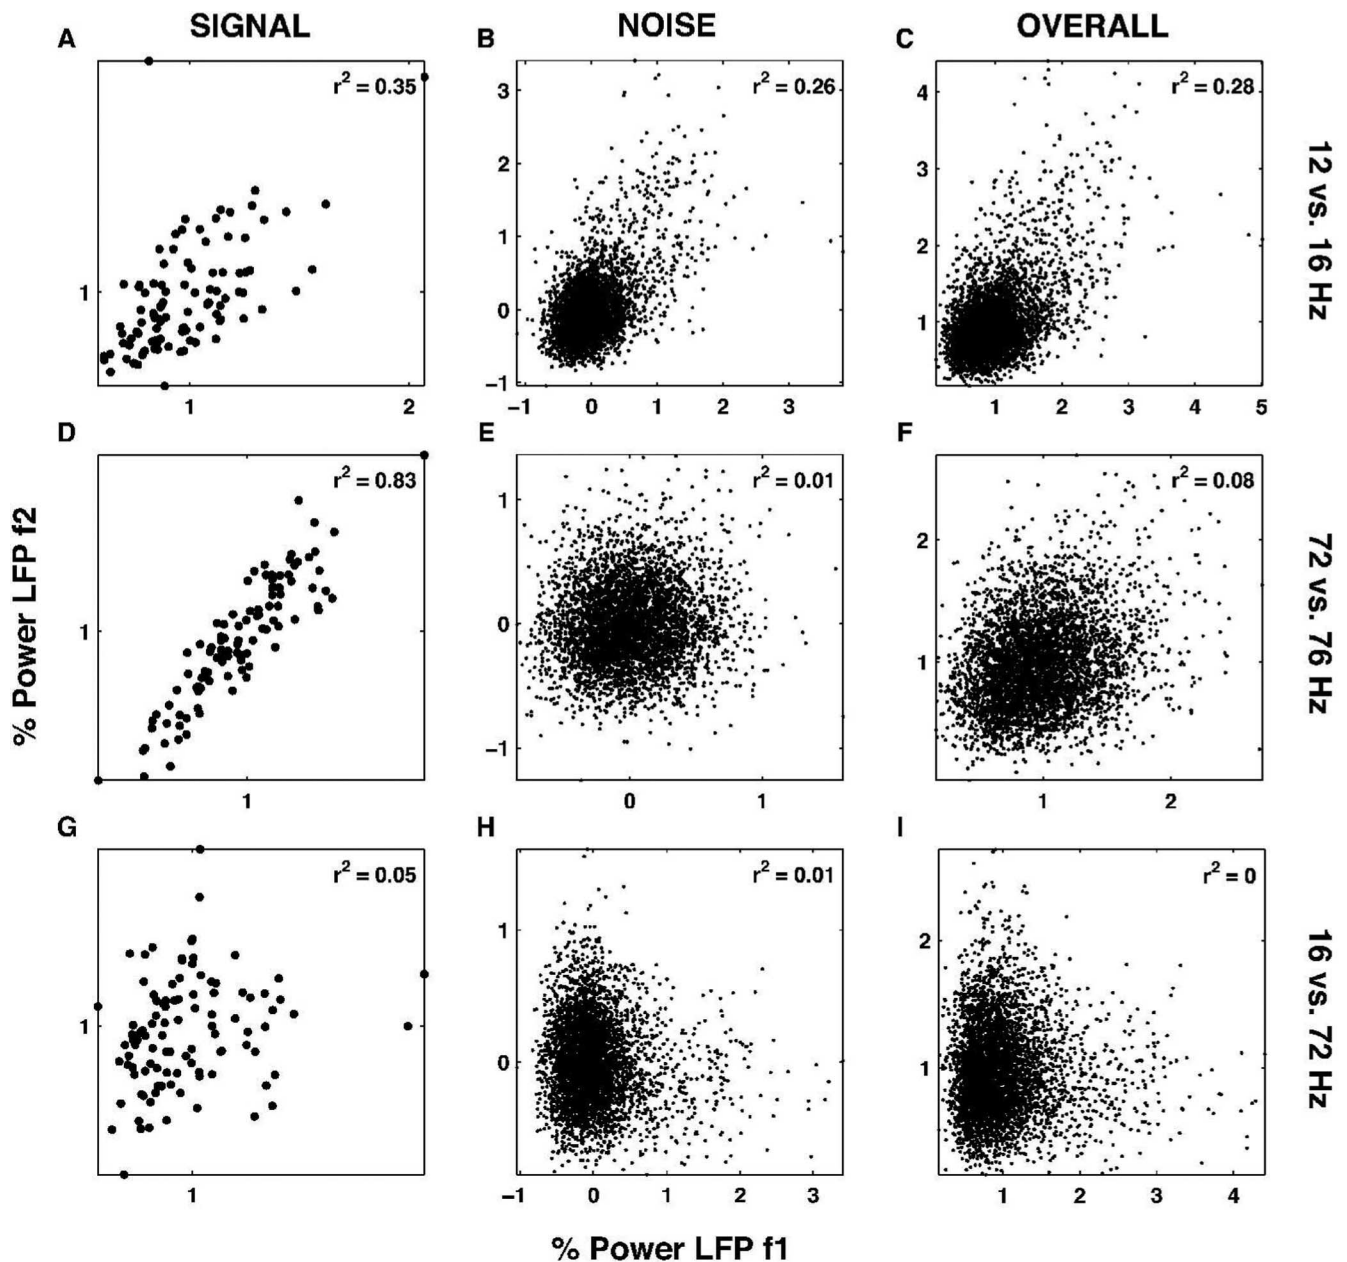


Figure 6. Signal, noise, and overall correlations of the power across different LFP bands. *A–I* report data from an example electrode (no. 7, monkey d04nm1). *A*, *D*, and *G* illustrate the amount of signal correlation by showing scatter plots of the trial-averaged power response at different frequency pairs (12 vs 16, 72 vs 76, and 16 vs 72 Hz, respectively) in each 2.048-s-long window of the movie presentation. *B*, *E*, and *H* (again corresponding to 12 vs 16, 72 vs 76, and 16 vs 72 Hz, respectively) illustrate noise correlation by showing the scatterplot of the fluctuations (around the mean across trials in the same time window) of the power in each trial and window. *C*, *F*, and *I* illustrate the amount of overall correlation by showing the scatterplot of the power in each trial and window at two different frequency pairs (12 vs 16, 72 vs 76, and 16 vs 72 Hz, respectively). The insets in each panel report the r^2 values of the linear fit of each scatterplot. For clarity of illustration, the power in each frequency has been normalized to 1 when summed over all windows and trials.

demonstrate the stimulus information at high frequencies in the spike spectrogram is to a substantial degree the same information conveyed by spike count.

Redundancy and independence across different LFP and spiking frequencies

We now investigate whether LFPs and spikes convey common stimulus information at particular frequencies. We start by considering how best to extract information about the stimulus by combining the observation of one spike frequency f_1 and one LFP frequency f_2 . Figure 8A shows that the joint information that

could be extracted from the combined observation of the LFP and spike spectra was relatively constant across spiking frequencies. This is because, as explained above, the power at all spike frequencies mostly reflects the spike count. The joint information $I(S; R_{f_1} R_{f_2})$ obtained is maximal (0.57 ± 0.03 bits) when pairing a low-frequency LFP power (4 Hz) to the lowest frequency spike power (1 Hz). This value was approximately equal to the sum of the information conveyed by the 4 Hz LFP and the 1 Hz spike powers (0.22 and 0.35 bits, respectively). A local maximum of lesser size (0.52 ± 0.03 bits) was reached (Fig. 8A) when combining a high-gamma LFP power (60–100 Hz) with a very low-

frequency spike power (1–4 Hz). The joint information carried the low-frequency spike power and the high-gamma LFP power was less than the sum of the two corresponding information values.

We now investigate whether there is any sharing of signal and noise correlations between different LFP and spiking frequencies f_1 and f_2 . First, we found that LFP frequencies <40 Hz were very independent from the spike spectra. There was very little signal, noise, or overall correlation between <40 Hz LFP oscillations and spikes. This means that these low-frequency LFP oscillations are totally decoupled from spikes under natural stimulation condition and thus convey independent information about the stimulus. We found that LFP frequencies >40–50 Hz showed positive signal and noise correlations with the spikes. The informative high-gamma LFP fluctuations had a sizable signal correlation of ~ 0.5 and noise correlation of 0.1 with spikes. The presence of positive signal and noise correlation explains why there is information redundancy between high-gamma LFPs and spikes.

Finally, it is possible that spikes and LFPs are mostly independent because LFPs reflect the contribution of a large population near the electrode, whereas the spikes only reflect the activity of a small number of neurons. The neurons from which we collect the spikes may happen by chance to be outliers that do not reflect the typical local spiking behavior. To address this concern, we considered the relationship between LFP and MUA recorded from the same electrode. The MUA signal that was extracted by our recordings reflects the weighted sum of extracellular spikes of all neurons sufficiently close to the tip of the electrode (Logothetis, 2003), whereas the spikes signal consists of detected spike arrival times of a small neural population with high enough signal-to-noise ratio (see Materials and Methods). We found (results not shown) that the signal and noise correlation structure between LFPs and MUA was almost identical with that found between LFPs and spikes and reported in Figure 8. Therefore, we can rule out the concern that LFPs and spikes were mostly independent because the extracted spikes were not representative of the overall spiking activity around the tip of the electrode.

Discussion

We investigated the relationship between the visual information carried by different frequencies of LFPs and spikes. Characterizing these relationships is useful to understand which neural signals can optimally communicate with brain-machine interfaces (Donoghue et al., 1998; Andersen et al., 2004), and to better interpret the blood oxygenation level-dependent response, which, under many conditions, correlates with dif-

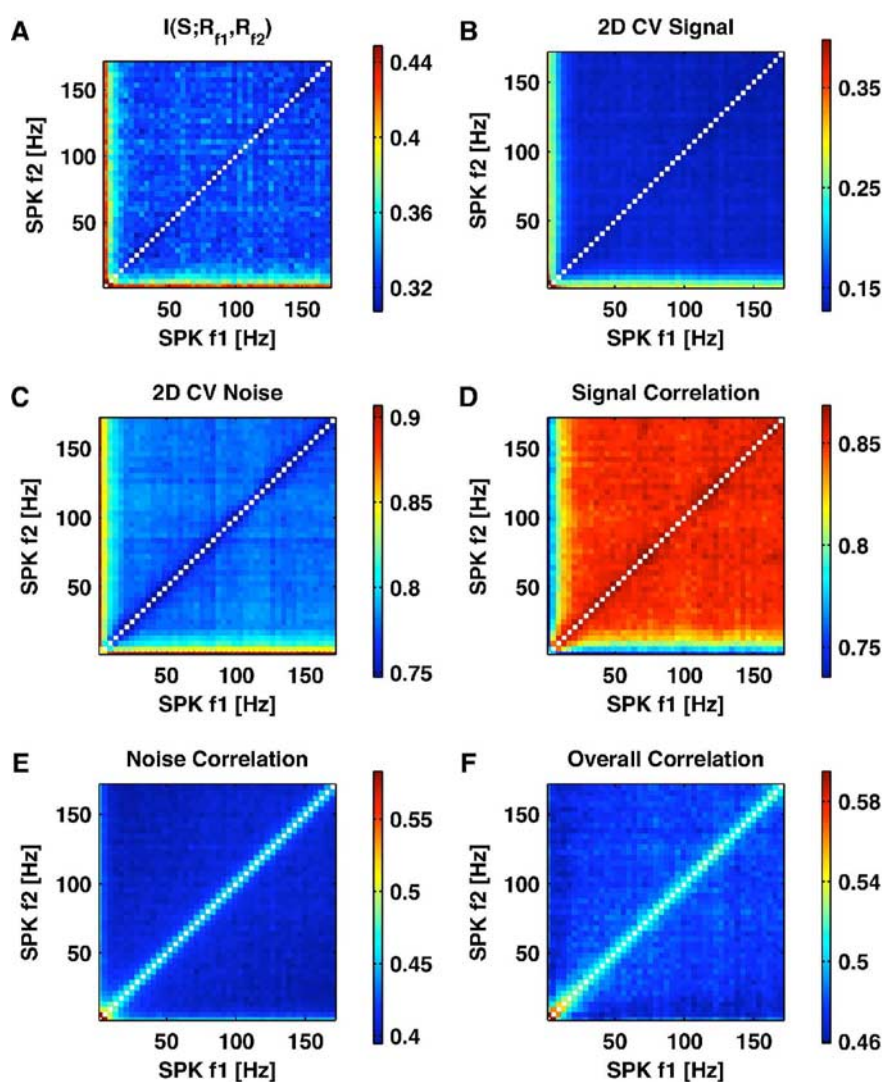


Figure 7. Information and correlations of pairs of different spike frequencies. **A**, The information about the movie $I(S; R_{f_1}, R_{f_2})$ that can be extracted from two joint spike frequencies f_1 and f_2 . **B**, The two-dimensional signal CV quantifying the stimulus modulation of the trial-averaged power of two joint spike frequencies f_1 and f_2 . **C**, The two-dimensional noise CV quantifying the reliability across trials of the power of two joint spike frequencies f_1 and f_2 . **D**, The signal correlation between the trial-averaged powers observed at two different spike frequencies f_1 and f_2 . **E**, The noise correlation between the spike powers at frequencies f_1 and f_2 . **F**, The overall correlation between the spike power at frequencies f_1 and f_2 . Results in all panels report data collected during movie presentation and show the average over the entire dataset. SPK, Spike.

ferent aspects of both spikes and LFPs (Logothetis et al., 2001; Kayser et al., 2004; Niessing et al., 2005; Nir et al., 2007).

Several recent studies have examined the tuning properties of LFPs to visual stimuli and their relationship to the tuning of spikes (Frien et al., 2000; Pesaran et al., 2002; Kayser and Konig, 2004; Henrie and Shapley, 2005; Kreiman et al., 2006; Liu and Newsome, 2006; Kraskov et al., 2007). Our study builds on this knowledge to add a number of critical new results. First, we characterized the tuning of LFPs and spikes to a color movie, which is a rich and naturalistic dynamic visual stimulus. Second, we quantified how much information about all stimulus features appearing in the movie is available, on a single trial, in each frequency of LFPs and spikes. Third, we determined in detail the signal and noise correlations existing between neural signals at different frequencies.

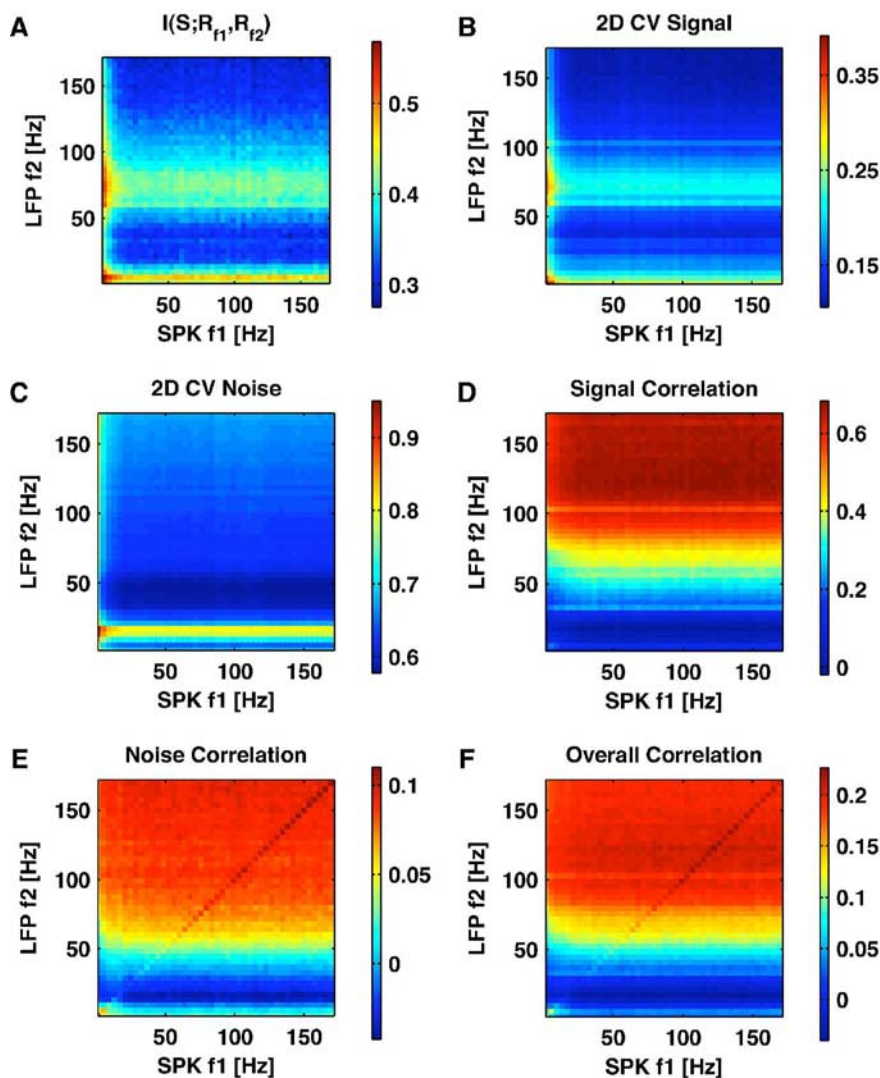


Figure 8. Information and correlations between LFP and spike frequencies. *A*, The information about the movie $I(S; R_{f_1}, R_{f_2})$ that can be extracted from the joint observation of a spike frequency f_1 and an LFP frequency f_2 . *B*, The two-dimensional signal CV quantifying the stimulus modulation of the trial-averaged power of a spike frequency f_1 and an LFP frequency f_2 . *C*, The two-dimensional noise CV quantifying the reliability across trials of the power of a spike and an LFP frequency. *D*, The signal correlation between the trial-averaged powers of a spike and an LFP frequency. *E*, The noise correlation between a spike frequency f_1 and an LFP frequency f_2 . *F*, The overall correlation between a spike and an LFP frequency. Results in all panels report data collected during movie presentation and show the average over the entire dataset. SPK, Spike.

Temporal scale of the LFP and spiking signal

We first determined the frequencies at which the power modulation of LFPs conveys maximal information about the movie. Consistent with previous studies using less rich visual stimuli (Pesaran et al., 2002; Henrie and Shapley, 2005; Liu and Newsome, 2006), we found that the power of LFPs in the high-gamma range (60–100 Hz) conveyed maximal visual information about movie stimuli. In contrast to most previous studies, however, we found that low-frequency LFPs (in the 1–8 Hz range) were equally highly informative about the movie. The previously reported lack of visual stimulus selectivity at low LFP frequencies was commonly ascribed to the hypothesis that cortex acts as capacitive filter, spatially blurring low-frequency signals, which travel longer distances than those at high frequencies (Bedard et al., 2006). The hypothesis of capacitive filtering proved wrong according to recent detailed measurements of the cortical impedance spectrum

(Logothetis et al., 2007). An alternative explanation for the discrepancy between the current findings and previous reports is the nature of the visual stimulation. Movies permit the assessment of stimulus encoding in a considerably more general way than receptive field-specific stimuli. LFPs might in principle be selective for stimulus features that were never explored in previous studies and are present in naturalistic movies. Because movies contain most power at low frequency, it is conceivable that some of the features for which LFPs are selective are characterized by slow fluctuations and thus are reflected in LFPs at low frequency. Additional studies, such as reverse correlations between stimulus and low-frequency power modulation, are needed to characterize what stimulus features are encoded by low-frequency LFPs during natural vision.

LFPs in the 12–24 Hz range did not increase the power during movie presentation and [unlike what was reported, e.g., in the motor cortex (Rubino et al., 2006)] were very unreliable and not very informative, suggesting that in our experimental conditions they mainly reflect a neuromodulatory input mostly unrelated to the stimulus. LFPs in the 24–40 Hz range did show an increase of power during movie presentation, but this visually induced power was not reliably modulated from movie scene to movie scene (low signal CV) and carried little information, suggesting that (although partly modulated by a sensory-related component) they are also primarily affected by neuromodulations.

In contrast to LFPs, changes in the power of the spike signal conveyed information only at low frequencies. The coexistence of precisely timed and reliable gamma LFP oscillations with unreliable spiking activity at the same frequency scales is possible when gamma oscillations are generated by strong recurrent inhibitory and excitatory cortical interactions (Brunel and Wang, 2003). That said, the absence of high-frequency spike rate power modulations is still fully compatible with the presence of spike timing codes (Hopfield, 1995; de Ruyter van Steveninck et al., 1997; Victor and Purpura, 1998; Golisch and Herz, 2005) such as latency codes (Panzeri et al., 2001) or temporal codes based on spike timing relative to LFP phases (Buzsaki, 2002; Mehta et al., 2002; Huxter et al., 2003; Montemurro et al., 2008).

Shared sources of noise and signal between different LFP and spike frequencies

We discovered three relationships between power modulation at different LFP frequencies, and between LFPs and spikes, that we found of particular interest. We now discuss their implication in terms of understanding the mechanisms underlying such signals.

First, we found that LFP frequencies <40 Hz do not share any

substantial signal or noise correlations with higher LFP frequencies or with spikes. This suggests that the power of LFP fluctuations <40 Hz reflect different neural processes from those giving rise to higher-frequency LFPs and to spikes, and that these processes are mostly decoupled in naturalistic stimulation conditions.

Second, we found that LFP bands in the low-frequency range (1–24 Hz) had the strongest noise correlations of the whole LFP frequency range studied, and had relatively little signal correlations. The comparably small amount of signal correlation (i.e., the weak similarity in stimulus selectivity) explains why, although they share a substantial amount of stimulus unrelated variability, the joint information carried by two low LFP frequencies is almost equal to the sum of the two informations carried by the individual frequencies. This result implies that the separation of the 1–24 Hz frequency range into smaller frequency bands, such as the delta, alpha, and theta bands used in the EEG literature, is indeed useful, because each such band adds almost independent information about the external correlates. However, it is important to note that the low amount of redundant information between bands is attributable to the relatively small signal correlation, and not because these bands are truly uncorrelated and decoupled phenomena. The noise correlations in the 1–24 LFP range were present both during movie stimulation and during spontaneous activity. One potential explanation of this behavior is that low-frequency LFPs in the region 1–24 Hz are generated in part by some common ascending pathways (because shared inputs are able to generate both signal and noise correlations), and especially share a strong common source of a noise unrelated to the stimulus, for example, a common source of neuromodulation (which would more easily account for the presence of noise correlations in absence of visual stimulation and in absence of a very strong signal correlation). The sharing of both signal and noise correlation implies that the low-frequency LFP bands do not originate from different types of neural events triggered by different processing pathways, even if they add up independent information about the external correlates.

Third, we found that LFP in the 60–100 Hz high-gamma band shared very little noise correlation during visual stimulation, but shared the highest observed signal correlation across all LFP frequencies. The 60–100 Hz LFP band also had the highest proportional power increase during visual stimulation, and also had a substantial signal correlation with spikes. This latter finding is consistent with the significant but loose relationship between auditory cortex spikes and gamma power recently reported by Nir et al. (2007). All these facts are consistent with the view that high LFP frequency signals might arise within the same localized cortical cluster: namely, they may originate when a sensory input activates inhibitory or excitatory–inhibitory recurrent loops between local cortical populations (Brunel and Wang, 2003; Fries et al., 2007). The high signal correlation between 60 and 100 Hz LFP frequencies, and their overall power increase during stimulus presentation, may indicate that all fluctuations in this range are generated by such local recurrent dynamics driven by the same external stimuli. The relatively small amount of noise correlations between 60 and 100 Hz LFPs indicates that this variability is not simply a global effect on all frequencies in this range, such as trial-to-trial variation in a single neuromodulation factor, but more likely originates from complex network interactions. The fact that the signal correlation between 60 and 100 Hz LFPs and spikes is strong may arise because a local stimulus-related input would affect both the amount of spiking activity of pyramidal neurons and

the engagement of the recurrent loop in the region around the tip of the electrode, which in turn reflects into the high-gamma LFP power. The same local generation mechanism is also consistent with our finding that there is only a partial relationship between the stimulus tuning of high-gamma LFPs and spikes, because LFPs may capture factors, such as periodic firing of interneurons, which may not be directly accessible from the spikes of large pyramidal neurons.

References

- Andersen RA, Musallam S, Pesaran B (2004) Selecting the signals for a brain-machine interface. *Curr Opin Neurobiol* 14:720–726.
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366.
- Bair W, Koch C, Newsome WT, Britten K (1994) Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J Neurosci* 14:2870–2892.
- Basar E (1980) EEG-brain dynamics: relation between EEG and brain evoked potentials. Amsterdam, New York, Oxford: Elsevier/North Holland Biomedical.
- Bedard C, Kroger H, Destexhe A (2006) Model of low-pass filtering of local field potentials in brain tissue. *Phys Rev E* 73:051911.
- Brunel N, Wang XJ (2003) What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. *J Neurophysiol* 90:415–430.
- Buzsaki G (2002) Theta oscillations in the hippocampus. *Neuron* 33:325–340.
- Buzsaki G, Kandel A (1998) Somadendritic backpropagation of action potentials in cortical pyramidal cells of the awake rat. *J Neurophysiol* 79:1587–1591.
- de Ruyter van Steveninck RR, Lewen GD, Strong SP, Koberle R, Bialek W (1997) Reproducibility and variability in neural spike trains. *Science* 275:1805–1808.
- Donoghue JP, Sanes JN, Hatsopoulos NG, Gaal G (1998) Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. *J Neurophysiol* 79:159–173.
- Eckhorn R, Thomas U (1993) A new method for the insertion of multiple microprobes into neural and muscular tissue, including fiber electrodes, fine wires, needles and microsensors. *J Neurosci Methods* 49:175–179.
- 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. *Eur J Neurosci* 12:1453–1465.
- Fries P, Nikolic D, Singer W (2007) The gamma cycle. *Trends Neurosci* 30:309–316.
- Gawne TJ, Richmond BJ (1993) How independent are the messages carried by adjacent inferior temporal cortical neurons? *J Neurosci* 13:2758–2771.
- Gollisch T, Herz AVM (2005) Disentangling sub-millisecond processes within an auditory transduction chain. *PLoS Biol* 3:144–154.
- Granit R, Kernell D, Smith RS (1963) Delayed depolarization and the repetitive response to intracellular stimulation of mammalian motoneurons. *J Physiol (Lond)* 168:890–910.
- Gustafsson B (1984) Afterpotentials and transduction properties in different types of central neurones. *Arch Ital Biol* 122:17–30.
- Harada Y, Takahashi T (1983) The calcium component of the action potential in spinal motoneurons of the rat. *J Physiol (Lond)* 335:89–100.
- Henrie JA, Shapley R (2005) LFP power spectra in V1 cortex: the graded effect of stimulus contrast. *J Neurophysiol* 94:479–490.
- Hopfield JJ (1995) Pattern recognition computation using action potential timing for stimulus representation. *Nature* 376:33–36.
- Huxter J, Burgess N, O'Keefe J (2003) Independent rate and temporal coding in hippocampal pyramidal cells. *Nature* 425:828–832.
- Juergens E, Guettler A, Eckhorn R (1999) Visual stimulation elicits locked and induced gamma oscillations in monkey intracortical- and EEG-potentials, but not in human EEG. *Exp Brain Res* 129:247–259.
- Kamondi A, Acsády L, Wang XJ, Buzsáki G (1998) Theta oscillations in somata and dendrites of hippocampal pyramidal cells in vivo: activity-dependent phase-precession of action potentials. *Hippocampus* 8:244–261.
- Kayser C, König P (2004) Stimulus locking and feature selectivity prevail in

- complementary frequency ranges of V1 local field potentials. *Eur J Neurosci* 19:485–489.
- Kayser C, Kim M, Ugurbil K, Kim DS, Konig P (2004) A comparison of hemodynamic and neural responses in cat visual cortex using complex stimuli. *Cereb Cortex* 14:881–891.
- Kraskov A, Quiroga RQ, Reddy L, Fried I, Koch C (2007) Local field potentials and spikes in the human medial temporal image lobe are selective to image category. *J Cogn Neurosci* 19:479–492.
- Kreiman G, Hung CP, Kraskov A, Quiroga RQ, Poggio T, DiCarlo JJ, Poggio T, DiCarlo JJ (2006) Object selectivity of local field potentials and spikes in the macaque inferior temporal cortex. *Neuron* 49:433–445.
- Lindsley DB, Wicke JD (1974) The electroencephalogram: autonomous electrical activity in man and animals. In: *Electroencephalography and human brain potentials* (Thomson RF, Patterson MM, eds), pp 3–83. New York: Academic.
- Liu J, Newsome WT (2006) Local field potential in cortical area MT: stimulus tuning and behavioral correlations. *J Neurosci* 26:7779–7790.
- Logothetis NK (2003) The underpinnings of the BOLD functional magnetic resonance imaging signal. *J Neurosci* 23:3963–3971.
- Logothetis NK, Wandell BA (2004) Interpreting the BOLD signal. *Annu Rev Physiol* 66:735–769.
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150–157.
- Logothetis NK, Merkle H, Augath M, Trinath T, Ugurbil K (2002) Ultra high-resolution fMRI in monkeys with implanted RF coils. *Neuron* 35:227–242.
- Logothetis NK, Kayser C, Oeltermann A (2007) In vivo measurement of cortical impedance spectrum in monkeys: implications for signal propagation. *Neuron* 55:809–823.
- Mehta MR, Lee AK, Wilson MA (2002) Role of experience and oscillations in transforming a rate code into a temporal code. *Nature* 417:741–746.
- Mitzdorf U (1987) Properties of the evoked potential generators: current source-density analysis of visually evoked potentials in the cat cortex. *Int J Neurosci* 33:33–59.
- Montemurro MA, Senatore R, Panzeri S (2007) Tight data-robust bounds to mutual information combining shuffling and model selection techniques. *Neural Comput* 19:2913–2957.
- 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.
- Niessing J, Ebisch B, Schmidt KE, Niessing M, Singer W, Galuske RAW (2005) Hemodynamic signals correlate tightly with synchronized gamma oscillations. *Science* 309:948–951.
- 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.
- Panzeri S, Schultz SR, Treves A, Rolls ET (1999) Correlations and the encoding of information in the nervous system. *Proc Biol Sci* 266:1001–1012.
- Panzeri S, Petersen RS, Schultz SR, Lebedev M, Diamond ME (2001) The role of spike timing in the coding of stimulus location in rat somatosensory cortex. *Neuron* 29:769–777.
- Panzeri S, Senatore R, Montemurro MA, Petersen RS (2007) Correcting for the sampling bias problem in spike train information measures. *J Neurophysiol* 98:1064–1072.
- Percival DB, Walden AT (1993) *Spectral analysis for physical applications: multitaper and conventional univariate techniques*. Cambridge, UK: Cambridge UP.
- Pesaran B, Pezaris JS, Sahani M, Mitra PP, Andersen RA (2002) Temporal structure in neuronal activity during working memory in macaque parietal cortex. *Nat Neurosci* 5:805–811.
- Pola G, Thiele A, Hoffmann KP, Panzeri S (2003) An exact method to quantify the information transmitted by different mechanisms of correlational coding. *Network* 14:35–60.
- Rubino D, Robbins KA, Hatsopoulos NG (2006) Propagating waves mediate information transfer in the motor cortex. *Nat Neurosci* 9:1549–1557.
- Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 20:5503–5515.
- Schneidman E, Bialek W, Berry MJ (2003) Synergy, redundancy, and independence in population codes. *J Neurosci* 23:11539–11553.
- Shannon CE (1948) A mathematical theory of communication. *Bell System Tech J* 27:379–423.
- Steriade M (1991) Alertness, quiet sleep, dreaming. In: *Cerebral cortex*, pp 279–357. New York, London: Plenum.
- Steriade M, Hobson J (1976) Neuronal activity during the sleep-waking cycle. *Prog Neurobiol* 6:155–376.
- Steriade M, McCormick DA, Sejnowski TJ (1993) Thalamocortical oscillations in the sleeping and aroused brain. *Science* 262:679–685.
- Strong SP, Koberle R, de Ruyter van Steveninck RR, Bialek W (1998) Entropy and information in neural spike trains. *Phys Rev Lett* 80:197–200.
- Victor JD, Purpura KP (1998) Spatial phase and the temporal structure of the response to gratings in v1. *J Neurophysiol* 80:554–571.