Title
State Dependence of Noise Correlations in Macaque Primary Visual Cortex

Journal
Neuron, 82(1)

ISSN
0896-6273

Authors
Ecker, Alexander S
Berens, Philipp
Cotton, R James
et al.

Publication Date
2014-04-01

DOI
10.1016/j.neuron.2014.02.006

Peer reviewed

Published in final edited form as:
Neuron. 2014 April 2; 82(1): 235–248. doi:10.1016/j.neuron.2014.02.006.

State dependence of noise correlations in macaque primary visual cortex

Alexander S. Ecker[1,2,3,5], **Philipp Berens**[1,2,3], **R. James Cotton**[1], **Manivannan Subramaniyan**[1], **George H. Denfield**[1], **Cathryn R. Cadwell**[1], **Stelios M. Smirnakis**[1,4], **Matthias Bethge**[2,3,5], and **Andreas S. Tolias**[1,3,6]

[1]Department of Neuroscience, Baylor College of Medicine, Houston, TX, USA

[2]Werner Reichardt Centre for Integrative Neuroscience and Institute of Theoretical Physics, University of Tübingen, Germany

[3]Bernstein Centre for Computational Neuroscience, Tübingen, Germany

[4]Department of Neurology, Baylor College of Medicine, Houston, TX, USA

[5]Max Planck Institute for Biological Cybernetics, Tübingen, Germany

[6]Department of Computational and Applied Mathematics, Rice University, Houston, TX, USA

Abstract

Shared, trial-to-trial variability in neuronal populations has a strong impact on the accuracy of information processing in the brain. Estimates of the level of such noise correlations are diverse, ranging from 0.01 to 0.4, with little consensus on which factors account for these differences. Here we addressed one important factor that varied across studies, asking how anesthesia affects the population activity structure in macaque primary visual cortex. We found that under opioid anesthesia, activity was dominated by strong coordinated fluctuations on a timescale of 1–2 Hz, which were mostly absent in awake, fixating monkeys. Accounting for these global fluctuations markedly reduced correlations under anesthesia, matching those observed during wakefulness and reconciling earlier studies conducted under anesthesia and in awake animals. Our results show that internal signals, such as brain state transitions under anesthesia, can induce noise correlations, but can also be estimated and accounted for based on neuronal population activity.

Keywords

Noise correlations; Anesthesia; Opioids; Gaussian Process Factor Analysis; Primary Visual Cortex; V1; Macaque

Corresponding authors: Alexander S. Ecker, Centre for Integrative Neuroscience, Otfried-Müller-Str. 25, 72076 Tübingen, Germany, alexander.ecker@uni-tuebingen.de, Phone: +49-7071-2988908. Andreas S. Tolias, Baylor College of Medicine, Department of Neuroscience, One Baylor Plaza, S. 553, Houston, TX 77030, USA, astolias@bcm.edu, Phone: +1-713-798-4071.

Introduction

A ubiquitous property of cortical neurons is their high degree of response variability (Softky & Koch 1993). Since repeated presentations of the same stimulus never elicit the same response twice, an accurate representation of the stimulus can be obtained only by considering the joint response profile of populations of neurons. The accuracy of such a population code strongly depends on neuronal correlations (Averbeck et al. 2006; Zohary et al. 1994; Abbott & Dayan 1999; Sompolinsky et al. 2001). Specifically, noise correlations, which express the amount of covariability in the trial-to-trial fluctuations of responses of two neurons to repeated presentations of the same stimulus, are central to such questions of coding accuracy.

In recent years, both the level and the origin of such noise correlations have been subject to debate. While it was originally suggested that noise correlations arise due to shared sensory noise arising in the afferent sensory pathway (Zohary et al. 1994; Shadlen & Newsome 1998), more recent studies suggest that they in fact represent meaningful top-down signals generated internally to the brain (Cohen & Newsome 2008; Nienborg & Cumming 2009; Ecker et al. 2010). Moreover, the observed level of correlations varies greatly between studies, with average values ranging from 0.01 to 0.4 (Bach & Krüger 1986; Zohary et al. 1994; Gawne & Richmond 1993; Gawne et al. 1996; Bair et al. 2001; Kohn & Smith 2005; Gutnisky & Dragoi 2008; Smith & Kohn 2008; Cohen & Newsome 2008; Mitchell et al. 2009; Cohen & Maunsell 2009; Ecker et al. 2010; Hansen et al. 2012; Smith et al. 2013; Smith & Sommer 2013; Herrero et al. 2013). It has recently been suggested that much of the differences between studies may be accounted for by differences in firing rates (Cohen & Kohn 2011). However, there are striking differences in correlations even between studies conducted in the same brain area with similar stimuli and similar firing rates (e.g. Smith & Kohn 2008; Ecker et al. 2010), suggesting that the firing rate dependence is insufficient to explain the variability across studies and other factors need to be taken into account as well.

One such factor that varies across studies is anesthesia. It constitutes a drastic alteration of global brain state, the mechanisms of which are only partly understood and depend on drugs that are used (Campagna et al. 2003). One of the most striking features of anesthesia, also observed during natural deep sleep, are strong slow-wave oscillations in the electroencephalogram (EEG) at frequencies below 2 Hz (Steriade et al. 1993). Many commonly used anesthetics, such as isoflurane, urethane and ketamine, substantially alter neural activity by suppressing sensory responses and increasing response latencies (Angel 1993; Drummond 2000; Chi & Field 1986; Kohn et al. 2009) as well as inducing so-called up and down states (Renart et al. 2010; Constantinople & Bruno 2011; Harris & Thiele 2011). Some neuroscientists resort to opioids, such as fentanyl or sufentanil (Kohn & Smith 2005; Smith & Kohn 2008; Reich et al. 2001), which are believed to affect neural activity in less dramatic ways (Loughnan et al. 1987; Schwender et al. 1993; Drummond 2000; Constantinople & Bruno 2011). However, although opioids seem to have a number of advantages over other drugs, they have similarly been shown to affect neural response properties (Schwender et al. 1993) and induce low-frequency oscillations (Bowdle & Ward 1989).

To shed light on how opioids modify the structure of neural population activity, we measured noise correlations in primary visual cortex of anesthetized and awake monkeys using identical recording techniques. Under anesthesia we observed periods of almost complete silence across the population as well as periods of very strong activity. These periods lasted for a few hundred milliseconds, arose spontaneously and were not linked to the visual stimulus. They resembled up and down states commonly observed using non-opioid anesthetics (Renart et al. 2010; Constantinople & Bruno 2011; Harris & Thiele 2011) and their characteristic frequency was comparable to slow-wave oscillations in the EEG (Steriade et al. 1993). Interestingly, they could be almost completely accounted for by a latent variable model of the population activity with a single latent variable indicating the network state. When we conditioned on this latent variable, the magnitude and structure of noise correlations under anesthesia were almost indistinguishable from those we observed previously in awake monkeys (Ecker et al. 2010).

Our results show that spontaneous transitions in network state under anesthesia induce noise correlations between neurons. These transitions are absent in awake, fixating monkeys. This indicates a clear qualitative difference between the two states despite similar firing rates. Thus, anesthesia is an important, but often neglected factor accounting for differences between studies that cannot be explained by firing rates, as suggested previously (Cohen & Kohn 2011).

Results

First and second order statistics of neuronal responses

We recorded the spiking activity of populations of neurons in primary visual cortex of awake and anesthetized macaque monkeys. We recorded from 487 neurons in two awake monkeys and 636 neurons in three anesthetized monkeys. Our dataset consists of 58 recording sessions (31 awake, 27 anesthetized), each containing 10–42 simultaneously recorded cells (medians: 15 awake, 23 anesthetized). The awake dataset is a subset of previously published data (Ecker et al. 2010, see Experimental Procedures for details). We presented sinusoidal gratings covering the receptive fields of all recorded neurons. Gratings were drifting, except in 14 of the awake sessions where static gratings were shown.

As expected, neurons in V1 of awake monkeys were robustly driven by the grating stimulus (Figure 1A) and the vast majority of cells were tuned to orientation (Figure 1B; for this example session: 27/29 cells; overall 82% or 400/487 cells at $p < 0.01$, permutation test, not corrected for multiple testing). The same was true for anesthetized recordings (Figure 1C, D), where an even larger fraction of cells was tuned (example session: all 44 cells; overall 92% or 586/636 cells tuned at $p < 0.01$), probably reflecting the fact that anesthetized recordings on average contained larger amounts of data. Thus, when averaging spike trains across multiple trials, responses recorded during wakefulness and under anesthesia were qualitatively similar in the sense that a large fraction of cells was robustly tuned to orientation.

We noticed, however, that anesthetized responses appeared noisier than those recorded during wakefulness (Figure 1A, C). To test whether this impression was true at the

population level we computed the Fano factors (variance of the response divided by its mean) for all recorded neurons. Indeed, response variability was roughly twice as large under anesthesia as during wakefulness (Figure 2A; average F = 2.2 vs. 1.2, respectively; $p < 10^{-15}$, Wilcoxon rank sum test). This was not due to systematic differences in firing rates between wakefulness and anesthesia, as it was true for the entire range of firing rates (Figure 2B).

This increased trial-to-trial variability could be a single-neuron effect, where the anesthetic causes individual neurons to fire more randomly, or a population effect, where groups of neurons are co-modulated by a common source present only under anesthesia. While the former would add independent noise and manifest itself primarily in increased variances (and Fano factors), the latter would also give rise to elevated noise correlations. Indeed, the average level of correlations was roughly six times higher under anesthesia than during wakefulness (Figure 2C; 0.05 vs. 0.008, respectively; $p < 10^{-15}$, Wilcoxon rank sum test, 8012 vs. 3878 pairs). Again, this difference was present at the full range of firing rates and most prominent for pairs of cells with high rates (Figure 2D).

State fluctuations under anesthesia

Our data seem to argue for a population level effect of anesthesia, where many neurons are modulated simultaneously on a trial-to-trial basis. Indeed, population raster plots showing the activity of all simultaneously recorded neurons for a given trial revealed periods of almost complete silence as well as periods of vigorous activity (Figure 3C, see e.g. trials 2–4). The transitions between such periods seemed to arise spontaneously and were not linked to the stimulus, suggesting that at least part of the increased variability was caused by a common noise source.

To characterize this common source of variability in more detail, we used a recently developed latent variable model called Gaussian Process Factor Analysis (GPFA, Figure 3A and Experimental Procedures for details) (Yu et al. 2009). The GPFA model promises to be a good candidate for capturing the phenomena observed here as it seeks to describe the correlations in the data by a low-dimensional state variable, which evolves smoothly in time and affects each neuron's firing rate linearly. We use the GPFA model to represent the fluctuations around the stimulus-driven response (noise correlations):

$$r_k(t) = f_k(s(t)) + c_k x(t) + \eta. \quad (1)$$

Here, $f_k(s(t))$ is the time-resolved tuning curve of neuron k, which captures the stimulus-induced response dynamics; $x(t)$ is the network state, which is a one-dimensional function of time; c_k is the weight that determines how x affects the neuron's response; and η is independent Gaussian noise. The network state x has a smooth autocorrelation function with timescale τ (Figure 3A and Experimental Procedures).

Using such a latent variable model affords several advantages over the traditional approach of computing pairwise correlations and analyzing their relationship to other quantities such as signal correlations or distance between neurons. First, the number of parameters that need

to be estimated is substantially lower than when estimating the full correlation matrix. Second, if there are processes contributing to the observed correlations that affect many neurons at the same time, they can be estimated more efficiently and their timescale can be extracted simultaneously.

The GPFA model with a single state variable captured the structure and dynamics of the population response under anesthesia well. Visually, the estimate of the network state corresponded well to the apparent on and off periods (Figure 3C). We quantified how much explanatory power the network state variable has under the two different brain states by computing the fraction of variance explained (see Experimental Procedures for details) on a separate subset of the data not used for fitting the model. In the awake dataset the state variable explained on average less than 5% of the variance (Figure 3D, F). Strikingly, under anesthesia up to 40% of individual cells' variances were explained by network state (Figure 3E, F). To ensure that this effect was not due to longer trials in our anesthetized experiments (2 s anesthetized vs. 500 ms awake), we repeated the analysis on the anesthetized data using only the first 500 ms of the response (Figure 3F, dashed line), which reproduced the result obtained with the full response. Generally, the fraction of variance explained was substantially higher for cells with high firing rates (Figure 3F) and increased with the size of the window over which spikes were counted (Figure 3G). This effect was particularly strong under anesthesia, but much less so during wakefulness.

To gain insights into the structure of variability induced by the network state variable, we analyzed the key parameters of the model: weights and timescale. The weight of a cell tells us how the network state affects its firing rate. If all cells are co-modulated in the same direction we expect mostly positive weights and, thus, positive correlations between cells. If, on the other hand, some cells are enhanced (positive weights) while others are suppressed (negative weights) we expect mostly positive correlations within each group and negative correlations across groups. During wakefulness the weights were mostly distributed around zero (Figure 4A, C; 65% positive) while during anesthesia most weights were positive (Figure 4B, D; 88% positive). Note, though, that there is an ambiguity in the GPFA model: one can always flip the sign of all weights without changing the model, by simply flipping the sign of the latent variable (see Eq. 1). By convention, we set the sign such that the majority of weights for each model are positive. We therefore expect a fraction greater than 50% to have positive sign, even in the absence of any effect (bootstrap 95% confidence intervals under the null hypothesis: awake 61.6–62.9% positive weights, anesthetized 59.2–60.4%). Thus, although it was significant ($p = 2 \times 10^{-9}$), only marginally more neurons than expected by chance had positive weights during wakefulness. Together with the finding above that the model explained very little variance, this indicates that there were no strong state fluctuations in our data during wakefulness. Under anesthesia, in contrast, the weights were mostly positive ($p < 10^{-15}$), indicating that the firing rates of most cells were co-modulated by a common term, which presumably caused the elevated correlations observed above (Figure 2B, C; we will quantify below what fraction of the correlations is accounted for by the network state variable).

The inferred timescale can help us to constrain our hypotheses on the origin of the observed correlations. If the common noise was due to shared sensory noise (Zohary et al. 1994;

Shadlen & Newsome 1998) its time constant should be relatively small, corresponding roughly to the membrane time constants of the post-synaptic cells (10–50 ms; Mason et al. 1991; Shadlen & Newsome 1998). On the other hand, intrinsically generated up and down states, which have been observed with many non-opioid anesthetics, are much slower (< 2 Hz; Renart et al. 2010; Constantinople & Bruno 2011; Haider et al. 2013). More consistent with the latter hypothesis, the timescale of the network state dynamics during anesthesia was relatively slow. The median width of the Gaussian temporal kernel was 207 ms (Figure 4F). In the frequency domain this corresponds to a lowpass cutoff frequency of 2.35 Hz (at −40 dB attenuation). This estimate of the timescale appears somewhat higher than that previously reported for anesthetized monkey V1 (Smith & Kohn 2008). However, this difference is caused by what appears to be a bias in their method of estimating the timescale, rather than reflecting a discrepancy between the two datasets (performing the same analysis as they did showed that our dataset is consistent with theirs; see Supplemental Material for an in-depth discussion of this issue). During wakefulness, in contrast, a large fraction of timescale values were around 800 ms (Figure 4E; median 688 ms), which is substantially longer than a single trial (500 ms). As the model does not take into account correlations of the network state across trials, this indicates that the network state was essentially constant within a trial. Thus, the strongest common modulations the model picked up during wakefulness were, in addition to being much weaker, substantially slower than the state fluctuations we observed under anesthesia.

We next turned to the pairwise correlation structure and asked to what extent it was explained by the network state fluctuations. The raw correlation structure under anesthesia resembled that in previous reports of anesthetized monkey V1 (Kohn & Smith 2005; Smith & Kohn 2008). Raw noise correlations were strongest for pairs with high firing rates (Figure 5A) (see also Smith & Sommer 2013) and high signal correlations (Figure 5B). Moreover, they decreased significantly with the spatial separation between cells (Figure 5C). To determine to what degree the GPFA model accounted for this correlation structure, we computed the residual correlations after accounting for the network state. This can be thought of as computing correlations by not only conditioning on the stimulus but also on the network state (see e.g. Renart et al. 2010). We found that the network state explained most of the difference in the magnitude and structure of noise correlations between wakefulness and anesthesia. The residual correlation structure under anesthesia resembled the raw correlation structure during wakefulness remarkably well: except for pairs recorded on the same tetrode the differences were within the margin of error (Figure 5A–C). For pairs recorded on the same tetrode the residual correlations under anesthesia were significantly higher than during wakefulness (Figure 5C, see discussion). Accounting for network state did not alter the correlation structure during wakefulness. This finding was expected due to the low fraction of variance captured by the model during wakefulness (Figure 3F).

Model of state fluctuations as common gain

The analysis of residual correlations showed that the correlation structure changed when accounting for network state: the firing rate dependence was nearly abolished (Figure 5A) and both the relation with signal correlations and that with distance were weakened substantially (Figure 5B, C). This may seem counter-intuitive at first, since all neurons are

modulated by the same common network state variable and, thus, one may expect a uniform effect on all neurons. However, since the network state can affect different neurons with different weights and those weights may depend on the stimulus, network state fluctuations can induce a non-uniform correlation structure. In our data, the weights were positively correlated with firing rates (data not shown), indicating that the network state acted as a common gain, modulating each neuron's firing rate multiplicatively.

To understand how such fluctuations in common gain would affect the correlation structure, we considered a simple network model: the firing rate of each neuron was determined by its tuning curve times a common gain and neurons spiked according to independent, inhomogeneous Poisson processes (Figure 6A, see Experimental Procedures for details). The gain term was fluctuating randomly with temporal correlations matching that in the data (~200 ms). This simple model was able to reproduce both the firing rate dependence of noise correlations in our data as well as their dependence on signal correlations quite naturally (Figure 6B, C). To capture the spatial dependence of correlations, we would have to include spatial structure, e.g. by replacing the global gain by one that can vary across space with a certain correlation structure. However, we do not pursue the question in more detail here, since the main point of the model is to illustrate that very simple mechanisms can cause remarkably non-uniform correlation structures.

Spontaneous activity

We next asked whether the state fluctuations observed under anesthesia were also present during spontaneous activity in the absence of visual stimulation. To address this question, we analyzed the blank periods between subsequent stimulus presentations. The results essentially mirrored those obtained during visual stimulation (Figure 7). Variance explained increased with both firing rates and the size of integration window (Figure 7A–C). Weights were almost exclusively positive (96%, Figure 7D, E) and the timescale of the network state was comparable to that during visual stimulation (Figure 7F; median: 181 ms; cutoff frequency: 2.75 Hz). As for the evoked responses, residual correlations after accounting for network state were profoundly reduced (Figure 7H–J).

Local field potential as a predictor of global network state

We showed that under anesthesia most neurons are affected in a similar way by the network state and this network state can change on a timescale of a few hundred milliseconds. If the effect is as global as it appears, we should find its signature in more global measures of neural activity such as the local field potential (LFP). We thus asked whether the low-frequency range of the LFP correlated with the network state we inferred above. This was indeed the case for all three anesthetized but for none of the awake animals (Figure 8A, B). The magnitude of the correlation was strongest at zero time lag and had additional peaks/troughs of opposite sign at time lags of ±500 ms between LFP and inferred network state.

If the low-frequency range of the LFP is correlated with the network state, it should be possible to use it to predict the trial-to-trial variability observed under anesthesia. To verify this, we followed the approach taken by Kelly et al. (2010) and fitted a generalized linear

model (GLM) with the lowpass-filtered LFP as input (see Experimental Procedures for details):

$$\mu(t)=\exp(a(t)+bu(t)) \quad (2)$$

Here $\mu(t)$ is the firing rate, $a(t)$ the stimulus response (PSTH) and $u(t)$ the LFP, all of which are functions of time. The linear weight b determines by how much a change in the LFP affects the firing rate of the neuron. During wakefulness the LFP weights were distributed mostly around zero (Figure 8C), whereas under anesthesia they were mostly negative (Figure 8D).

In summary, the network state we inferred above in an unsupervised way from spiking data alone (GPFA model) has its physiological counterpart in the low-frequency oscillations in the LFP. Both the low-frequency oscillations and the apparent network state fluctuations in the spiking activity of local populations are pronounced under anesthesia but relatively small, if not absent, during awake fixation.

Finally, our analysis so far has focused on comparing wakefulness and anesthesia using different cells recorded in different animals. However, anesthesia has multiple different stages, with light anesthesia being characterized by relatively desynchronized EEG activity, whereas deep anesthesia displays strong, coherent network oscillations. We therefore asked whether we could use the LFP to find evidence for slow changes in brain state (depth of anesthesia) within recording sessions. Indeed, in many sessions we observed slow changes in LFP power in a low-frequency range and sometimes in the gamma range (Figure 9A, B). To quantify these changes we computed an LFP power ratio in windows of approximately 90 seconds (power at 0.5–2 Hz divided by that in the gamma band, 30–70 Hz) (Goard & Dan 2009), which we used as a proxy for depth of anesthesia. This power ratio displayed changes on timescales of a few minutes up to half an hour and longer (Figure 9C, D, black lines). Remarkably, the time-resolved LFP power ratio was tracked very closely by the total correlation in the network as measured by the variance of the network state variable inferred by the GPFA model (Figure 9C, D, red lines). Across all sessions, the LFP power ratio and the overall level of correlations were significantly correlated (Figure 9E, Spearman's ρ = 0.42, $p < 10^{-15}$). This correlation was positive and significant in 19/27 individual sessions ($p < 0.05$, uncorrected). Thus, the degree of network-wide correlations varied within a recording session in the same cells over the course of several minutes and correlated well with more traditional, LFP- or EEG-based measures of brain state or depth of anesthesia.

Discussion

State fluctuations under opioids

We demonstrated a striking feature of cortical activity under opioid anesthesia that had previously not been appreciated: neurons undergo spontaneous coordinated transitions between states of almost complete silence, highly elevated levels of activity, and intermediate levels of activity. These state transitions resemble up and down states, which have been described previously for other, non-opioid anesthetics (Steriade et al. 1993;

Renart et al. 2010; Constantinople & Bruno 2011), and they occur on a timescale of several hundred milliseconds. In addition, the strength of these state fluctuations can change slowly over several minutes, which may reflect slow changes in the depth of anesthesia.

Although the effect of opioid anesthetics may be less dramatic than that of non-opioids such as urethane, isoflurane or ketamine (Constantinople & Bruno 2011; Movshon et al. 2003; Kohn et al. 2009; Smith & Sommer 2013), it should be emphasized that they still have a substantial effect on neural responses, explaining on average more than one third of the variance of cells firing at rates of more than 10 spikes/s (Figure 3). Since the effect is largely common to all cells within a few millimeters of cortex, it becomes particularly evident when considering populations of simultaneously recorded neurons and substantially biases the structure of noise correlations compared with awake recordings.

We are aware of two reports that directly addressed the effect of opioids and found no differences to the awake state (Loughnan et al. 1987; Constantinople & Bruno 2011). Although they may superficially appear at odds with our results, this is not the case. One study measured the average sensory-evoked EEG response in humans (Loughnan et al. 1987) and found no difference between anesthetized and awake subjects. While this finding is consistent with our results that sensory responses were intact, it does not rule out spontaneous state transitions, as those would have been averaged out. The other study measured membrane potential fluctuations in single neurons (Constantinople & Bruno 2011). It is possible that opioids act more subtly than other anesthetics, not inducing the bimodal distribution of membrane potentials that typically characterizes up and down states (Petersen et al. 2003; Constantinople & Bruno 2011), but nevertheless leading to global fluctuations in spiking output that are strong enough to be picked up when recording populations of neurons simultaneously. Another important point to be noted is that the two studies cited above were conducted under much lighter anesthesia. The fentanyl doses used (3 μg/kg bolus and 10 μg/kg/h, respectively) were substantially lower than the minimum equivalent sufentanil dose used in acute primate experiments (our study and e.g. Smith & Kohn 2008; Kelly et al. 2010: 4–15 μg/kg/h sufentanil, equivalent to 40–150 μg/kg/h fentanyl). Thus, the differences in depth of anesthesia, different measures of neural activity, or differences between species could account for the differences between these studies and ours.

State fluctuations during wakefulness

State transitions similar to those we observed under anesthesia have been observed in rodents also during wakefulness. Poulet & Petersen (2008) found that periods of inactivity (termed quiet wakefulness) resembled the anesthetized state. Both, the intracellular membrane potentials and the LFP displayed increased power in the low frequencies, similar to our and other labs' findings under anesthesia, and spikes were tightly locked to those oscillations. During periods of active whisking, in contrast, somatosensory cortex was in a desynchronized state that resembled our awake results. In addition, Niell & Stryker (2010) showed that the firing rates of neurons in primary visual cortex of mice depend strongly on whether the mouse is still or running on a treadmill. Although they did not explicitly test whether response variability or properties of the LFP were different between the two states,

their findings support the idea that the cortex can switch between different states of activity during wakefulness.

To our knowledge, such state fluctuations have not been observed in the visual system of awake, behaving primates. One could speculate that this is due to a species difference between rodents and primates. However, it seems more likely that we did not observe such quiet states during our awake experiments because the monkeys had to actively initiate a trial by moving their eyes to the fixation spot and maintain fixation throughout the trial, actively suppressing their natural reflex to move the eyes several times per second. This required oculomotor action before and during the stimulus could trigger an active state similar to whisking or running in rats and mice.

This action to initiate a trial may be an important difference between experiments in the visual system of awake monkeys and rodents. Unlike with monkeys, in most studies of the rodent visual system the animals do not have to actively initiate a trial, but stimuli are presented periodically. To obtain a similar level of control over the brain state, one would have to either infer it post-hoc from recordings of locomotion, eye or whisker movements or – as we did in this study – directly from neuronal population activity. Since this is not usually done (but see, e.g. Poulet & Petersen 2008; Niell & Stryker 2010), many datasets collected in awake rodent visual cortex are likely to contain a mixture of brain states. We, therefore, do not expect large differences between wakefulness and anesthesia in such cases, a hypothesis corroborated by a recent study of noise correlations in mouse V1 (Denman & Contreras 2013).

Role of firing rates

Could the difference between our awake and anesthetized data be attributed to factors other than anesthesia? It has been suggested that the low correlations we measured in awake monkeys were a result of unusually low firing rates (Cohen & Kohn 2011). However, this is not a viable explanation since firing rates were similar in our awake and anesthetized recordings and they were comparable to (in fact, slightly higher than) those reported by other labs using similar stimuli in the same cortical area as in our present study (Smith & Kohn 2008: 3.4 spikes/s; this study, awake: 5.4 spikes/s, anesthetized: 5.0 spikes/s). In addition, the difference between awake and anesthetized correlations was evident at the full range of firing rates (Figure 2) and the firing-rate dependence of correlations in our anesthetized dataset resembled that reported by other labs in anesthetized monkey V1 (Smith & Kohn 2008; Smith & Sommer 2013). In summary, while firing rates certainly contribute to differences between studies to some extent, they cannot account for the difference between wakefulness and anesthesia.

Role of cortical layers

Recent studies suggest that noise correlations are low in the granular layers of V1, raising the possibility that our awake recordings were mostly restricted to those layers (Hansen et al. 2012; Smith et al. 2013). If this was the case, the effects we describe in this study could be caused by laminar differences, rather than reflecting a difference between wakefulness and anesthesia. Based on our data we cannot rule out this possibility entirely, but a number of

observations argue against it. Although in our anesthetized experiments we recorded throughout all layers and tried to adjust all tetrodes to a similar depth for each recording, we were unable to identify the region of low correlations. This is most likely owed to the limitations of our experimental approach. Tetrodes have a blunt tip, presumably causing more tissue displacement than single electrodes with small beveled tips, making the point of entry into the brain a poor reference to estimate laminar location. In addition, we did not reach white matter with all tetrodes before the end of the experiment, precluding the use of white matter as a reference. Furthermore, tetrodes have much lower impedances than single electrodes. Therefore they probably sample cells from a larger volume. We thus expect considerable variability in both the laminar location of the tetrodes and our estimates thereof. Since the region of low correlations reported previously (Hansen et al. 2012; Smith et al. 2013) is a narrow strip of 200–300 μm, it may not be surprising that we were unable to identify it. However, for the same reasons it seems implausible that laminar variation should explain the low correlations we observed during wakefulness. For this to be the case, most of our tetrodes should have been located in exactly this narrow region. Yet, unlike in our anesthetized experiments, we neither adjusted the tetrodes together nor did we target any specific layer, but instead adjusted each tetrode to a position where it isolated cells. In addition, between awake recording sessions we sometimes adjusted the tetrodes, in total by up to 600 μm (median: ~300μm) between the first and the last experiment. As a result, we should either have been able to localize the region of low correlations during anesthesia or we likely recorded from outside it as well during wakefulness, suggesting that the effect we describe is not simply explained by laminar differences.

Relation to other studies of noise correlations in the primate visual system

By modeling the state fluctuations under anesthesia with a latent variable model (GPFA) we recovered the residual correlation structure, which was remarkably similar to that observed in the awake monkey. This finding reconciles the results of previous studies conducted in V1 under anesthesia with our awake, fixating data (Ecker et al. 2010). The raw correlation structure we observed under anesthesia is entirely consistent with previous reports using the same preparation (Kohn & Smith 2005; Smith & Kohn 2008). The higher average level of correlations during anesthesia (Reich et al. 2001; Kohn & Smith 2005; Smith & Kohn 2008) is accounted for by the one-dimensional network state variable. The LFP can be used to predict some of these state fluctuations under anesthesia, which has been reported previously (Kelly et al. 2010). Finally, another study characterizing higher-order correlations in anesthetized monkey V1 (Ohiorhenuan et al. 2010) reports an excess probability of silence in triplets of neurons, suggesting that the periods of almost complete silence we observe are also present in other anesthetized preparations.

Some discrepancies remain between the papers discussed above (Kohn & Smith 2005; Smith & Kohn 2008; Reich et al. 2001; Ohiorhenuan et al. 2010; Ecker et al. 2010) and some other studies. For instance, some authors report substantially higher noise correlations in awake monkey V1 (0.2–0.4; Gutnisky & Dragoi 2008; Hansen et al. 2012; Herrero et al. 2013) than we did (Ecker et al. 2010). Note that in addition to substantially higher average firing rates, these studies typically also observed relatively high Fano factors ($F > 2$; Gutnisky and Dragoi, 2008; Herrero et al., 2013; our awake data: average $F = 1.1$, Figure

2A), indicating that either different cell populations were sampled or additional confounding factors were present that were not accounted for (e.g. as argued in Ecker et al. 2010; Bair et al. 2001). For instance, accounting for eye movements reduced the correlations by almost 50% in Hansen et al. 2012 (their Supplementary Figure 3).

Correlations between nearby neurons

Similar to other authors (e.g. Smith & Kohn 2008; Cohen & Maunsell 2009), we focused mainly on pairs recorded by different electrodes. For such pairs, accounting for the network state under anesthesia reduced the noise correlations consistently below 0.01, similar to the level observed during wakefulness. However, a notable observation we made was that residual correlations between pairs recorded by the same tetrode were still higher under anesthesia than during wakefulness (Figure 5C). This could reflect an additional, more local contribution of anesthesia that was not captured by the single latent variable in our model. Alternatively, there may be some degree of heterogeneity in the local connectivity, which gives rise to different levels of correlation depending on where one records from (e.g. close to pinwheels vs. linear zones, or differences between layers). Indeed, when we re-analyzed the awake data, focusing on pairs recorded by the same tetrode, we observed some differences between the two monkeys. In one monkey, signal correlations for pairs recorded on the same tetrode were close to zero (average 0.025) and so were the noise correlations (average 0.006), while in the other monkey signal correlations were positive (0.24) and noise correlations were somewhat higher as well (0.045). The latter is more consistent with the anesthetized results (average signal correlations: 0.17, average residual noise correlations: 0.065). It is possible that we sampled cells in a more unbiased fashion in our anesthetized experiments, in which we recorded from more monkeys and more individual tetrodes than in our awake dataset. To reach a definite conclusion regarding the structure and level of correlations for neurons separated by less than 200 μm and to resolve the potential contribution of cortical layers, more extensive future experiments with high-density laminar probes (Blanche et al. 2005) are needed.

Conclusions and future directions

Most of what we know today about the early visual system we learned through studies in anesthetized animals (e.g. Hubel & Wiesel 1968; Zeki 1974; De Valois, William Yund, et al. 1982; De Valois, Albrecht, et al. 1982; Movshon et al. 1985; Carandini et al. 1997). Without doubt, the acute anesthetized preparation is an extremely valuable tool, which offers many advantages for studying the early visual system (no training of animals, no issues due to eye movements/microsaccades, longer experiments with more trials, etc.). For instance, receptive fields or tuning curves can be measured under anesthesia just as well as in the awake animal.

More recently, however, many groups have started to characterize the joint activity patterns of pairs and groups of neurons (Zohary et al. 1994; Gawne & Richmond 1993; Gawne et al. 1996; Bair et al. 2001; Reich et al. 2001; Kohn & Smith 2005; Smith & Kohn 2008; Gutnisky & Dragoi 2008; Ecker et al. 2010; Berens et al. 2012) and both the origin and the implications of neuronal correlations have been of great interest (Zohary et al. 1994; Shadlen & Newsome 1998; Abbott & Dayan 1999; Sompolinsky et al. 2001; Averbeck et al.

2006; Cohen & Newsome 2008; Josic et al. 2009; Nienborg & Cumming 2009; Ecker et al. 2011). For these studies it is important to distinguish between different sources of correlation: if the network transitions from one state to another such widely distributed dynamics can quickly become the dominant source of (co-)variance. However, if such state transitions do not occur in alert animals paying attention to or interacting with their environment, the functional relevance of these correlations may be very different from those originating from shared input in the feed-forward signal chain of upstream neurons. Thus, one should be aware of possible state fluctuations and, if necessary, take them into account. While some authors have done so by considering only data during those periods where the brain was in a certain state (e.g. Renart et al. 2010) or incorporating global signals such as the LFP directly into the response model (Kelly et al. 2010), our study showed that in some situations the network state may also be inferred directly from population data using a latent variable model (Figure 3–5).

Latent variable models like the one we used in this study (GPFA, Yu et al. 2009; see also Macke et al. 2011; Buesing et al. 2012) are powerful tools for future studies of neuronal population activity. In light of current and future technological developments, the number of neurons that can be monitored simultaneously will increase substantially. The amount of time that can be used to collect data, however, is and remains limited by experimental and ethical constraints. Thus, an accurate characterization of the joint population response will be feasible only if much of the variability is restricted to a relatively low-dimensional subspace. Fortunately, this is very likely to be the case if our original hypothesis is correct and most of the correlations observed in awake animals are driven by unobserved internal signals rather than by shared sensory noise (Ecker et al. 2010). In this case, latent variable models will not only afford a parsimonious statistical description of neuronal population data, but they may also provide us with a method to read out internal signals, such as the focus of attention (Cohen & Maunsell 2010), task strategies or many more, in real time on a trial-by-trial basis.

Experimental Procedures

Electrophysiology in awake monkeys

We recorded from two adult, male rhesus monkeys (*macaca mulatta*) using chronically implanted tetrode arrays. The awake dataset used in this study is a subset of a dataset analyzed previously (Ecker et al. 2010; Berens et al. 2012, see below for inclusion criteria). Surgical methods and recording protocol for our awake experiments have been described previously (Tolias et al. 2007; Ecker et al. 2010).

Electrophysiology in anesthetized monkeys

In acute experiments lasting 4–5 days, we recorded from three adult, male rhesus monkeys (*macaca mulatta*) using the same 24-tetrode arrays as in the awake recordings. Surgical details are described in the Supplemental Experimental Procedures. Prior to each set of recordings, all tetrodes were adjusted to a new target depth approximately 200 μm deeper than the previous one. The exact amount of adjustment varied by tetrode, leaving tetrodes (if possible) at a position where cells could be isolated. Throughout the experiments anesthesia

was maintained by intravenous infusion of sufentanil (4–15 μg/kg/hr; protocol similar to Kohn & Smith 2005; Smith & Kohn 2008). Animals were paralyzed using vecuronium bromide by intravenous infusion (100 μg/kg/hr). The pupils were dilated by topical application of cyclopentolate. Refraction was provided by contact lenses. Stimuli were presented monocularly; the other eye was closed and covered. The open eye was kept irrigated using saline. Vital signs (ECG, heart rate, respiratory rate and volume, blood pressure, temperature, CO_2, O_2, $\%SpO_2$) were monitored continuously. All experimental procedures complied with guidelines approved by the Baylor College of Medicine Institutional Animal Care and Use Committee (IACUC).

Visual stimuli/behavioral paradigm

Visual stimuli were drifting gratings (16 different directions of motion) under a circular aperture presented at full contrast on gray background using the Psychophysics toolbox for Matlab (Brainard 1997). In a subset of awake experiments stimuli were static gratings (eight orientations), partly at lower contrasts (see Ecker et al. 2010 for details). Because of space constraints in the anesthetized setup, we used an LCD monitor running at a refresh rate of 60 Hz and positioned at a distance of 55 cm to the eye during our anesthetized experiments. The stimuli for awake monkeys were presented on CRT monitors running at 100 Hz and positioned at a distance of 100 cm. To address concerns about low firing rates in our data raised previously (Cohen & Kohn 2011), we reduced the size of the stimuli during the anesthetized experiments to 2–3° in diameter, compared with 4° in awake experiments. We ensured that the gratings covered the receptive fields of all neurons by mapping multi unit receptive fields of most tetrodes manually before each recording session. Temporal frequency was 3.4 cycles/sec for all sessions. Spatial frequency varied between 3–6 cycles/deg, roughly matching the preferences of the recorded cells due to some variability in eccentricity of recording locations (estimated between 1–4° from the fovea). Stimulus conditions were randomized in blocks of 16 trials to ensure a balanced number of repetitions.

In awake experiments trials were initiated by a sound and the appearance of a fixation target (~0.15°). After the monkey fixated for 300 ms, the stimulus was shown for 500 ms and the monkey had to fixate for another 300 ms. Monkeys were required to fixate within a radius of 0.5–1° but typically fixated much more accurately, as revealed by offline analysis. Monkeys were rewarded by a drop of juice upon completion of a successful trial.

In anesthetized experiments stimuli were shown for 2 s, separated by blank periods with a gray screen lasting approximately 1.1–1.6 s (randomly drawn from a uniform distribution).

Spike detection and sorting

Our data processing methods are based on previously published work (Tolias et al. 2007) but have been revised since the original report. A detailed description can be found in the Supplemental Experimental Procedures. Briefly, spikes were detected offline when the signal on any of the four channels crossed a threshold of five times the SD of the noise. After spike alignment, we extracted the first three principal components on each channel, resulting in a 12-dimensional feature vector used for spike sorting. To deal with waveform

drift, we fit a mixture model that uses Kalman filters to track changing cluster means over time (Calabrese & Paninski 2011). Single unit isolation was assessed quantitatively using the mixture model. Since the focus of this paper is on global fluctuations that are distributed among many tetrodes, spike-sorting errors are unlikely to play an important role; they would affect primarily pairs recorded by the same tetrode (Ecker et al. 2010). Therefore, we included all units flagged as single units in the analysis to increase statistical power. The sum of the false positive rate and the false negative rate was less than 10% for 62% of the single units in our dataset and less than 20% for 83% of the single units (awake: 63% and 82%, anesthetized: 61% and 83%).

Dataset and inclusion criteria

We recorded from two awake and three anesthetized monkeys, a total of 46 and 30 recording sessions, respectively. We included recording sessions where gratings were shown for at least 500 ms per trial, at least 20 trials per condition, and at least 10 single units with stable firing rates were recorded. Firing rate stability was assessed by computing the long-term component of the trial-autocorrelogram (Bair et al. 2001), which we estimated by taking a weighted average (Gaussian window with SD of eight trials) around zero, excluding the bin at zero lag (which is one by definition). Units were considered stable if the long-term component of the trial-autocorrelation was less than 0.1. These criteria resulted in 31 awake and 27 anesthetized recording sessions with 487 and 636 single units, respectively. The stability criterion was important since the anesthetized experiments were performed acutely and tetrodes were adjusted every 8–10 hours (see Supplemental Experimental Procedures for a discussion of this issue).

Data analysis/availability of code and data

Data analysis was done in Matlab using a data analysis framework with MySQL database backend (DataJoint: https://github.com/datajoint; D. Yatsenko, Tolias Lab, Baylor College of Medicine). The complete dataset and code used for data processing, data analysis and creating the figures in this article are available at http://toliaslab.org/publications/ecker-et-al-2014/.

Orientation tuning

We assessed the significance of orientation tuning by a permutation test. We first extracted the magnitude of the second Fourier component (i.e. orientation) by projecting the vector of average responses for each orientation onto a complex exponential with two cycles:

$$q=\sum_{k=1}^{16}\langle r\rangle_k \exp(\pi i k/4), \quad (3)$$

where $\langle r\rangle_k$ is the average response to the k^{th} direction of motion. We compared $|q|$ to its null distribution, which we obtained by shuffling the trial labels. We ran 1000 iterations of the shuffling procedure and used the fraction of runs with $|q|$ greater than that observed in the real data as the p value.

Fano factors/noise correlation analysis

Fano factors and noise correlations were computed on the first 500 ms of the response for both awake and anesthetized experiments. Fano factors were computed as the variance of the spike count divided by its mean. Noise correlations were computed as the Pearson correlation coefficient of two neurons' responses to identical repetitions of the same stimulus condition, averaged (for each pair) over all stimulus conditions with non-zero firing rates in both neurons.

Gaussian Process Factor Analysis (GPFA)

A detailed description of the GPFA model and the derivation of the Expectation Maximization (EM) algorithm to fit it can be found in Yu et al. (2009). Here we describe only the key points.

The GPFA model is described in the main text (Eq. 1 and Figure 3). We extracted spike counts in each trial during the stimulus period in T non-overlapping bins of 100 ms starting 30 ms after stimulus onset (awake: $T = 5$, anesthetized: $T = 20$). We square-root transformed spike counts to stabilize the variances (Yu et al. 2009). Before fitting the model we subtracted the average across trials for each stimulus condition and time bin. This procedure removes systematic contributions by the stimulus and, thus, the model explains only the trial-to-trial variability. Note that in this case both the network state x and the observed (transformed) spike counts $\boldsymbol{y}$ have zero mean (over trials) in each bin. The noise covariance under this model is given by

$$\mathrm{Cov}[\boldsymbol{y}] = \boldsymbol{c}\boldsymbol{c}' + R, \quad (4)$$

where the prime ($'$) denotes the transpose, $\boldsymbol{y}$ are the square-root-transformed and mean-subtracted spike counts, $\boldsymbol{c}$ is a vector of linear weights mapping network state to firing rate, and R is a diagonal matrix of residual (independent) variances. We fitted the model for each stimulus condition independently to allow the weights to depend on the stimulus (this was indeed the case: weights increased with firing rates, which was reflected in both the increase of correlations and variance explained with firing rates, Figures 2 and 3). Units were included in the model in all stimulus conditions where they fired at least 0.5 spikes/s during the stimulus period.

The network state x was assumed to evolve smoothly in time. This was achieved by modeling its temporal correlations by a Gaussian kernel:

$$K_{ij} \equiv \mathrm{Cov}[x(t_i), x(t_j)] = \exp\left(-\frac{(t_i - t_j)^2}{2\tau^2}\right) \quad (5)$$

To keep the algorithm computationally tractable we set temporal correlations in network state extending across trials to zero.

To evaluate the fraction of variance explained (VE; Figure 3) and the residual correlations (Figure 5) we used an independent test set that had not been used for fitting the model. Training and test set consisted of the first and second half of the data (and vice versa; i.e. two-fold cross-validation). We fit the model on spike counts in 100-ms windows, but residual correlations and VE can also be evaluated for larger counting windows by summing up variances and (temporal) covariances over several time bins. VE (Figure 3) and residual noise correlations (Figure 5) were calculated for 500 ms windows since this was the maximum available in the awake dataset. For details on how to compute VE and residual correlations see Supplemental Experimental Procedures.

Model of common gain modulation

The model population (Figure 6) consisted of 64 neurons with uniformly spaced preferred orientations and von Mises tuning curves given by

$$f_k(\theta)=\exp(\kappa\cos(2(\theta-\varphi_k))+\alpha), \quad (6)$$

where $\boldsymbol{\varphi}_k$ is the preferred orientation, $k = 2$ and $\boldsymbol{\alpha} = 1.8$, resulting in a bandwidth of ~25° (half-width at half-maximum) and a peak firing rate of 45 spikes/s. The firing rate of each neuron was determined by the product of its tuning curve and the value of the common gain:

$$\mu_k(t)=g(t)\cdot f_k(\theta) \quad (7)$$

The gain had $E[g] = 1$ and its temporal autocorrelation was a Gaussian kernel

$$K_{jk}\equiv \mathrm{Cov}[g(t_j),g(t_k)]=\sigma^2\exp\left(-\frac{(t_j-t_k)^2}{2\tau^2}\right) \quad (8)$$

with $\boldsymbol{\sigma} = 0.15$ and $\boldsymbol{\tau} = 200$ ms. We sampled independent Poisson spike counts from the given rates $\boldsymbol{\mu}(t)$. As for the data, we used bins of 100 ms and computed correlations in bins of 500 ms.

Analysis of spontaneous activity under anesthesia

For the analysis of spontaneous activity (Figure 7) we used the blank periods between two subsequent stimuli. We analyzed segments of 1-s duration starting 200 ms after the end of the stimulus (to avoid contamination by off responses to the stimulus). Approximately 75% of the blank periods were long enough to be included given these criteria, resulting in an average of 1188 ‘trials’ (min: 1148, max: 1225).

Generalized Linear Model accounting for network state

Following Kelly et al. (2010) we fitted a Generalized Linear Model (GLM) with the lowpass-filtered LFP as inputs (Figure 8). The model is defined in Eq. 2 in the main text. As for the GPFA model above we used spike counts in 100-ms bins and fitted the model independently for each stimulus condition. The contribution of the stimulus was captured by the parameter $a(t)$, which represents the PSTH. The LFP predictor $u(t)$ was the bandpass-

filtered (0.5–10 Hz) LFP. We averaged the LFP over all tetrodes that recorded at least one single unit in this session and subtracted the average stimulus-evoked response. The latter ensured that LFP weights captured only fluctuations around the average response to the stimulus. For analysis of the weights (Figure 8C, D) we averaged the weights of each neuron across all conditions in which it was included (firing rate > 0.5 spikes/s). The cross-correlation between LFP and network state estimated by GPFA (Figure 8A, B) was computed by first subtracting the average of each measure within each trial (i.e. it is the correlation of the fluctuations within trials rather than across trials).

Analysis of depth of anesthesia

To assess slow changes in brain state we performed spectral analysis on the LFP (Figure 9). We averaged the LFP across all tetrodes that recorded at least one single unit in this session and computed the spectrogram using 200 overlapping windows (16 trials or ~1 min per window, 50% overlap). The spectrogram was computed on the continuous LFP trace including both evoked and spontaneous activity; no average stimulus response was subtracted. Following Goard & Dan (2009) we computed a power ratio to assess brain state. The power ratio was defined as the power in the low-frequency band (0.5–2 Hz) divided by that in the gamma band (30–70 Hz). To quantify the overall correlation in the network, we computed the variance of the network state variable inferred by the GPFA model in the same windows as we used for the spectral analysis above. For the population analysis (Figure 9E) we used 20 non-overlapping windows to quantify both the power ratio and the overall correlation. This smaller number was chosen as a trade-off between temporal resolution and reducing noise by including more data.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank Dennis Murray, Tori J. Shinn, Allison N. Laudano, Amy M. Morgan, Jessica L. Rudy and Wangchen Wang for technical assistance and help with data collection; Byron Yu and John P. Cunningham for sharing their GPFA code; Sven Simon and Rahul G. Baijal for discussion and consultation with anesthesia; Dimitri Yatsenko for discussion and the development of Data-Joint; Ralf M. Haefner for comments on the manuscript.

This work was supported by grants NEI R01-EY018847, NEI P30-EY002520-33 and the NIH-Pioneer award DP1-OD008301 to A.S.T., the McKnight Scholar Award to A.S.T; the Bern-stein Center for Computational Neuroscience (FKZ 01GQ1002); the German Excellency Initiative through the Centre for Integrative Neuroscience Tübingen (EXC307); an HHMI Early Career Award to S.S. and NEI R01-EY019272 to S.S.

References

Abbott LF, Dayan P. The Effect of Correlated Variability on the Accuracy of a Population Code. Neural Computation. 1999; 11(1):91–101. [PubMed: 9950724]

Angel A. Central Neuronal Pathways and the Process of Anaesthesia. British Journal of Anaesthesia. 1993; 71(1):148–163. [PubMed: 8102065]

Averbeck BB, Latham PE, Pouget A. Neural correlations, population coding and computation. Nat Rev Neurosci. 2006; 7(5):358–366. [PubMed: 16760916]

Bach M, Krüger J. Correlated neuronal variability in monkey visual cortex revealed by a multi-microelectrode. Experimental Brain Research. 1986; 61(3):451–456. [PubMed: 3956608]

Bair W, Zohary E, Newsome WT. Correlated Firing in Macaque Visual Area MT: Time Scales and Relationship to Behavior. The Journal of Neuroscience. 2001; 21(5):1676–1697. [PubMed: 11222658]

Berens P, et al. A Fast and Simple Population Code for Orientation in Primate V1. The Journal of Neuroscience. 2012; 32(31):10618–10626. [PubMed: 22855811]

Blanche TJ, et al. Polytrodes: High-Density Silicon Electrode Arrays for Large-Scale Multiunit Recording. J Neurophysiol. 2005; 93(5):2987–3000. [PubMed: 15548620]

Bowdle TA, Ward RJ. Induction of anesthesia with small doses of sufentanil or fentanyl: dose versus EEG response, speed of onset, and thiopental requirement. Anesthesiology. 1989; 70(1):26–30. [PubMed: 2521435]

Brainard DH. The Psychophysics Toolbox. Spatial Vision. 1997; 10(4):433–436. [PubMed: 9176952]

Buesing L, Macke J, Sahani M. Spectral learning of linear dynamics from generalised-linear observations with application to neural population data. In. Advances in Neural Information Processing Systems. 2012; 25:1691–1699.

Calabrese A, Paninski L. Kalman filter mixture model for spike sorting of non-stationary data. Journal of Neuroscience Methods. 2011; 196(1):159–169. [PubMed: 21182868]

Campagna JA, Miller KW, Forman SA. Mechanisms of Actions of Inhaled Anesthetics. New England Journal of Medicine. 2003; 348(21):2110–2124. [PubMed: 12761368]

Carandini M, Heeger DJ, Movshon JA. Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex. The Journal of Neuroscience. 1997; 17(21):8621–8644. [PubMed: 9334433]

Chi OZ, Field C. Effects of Isoflurane on Visual Evoked Potentials in Humans. Anesthesiology. 1986; 65(3):328–330. [PubMed: 3752582]

Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. Nature Neuroscience. 2011; 14(7):811–819.

Cohen MR, Maunsell JHR. A Neuronal Population Measure of Attention Predicts Behavioral Performance on Individual Trials. The Journal of Neuroscience. 2010; 30(45):15241–15253. [PubMed: 21068329]

Cohen MR, Maunsell JHR. Attention improves performance primarily by reducing interneuronal correlations. Nature Neuroscience. 2009; 12(12):1594–1600.

Cohen MR, Newsome WT. Context-Dependent Changes in Functional Circuitry in Visual Area MT. Neuron. 2008; 60(1):162–173. [PubMed: 18940596]

Constantinople CM, Bruno RM. Effects and Mechanisms of Wakefulness on Local Cortical Networks. Neuron. 2011; 69(6):1061–1068. [PubMed: 21435553]

Denman DJ, Contreras D. The Structure of Pairwise Correlation in Mouse Primary Visual Cortex Reveals Functional Organization in the Absence of an Orientation Map. Cerebral Cortex. 2013

Drummond JCMD. Monitoring Depth of Anesthesia: With Emphasis on the Application of the Bispectral Index and the Middle Latency Auditory Evoked Response to the Prevention of Recall. Anesthesiology. 2000 Sep; 93(3):876–882. 2000. [PubMed: 10969323]

Ecker AS, et al. Decorrelated Neuronal Firing in Cortical Microcircuits. Science. 2010; 327(5965): 584–587. [PubMed: 20110506]

Ecker AS, et al. The Effect of Noise Correlations in Populations of Diversely Tuned Neurons. The Journal of Neuroscience. 2011; 31(40):14272–14283. [PubMed: 21976512]

Gawne TJ, et al. Adjacent Visual Cortical Complex Cells Share About 20% of Their Stimulus-Related Information. Cerebral Cortex. 1996; 6(3):482–489. [PubMed: 8670673]

Gawne TJ, Richmond BJ. How independent are the messages carried by adjacent inferior temporal cortical neurons? The Journal of Neuroscience. 1993; 13(7):2758–2771. [PubMed: 8331371]

Goard M, Dan Y. Basal forebrain activation enhances cortical coding of natural scenes. Nature Neuroscience. 2009; 12(11):1444–1449.

Gutnisky DA, Dragoi V. Adaptive coding of visual information in neural populations. Nature. 2008; 452(7184):220–224. [PubMed: 18337822]

Haider B, Häusser M, Carandini M. Inhibition dominates sensory responses in the awake cortex. Nature. 2013; 493(7430):97–100. [PubMed: 23172139]

Hansen BJ, Chelaru MI, Dragoi V. Correlated Variability in Laminar Cortical Circuits. Neuron. 2012; 76(3):590–602. [PubMed: 23141070]

Harris KD, Thiele A. Cortical state and attention. Nat Rev Neurosci. 2011; 12(9):509–523. [PubMed: 21829219]

Herrero JL, et al. Attention-Induced Variance and Noise Correlation Reduction in Macaque V1 Is Mediated by NMDA Receptors. Neuron. 2013; 78(4):729–739. [PubMed: 23719166]

Hubel DH, Wiesel TN. Receptive fields and functional architecture of monkey striate cortex. The Journal of physiology. 1968; 195(1):215–243. [PubMed: 4966457]

Josic K, et al. Stimulus-Dependent Correlations and Population Codes. Neural Computation. 2009; 21(10):2774–2804. [PubMed: 19635014]

Kelly RC, et al. Local field potentials indicate network state and account for neuronal response variability. Journal of Computational Neuroscience. 2010; 29(3):567–579. [PubMed: 20094906]

Kohn A, Smith MA. Stimulus Dependence of Neuronal Correlation in Primary Visual Cortex of the Macaque. The Journal of Neuroscience. 2005; 25(14):3661–3673. [PubMed: 15814797]

Kohn A, Zandvakili A, Smith MA. Correlations and brain states: from electrophysiology to functional imaging. Current Opinion in Neurobiology. 2009; 19(4):434–438. [PubMed: 19608406]

Loughnan BL, et al. Evoked potentials following diazepam or fentanyl. Anaesthesia. 1987; 42(2):195–198. [PubMed: 3826596]

Macke JH, et al. Empirical models of spiking in neural populations. Advances in neural information processing systems. 2011; 24:13501358.

Mason A, Nicoll A, Stratford K. Synaptic transmission between individual pyramidal neurons of the rat visual cortex in vitro. The Journal of Neuroscience. 1991; 11(1):72–84. [PubMed: 1846012]

Mitchell JF, Sundberg KA, Reynolds JH. Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4. Neuron. 2009; 63(6):879–888. [PubMed: 19778515]

Movshon JA, et al. Cortical responses to visual motion in alert and anesthetized monkeys. Nature Neuroscience. 2003; 6(1):3–3.

Movshon JA, et al. The analysis of moving visual patterns. Pattern recognition mechanisms. 1985; 54:117–151.

Niell CM, Stryker MP. Modulation of Visual Responses by Behavioral State in Mouse Visual Cortex. Neuron. 2010; 65(4):472–479. [PubMed: 20188652]

Nienborg H, Cumming BG. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature. 2009; 459(7243):89–92. [PubMed: 19270683]

Ohiorhenuan IE, et al. Sparse coding and high-order correlations in fine-scale cortical networks. Nature. 2010; 466(7306):617–621. [PubMed: 20601940]

Petersen CCH, et al. Interaction of sensory responses with spontaneous depolarization in layer 2/3 barrel cortex. Proceedings of the National Academy of Sciences. 2003; 100(23):13638–13643.

Poulet JFA, Petersen CCH. Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice. Nature. 2008; 454(7206):881–885. [PubMed: 18633351]

Reich DS, Mechler F, Victor JD. Independent and Redundant Information in Nearby Cortical Neurons. Science. 2001; 294(5551):2566–2568. [PubMed: 11752580]

Renart A, et al. The Asynchronous State in Cortical Circuits. Science. 2010; 327(5965):587–590. [PubMed: 20110507]

Schwender D, et al. Effects of Increasing Doses of Alfentanil, Fentanyl and Morphine on Mid-Latency Auditory Evoked Potentials. British Journal of Anaesthesia. 1993; 71(5):622–628. [PubMed: 8251268]

Shadlen MN, Newsome WT. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. The Journal of Neuroscience. 1998; 18(10): 3870–3896. [PubMed: 9570816]

Smith MA, et al. Laminar dependence of neuronal correlations in visual cortex. Journal of Neurophysiology. 2013; 109(4):940–947. [PubMed: 23197461]

Smith MA, Kohn A. Spatial and Temporal Scales of Neuronal Correlation in Primary Visual Cortex. J Neurosci. 2008; 28(48):12591–12603. [PubMed: 19036953]

Smith MA, Sommer MA. Spatial and Temporal Scales of Neuronal Correlation in Visual Area V4. The Journal of Neuroscience. 2013; 33(12):5422–5432. [PubMed: 23516307]

Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. The Journal of Neuroscience. 1993; 13(1):334–350. [PubMed: 8423479]

Sompolinsky H, et al. Population coding in neuronal systems with correlated noise. Physical Review E. 2001; 64(5):051904.

Steriade M, McCormick D, Sejnowski T. Thalamocortical oscillations in the sleeping and aroused brain. Science. 1993; 262(5134):679–685. [PubMed: 8235588]

Tolias AS, et al. Recording Chronically From the Same Neurons in Awake, Behaving Primates. Journal of Neurophysiology. 2007; 98(6):3780–3790. [PubMed: 17942615]

De Valois RL, Albrecht DG, Thorell LG. Spatial frequency selectivity of cells in macaque visual cortex. Vision Research. 1982; 22(5):545–559. [PubMed: 7112954]

De Valois RL, William Yund E, Hepler N. The orientation and direction selectivity of cells in macaque visual cortex. Vision Research. 1982; 22(5):531–544. [PubMed: 7112953]

Yu BM, et al. Gaussian-Process Factor Analysis for Low-Dimensional Single-Trial Analysis of Neural Population Activity. Journal of Neurophysiology. 2009; 102(1):614–635. [PubMed: 19357332]

Zeki SM. Cells responding to changing image size and disparity in the cortex of the rhesus monkey. The Journal of physiology. 1974; 242(3):827–841. [PubMed: 4449056]

Zohary E, Shadlen MN, Newsome WT. Correlated neuronal discharge rate and its implications for psychophysical performance. Nature. 1994; 370(6485):140–143. [PubMed: 8022482]

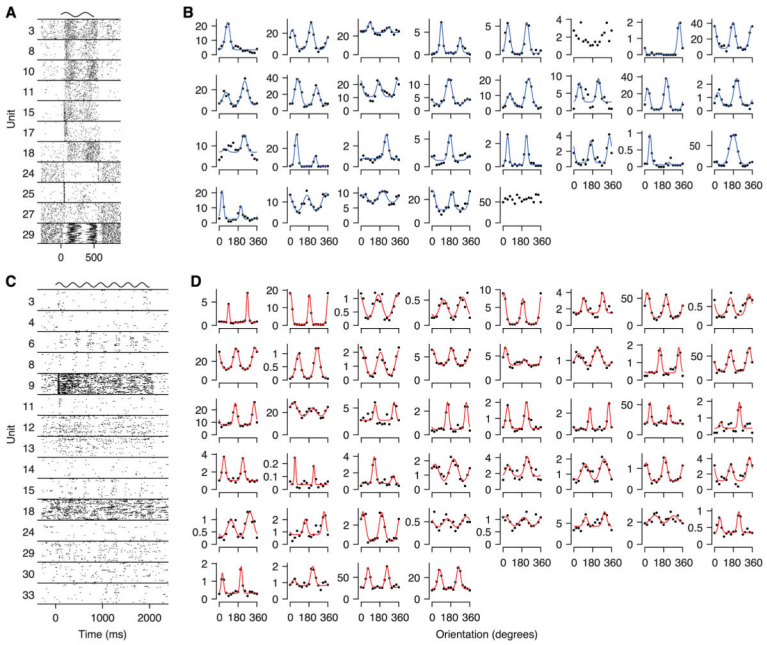

Figure 1. Recordings of population activity in V1

A, Spike rasters for a subset of the neurons recorded in one example session during wakefulness. The sinusoid at the top indicates the stimulus duration (500 ms) and its temporal frequency. Numbers: neuron numbers in B, counted from left to right, top to bottom. **B**, Tuning curves for all neurons in the same session as in A. Solid lines: least squares fit, shown only for cells significantly tuned to orientation (27/29 cells at $p < 0.01$, non-corrected). **C**, Spike rasters during anesthesia (as in A). **D**, Tuning curves (as in B; all 44 neurons significantly tuned at $p < 0.001$).

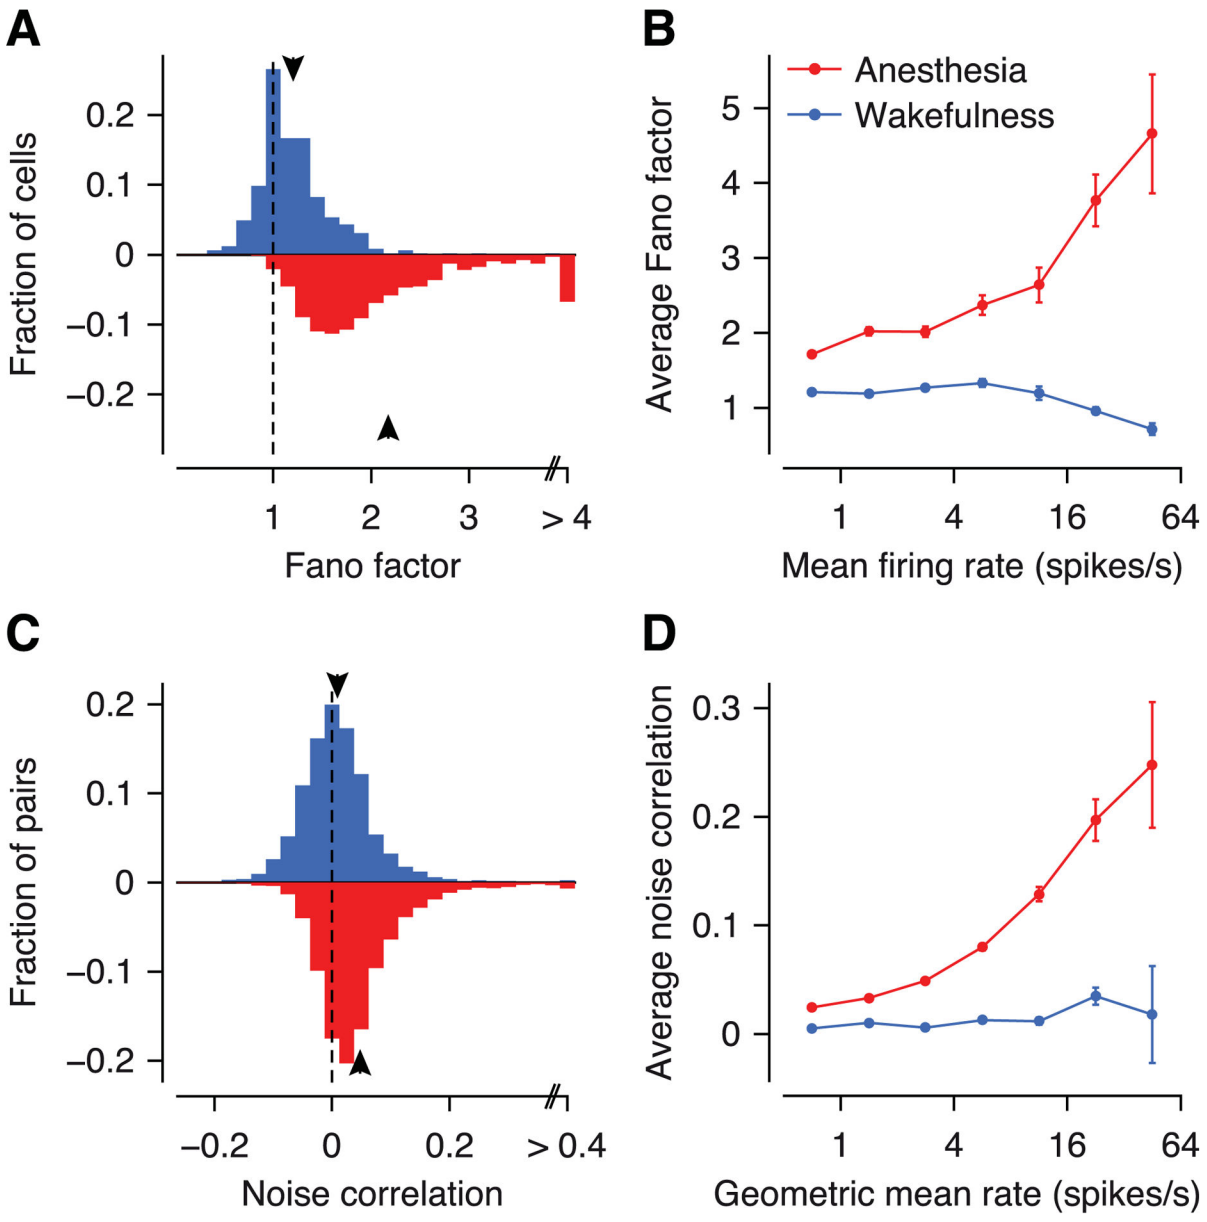

Figure 2. Fano factors and noise correlations during wakefulness (blue) and anesthesia (red)

A, Distribution of Fano factors. Arrows: means. **B**, Dependence of Fano factors on firing rates. Error bars: SEM. **C**, Distribution of noise correlations. **D**, Dependence of noise correlations on geometric mean firing rates.

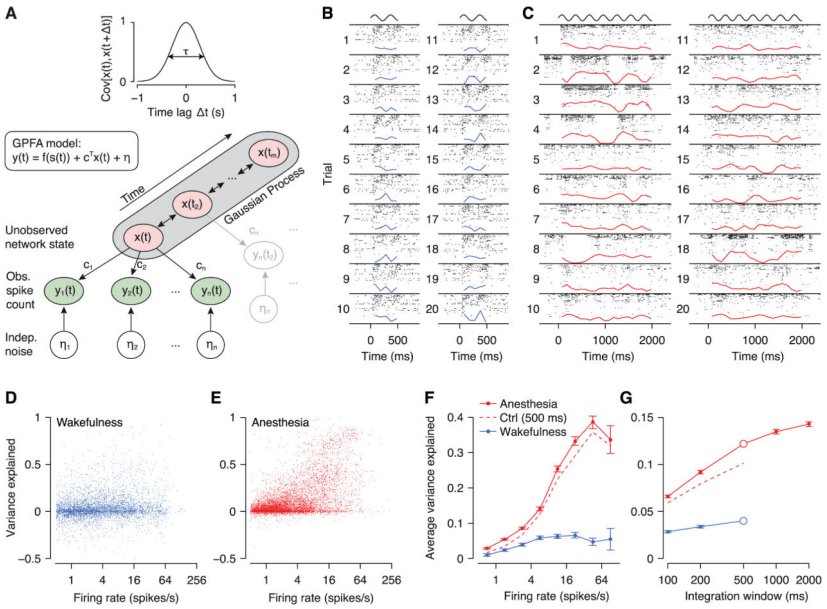

Figure 3. Gaussian Process Factor Analysis (GPFA)

A, Schematic of the GPFA model. Spike count variability is generated by an unobserved (one-dimensional) network state (x) linearly driving neural activity (weights *c*) plus independent noise (η). The network state evolves smoothly in time, which is modeled by a Gaussian Process with temporal covariance shown at the top (correlation timescale **τ** is learned from the data). **B**, Population rasters for an example session recorded in an awake animal. Each numbered row shows the rasters of all recorded neurons during a single trial. All trials were under identical stimulus conditions (500 ms drifting grating, indicated by sine wave at the top). Blue line: estimate of the network state (x). The visible rate modulations are locked to the phase of the stimulus, but not to the estimated network state (which in this case had very little explanatory power). **C**, As in B, but under anesthesia. The estimated network state captures the population rate dynamics very well (see e.g. trials 1–4), but is unrelated to the stimulus (stimulus duration: 2 s). **D**, Scatter plot of variance explained (VE) vs. firing rate during wakefulness. Each dot is a single neuron under one stimulus condition. VE is computed in 500-ms windows. **E**, As in D, but under anesthesia. **F**, Binned and averaged representation of D and E. Error bars: SEM. Dashed lines: model fit on anesthetized data using only the first 500 ms of each trial for better comparison with awake data (error bars omitted for clarity; they were comparable to those for the solid red line). **G**, Average VE versus size of integration window. Open circles: 500 ms window, which was used for panels D–F. Dashed line: control analysis as in panel F.

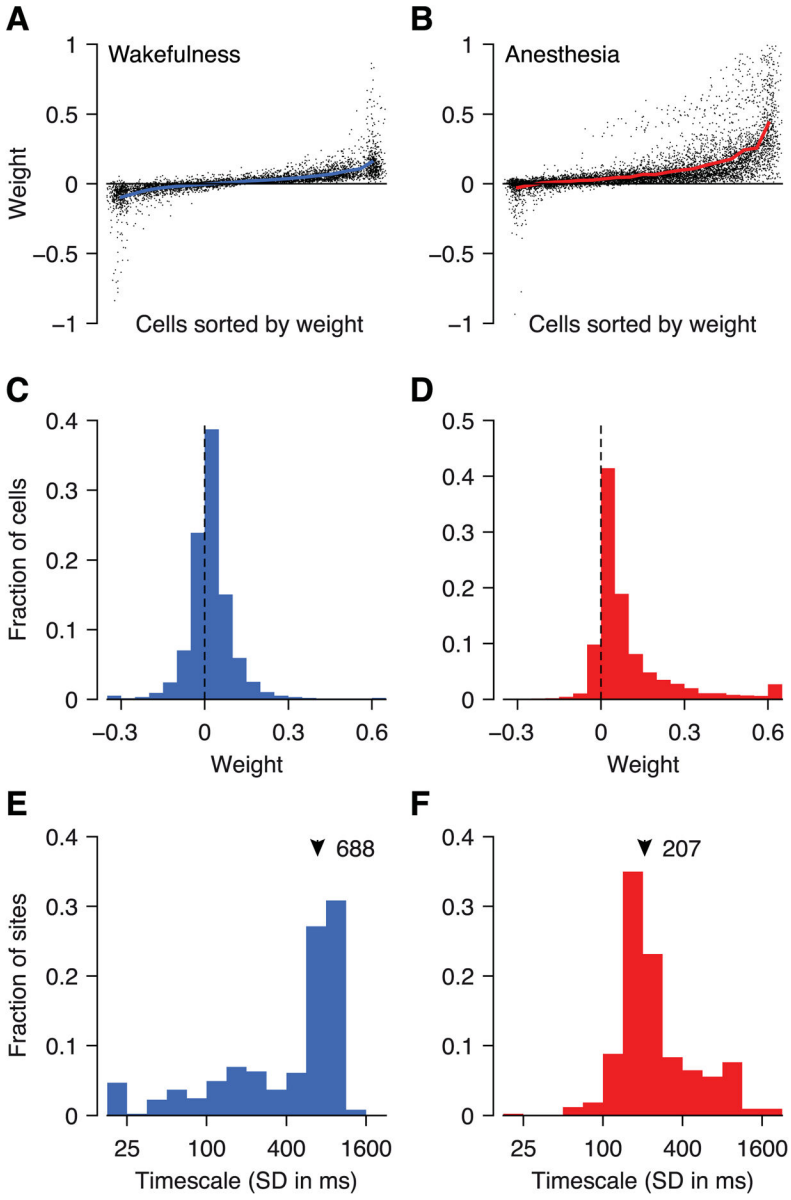

Figure 4. GPFA model parameters

Distribution of weights (variable ***c***, Eq. 1) during wakefulness (**A**, **C**) and under anesthesia (**B**, **D**). Timescale of network state dynamics during wakefulness (**E**) and under anesthesia (**F**). The timescale is the parameter (**τ**) in the squared-exponential temporal correlation function of the latent variable (*x*) in the GPFA model.

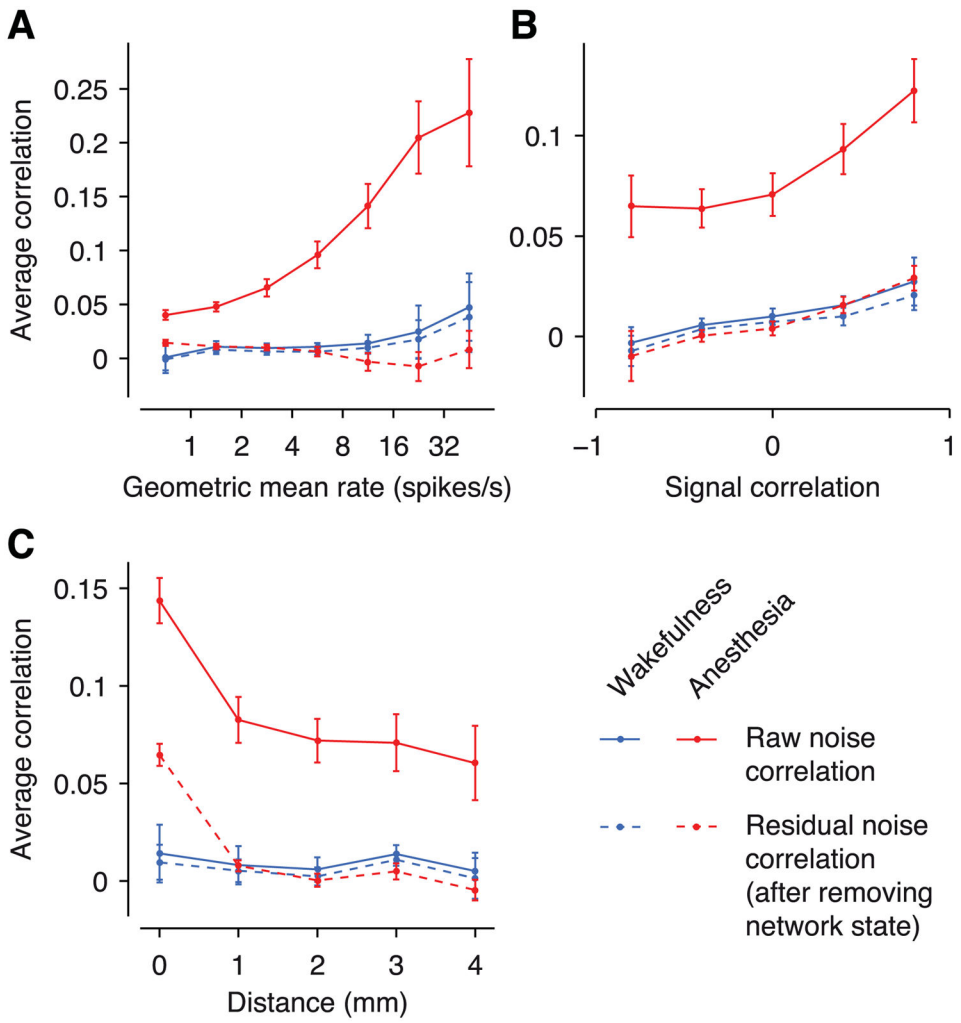

Figure 5. Accounting for network state reduces noise correlations under anesthesia

Raw (solid lines) and residual (after accounting for network state; dashed lines) noise correlations during wakefulness (blue) and under anesthesia (red). Dependence on firing rates (**A**), signal correlations (**B**) and distance between cells (**C**). Raw correlations in panel A are as in Figure 2D, except that here the model is fit for each condition separately. Error bars: SEM.

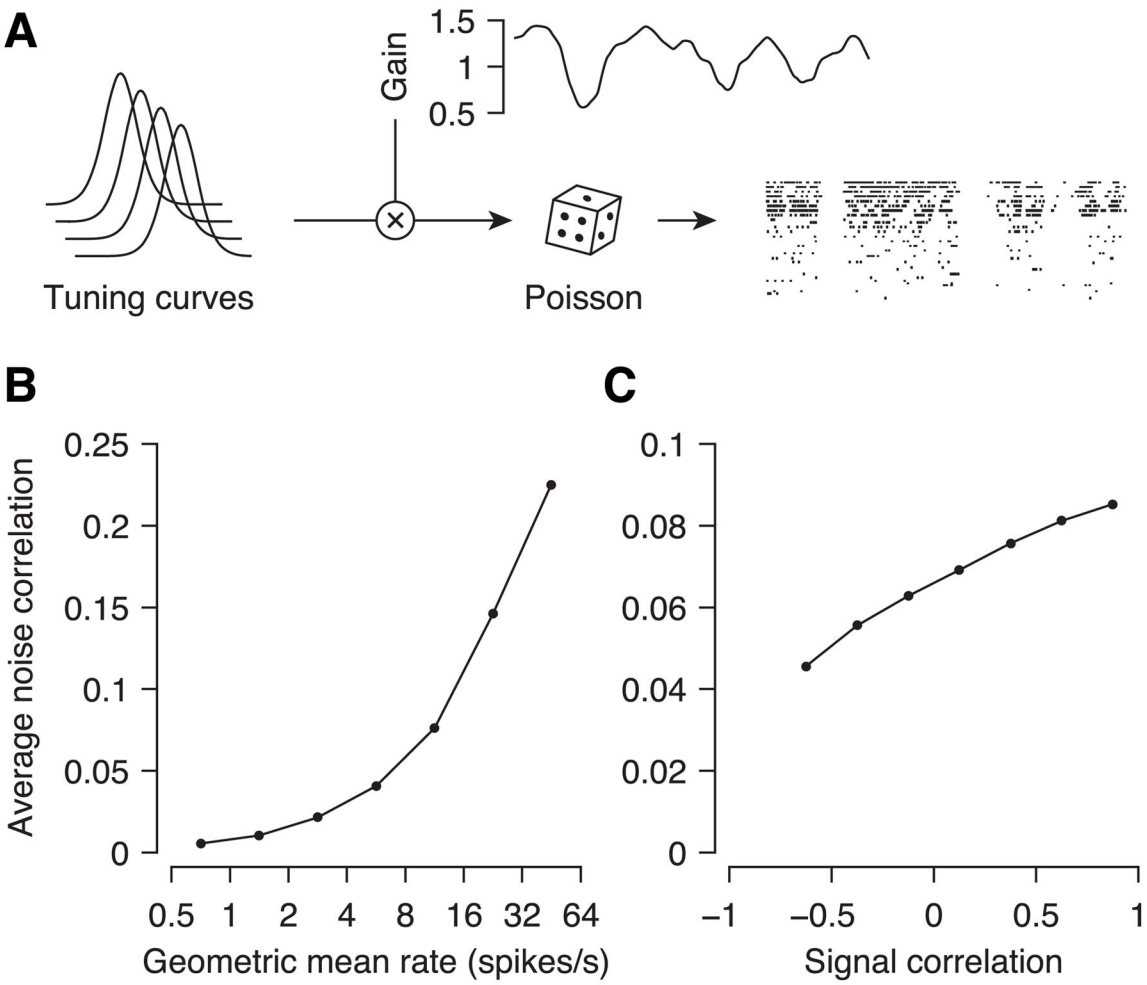

Figure 6. Model of state fluctuations as common fluctuations in excitability

A, Illustration of the model. Cells have tuning curves with identical shapes and regularly spaced preferred orientations. Each cell's firing rate is given by the tuning curve multiplied by the common gain, which changes slowly as in our data. Spikes are generated by independent inhomogeneous Poisson processes with the given rates. The resulting noise correlations increase with firing rates (**B**) and signal correlations (**C**), as in the data.

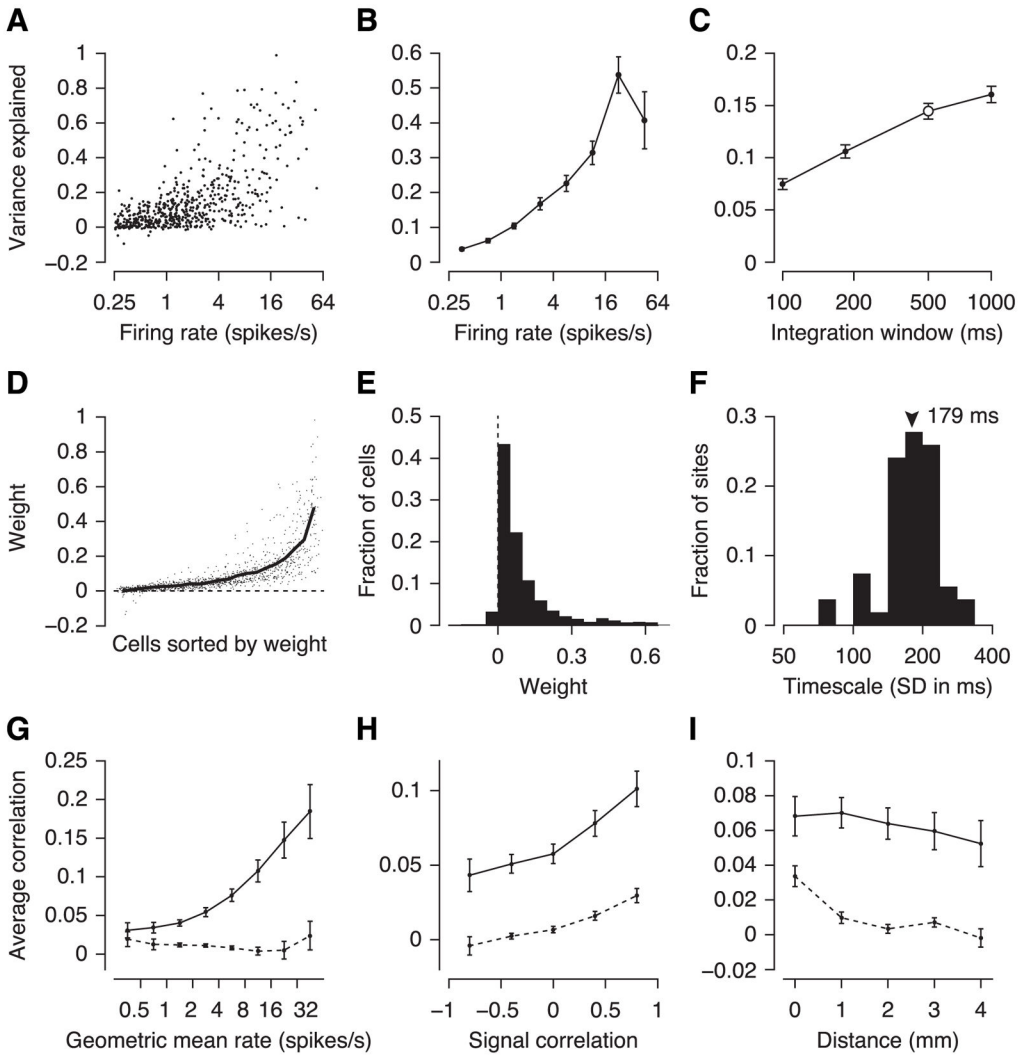

Figure 7. GPFA model during spontaneous activity under anesthesia

A–B, Variance explained (VE) versus firing rates (as in Figure 3D–F). **C,** VE vs. integration time (as in Figure 3G). **D–E**, Distribution of weights (as in Figure 4B, D). **F,** Distribution of timescales (as in Figure 4G). **G–I**, Residual correlations versus firing rate, signal correlation and distance, respectively (as in Figure 5A–C).

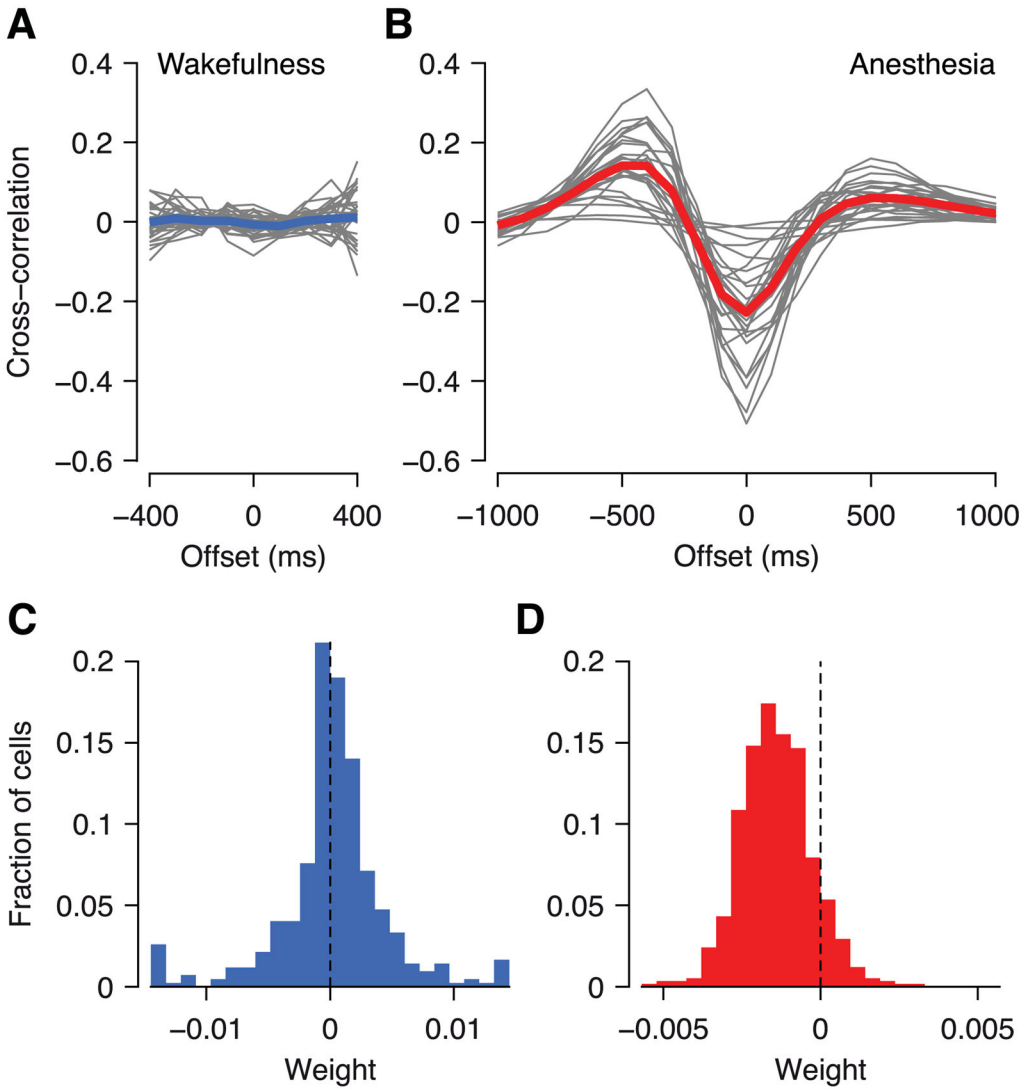

Figure 8. Local field potential is correlated with inferred network state and predicts trial-to-trial variability under anesthesia but not during wakefulness

A, Cross-correlation between low-frequency LFP (0.5–10 Hz) and network state inferred by GPFA model during wakefulness. Gray lines: individual sessions; blue line: average across all sessions. **B,** As in A, but under anesthesia. **C**, Distribution of LFP weights in Generalized Linear Model taking stimulus and LFP into account; during wakefulness. **D,** As in C, but under anesthesia.

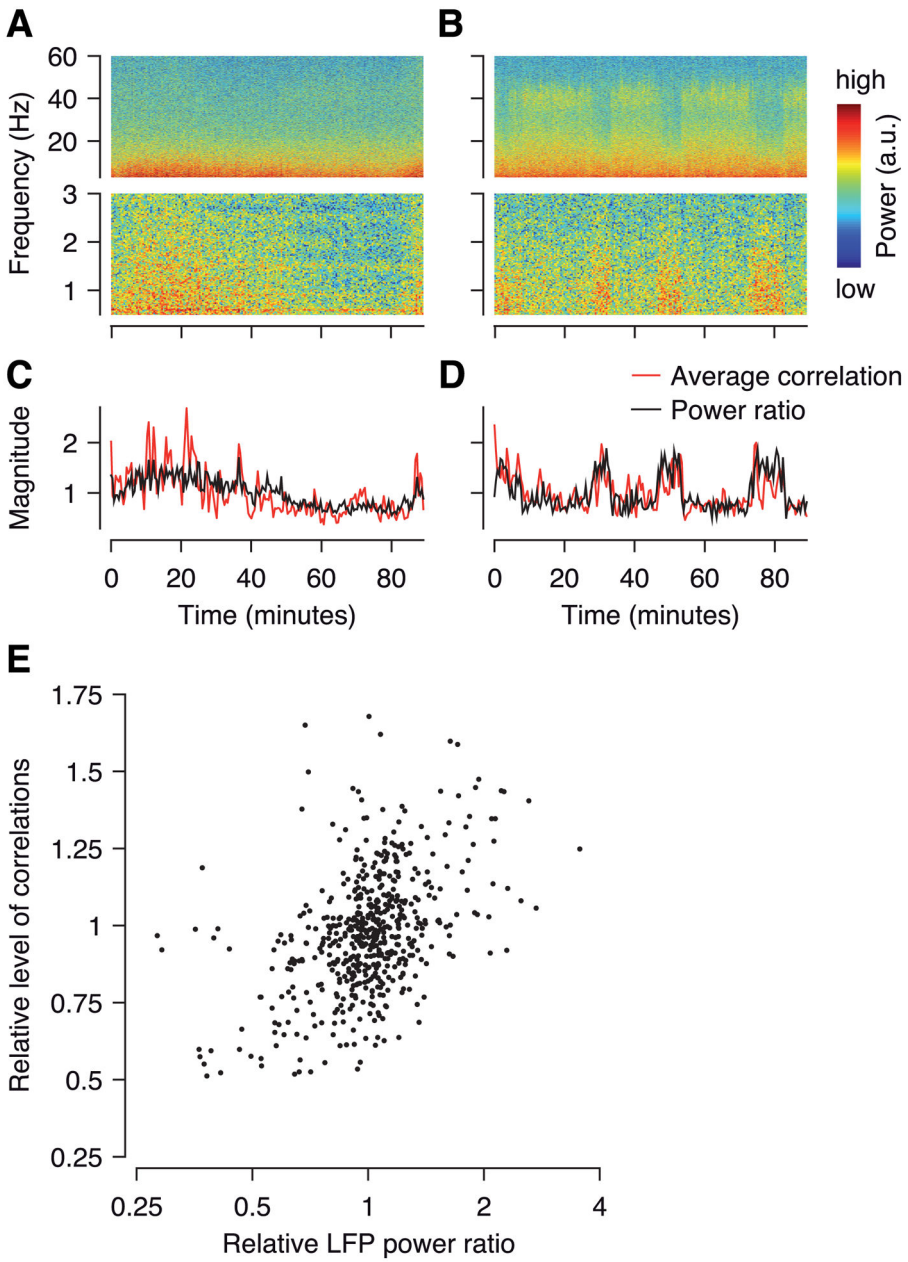

Figure 9. LFP power ratio correlates with overall level of noise correlations

A–B, Spectrogram of LFP over the course of two example recordings (~90 minutes). **C–D,** LFP power ratio (black line, power in 0.5–2 Hz band divided by that in the gamma band, 30–70 Hz) and average level of correlations (red line, variance of the network state inferred by GPFA) for the same sessions. Both quantities are normalized by the session average. **E,** Population analysis. LFP power ratio versus overall correlation (variance of network state inferred by GPFA) in 20 separate blocks per recording (27 recordings in total, i.e. $n = 540$). Both quantities normalized by the session average for each session. One outlier cropped for clarity.