

Variability and Correlations in Primary Visual Cortical Neurons Driven by Fixational Eye Movements

 James M. McFarland,¹ Bruce G. Cumming,² and  Daniel A. Butts¹

¹Department of Biology and Program in Neuroscience and Cognitive Science, University of Maryland, College Park, Maryland 20815, and ²Laboratory of Sensorimotor Research, National Eye Institute, National Institutes of Health, Bethesda, Maryland 20892

The ability to distinguish between elements of a sensory neuron's activity that are stimulus independent versus driven by the stimulus is critical for addressing many questions in systems neuroscience. This is typically accomplished by measuring neural responses to repeated presentations of identical stimuli and identifying the trial-variable components of the response as noise. In awake primates, however, small “fixational” eye movements (FEMs) introduce uncontrolled trial-to-trial differences in the visual stimulus itself, potentially confounding this distinction. Here, we describe novel analytical methods that directly quantify the stimulus-driven and stimulus-independent components of visual neuron responses in the presence of FEMs. We apply this approach, combined with precise model-based eye tracking, to recordings from primary visual cortex (V1), finding that standard approaches that ignore FEMs typically miss more than half of the stimulus-driven neural response variance, creating substantial biases in measures of response reliability. We show that these effects are likely not isolated to the particular experimental conditions used here, such as the choice of visual stimulus or spike measurement time window, and thus will be a more general problem for V1 recordings in awake primates. We also demonstrate that measurements of the stimulus-driven and stimulus-independent correlations among pairs of V1 neurons can be greatly biased by FEMs. These results thus illustrate the potentially dramatic impact of FEMs on measures of signal and noise in visual neuron activity and also demonstrate a novel approach for controlling for these eye-movement-induced effects.

Key words: fixational eye movements; noise correlations; V1; variability

Significance Statement

Distinguishing between the signal and noise in a sensory neuron's activity is typically accomplished by measuring neural responses to repeated presentations of an identical stimulus. For recordings from the visual cortex of awake animals, small “fixational” eye movements (FEMs) inevitably introduce trial-to-trial variability in the visual stimulus, potentially confounding such measures. Here, we show that FEMs often have a dramatic impact on several important measures of response variability for neurons in primary visual cortex. We also present an analytical approach for quantifying signal and noise in visual neuron activity in the presence of FEMs. These results thus highlight the importance of controlling for FEMs in studies of visual neuron function, and demonstrate novel methods for doing so.

Introduction

Understanding which aspects of a sensory neuron's activity represent features of the stimulus and which constitute “noise” with respect to the stimulus is central to many fundamental questions

in systems neuroscience. In particular, measures of neuronal variability are required for determining how reliably sensory neurons encode stimuli, and understanding how such variability is correlated across populations of neurons has important implications for theories of neural coding (Tolhurst et al., 1983; Shadlen and Newsome, 1998; Abbott and Dayan, 1999; Averbeck et al., 2006). Furthermore, determining how experimental manipulations, as well as changes in internal state (e.g., attention, arousal), affect sensory coding generally depends on the ability to accurately measure the components of neural activity that are reliably driven by the stimulus.

Measurements of the “signal” and “noise” components of a neuron's response typically rely on making repeated presentations of a set of stimuli. The across-trial average response (peristimulus time histogram, PSTH) captures the component of the

Received Dec. 31, 2015; revised March 30, 2016; accepted April 20, 2016.

Author contributions: J.M.M., B.G.C., and D.A.B. designed research; J.M.M. and B.G.C. performed research; J.M.M. analyzed data; J.M.M. and D.A.B. wrote the paper.

This work was supported by the National Science Foundation Grant IIS-1350990 (D.A.B.) and the National Institutes of Health Grant F32-EY-023921 (J.M.M.) and National Eye Institute Intramural Research Program Grant (B.G.C.). We thank D. Parker and I. Bunea for providing excellent animal care and to M. Whiteway for helpful comments on the manuscript.

The authors declare no competing financial interests.

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

DOI:10.1523/JNEUROSCI.4660-15.2016

Copyright © 2016 the authors 0270-6474/16/366225-17\$15.00/0

neuron's activity driven by the stimulus, whereas across-trial variability reflects the stimulus-independent component, typically treated as noise. In the visual cortex of awake animals, however, such measurements can be confounded by the presence of fixational eye movements (FEMs), which are composed of rapid jumps known as microsaccades and slower fixational drift (Martinez-Conde et al., 2004). Therefore, even for animals that are well trained to maintain precise fixation, FEMs will generate uncontrolled trial-to-trial variability in the visual stimulus on the retina. Because the receptive fields (RFs) of neurons in primary visual cortex (V1) are fixed in retinotopic coordinates (Gur and Snodderly, 1987, 1997), the responses of V1 neurons will exhibit additional trial-to-trial variability reflecting variation in the position of the stimulus on the retina. Indeed, previous studies have shown that measurements of V1 response reliability in awake animals can be greatly biased by FEMs (Gur and Snodderly, 1987, 1997, 2006).

In general, however, few studies have attempted to address this issue, largely because of technical limitations in measuring an animal's eye position with sufficient accuracy to track FEMs during electrophysiological recordings. In particular, standard eye-tracking hardware (such as scleral search coils and video trackers) typically have accuracies of $\sim 0.1^\circ$ (Read and Cumming, 2003; Kimmel et al., 2012; McFarland et al., 2014), which is approximately the magnitude of the FEMs themselves. Furthermore, even with accurate measurements of eye position, without the ability to present the same stimulus over many trials, distinguishing between stimulus-driven and stimulus-independent response variability remains a challenging problem. Therefore, for the most part, the effects of FEMs remain overlooked, presumably under the assumption that they are small, at least for neurons with RFs outside of the high-resolution central portion of the visual field (the fovea).

Here, we address this issue using several analytical innovations. First, we overcome the limited accuracy of conventional eye-tracking approaches by using a recently developed method for inferring an animal's eye position from the recorded activity of populations of V1 neurons (McFarland et al., 2014). Second, we develop general analytical methods for quantifying stimulus-driven and stimulus-independent components of V1 activity in the presence of measured FEMs. Importantly, our approach makes minimal assumptions about the stimulus tuning of V1 neurons or the structure of noise in their spiking activity (Amarasingham et al., 2015).

Applying this combination of analytical methods, we find that FEMs typically have large, and in many cases quite dramatic, effects. On average, about half of the stimulus-driven response variance of parafoveal V1 neurons (i.e., those recorded in typical experiments) is missed by the across-trial averages (PSTHs). This stimulus-driven component is misattributed to "noise," inflating measures of response variability such as the Fano factor (FF). We also describe a simple theoretical framework demonstrating that the effects we observe are not specific to the visual stimulus used in our study and are likely to remain substantial across a wide range of commonly used visual stimuli, as well as when measuring spiking responses over much larger time windows. Finally, we apply our method to quantify the effects of FEMs on estimates of correlations between pairs of simultaneously recorded V1 neurons. We find that FEM-induced biases lead to overestimates of the magnitude of stimulus-independent correlations and thus will distort typical measurements of "signal" and "noise" correlations.

Materials and Methods

Electrophysiology. Multielectrode recordings were made from V1 of two awake, head-restrained male rhesus macaques (*Macaca mulatta*; 13–15 years old) trained to fixate for fluid reward. A head-restraining post, recording chamber, and scleral search coils were implanted under general anesthesia and sterile conditions (Read and Cumming, 2003). A custom microdrive was used to introduce linear electrode arrays (U-probe or V-probe; Plexon; 24 contacts, 50 μm spacing) on each recording day. Eye position was monitored continuously using scleral search coils and sampled at 600 Hz. Stimuli were displayed on cathode ray tube monitors (100 Hz refresh rate; 1280 \times 1024 resolution) subtending 24.3° \times 19.3° of visual angle and viewed through a mirror haploscope. All protocols were approved by the Institutional Animal Care and Use Committee and complied with Public Health Service policy on the humane care and use of laboratory animals.

Broadband extracellular signals were sampled continuously at 40 kHz and stored to disk. Spikes were detected using a voltage threshold applied to the high-pass-filtered signals (low cutoff frequency of 100 Hz). Thresholds were initially set to yield an event rate of 50 Hz and were subsequently lowered if needed to capture any putative single-unit spiking. Spike sorting was performed offline using custom software. Briefly, spike clusters were modeled using Gaussian mixture distributions that were based on several different spike features, including principal components, voltage samples, and template scores. The features providing the best cluster separation were used to cluster single units. Cluster quality was quantified using a variety of measures including "L-ratio," "isolation distance" (Schmitzer-Torbert et al., 2005), and a variant of "D-prime." Only spike clusters that were well isolated using these measures—confirmed by visual inspection—were used for analysis. We verified that our results were not sensitive to cluster isolation criteria.

Behavioral task and visual stimuli. During recordings, the animals performed a simple fixation task in which they were required to maintain gaze within a small window (between 1° and 1.4° wide) around a fixation target to obtain a liquid reward after each completed 4 s trial. A "ternary bar noise" stimulus consisting of random patterns of black, white, and gray bars (gray matching the mean luminance of the screen) was used. Bars stretched the length of the monitor and spanned a region from 1.2° to 6° wide, centered on the neurons' RFs. The width of individual bars ranged from 0.038° to 0.1°, depending on the RF sizes of the recorded neurons. Bar patterns were displayed at 100 Hz and were uncorrelated in space and time. The probability of a given bar being nongray (i.e., black or white) was set to provide either a sparse luminance distribution (88% gray) or a dense distribution (33% gray), yielding similar results in both cases. The orientation of the bars was chosen to correspond to the preferred orientation of the majority of units recorded in a given session. In cases in which the neurons did not have a consistent preferred orientation, we performed recordings using two orthogonal bar orientations.

For most trials, the stimulus consisted of a unique sequence of bar patterns. A subset of "repeat" trials in which identical "frozen noise" sequences were presented was interleaved randomly. Either one or two unique frozen noise sequences (each spanning the entire 4 s trial) were presented between 26 and 414 total repeat trials (average, 233 trials; $n = 22$ recording sessions) in each experiment.

In some experiments (for purposes of a separate study), the fixation target made periodic jumps (3–4° amplitude) parallel to the orientation of the bar stimuli, requiring the animal to make "guided saccades" to maintain fixation on the target (McFarland et al., 2015). Furthermore, in some trials, histogram-equalized natural images (van Hateren and van der Schaaf, 1998) were displayed in the background. The animals' overall fixation accuracy was slightly higher in the "guided saccade" trials (median difference in eye position SD: 3%; $p = 0.055$; Wilcoxon signed-rank test; $n = 8$ experiments), as well as when there was an image background (median difference: 15%; $p = 0.0078$; $n = 8$). Because these differences were small, we pooled together these conditions for all analyses. Note that, in both cases, fixation accuracy tended to be lower during the basic "fixation with gray background" condition so our reported effects of FEMs would be larger if we isolated our analyses to only those trials.

Criteria for data selection. The first 200 ms and last 50 ms of each trial were excluded from analysis to minimize the impact of onset transients and fixation breaks. Only single units that had an average firing rate of at least 5 spikes/s during repeat trials, and for which the neuron was isolated for at least 25 repeat trials, were used. Similarly, for analysis of covariance, only pairs of neurons that were simultaneously isolated for at least 25 repeat trials were included. For neurons that were recorded with multiple stimulus orientations, the stimulus orientation that provided the better stimulus-processing model (larger log-likelihood improvement relative to the null model) was used. The results were verified as being robust to the precise choices of these selection criteria.

Estimation of stimulus-processing models. A recently developed modeling framework was used to describe each neuron's stimulus processing as a second-order linear-nonlinear cascade (LNLN) (Butts et al., 2011; Park and Pillow, 2011; Vintch et al., 2012; McFarland et al., 2013) following the approach described in detail previously (McFarland et al., 2013). Specifically, a neuron's predicted spike rate $r(t)$ is given by a sum over LN subunits, followed by a spiking nonlinearity function:

$$r(t) = F \left[\sum_j w_j f_j[\mathbf{k}_j s(t)] \right] = F[g(t)] \quad (M1)$$

where $g(t)$ is the "generating signal," $s(t)$ is the retinal stimulus at time t (with relevant history "time-delay-embedded"), the \mathbf{k}_j are a set of stimulus filters, with corresponding static nonlinearities $f_j(\cdot)$, the w_j are coefficients that determine whether each input is excitatory or inhibitory (constrained to be ± 1), and $F[\cdot]$ is the spiking nonlinearity function. A spiking nonlinearity of the following form was used:

$$F[x; \alpha, \beta, \theta] = \alpha \log[1 + e^{\beta(x-\theta)}] \quad (M2)$$

Assuming that the neuron's spike counts $R_{\text{obs}}(t)$ are described by a conditionally inhomogeneous Poisson count process, the log-likelihood LL of the model is given by:

$$LL = \sum_t [R_{\text{obs}}(t) \log[r(t)] - r(t)] + C \quad (M3)$$

where C is independent of the predicted rate $r(t)$ and is thus unaffected by changes in model parameters. Estimation of model parameters is accomplished by direct maximization of the LL (McFarland et al., 2013).

Although the set of functions $f_j(\cdot)$ can be inferred from the data (McFarland et al., 2013), in this case, the $f_j(\cdot)$ were limited to linear and squared functions, which represent a quadratic approximation of the neuron's stimulus-response function that provides a robust description of the stimulus processing of V1 simple and complex cells (Park and Pillow, 2011; McFarland et al., 2013).

To determine the number of excitatory and inhibitory quadratic subunits, a sequence of such models was fitted with increasing numbers of subunits. We stopped adding subunits when the cross-validated model LL (using a randomly selected subset of 20% of trials for cross-validation) no longer improved, allowing up to a maximum of four excitatory and four inhibitory "squared" subunits in addition to the linear filter. Note that the repeat trials were not used for estimating parameters of the stimulus-processing models.

To mitigate the possibility of overfitting, spatiotemporal smoothness and sparseness regularization were incorporated when estimating the stimulus filters. The magnitude of the smoothness regularization penalty for each unit was selected automatically by searching a grid spanning four orders of magnitude and selecting the value that provided the largest cross-validated LL (using the same 20% of cross-validation trials). In addition, a separate sparseness regularization parameter, which was set by hand, was also used. Details regarding stimulus filter regularization, as well as methods for finding maximum likelihood parameter estimates, were described previously (McFarland et al., 2013).

Estimation of RF tuning properties. The estimated stimulus-processing models were used to derive several measures of stimulus tuning for each neuron. The spatial location of each neuron's RF, as well as its width, were estimated as follows. First, the "spatial profile" of each stimulus filter was obtained by computing the SD of its coefficients across time lags at each spatial position. The average of these spatial profiles across a

neuron's stimulus filters (weighted by the relative contribution of each filter to the overall predicted firing rate) was then computed. Finally, a Gaussian function was fitted to the average spatial profile (minimizing squared error) given by the following:

$$f(x; \alpha, \mu, \sigma, c) = c + \alpha e^{-(x-\mu)^2/2\sigma^2} \quad (M4)$$

The RF location and RF width were then defined to be μ and 2σ , respectively.

Each neuron's sensitivity to the spatial phase of the stimulus was measured according to the following:

$$\frac{\sum_t [r(\mathbf{s}(t)) - r(-\mathbf{s}(t))]^2}{\sum_t [r(\mathbf{s}(t)) - \bar{r}]^2} \quad (M5)$$

where $r(\mathbf{s}(t))$ is the neuron's predicted firing rate in response to the stimulus at time t , and $r(-\mathbf{s}(t))$ is the predicted firing rate in response to the polarity-reversed stimulus. This quantity thus measures a neuron's average firing rate modulation under polarity reversal of the stimulus relative to its firing rate modulation across the ensemble of stimuli.

Model-based eye tracking. To overcome the limitations of hardware-based eye-trackers, a recently developed model-based eye-tracking method capable of accurately tracking FEMs was used. Details of this procedure have been described previously (McFarland et al., 2014) and are summarized here. This method uses probabilistic models of each neuron's stimulus processing to infer the most likely sequence of eye positions given the spiking activity recorded from a population of V1 neurons (using multi-unit and single-unit activity). Although there is typically only a small amount of information about eye position provided by a single neuron's activity at a given time, such information can be integrated across a population of neurons over time, resulting in accurate estimates of the animal's eye position with good spatial and temporal resolution. Although the accuracy of eye position estimates depends on the size and composition of the recorded neural population, we previously showed that reliable estimates can be obtained even when using only multi-unit signals from a 24-channel laminar probe (McFarland et al., 2014).

Eye movements were modeled as consisting of slow drifts by putting a zero-mean Gaussian prior on the instantaneous eye velocity except during saccades, when the eye position was allowed to change rapidly (McFarland et al., 2014). Saccades were detected using the scleral search coils by finding times when the instantaneous eye velocity exceeded 10°/s, with the onset/termination of saccades defined as the time when eye velocity first/last exceeded 3°/s. Microsaccades were defined as saccades with an amplitude of $<1^\circ$. Additional information from the eye-coil signals can also be incorporated into the eye position inference procedure to further improve its accuracy. In particular, coil measurements of short-timescale changes in eye position, such as displacements during microsaccades and measured fixational drift velocities, avoid biases created by slow drifts in the eye coil signals (McFarland et al., 2014). These additional coil signals were only used for one of the two animals, which was found to have more reliable eye coil signals. For the other animal, the eye coils were only used to detect saccade timing. Absolute position information from the coils was not used from either animal.

For purposes of eye tracking, quadratic models were used to describe each neuron's stimulus processing, consisting of a linear filter and two excitatory squared filters, such that each neuron's firing rate at time t was given by:

$$r(t) = F \left[\sum_i f_i[\mathbf{k}_i s(t)] + \sum_\tau c_\tau I_\tau(t) + \gamma p(t) + \sum_b \theta_b I_b(t) \right] \quad (M6)$$

In addition to modeling the dependence of each neuron's firing rate on the stimulus (left-most terms, matching those in Eq. M1 above), several additional covariates were incorporated into the models to improve eye-tracking performance (McFarland et al., 2014). Specifically, a set of parameters (θ_b) were included to capture slow variations in average spike

rates across separate 60–90 trial-recording “blocks,” coefficients c_r to capture “extra-retinal” perisaccadic firing rate modulation (McFarland et al., 2015), and a term γ to describe each neuron’s coupling to the population average firing rate $p(t)$. The $I_a(t)$ here are indicator functions that take a value of 1 if a saccade/microsaccade occurred with latency τ at time t and are 0 otherwise, such that the c_r coefficients describe the time course of perisaccadic modulation. Similarly, the $I_b(t)$ are indicator functions encoding the recording block. The population rate $p(t)$ was estimated by smoothing the binned spike counts of each multi-unit signal with a Gaussian kernel ($\sigma = 50$ ms), z-scoring the smoothed firing rate functions, and then averaging them across all multi-units. This additional covariate helped to prevent overfitting of eye position variability to shared noise in some recordings. Incorporating these additional components into the models can help to improve the performance of the eye tracking, although it is not necessary for achieving accurate eye position estimates. Note that multi-unit activity was only used for the purposes of eye tracking, not for the analysis of responses presented in the Results.

Learning of the model parameters and inference of the posterior distribution over eye position at each time was achieved by iterative optimization. Eye position was treated as a discrete variable, with the final inference performed over a grid of possible positions with spacing that was between 4 and 8 times finer than the size of the individual bar stimuli (i.e., 0.01–0.02° spacing). Although our inference procedure provided estimates of the full posterior distribution of eye positions at each time, point estimates (the posterior mean) of eye position were used for all analysis.

As with the subsequent characterization of each neuron’s stimulus-processing models (described above), the stimulus repeat trials were not used to estimate model parameters during the eye position inference procedure. Furthermore, a leave-one-out cross-validation procedure was used so that, when the activity of a given neuron was analyzed, eye positions inferred from the population excluding that unit were used (McFarland et al., 2014). When analyzing the covariance between pairs of neurons (see Fig. 9), the average of covariances estimated using each neuron’s leave-one-out eye position estimate was used. This greatly reduced the computational load compared with estimating separate eye position sequences for each possible pair of neurons. Any effects of overfitting of eye position variability were verified to be minimal (e.g., estimates of firing rate variance without using leave-one-out eye position signals were typically only slightly higher than estimates obtained with leave-one-out eye position signals; median difference = 1.6%, interquartile range 0–7.8%).

Estimating firing rate variance in the presence of FEMs. The core elements of the methods used for estimating stimulus-driven firing rate variance (with and without correcting for FEMs) are presented in the Results section (“Estimating stimulus-driven and stimulus-independent response variance in the presence of FEMs”). Here, we describe additional details, and provide a derivation for the FEM-corrected stimulus-driven variance $\text{Var}_{i,j}[r^i(t)]$ (Eq. 8).

First, to estimate the “PSTH variance” $\text{Var}_i[E_i[r^i(t)]]$ while avoiding biases due to sampling noise with finite trials (Sahani and Linden, 2003), the average product of spiking responses to the frozen noise stimulus across all pairs of unequal trials at each time was measured and then averaged across time (Eq. 6).

Our approach for estimating the FEM-corrected stimulus-driven variance $\text{Var}_{i,j}[r^i(t)]$ (Eq. 8) is analogous. Let $r(\mathbf{e}, t) = E_i[Y(t)|\mathbf{e}]$ be the expectation (average) of the spike count Y at time t given a particular sequence of eye positions, where the vector \mathbf{e} represents the eye position at time t , as well as L previous time samples: $\mathbf{e} = [e(t), e(t-1), \dots, e(t-L)]$. If it is assumed that the distribution of eye positions $p(\mathbf{e})$ is independent of the stimulus (and thus the same across time points), then the stimulus-driven response variance can be written as follows:

$$\text{Var}_{i,j}[r^i(t)] = E_i[E_{p(\mathbf{e})}[r^2(\mathbf{e}, t)]] - \bar{r}^2 \quad (\text{M7})$$

This quantity can be estimated by computing the product of spiking responses at time t across all pairs of unequal trials where the difference in eye positions Δe_{ij} between trials i and j was less than some (sufficiently small) threshold ε , and then averaging across time points:

$$\text{Var}_{i,j}[r^i(t)] = \langle \langle Y^i(t)Y^j(t) | \Delta e_{ij} < \varepsilon \rangle_{i \neq j} \rangle_t - \bar{Y}^2 \quad (\text{M8})$$

Estimation of Δe_{ij} is described further below. To show that Equation M8 provides an unbiased estimate of the stimulus-driven response variance, note that the spiking response at time t on the i^{th} trial can be written as follows: $Y^i(t) = r(\mathbf{e}^i, t) + \eta^i(t)$, where $\eta^i(t)$ is the “noise term.” Assuming that the noise terms on trials i and j are independent, Equation M8 then becomes:

$$\text{Var}_{i,j}[r^i(t)] = \langle \langle r(\mathbf{e}^i, t)r(\mathbf{e}^j, t) | \Delta e_{ij} < \varepsilon \rangle_{i \neq j} \rangle_t - \bar{Y}^2 \quad (\text{M9})$$

In the limit as $\varepsilon \rightarrow 0$, this then gives the average of $r^2(\mathbf{e}, t)$ over eye positions and time minus the square of the average rate, which is equal to the rate variance, as desired.

One subtle point about the above derivation is that, by restricting analysis to trial pairs where $\Delta e_{ij} \approx 0$, Equation M9 gives an estimate of $E[r^2(\mathbf{e}, t)]$ under the conditional eye position distribution $p(\mathbf{e} | \Delta e \approx 0)$ rather than $p(\mathbf{e})$. However, for a stimulus that is statistically invariant to spatial translations (such as used in our study), the expectation of $r^2(\mathbf{e}, t)$ with respect to any distribution over \mathbf{e} will be the same when averaged over a sufficiently large sample of the stimulus ensemble (i.e., averaged over a sufficient length of time). Therefore, in this case, sampling $r^2(\mathbf{e}, t)$ across eye positions at a given time point is equivalent, on average, to sampling $r^2(\mathbf{e}, t)$ across time points at a given eye position.

To measure the difference $\Delta e_{ij}(t)$ between the eye position trajectories at time t on trials i and j , the sequences of eye positions on each trial were compared over the range of times relevant for generating the neuron’s response to the stimulus in time bin t . Specifically, the sequence of eye positions within the current time bin plus the preceding 50 ms was used and then this window was shifted backward in time by 30 ms to account for V1 response delays. Therefore, eye position sequences spanning from the 80 ms before the start of the spike count window until 30 ms before the end of the window were used. The L2 norm of the vector difference of these eye position sequences (normalized by the number of time dimensions T in each vector $\mathbf{e}(t)$) was then used to compute $\Delta e_{ij}(t)$:

$$\Delta e_{ij}(t) = \sqrt{\frac{1}{T} \sum_{\tau=\tau_1}^{\tau_2} [e^i(t-\tau) - e^j(t-\tau)]^2} \quad (\text{M10})$$

Several methods were tested for estimating the limit of $\langle Y^i Y^j | \Delta e_{ij} \rangle_{i \neq j}$ as Δe_{ij} goes to zero, including fitting spline regression models. A simple procedure of computing $\langle Y^i Y^j \rangle$ across all trial pairs where Δe_{ij} was below a fixed threshold ε produced reliable results without the need for data-dependent parameter selection (such as a choice of spline basis). $\varepsilon = 0.01^\circ$ was used for all analysis, although results were largely insensitive to precise selection of this threshold. Although our models of neural stimulus processing (McFarland et al., 2013) and model-based eye-tracking method (McFarland et al., 2014) involve more sophisticated parameter estimation procedures, our primary analysis of eye-movement-related response variability primarily depends on selection of this single parameter.

To measure uncertainty in our estimates of the stimulus-driven rate variance (using Eq. 8), a bootstrap resampling procedure was used. Specifically, the set of trial pairs with $\Delta e_{ij} < \varepsilon$ was resampled (with replacement) using 5000 bootstrap samples. The bootstrap distribution was then used to estimate the CV of the stimulus-driven rate variance for each neuron to generate the data shown in Figure 8B. When computing rate variance as a function of Δe_{ij} (see Fig. 2E), 100 equipopulated bins of Δe_{ij} were used.

Model-based estimates of response variability. To validate our nonparametric estimates of α (see Fig. 4A), the stimulus-processing models estimated for each neuron were used to compute predicted firing rates on each trial using the same set of stimulus repeat trials used for nonparametric estimates. α was then computed simply as the fraction of the overall (model-predicted) firing rate variance arising from trial-to-trial variability, which was necessarily due to fluctuations in eye position.

When using the stimulus-processing models to estimate α under different conditions (e.g., as a function of time bin size; see Fig. 8), a slightly different approach was used that provided a better approximation of neural response statistics with respect to the entire stimulus ensemble

(rather than the particular stimulus sequences used for repeat trials). Specifically, the predicted firing rate of each model neuron in response to the full set of nonrepeating stimulus trials given the estimated eye positions was computed. The neuron's firing rate in response to the same sequence of stimuli was then computed, but the trial labels of the eye position data were randomly shuffled. An unbiased estimate of the model neurons' "PSTH variance" is then given by the covariance of the shuffled and unshuffled responses:

$$\text{Var}_i[E_i[r^i(t)]] = \text{Cov}_i[r(t), r_{\text{shuff}}(t)] \quad (\text{M11})$$

where $r_{\text{shuff}}(t)$ is the model-predicted firing rate after trial shuffling of eye position data. This method provided very similar results to computing α directly from model-predicted firing rates on the set of repeat trials, but has the advantage of not being limited to the repeat trials.

Analytical approximation for α . For the simulations shown in Figures 5 and 6, illustrating the analytical approximation of α , simple model neurons with Gabor spatial filters (preferred spatial frequency: $1/\lambda = 2$ cycles/° unless otherwise stated, and RF width $\sigma = 0.5\lambda$) were used. The model complex cell consisted of a quadrature pair of Gabor filters with outputs that were squared and summed together. For both the simple and complex cell models, the filter outputs were then passed through a rectifying spiking nonlinearity (Eq. M2) to generate simulated firing rates. For Figures 5, A–E, and 6, these model neurons were presented with a one-dimensional "ternary" white noise stimulus similar to the stimulus used in our experimental recordings. For Figure 5F, colored Gaussian noise was used. For simplicity, in these simulations, a Gaussian eye position distribution was used, setting the SD to the median value measured in our experiments (0.11°) unless otherwise stated.

The stimulus-processing models for each neuron were also used to estimate α using the analytical approximation of Equation 9. In this case, the same random bar stimulus used experimentally for each recorded neuron, as well as the measured distribution of eye positions, were used. These "analytical" estimates of α were then compared with "direct" estimates using Equation M11.

Estimation of cross-correlograms. To compute cross-correlograms between pairs of neurons, the array of spike counts across neurons and trials were time-shifted by an amount τ using padding with NaNs to handle the trial boundaries. The stimulus-driven component of the cross-correlogram between neurons m and n was estimated according to the following:

$$\text{Cov}_{i,t}[r_m^i(t), r_n^i(t + \tau)] \approx \langle \langle Y_m^i(t)Y_n^i(t + \tau) \mid \Delta e_{ij}(t) < \varepsilon \rangle_{i \neq j} \rangle_t - \bar{Y}_m \bar{Y}_n \quad (\text{M12})$$

When estimating this function at nonzero time lags τ , eye position trajectories were simply compared between pairs of trials at time t (ignoring the changes in eye position between t and $t + \tau$). Similar results were found, however, when comparing eye position trajectories over larger time windows.

Cross-correlograms were normalized based on the measured total response variances of each neuron:

$$\frac{\text{Cov}_{i,t}[r_m^i(t), r_n^i(t + \tau)]}{\sqrt{\text{Var}_{i,t}[Y_m^i(t)]\text{Var}_{i,t}[Y_n^i(t)]}} \quad (\text{M13})$$

Such normalization produces numbers bounded between -1 and 1 that are more comparable across pairs of neurons. This same normalization was applied when estimating both the stimulus-driven and stimulus-independent components of the cross-correlogram.

To obtain single estimates of the strengths of stimulus-driven and stimulus-independent correlations between each pair of neurons, the time lag in which the magnitude of the cross-correlogram was maximal (searching over the central ± 30 ms) was calculated and then the normalized values of each component of the cross-correlogram at that time lag were found. For analysis of the timescale of stimulus-independent correlations (see Fig. 9D), the set of neuron pairs with strong positive correlations was identified by dividing neuron pairs into two groups using a median split applied to the peak value of the normalized cross-correlograms. This set of neuron pairs were then further divided into two groups based on the magnitude of FEM-corrected stimulus-independent

correlations (again using a median split). The timescale of stimulus-independent correlations was measured using the full peak width at one-quarter of the maximal correlation value.

Statistical methods. Because the distribution of eye positions was heavy tailed, the SD of eye positions was measured using a robust estimator based on the median absolute deviation, given by the following:

$$\hat{\sigma} = 1.48 \text{ median}(|x_i - \text{median}_j(x_j)|) \quad (\text{M14})$$

where the scaling factor of 1.48 provides an estimate of variability that converges to the SD for a normal distribution.

Correlations and associated p -values were computed using the non-parametric Spearman's rank correlation (using Matlab's *corr* function).

Additional analysis details. Unless otherwise specified, spikes were binned at 10 ms resolution. Blinks were identified by finding instances in which any component of the instantaneous eye velocity exceeded a threshold (three times the median speed) for at least 100 ms and where the overall change in eye position was small. Data occurring from 30 ms after the onset of a blink to 100 ms after the termination of the blink were excluded from all analysis of trial-to-trial variability (though such exclusion had little impact on the results). Exclusion of data after detected microsaccades also had little impact on the results (and this was not done for the analysis presented).

During some experiments, problems with the stimulus display caused individual stimulus patterns to occasionally be presented for multiple video frames. On stimulus repeat trials, such repeated frames caused a misalignment of the stimulus sequence for the remainder of the trial. In these cases, the data were time shifted to realign neural responses to the displayed stimuli across repeat trials and 100 ms of data following any repeated stimulus frame were excluded. These video frame repeats were relatively rare (occurring in, on average, 2% of repeat trials) and similar results were obtained when excluding trials with any repeat frames from analysis.

Results

Effects of fixational eye movements on estimates of trial-to-trial variability

The standard approach for segregating the stimulus-driven and stimulus-independent components of a sensory neuron's activity is to measure the neuron's response to repeated presentations of a set of stimuli $\{s_j\}$. If we define $Y^i(s_j)$ as the spike count over some time interval in response to the j^{th} stimulus on the i^{th} trial, then the stimulus-driven component is typically measured by the across-trial average response, or firing rate: $r(s_j) = E_i[Y^i(s_j)]$. Conversely, the across-trial variability, measured by $\text{Var}_i[Y^i(s_j)]$, is assumed to represent the stimulus-independent component of the neuron's activity. The "law of total variance" provides a general means of decomposing a neuron's total response variance into a sum of this stimulus-driven component (the variance of the firing rate across stimuli) and the stimulus-independent component (the average variance of the response to a fixed stimulus):

$$\text{Var}_{i,j}[Y^i(s_j)] = \text{Var}_j[r(s_j)] + E_j[\text{Var}_i[Y^i(s_j)]] \quad (1)$$

where $\text{Var}_{i,j}[\cdot]$ indicates the total variance computed across both trials and stimuli. In the context of a time-varying stimulus presented over multiple trials, we can replace the stimulus index j with an index running over time:

$$\text{Var}_{i,t}[Y^i(t)] = \text{Var}_t[r(t)] + E_t[\text{Var}_i[Y^i(t)]] \quad (2)$$

where the time-varying firing rate $r(t) = E_i[Y^i(t)]$ is estimated by the PSTH.

In the presence of fixational eye movements (FEMs), however, Equation 2 no longer provides a segregation of stimulus-driven and stimulus-independent response variance because the stimulus on the retina will itself vary from trial to trial (Fig. 1A). In this case, the variance of the PSTH will underestimate the neuron's

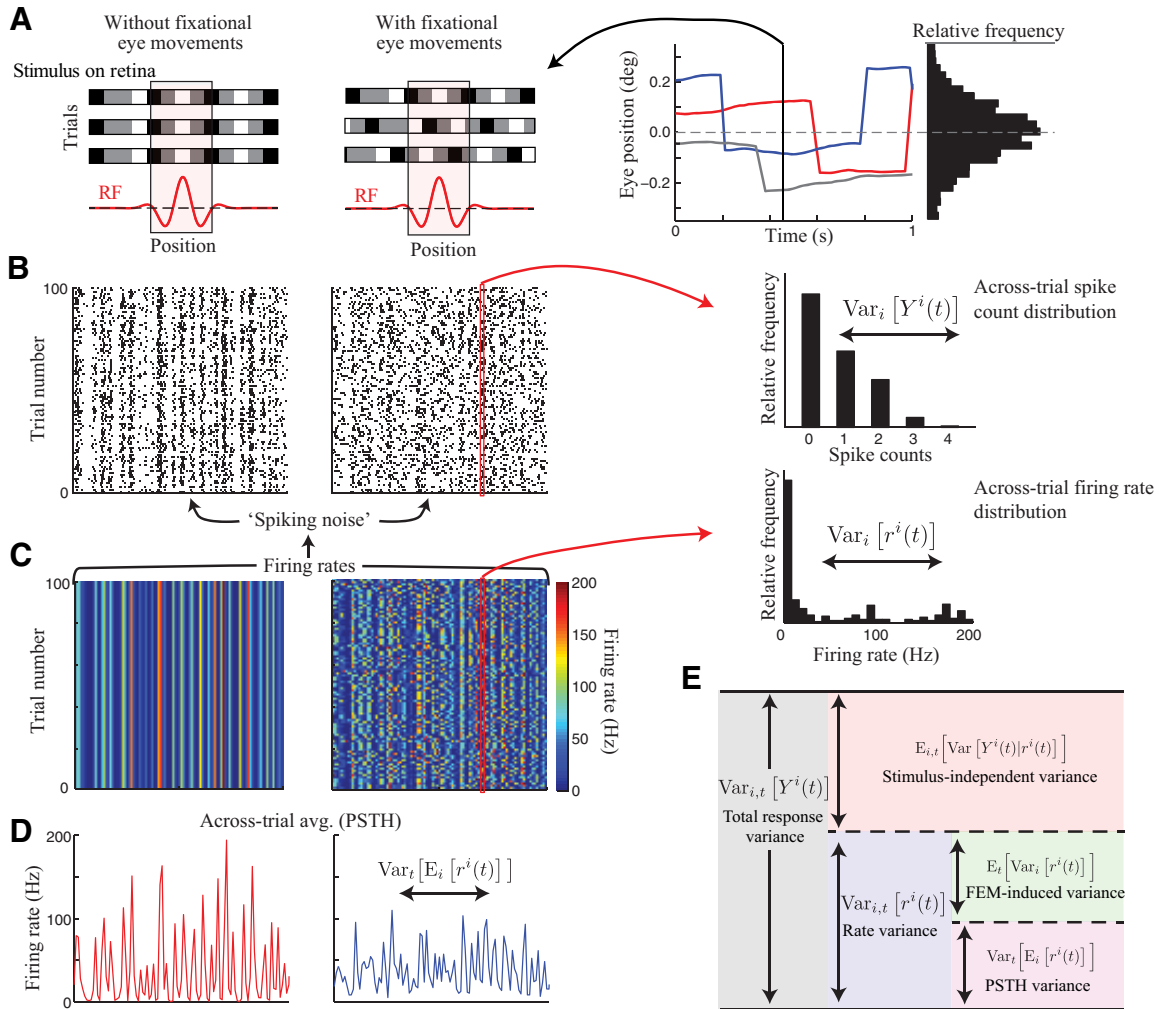


Figure 1. Decomposing sources of neural response variability. **A**, Schematic illustration of a single frame of the visual stimulus, presented across multiple trials (different rows). The spatial profile of an example neuron's RF (red trace) is shown below. Without FEMs (left), the stimulus in the RF is identical across trials, whereas FEMs serve to "jitter" the position of the visual stimulus relative to the RF across trials. Example eye position trajectories across three trials are shown at right. **B**, Simulated spike rasters of a model neuron in response to repeated presentations of a "frozen noise" stimulus showing that FEMs produce increased trial-to-trial variability (right) compared with the neuron's response in the absence of eye movements (left). Inset at right shows the across-trial distribution of spike counts over the time window indicated by the red rectangle. **C**, Stimulus-driven firing rate on each trial from which the spiking data in **B** were generated (as a rate-modulated Poisson process). Note that this is by definition constant across trials in which FEMs are absent (left), although it exhibits substantial across-trial variability in the presence of eye movements (right). Inset at right shows the across-trial distribution of firing rates at the example time point (red rectangle). **D**, The across-trial average response ("PSTH") is less variable across different time points of the stimulus in the presence of FEMs (right) because it averages together responses to different retinal stimuli in the neuron's RF. **E**, Schematic illustrating the decomposition of total spike count variance into its constituent components: stimulus-driven (blue region) and stimulus-independent (red) variance. The stimulus-driven variance is then further decomposed into FEM-induced across-trial variability (green) and across-time variability in the PSTH (magenta).

true stimulus-driven response variance because it does not take into account FEM-induced trial-to-trial variability in the firing rate $r^i(t)$ (Fig. 1B–D). We can quantify this bias by again using the law of total variance to write the total firing rate variance as a sum of FEM-induced across-trial variability and variability in the PSTH over time:

$$\text{Var}_{i,t}[r^i(t)] = \text{Var}_i[\text{E}_i[r^i(t)]] + \text{E}_i[\text{Var}_i[r^i(t)]] \quad (3)$$

Given that $\text{E}_i[r^i(t)] = \text{E}_i[Y^i(t)]$, the measured PSTH variance is only the first term on the right side of Equation 3, thus underestimating the true stimulus-driven response variance by an amount equal to $\text{E}_i[\text{Var}_i[r^i(t)]]$. Furthermore, this FEM-induced across-trial firing rate variability will increase the measured across-trial response variability, inflating estimates of the stimulus-independent response variance. In fact, because these two sources of variance must sum to equal the total response

variance (Eq. 2), FEMs will produce equal and opposite biases in estimates of each (Fig. 1E).

To measure the magnitude of these effects in V1 recordings, we used the fraction of the true stimulus-driven response variance that is captured by the measured PSTH variance as follows:

$$\alpha = \frac{\text{Var}_t[\text{E}_i[r^i(t)]]}{\text{Var}_{i,t}[r^i(t)]} \quad (4)$$

Therefore, α is a measure of a neuron's sensitivity to FEMs that ranges between 0 and 1, with a value of 1 indicating that FEMs have no effect on measures of the stimulus-driven variance and a value of 0 indicating that all of the neuron's stimulus-driven response appears as trial-to-trial variability (with the measured PSTH showing no stimulus-driven response variability as a result).

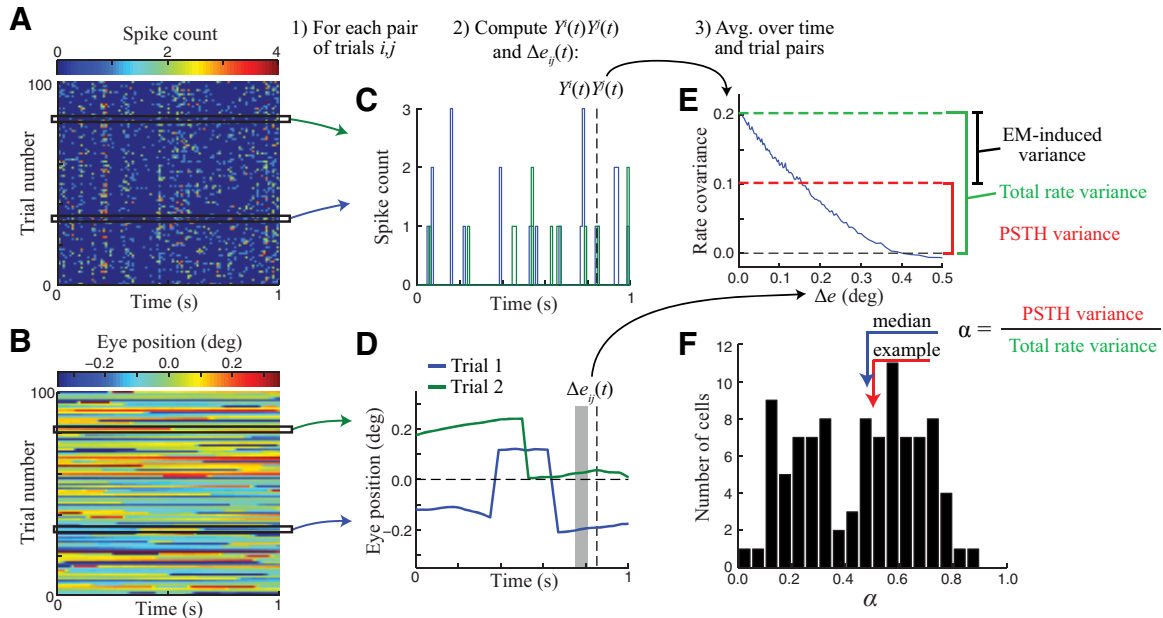


Figure 2. Nonparametric estimation of eye-movement induced variability. **A**, Spike counts for an example neuron (in 10 ms time bins) across trials in response to repeated presentations of a frozen noise stimulus. **B**, Corresponding eye position data for each trial. **C**, Spiking responses are compared at each time point across all pairs of trials (shown for the example trial pair indicated by the black horizontal rectangles in **A** and **B**). **D**, Eye position trajectories for the same two example trials. The sequence of eye positions relevant for determining the neuron's response at a given time (vertical dashed line) is highlighted by the gray region (see Materials and Methods). **E**, Average product of spiking responses is computed across time points and trial pairs as a function of the similarity of eye trajectories Δe . In the limit as Δe goes to 0, this function provides an estimate of the firing rate variance (Eq. 8; green dashed line). The PSTH-based estimate of the neuron's stimulus-driven response variance (Eq. 6) is shown as the red dashed line for comparison. For this neuron, the fraction of firing rate variance captured by the PSTH (α) was 0.50. **F**, Across the population of neurons ($n = 97$), α values varied widely (median 0.48, interquartile range: 0.27–0.64).

Estimating stimulus-driven and stimulus-independent response variance in the presence of FEMs

Given that we only ever observe a neuron's spiking activity and never have direct access to its underlying firing rate on single trials $r^i(t)$, the question remains how we can estimate the stimulus-driven response variance given measurements of the spike count $Y^i(t)$ and the sequence of eye positions $e^i(t)$ on each trial. In principle, one could construct estimates of the neuron's average response as a function of both time and the animal's eye position $r(t, \mathbf{e})$. Given this “joint” response function, one could then compute the stimulus-driven variance directly:

$$\text{Var}_{i,t}[r^i(t)] = \text{Var}_{p(\mathbf{e})}[r(t, \mathbf{e})] \quad (5)$$

where $p(\mathbf{e})$ is the distribution of eye positions over the experiment. However, while $r(t, \mathbf{e})$ can in principle be estimated using nonparametric regression methods (Rad and Paninski, 2010; Park et al., 2011), such methods generally require a large number of trials to produce reliable results, depending on how sensitive $r(t, \mathbf{e})$ is to eye position. Furthermore, any variance in the estimation of $r(t, \mathbf{e})$ will produce systematic biases in estimates of the firing rate variance (Eq. 5), making this approach highly sensitive to the binning/regularization procedures used.

We thus developed a procedure for directly estimating a neuron's stimulus-driven and stimulus-independent response variance in the presence of FEMs that circumvents the above limitations. Our approach is analogous to estimating the PSTH variance using the covariance of spiking responses on unequal trials (Perkel et al., 1967; Brody, 1999):

$$\text{Var}_i[E_i[r^i(t)]] = \langle\langle Y^i(t)Y^j(t) \rangle\rangle_{i \neq j} - \bar{Y}^2 \quad (6)$$

where $\langle\cdot\rangle_t$ indicates the empirical average over time and \bar{Y} is the overall average response. We extend this method to compute the

true stimulus-driven response variance in the presence of FEMs by using only data from when the sequences of eye positions on trials i and j at time t were sufficiently similar (Fig. 2A–E). The key advantage of this approach is that it allows us to estimate the stimulus-driven variance directly from the data without having to construct data-intensive (and potentially biased) estimates of the firing rate function $r(t, \mathbf{e})$ explicitly.

Specifically, we first define a measure of the similarity of eye position trajectories on a pair of trials at time t as $\Delta e_{ij}(t) = |\mathbf{e}^i(t) - \mathbf{e}^j(t)|$ (i.e., the magnitude of the difference in eye position trajectories over a chosen time window; Fig. 2D; see Materials and Methods). We can then obtain an unbiased estimate of $\text{Var}_{i,t}[r^i(t)]$ (derived in the Materials and Methods), given by:

$$\text{Var}_{i,t}[r^i(t)] = \langle\langle Y^i(t)Y^j(t) \mid \Delta e_{ij} = 0 \rangle\rangle_{i \neq j} - \bar{Y}^2 \quad (7)$$

Given that $e(t)$ is a continuous variable, we will not observe precisely the same sequence of eye positions on different trials. However, assuming that $r(t, \mathbf{e})$ is a smoothly varying function of e , we can estimate the firing rate variance by taking the limit of the conditional average in Equation 7 as Δe approaches zero (Fig. 2E). We approximate this limit by using a sufficiently small threshold ε on the similarity of eye position trajectories between pairs of trials (see Materials and Methods), giving:

$$\text{Var}_{i,t}[r^i(t)] = \langle\langle Y^i(t)Y^j(t) \mid \Delta e_{ij} < \varepsilon \rangle\rangle_{i \neq j} - \bar{Y}^2 \quad (8)$$

We can then directly estimate the fraction of the stimulus-driven variance captured by the PSTH variance (α) as the ratio of Equations 6 and 8.

Critically, this method does not make any assumptions about the form of the noise distribution (e.g., Poisson) nor its independence across time bins (Amarasingham et al., 2015); it only

requires that the noise is independent on different trials. Furthermore, it does not assume a particular functional form of the neurons' response to different stimuli, only that the firing rate is a smooth function of eye position [i.e., slowly varying relative to the sampling of the eye position distribution $p(e)$].

Large fraction of V1 response variability driven by FEMs

To evaluate the effects of FEMs on estimates of V1 response variability, we recorded the activity of V1 neurons in two monkeys using linear multielectrode arrays. The animals performed a simple fixation task while they were presented with a "one-dimensional" dynamic noise stimulus, with stimulus frames drawn independently at 100 Hz. Each stimulus frame consisted of a random pattern of black, white, and gray bars, with the bar orientation selected to match the preferred orientations of the recorded neurons (see Materials and Methods; McFarland et al., 2014; McFarland et al., 2015). Trial-to-trial variability was then evaluated using repeated presentations of 4-s-long "frozen noise" sequences (randomly selected from the same stimulus ensemble).

We measured the animals' eye position directly using standard methods (scleral search coils); however, the accuracy of such measurements is not sufficient for tracking FEMs (Read and Cumming, 2003; Kimmel et al., 2012; McFarland et al., 2014). Therefore, we used a recently developed model-based eye-tracking method that allowed us to infer the animals' eye position offline with approximately arc-minute accuracy using the activity of recorded V1 neurons (McFarland et al., 2014). This method uses models of each neuron's response to the stimulus to infer the most likely eye position as a function of time. To avoid any biases due to overfitting of eye position variability to the observed neural activity, we always used eye position signals inferred from the population excluding the particular neuron under study (see Materials and Methods). Note that, because the stimulus only varied along one spatial dimension, neural responses were only sensitive to one dimension of eye position, simplifying both the eye position inference and subsequent analysis.

We thus used this precise eye-tracking method to compare eye trajectories between different trials, allowing for application of the nonparametric approach described above. We found that estimated values of α varied widely across the population of recorded neurons ($n = 97$), with a median value of 0.48 (Fig. 2F). This means that standard methods for quantifying stimulus-driven and stimulus-independent response variance (i.e., computing the across-trial mean and variance) would typically miss about half of the true stimulus-driven response variance, mistakenly attributing that FEM-induced variability to stimulus-independent "noise." Furthermore, we found that, for many neurons the effects of FEMs were even more dramatic, with the across-trial average response (PSTH) almost completely failing to capture the neurons' true stimulus-driven response (e.g., 24% of neurons had α values < 0.25).

Dependence of α on V1 neuron stimulus tuning

Intuitively, we would expect the biases created by FEMs to depend largely on how sensitive a given neuron is to fine spatial features of the stimulus. For example, neurons with smaller RFs and/or those that are more sensitive to the spatial phase of the stimulus (i.e., more "simple-cell like") should have smaller values of α . We tested this intuition directly by estimating detailed stimulus-processing models for each neuron (Fig. 3A; Butts et al., 2011; Park and Pillow, 2011; McFarland et al., 2013), using these models to derive a number of measures of stimulus tuning, including RF width and spatial phase sensitivity (see Materials and

Methods). As expected, neurons with smaller RFs (Fig. 3B) and those that were more sensitive to spatial phase (Fig. 3C) tended to have smaller values of α . In addition, we found that both α ($p = 0.014$) and spatial phase ($p = 7.3 \times 10^{-3}$) were bimodally distributed (using Hartigan's dip statistic with the Monte Carlo method of Mechler and Ringach, 2002; $N = 10^5$ samples), potentially reflecting a division between "simple" and "complex" cells.

One might expect that such small values of α would be limited to more central RFs, where V1 RFs tend to be smaller. However, there was substantial variability across neurons at a given eccentricity in both the RF size (Fig. 3D), as well as α (Fig. 3E). This large variability in RF width at a given eccentricity is consistent with previous work using similar methods to map V1 RFs (Pack et al., 2006). Therefore, the effects of FEMs on measures of response variability were not isolated to neurons representing the center of gaze (i.e., in and immediately surrounding the fovea), but were often quite large at parafoveal eccentricities, where V1 neurons are most often recorded. Sensitivity to FEMs remained substantial even for neurons with RFs at an eccentricity of about 10° (recorded from the calcarine sulcus), with 3/12 such neurons having α values < 0.25 (Fig. 3E).

These stimulus-processing models could also be used to produce a second, independent measurement of α . Specifically, we used the models to predict the neurons' firing rate on each trial $r^i(t)$ given the measured eye positions (and resulting stimulus on the retina), and then computed α directly from the model-predicted $r^i(t)$. We found that model-based estimates were in very good agreement with our nonparametric estimates (Fig. 4A; Spearman's $\rho = 0.92$; $p = 7.5 \times 10^{-40}$). This model-based approach to estimating α also has the advantage that it does not require stimulus-repeat trials and can also be applied with simulated data to explore the effects of FEMs under a variety of conditions.

The close agreement between model-based and nonparametric estimates of α is perhaps surprising given that even the most sophisticated stimulus-processing models are known to provide substantial underestimates of the stimulus-driven response variance of V1 neurons (David et al., 2004). Indeed, our models generally captured only about half of the firing rate variance identified using the nonparametric estimates (median = 55%; Fig. 4B). However, because α is based on a ratio of firing rate variances, it is much less sensitive to the overall proportion of rate variance captured by the models. As discussed further below, accurate estimates of α in fact depend on how well a given model describes a neuron's relative sensitivity to spatial translations of the stimulus.

Analytical approximation for FEM-induced biases

To better understand the factors determining FEM-induced biases, we derived an analytic approximation for α , the proportion of stimulus-driven variance captured by the measured PSTH variance. This approximation is based on taking a neuron's stimulus-driven firing rate, typically measured as a function of time, and representing it as a function of the animal's gaze position x , allowing us to express α in terms of two relatively simple components: the neuron's sensitivity to spatial translations of the stimulus (such as those induced by FEMs) and the distribution of eye positions $p(e)$ sampled over an experiment.

We illustrate this method in Figure 5, A–C, using idealized model simple and complex cells presented with a random bar stimulus (see Materials and Methods). For a time-varying stimulus that is statistically invariant to spatial translations (such as noise stimuli and drifting gratings), we can represent the entire

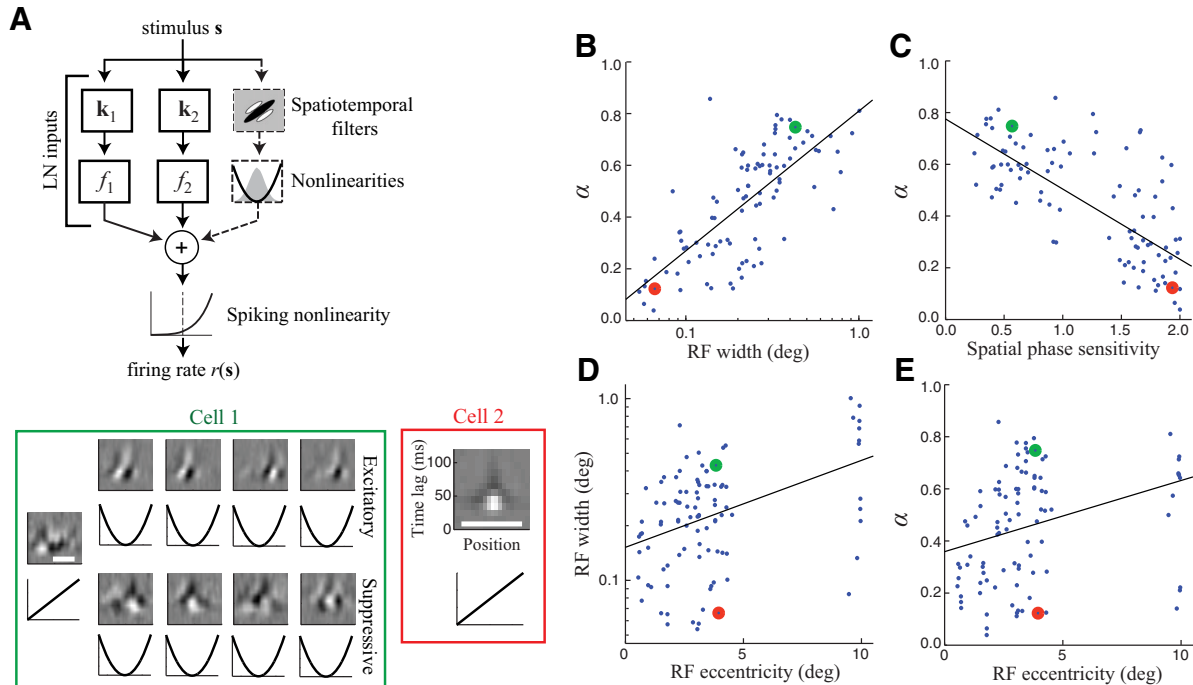


Figure 3. Dependence of α on RF properties. **A**, Top, Models of V1 neuron stimulus processing consisted of a set of “linear-nonlinear subunits,” each containing a spatiotemporal stimulus filter and a static nonlinearity. A final spiking nonlinearity was then used to map the summed output of these subunits to a predicted firing rate. Bottom, Models estimated for two example neurons recorded simultaneously (at the same eccentricity). Spatiotemporal filters are shown above (white scale bars are 0.4°), with corresponding nonlinearities of each subunit below. The model of cell 2 (red, right) consists of a single linear filter (i.e., simple-cell-like) with a small RF, making it very sensitive to FEM-induced variability ($\alpha = 0.12$). In contrast, the complex-cell-like model for cell 1 (green, left) consists of a number of excitatory and suppressive squared subunits, each with a larger RF, making it relatively insensitive to FEMs ($\alpha = 0.75$). **B**, Across the population of neurons ($n = 97$), α was strongly correlated with RF width (Spearman’s $\rho = 0.75, p \sim 0$). Black line shows a linear fit. **C**, α was also strongly correlated with a neuron’s spatial phase sensitivity ($\rho = -0.69, p \sim 0$). **D**, V1 neurons exhibited a wide range of RF properties at a given eccentricity. For example, RF width was relatively weakly related to RF eccentricity ($\rho = 0.31, p = 1.9 \times 10^{-3}$). **E**, Although α was systematically larger for neurons with RFs at larger eccentricities ($\rho = 0.37, p = 2.3 \times 10^{-4}$), there was a wide range of α values for neurons recorded at a given eccentricity, including those well outside of the fovea.

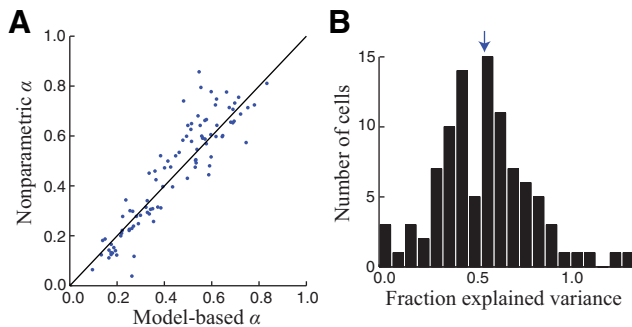


Figure 4. Model-based validation of α estimates. **A**, We used the estimated stimulus-processing model for each neuron to validate our nonparametric estimates of α . Specifically, we computed α directly from the model-predicted firing rates on each trial. These model-based estimates of α were in close agreement with nonparametric estimates across the population of neurons (Spearman’s $\rho = 0.92; p = 7.5 \times 10^{-40}$). **B**, This close agreement was despite the fact that the models typically captured about half of the stimulus-driven firing rate variance (median: 0.54, interquartile range: 0.37–0.68). Arrow indicates the median.

stimulus ensemble as a single function of space (with infinite spatial extent) such that the stimulus displayed at any time is located at some spatial position x (Fig. 5A). We can then compute a neuron’s stimulus-driven response variance as the variance of this firing rate function $r(x)$ over all spatial positions x (Fig. 5B). Similarly, at each position x , we can compute the neuron’s across-trial average response, measured under the distribution of eye positions $p(e)$ centered at x , as the following convolution: $E_{p(e)}[r(x)] = \int r(x)p(e - x)de$ (Fig. 5C),

assuming the distribution $p(e)$ is independent of the stimulus. Therefore, if we define \bar{r} as the mean-subtracted firing rate, we can compute α as follows:

$$\alpha = \frac{\text{Var}_i[E_i[r^i(t)]]}{\text{Var}_i[r^i(t)]} = \frac{\int [(p * \bar{r})(x)]^2 dx}{\int \bar{r}^2(x) dx} = \frac{\int |R(k)P(k)|^2 dk}{\int |R(k)|^2 dk} \tag{9}$$

where, in the last step, we have used Parseval’s theorem to represent these quantities in Fourier space so that the convolution of $r(x)$ and $p(e)$ can be written as a product of their spatial frequency spectra. Therefore, α is approximately determined by a weighted average of $R(k)$, the spatial frequency spectrum of $r(x)$, with the spectrum of the eye position distribution serving as the weighting function (Fig. 5D). Note that $R(k)$ is distinct from (although related to) a neuron’s spatial frequency tuning and characterizes its sensitivity to spatial translations of the stimulus.

For the model simple cell, the firing rate changes rapidly as a function of gaze position (Fig. 5B), so its spectrum $R(k)$ is largely concentrated at high spatial frequencies (Fig. 5D). The effect of FEMs [as determined by the spectrum of the eye position distribution $P(k)$] is to attenuate these high spatial frequencies in the PSTH, resulting in a high sensitivity to FEMs and a correspondingly small value of α (0.23). By comparison, the complex cell is much less sensitive to stimulus translations [i.e., the spectrum $R(k)$ is concentrated at lower spatial frequencies], resulting in a much higher value of α (0.85) despite the cell having the same preferred spatial frequency.

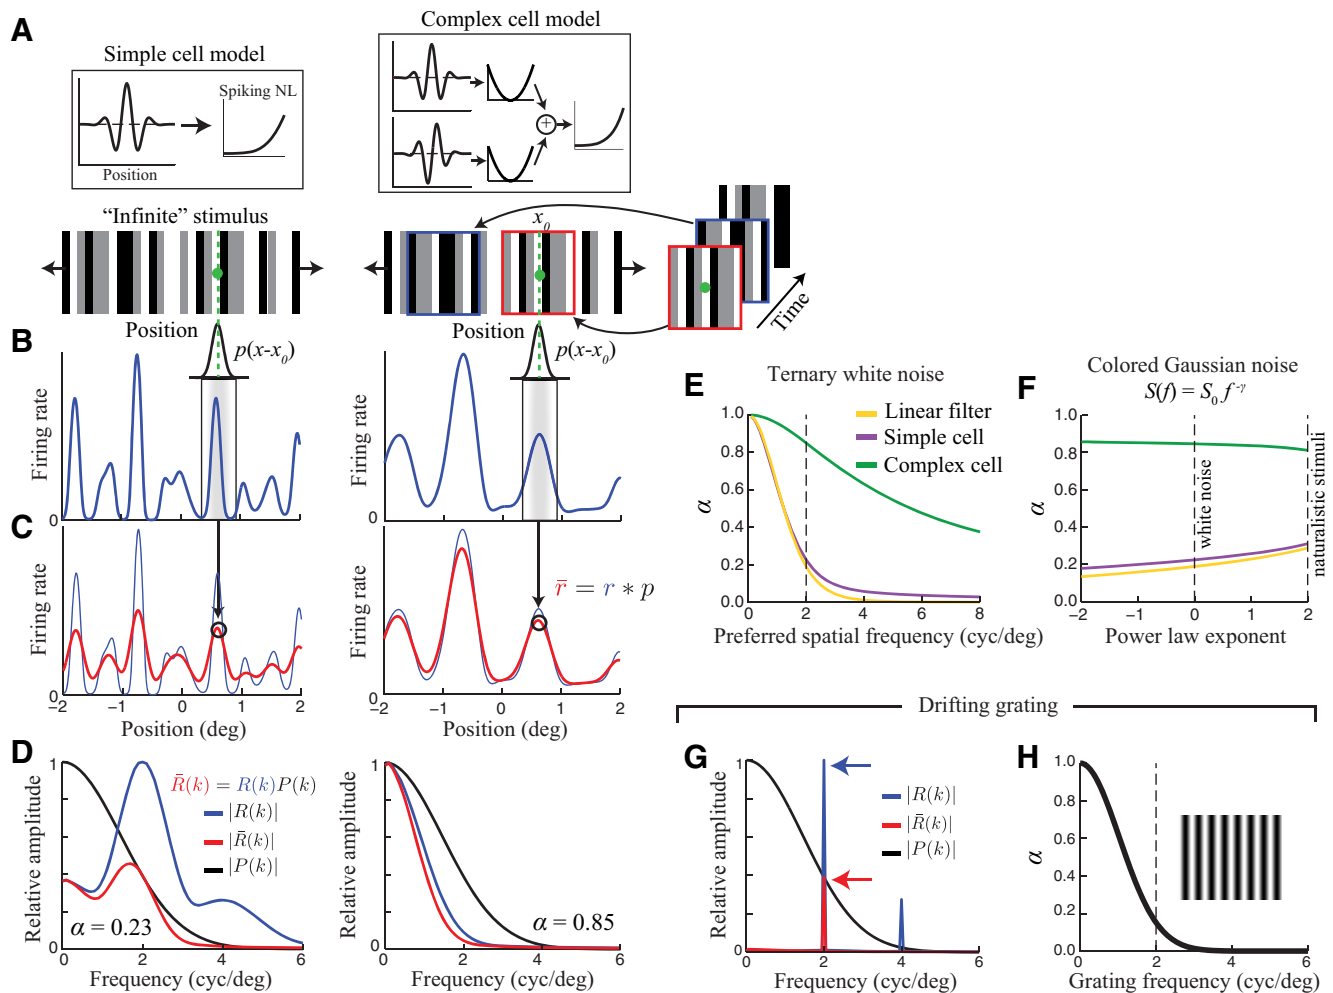


Figure 5. Analytical approximation of eye-movement-induced variability. **A**, To illustrate the analytical approximation of α , we consider model simple (left) and complex (right) cells, presented with a random bar stimulus (bottom). To simplify the calculation, we imagine the sequence of presented images (right) being compiled into a single visual stimulus with infinite spatial extent. **B**, Model neurons' firing rates can then be computed as a function of the animal's gaze position along the stimulus [$r(x)$; blue]. **C**, When the fixation target is located at a point x_0 , the across-trial average responses ("PSTHs"; $\bar{r}(x_0)$) measured under a distribution $p(e)$ of eye positions (black curves at top in **B**) are given by a local average of $r(x)$ weighted by the distribution $p(x - x_0)$. We can thus write $\bar{r}(x)$ as a convolution of $r(x)$ and $p(e)$. **D**, The spatial frequency spectrum of $\bar{r}(x)$ [$\bar{R}(k)$; red] can be computed by multiplying the spectrum of $r(x)$ [$R(k)$; blue] by the spectrum of $p(x)$ [$P(k)$; black]. α is then given by the ratio of the integrals of $\bar{R}^2(k)$ and $R^2(k)$ (Eq. 9). For the simple cell (left), the response spectrum is sharply peaked at 2 cycles/deg (the center frequency of the neuron's Gabor filter), where eye movements strongly attenuate the measured PSTH [i.e., $P(k)$ is small]. In contrast, for the complex cell, $R(k)$ is largely concentrated at low frequencies where $\bar{R}(k)$ is only moderately attenuated by $P(k)$. **E**, α is shown as a function of preferred spatial frequency for the simple (magenta) and complex (green) cell models in **A–D**, along with a corresponding linear Gabor filter (orange). The vertical dashed line indicates the preferred spatial frequency used in **A–D**. **F**, α is largely insensitive to the spatial frequency content of the stimulus. For the same three model neurons in **E** (with preferred SF = 2 cycles/deg), α was only slightly different for a stimulus with $1/f^2$ power spectrum (such as with natural images) compared with white noise. **G**, Spatial frequency spectra of $r(x)$ and $\bar{r}(x)$ for a drifting grating stimulus, showing strong attenuation of the PSTH at the spatial frequency of the grating (2 cycles/deg). **H**, For grating stimuli, α decreases sharply with the grating spatial frequency, reaching values near 0 for frequencies of just 2 cycles/deg (vertical dashed line).

Although these are highly simplified models, they serve to illustrate the range of FEM sensitivity likely to be exhibited by real V1 neurons having RFs of a given size, with the idealized simple/complex cells being most/least sensitive to FEMs, and most V1 neurons falling somewhere in between these extremes. Using this approach to analyze the dependence of α on RF properties, we find that α values can be very large even for neurons with relatively low preferred spatial frequencies (Fig. 5E), highlighting the potential impact of FEMs even for neurons with large RFs.

We also examined the dependence of α on the spatial frequency content of the stimulus using temporally uncorrelated stimuli with power-law spatial frequency spectra of the form $S(f) = Cf^{-\gamma}$, where the exponent γ controls the rate of increase/decrease in power with frequency. This analysis shows that a neuron's sensitivity to FEMs was similar over a range of relevant

spatial frequency spectra, including white noise (such as used in our experiments) and naturalistic spectra ($\gamma \approx 2$; Fig. 5F), being largely determined instead by the neuron's RF properties.

For the special case of a spatially periodic stimulus (such as a drifting grating), $r(x)$ is a periodic function of x and $R(k)$ can be approximated by a single delta function centered at the spatial frequency of the grating k_0 , ignoring contributions from higher spatial frequency harmonics which will be strongly attenuated by the spectrum of the eye position distribution $P(k)$ (Fig. 5G). The equation for α is then determined by a single value of the eye position spectrum: $\alpha \approx |P(k_0)|^2$. For realistic eye position distributions (with SD $\approx 0.1^\circ$), α is < 0.2 even at grating frequencies as low as 2 cycles/deg (Fig. 5H) because FEMs will alter the spatial phase of the retinal stimulus across trials. Note that for a grating stimulus, this approximation for α is independent of a neuron's RF properties including its spatial phase sensitivity, although in




Figure 6. Dependence of α on fixation accuracy. To illustrate the dependence of α on the animal’s fixation accuracy, we used the analytical framework developed above to compute α as a function of the SD of the eye position distribution (assuming a Gaussian distribution) under several different conditions using the same model neurons as in Figure 5. **A**, For drifting grating stimuli, α decayed rapidly with eye position SD for a range of different grating spatial frequencies (represented by different line thicknesses). Note that this result applies independent of the spatial phase sensitivity of the neuron. The vertical dashed line and the shaded region show the median and range respectively of the eye position SDs measured in our experiments. **B**, Similar to **A**, but for a spatial white noise stimulus. α is shown as a function of eye position SD for model simple (magenta) and complex (green) cells with several different preferred spatial frequencies (different line thicknesses).

the case of a perfect complex cell [where $r(x)$ is constant], the response variance itself will approach zero.

In the above derivation, we have ignored the temporal dynamics of eye movements, relying only on the marginal distribution of eye positions $p(e)$ to estimate FEM-induced biases. This is equivalent to assuming that the animal’s eye position is approximately constant over the timescale relevant for determining a neuron’s response (e.g., the ~ 50 ms integration time). For rapidly flashed, temporally uncorrelated stimuli (such as used in our experiments), the temporal dynamics of FEMs will inevitably have a negligible effect on measures of trial-to-trial variability because they do not significantly alter the spatiotemporal statistics of the retinal stimulus (McFarland et al., 2015). To illustrate this point, we compared estimates of α using Equation 9 with those based on direct calculation (see Materials and Methods), using the estimated stimulus-processing models applied to the same random bar stimulus and eye position trajectories as in our recordings. These two estimates of α were in close agreement (median absolute deviation = 3%), illustrating the accuracy of this approximation for rapidly flashed stimuli. When using slowly varying or static stimuli, the temporal dynamics of eye movements will introduce an additional source of trial-to-trial variability (Gur et al., 1997; Gur and Snodderly, 2006) that will only increase FEM-related biases.

Using this analytical framework, we can also evaluate how FEM-induced biases depend on the animal’s fixation accuracy via the distribution of eye positions $p(e)$. For both grating stimuli (Fig. 6A) and

white noise stimuli (Fig. 6B), α decreases sharply as a function of the SD of eye positions. Indeed, increasing the scale of the eye position distribution is equivalent to proportionately decreasing the spatial scale of the stimulus and RFs, such that the dependence of α on eye position SD is essentially the same as its dependence on a neuron’s RF size/preferred SF. These results highlight the fact that any experimental or “internal” factors that influence an animal’s fixation accuracy could produce significant changes in measures of stimulus-driven and stimulus-independent responses by altering the magnitude of FEM-induced biases.

Effects of FEMs on Fano factor (FF) estimates

Neural response reliability is most commonly measured using the FF, which is defined as the variance-to-mean ratio of the spike count Y in response to a stimulus \mathbf{s} as follows:

$$FF(\mathbf{s}) = \frac{\text{Var}_i[Y^i(\mathbf{s})]}{r(\mathbf{s})} \quad (10)$$

As defined above, this quantity depends on the stimulus \mathbf{s} , but a single FF is usually reported for a given neuron, representing the average value of this ratio over an ensemble of stimuli (Tolhurst et al., 1983; Carandini, 2004; Churchland et al., 2010), which we will refer to as the “average FF.”

As described above, using the variance of Y across repeated presentations of a stimulus to estimate the stimulus-independent response variance (the numerator in Eq. 10) will be systematically biased upward by the presence of FEM-induced trial-to-trial firing rate variability. Our nonparametric approach, however, can be used to obtain an unbiased estimate of the average stimulus-independent variance across an ensemble of stimuli by simply subtracting the estimated stimulus-driven variance (Eq. 8) from the measured total response variance (i.e., using the law of total variance):

$$E_j[\text{Var}_i[Y^i(\mathbf{s}_j)]] = \text{Var}_{i,j}[Y^i(\mathbf{s}_j)] - \text{Var}_j[r(\mathbf{s}_j)] \quad (11)$$

The average FF can then be approximated by the ratio of the average stimulus-independent variance to the average firing rate:

$$E_j[FF(\mathbf{s}_j)] = E_j\left[\frac{\text{Var}_i[Y^i(\mathbf{s}_j)]}{r(\mathbf{s}_j)}\right] \approx \frac{E_j[\text{Var}_i[Y^i(\mathbf{s}_j)]]}{E_j[r(\mathbf{s}_j)]} \quad (12)$$

This approximation depends on the stimulus-independent variance being proportional to the mean rate, which will be true for any rate-modulated renewal process (Nawrot et al., 2008; Churchland et al., 2011) and approximately holds in many cortical recordings (Tolhurst et al., 1983; Britten et al., 1993; McAdams and Maunsell, 1999).

Using our previously derived measure of the stimulus-driven rate variance (Eq. 8) to estimate the stimulus-independent variance for each neuron (Eq. 11), we calculated FEM-corrected FFs and compared them with uncorrected estimates in which we simply took the PSTH variance to represent the stimulus-driven component. We first analyzed these effects using the same time bin size (10 ms) as above, but consider the case of the larger time bins typically used to calculate FFs in the next section (below). For 10 ms time bins, we found that the effects of FEMs on FF estimates were typically fairly modest (median overestimate of 1.15-fold when ignoring FEMs). However, this bias varied widely across neurons and, in some cases, uncorrected estimates were >2 -fold larger than their FEM-corrected values (Fig. 7). Correcting for FEMs also altered the statistical comparison of our population data to the null hypothesis of Poisson variability ($FF = 1$).

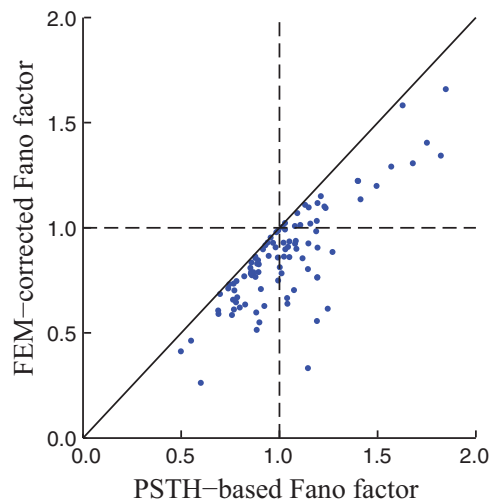


Figure 7. FEM-induced overestimates of FFs. Comparison of FF estimates using FEM-corrected methods (vertical axis) and standard methods that treat all across-trial variability as noise (i.e., “PSTH-based” estimates; horizontal axis). The effects of FEMs on FF estimates were typically fairly modest (median overestimate of 1.15-fold when ignoring FEMs), but were in some cases fairly dramatic. Overall, when ignoring FEMs, estimated FF values were not significantly different from 1 (median = 1.0, $p = 0.84$; Wilcoxon signed-rank test). After correcting for FEMs, however, they were significantly <1 (median = 0.86, $p = 1.3 \times 10^{-7}$). Geometric mean FFs were 1.0 and 0.83 without and with FEM corrections, respectively.

Namely, whereas PSTH-based estimates resulted in FFs that were not significantly different from 1 ($p = 0.84$; Wilcoxon signed-rank test), correcting for FEMs resulted in FFs that were significantly <1 (“sub-Poisson variability”; $p = 1.3 \times 10^{-7}$), emphasizing the importance of accounting for FEMs when making quantitative statements about response reliability (Gur et al., 1997; Gur and Snodderly, 2006).

Effects of spike-counting window on FEM-induced biases

Thus far, we have used a time window of 10 ms for binning spikes, corresponding to the frame duration of our dynamic noise stimulus. This choice of time resolution is useful because it adequately captures changes in the firing rate over time (driven by the stimulus). However, measures of V1 response variability are typically made using larger spike count windows of hundreds of milliseconds to several seconds (Tolhurst et al., 1983; Britten et al., 1993; Softky and Koch, 1993; Teich et al., 1996; Gur et al., 1997; Buracas et al., 1998; Kara et al., 2000; Gur and Snodderly, 2006; Schölvinc et al., 2015).

To understand how the effects of FEMs manifest with longer counting windows, we applied our nonparametric analysis (Eq. 8) to estimate α for a range of different time bin sizes (Fig. 8A). Surprisingly, this analysis revealed that α remains largely unchanged across a wide range of time bins, from 5 ms to several seconds. Nonparametric estimates of α became unreliable for bin sizes >100 –200 ms (Fig. 8B), because there were too few pairs of response windows $[Y^i(t), Y^j(t)]$ with sufficiently similar eye trajectories over the course of the time window. Therefore, we estimated the effects of FEMs for larger bin sizes using our stimulus-processing models (as in Fig. 4).

We also examined the effects of time bin size on FF estimates. Consistent with most previous reports (Teich et al., 1996; Buracas et al., 1998; Schölvinc et al., 2015), we found that FF estimates (uncorrected for FEMs) increased systematically with larger time bins, particularly for bin sizes >1 s (Fig. 8C). Conversely, FEM-corrected FF estimates remained largely unchanged, at least over

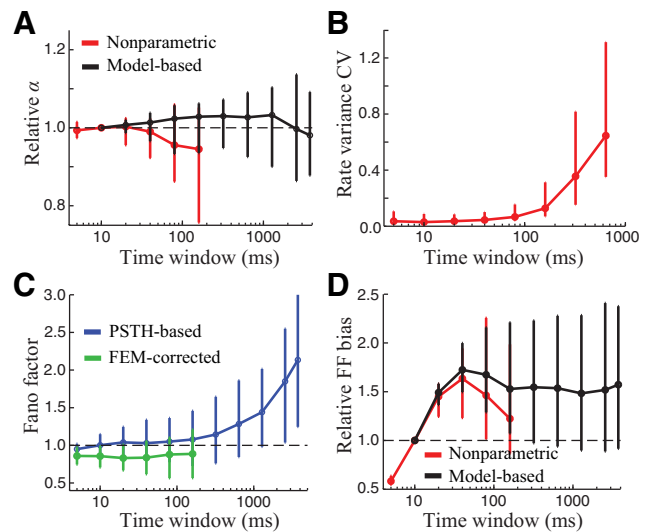


Figure 8. Dependence of FEM-induced biases on time bin size. **A**, α was largely insensitive to the time bin size used to count spikes. Model-based estimates of α (black) varied little across a wide range of time bin sizes (α values are normalized by their values at 10 ms bins). Nonparametric estimates of α (red) were similar over the range of time bins where reliable estimates could be obtained. Error bars depict the median and interquartile range across neurons. **B**, Nonparametric estimates of α became unreliable for time windows larger than ~ 100 ms, as shown by the rapid increase in the coefficient of variation (CV) of stimulus-driven response variance estimates (computed using bootstrap resampling, see Materials and Methods). **C**, Direct estimates of the FF (without correcting for FEMs; blue trace) increased substantially at larger bin sizes, whereas the FEM-corrected estimates (green) remained constant over the range of bin sizes where they could be reliably estimated. **D**, FEM-induced bias in FF estimates increased substantially for bin sizes up to ~ 50 ms (FF biases at each time window are normalized by their values at 10 ms bins, as in **A**). Both nonparametric (red) and model-based (black) estimates showed a similar trend in the FF bias for time windows up to ~ 100 ms. For larger bin sizes, the model-based estimates remained near their maximal values, even for time bins of several seconds.

the range of time bins where they could be reliably estimated (Fig. 8C). Indeed, the relative FEM-induced bias in FF estimates (which we define as the normalized difference between corrected and uncorrected estimates) increased substantially with bin size, reaching its maximal value for ~ 40 ms bins (Fig. 8D).

To better understand these FEM-induced biases in FF estimates, recall that, in the presence of FEMs, the across-trial response variance will overestimate the true stimulus-independent variance by an amount $E_i[\text{Var}_i[r^i(t)]]$. Therefore, estimates of the average FF will be biased upward by an amount equal to:

$$FF_{bias} = \frac{E_i[\text{Var}_i[r^i(t)]]}{\bar{Y}} = (1 - \alpha) \frac{\text{Var}_{i,i}[r^i(t)]}{\bar{Y}} \quad (13)$$

where, in the second step, we have used Equations 3 and 4 to write the FEM-induced bias in terms of α . In other words, FEMs will produce a bias given by the variance-to-mean ratio of the firing rate multiplied by the factor $1 - \alpha$.

Equation 13 makes clear that the FEM-induced bias in FF estimates is entirely a function of the neuron’s stimulus-driven firing rate (rather than the properties of its stimulus-independent “noise” distribution). Therefore, we can again use the stimulus-processing models for each neuron to estimate the magnitude of FEM-induced biases for time bin sizes >100 ms, where our nonparametric estimates are unreliable. This model-predicted dependence of FF bias on time bin size was consistent with our nonparametric estimates at short bin sizes (Fig. 8D) and remained largely stable (near its maximal value) for larger bin

sizes. Comparison of these results with those shown in Figure 8C suggests that the large increase in uncorrected FFs at large bin sizes reflects stimulus-independent variability occurring at slower timescales (Ecker et al., 2014; Goris et al., 2014) rather than effects of FEMs. Nevertheless, these results show that FEM-induced biases in measures of both the stimulus-driven and stimulus-independent response will remain significant even over substantially larger time windows.

Effects of FEMs on measures of correlation between neurons

Thus far, we have only considered the effects of FEMs on measures of trial-to-trial variability in single neurons; however, these effects will also be present when measuring the trial-to-trial “covariability” between pairs of neurons. Analogously to the single-neuron analysis, we can decompose the covariance between neurons m and n into stimulus-driven and stimulus-independent components:

$$\text{Cov}_{i,t}[Y_m^i(t), Y_n^i(t)] = \text{Cov}_i[r_m(t), r_n(t)] + E_i[\text{Cov}_i[Y_m^i(t), Y_n^i(t)]] \quad (14)$$

where Y_m and r_m represent the spike counts and firing rate of neuron m , respectively. The stimulus-driven covariance (first term on the right) is often estimated using a so-called “shuffle-corrector” (Perkel et al., 1967; Brody, 1999), which is essentially the covariance of the neurons’ PSTHs. The estimated stimulus-driven covariance is then subtracted from the total covariance to yield the stimulus-independent covariance (second term on the right in Eq. 14). Because the PSTHs fail to capture a significant portion of the stimulus-driven response in the presence of FEMs (as demonstrated above), this standard approach will lead to biased estimates of the stimulus-independent response covariance. Note that the stimulus-driven and stimulus-independent correlation terms defined above are conceptually similar to the so-called “signal” and “noise” correlations often studied in sensory neurons (Cohen and Kohn, 2011), although they differ in several important respects (see Discussion).

As with single-neuron variability, we can express the FEM-induced bias by decomposing the firing rate covariance into a sum of FEM-induced across-trial covariance, and the covariance of the PSTHs:

$$\text{Cov}_{i,t}[r_m^i(t), r_n^i(t)] = \text{Cov}_i[E_i[r_m^i(t)], E_i[r_n^i(t)]] + E_i[\text{Cov}_i[r_m^i(t), r_n^i(t)]] \quad (15)$$

Therefore, the PSTH covariance (and equivalently, the “shuffle-corrector”) will underestimate the true firing rate covariance by an amount: $E_i[\text{Cov}_i[r_m^i(t), r_n^i(t)]]$, producing a corresponding overestimate of the stimulus-independent covariance.

To estimate the FEM-corrected firing rate covariance between a pair of neurons, we can use an approach analogous to Equation 8: computing the covariance of the two neurons’ spiking responses on pairs of unequal trials with sufficiently similar eye position trajectories as follows:

$$\text{Cov}_{i,t}[r_m^i(t), r_n^i(t)] \approx \left\langle \left(Y_m^i(t) Y_n^i(t) \mid \Delta e_{ij}(t) < \varepsilon \right)_{i \neq j} \right\rangle_t - \bar{Y}_m \bar{Y}_n \quad (16)$$

Furthermore, by introducing a variable time lag between the responses of neurons m and n into the equations above (see Materials and Methods), we obtain the cross-correlogram and its

constituent stimulus-driven and stimulus-independent components (Brody, 1999).

Using this approach, we found that standard “PSTH-based” methods often greatly underestimate the stimulus-driven covariance (shown for four example pairs in Fig. 9A), producing biased estimates of the stimulus-independent covariance on the timescales of V1 stimulus processing (tens of milliseconds). To quantify these effects, we extracted the stimulus-driven and stimulus-independent components of the cross-correlation (at the time lag where the cross-correlogram amplitude was maximal; see Materials and Methods) with and without correcting for FEMs. Across all pairs of neurons ($n = 256$), correcting for FEMs produced estimates of stimulus-driven correlations that were substantially larger (~ 1.6 -fold) than PSTH-based estimates (Fig. 9B). As a result, FEM-corrected estimates of stimulus-independent correlations were 0.4-fold smaller than PSTH-based estimates (Fig. 9C). In a number of cases, FEM-induced across-trial variability accounted for nearly all of the apparent stimulus-independent correlation between a pair of neurons, such that their responses were essentially independent when properly conditioned on the stimulus (see first and third examples in Fig. 9A).

We found that, despite corrections for FEM, significant (generally positive) stimulus-independent correlations remained for a substantial fraction of neuron pairs (see second example in Fig. 9A). Indeed, median values of the stimulus-independent correlation were significantly positive both with (median = 0.0060; $p = 1.1 \times 10^{-16}$) and without (median = 0.013; $p = 3.7 \times 10^{-17}$) correcting for FEMs. These residual correlations were not simply due to a failure of our analysis to fully correct for FEM-induced biases because they tended to occur on slower timescales than stimulus-driven correlations (Fig. 9D). In fact, the width of the cross-correlogram peak between a pair of neurons was significantly correlated with the FEM-corrected magnitude of stimulus-independent correlations (Spearman’s $\rho = 0.33$; $p = 8.8 \times 10^{-8}$), suggesting that the remaining correlations were driven by slower processes than the neurons’ stimulus processing (Ecker et al., 2014; Goris et al., 2014).

Because FEM-induced biases in estimates of stimulus-independent correlations are fundamentally due to stimulus-driven trial-to-trial covariability, such biases could systematically affect measurements of the relationship between stimulus-driven and stimulus-independent correlations (typically measured in terms of the so-called “signal” and “noise” correlations; Averbek et al., 2006; Cohen and Kohn, 2011; Moreno-Bote et al., 2014). Indeed, we found that FEM-induced biases in estimates of stimulus-independent correlations were strongly dependent on the stimulus-driven correlation between a pair of neurons (Fig. 9E) such that neurons with more similar stimulus tuning were estimated to have larger stimulus-independent correlations. Given that this bias, $E_i[\text{Cov}_i[r_m^i(t), r_n^i(t)]]$, is entirely a function of the neurons’ stimulus-driven firing rates (Eq. 15), this may not seem surprising. However, such a relationship need not be the case in general because the components of Equation 15 can each be positive or negative (unlike single-neuron variances; Eq. 3). For example, a pair of neurons with partially overlapping RFs might have positively correlated firing rates overall yet exhibit negatively correlated across-trial firing rate variability as a result of eye movements shifting the stimulus between their RF peaks across trials (Fig. 9F). Indeed, the last example pair in Figure 9A demonstrates a case where the FEM-induced bias is not simply proportional to the stimulus-driven rate covariance.

Together, these results highlight the importance of carefully controlling for the effects of FEMs when analyzing

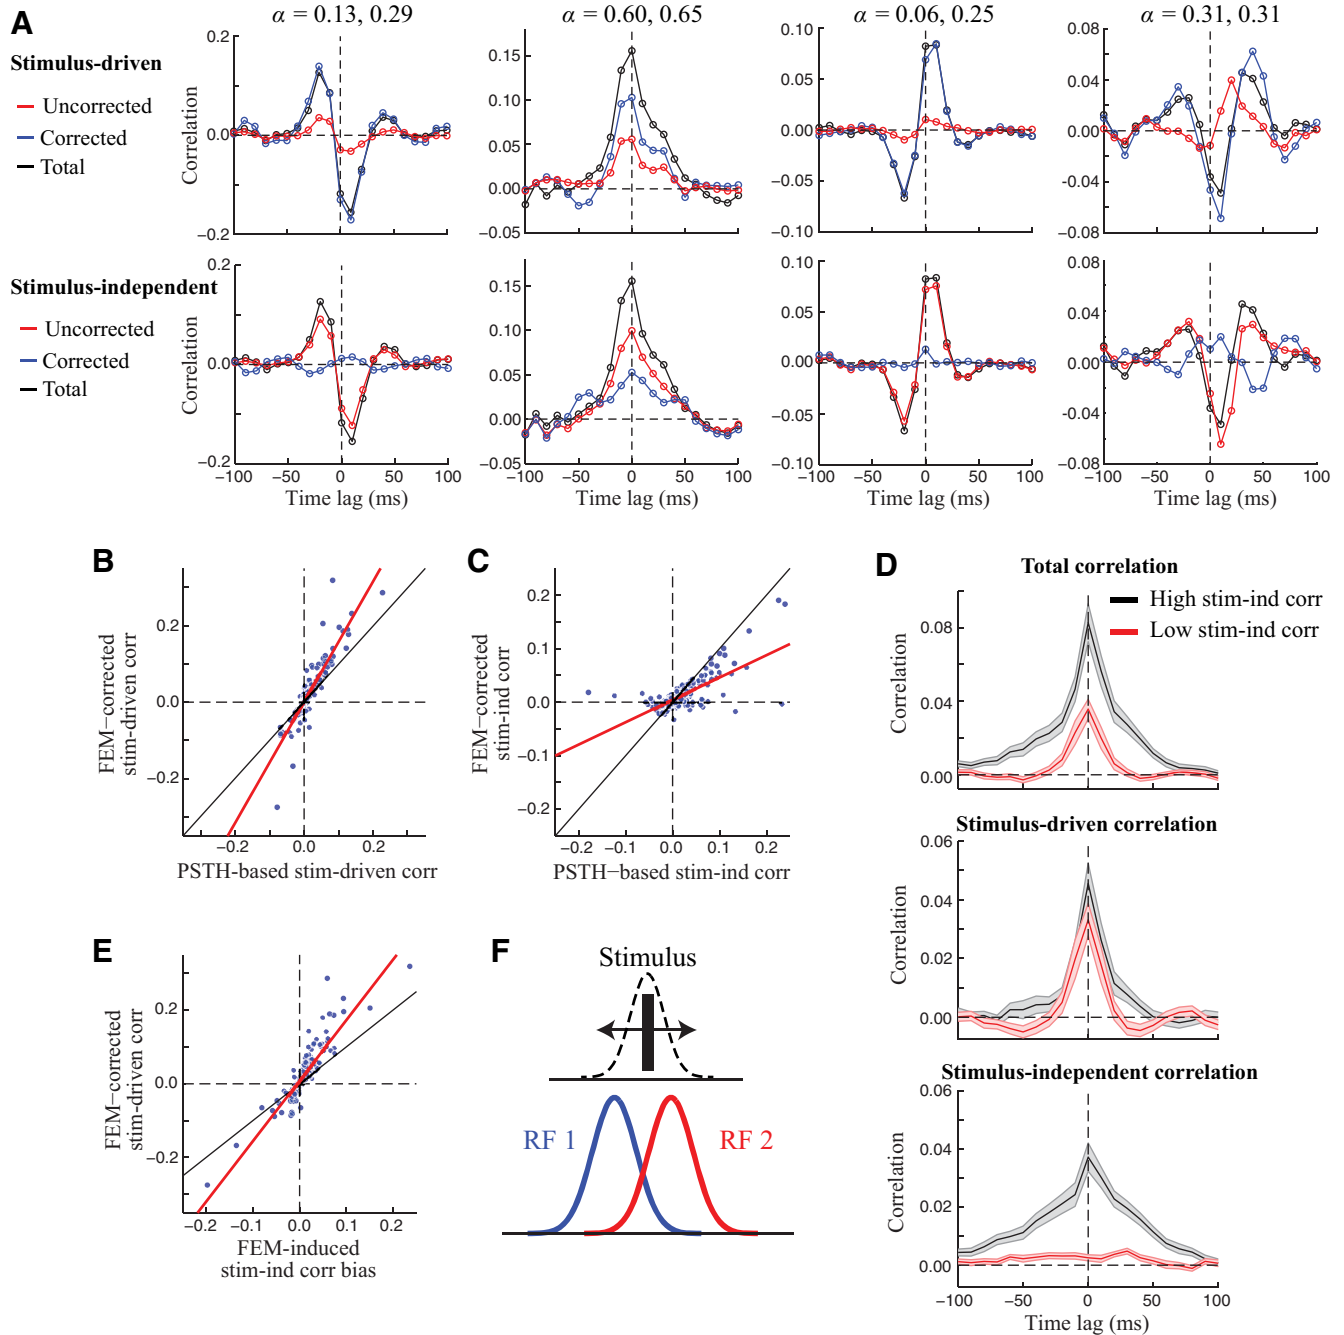


Figure 9. FEMs produce biased estimates of signal and noise correlations. **A**, Top, Raw cross-correlograms for four example neuron pairs (black), compared with FEM-corrected (blue) and PSTH-based (red) estimates of the stimulus-driven rate correlation. Bottom, Estimates of the stimulus-independent correlation with (blue) and without (red) correcting for FEMs are obtained by subtracting the stimulus-driven correlation from the total correlation (see Materials and Methods). In many cases, the uncorrected (PSTH-based) methods greatly underestimated the true firing rate comodulation, attributing it instead to correlated noise. α values for each neuron in the pair are listed above. **B**, Across all pairs of neurons ($n = 256$), FEM-corrected estimates of the stimulus-driven correlation (using the value at the time lag where the cross-correlogram amplitude was maximal) were approximately 1.6-fold larger than PSTH-based estimates (based on the slope of a linear fit; red line). Black line shows the diagonal. **C**, Correcting for FEMs reduced the magnitude of measured stimulus-independent correlations by more than half (slope of linear fit: 0.42). **D**, To illustrate the different timescales of stimulus-driven and stimulus-independent correlations, we performed a simple analysis in which we took neuron pairs with positive overall correlation and split them into two groups: those in which stimulus-independent correlations were large (black) and those in which they were small (red; see Materials and Methods). Both groups showed stimulus-driven correlations lasting several tens of ms (center), whereas stimulus-independent correlations remaining after correcting for FEMs (bottom) tended to occur on slower timescales. **E**, FEM-induced bias in estimates of the stimulus-independent correlation between a pair of neurons (i.e., the difference in corrected and uncorrected stimulus-independent correlation estimates) was strongly related to the (FEM-corrected) stimulus-driven correlation between the pair (slope of linear fit: 1.6). **F**, Schematic illustrating a simple scenario in which FEMs can introduce negatively correlated trial-to-trial variability between a pair of neurons with positive stimulus-driven correlation. The diagram shows the spatial profile of the RFs of two model neurons (red and blue curves). For a simple stimulus consisting of a single bar, FEMs will create trial-to-trial variability in the bar position relative to the neurons' RFs. Because the RFs are partially overlapping, this will produce negatively correlated variability, even though the overall "RF overlap" (and thus the overall stimulus-driven correlation) is still positive.

stimulus-independent “noise” correlations in V1, particularly when studying their relationship with measures of stimulus-driven correlations.

Discussion

Our results show that, in awake animals, FEMs generate trial-to-trial variability in the position of the visual stimulus on the retina that will tend to produce large biases in standard measures of “signal” and “noise” in V1 neuron activity. In particular, we found that the across-trial average response (i.e., the PSTH) typically captured less than half (and in many cases <20%) of the true stimulus-driven response variance. The remaining stimulus-driven activity necessarily appears as trial-to-trial variability in the neurons’ response and would typically be treated as stimulus-independent noise.

Precise measurement of these effects required several innovations. First, we developed a nonparametric procedure for estimating stimulus-driven firing rate modulation in the presence of FEMs based on computing the similarity of spiking responses on trials with similar eye positions. This method makes minimal assumptions about the nature of a neuron’s spiking “noise” or its stimulus tuning, and thus might be used more broadly to quantify variability associated with other experimentally uncontrolled factors such as the state of the cortical network (Arieli et al., 1996; Lakatos et al., 2005; Ecker et al., 2014; Schölvinck et al., 2015; Cui et al., 2016), or the animal’s attentional, arousal, or behavioral state (Saalmann et al., 2007; Cohen and Maunsell, 2009; Niell and Stryker, 2010; Constantinople and Bruno, 2011). Second, to overcome the limited accuracy of existing hardware-based eye-tracking methods, we used a recently developed model-based eye-tracking procedure, allowing us to infer the animal’s FEMs accurately using the recorded activity of populations of V1 neurons (McFarland et al., 2014). However, the analysis procedures presented here do not depend on the particular method of tracking eye position and could be used with sufficiently accurate hardware-based eye-tracking methods (Cornsweet and Crane, 1973; Roorda et al., 2002). We also obtained very similar estimates of FEM-induced biases using separate analyses based on detailed stimulus-processing models estimated for each neuron (Fig. 4A), which simultaneously validates both our nonparametric analyses and the model-based eye tracking.

We further demonstrated that FEM-induced biases are not particular to the experimental conditions studied here, and are likely to be a large factor in most studies of V1 in the awake animal. Critically, the effects of FEMs on measures of response variability were not isolated to neurons in the fovea, and were found across the full range of eccentricities we studied (out to $\sim 10^\circ$), which spans the eccentricities of typical V1 neuron recordings. Furthermore, the large biases created by FEMs are not unique to the particular visual stimuli used here. Although our model-based eye tracking is designed to work most effectively with rapidly changing stimuli that strongly modulate V1 neuron activity, we presented a simple analytical framework to illustrate that the biases created by FEMs are likely to remain substantial for the range of stimuli typically used in V1 experiments (i.e., drifting gratings, natural images, etc.). Furthermore, whereas most studies use larger time bins for measuring neural responses (i.e., hundreds of milliseconds to seconds) than the 10 ms bins used in most of our analyses, we showed that FEM-induced biases remain substantial even when using such larger spike-count windows.

The magnitude of FEM-induced biases was highly variable across neurons and was strongly dependent on several features of

the neurons’ stimulus tuning. Consistent with intuitive expectations, the width of a neuron’s RF and its sensitivity to the spatial phase of the stimulus were strongly predictive of the magnitude of FEM-induced biases. This suggests that great care must be made when comparing measures of response variability across groups of neurons where such stimulus tuning properties differ systematically (e.g., simple vs complex cells, comparisons across cortical layers, etc.). Even comparisons between different visual areas will potentially suffer from systematic biases related to the differential effects of FEMs on measures of response variability.

In addition to the neurons’ tuning properties, FEM-induced biases of course depend on the properties of the eye movements themselves, in particular on the animals’ fixation accuracy. Indeed, our analyses show that the effects of FEMs can be sensitive to even relatively small changes in fixation accuracy (Fig. 6). This suggests that any analysis relating V1 response properties (e.g., response amplitude, variability, and noise correlations) with experimental variables that are related systematically to an animal’s fixation accuracy could be strongly biased by the effects of FEMs. A number of factors often studied in relation to visual neuron activity could be associated with changes in fixation accuracy, including covert spatial attention (Hafed and Clark, 2002), arousal (Honda et al., 2013), task demands (Ko et al., 2010), behavioral training (Cherici et al., 2012), and even the properties of the visual stimulus itself (Snodderly, 1987). Therefore, particular care should be taken to control for FEM-related effects when studying how such factors relate to V1 activity.

Consistent with previous work (Gur et al., 1997; Gur and Snodderly, 2006), we found that estimates of V1 neuron response reliability (i.e., the Fano Factor) could be substantially biased by FEMs (Fig. 7). However, the work by Snodderly and colleagues examined a fundamentally different aspect of FEM-induced variability than what is shown here. Namely, by presenting a drifting bar stimulus and counting spikes over the duration of the trial, their results were largely insensitive to trial-to-trial variability in absolute eye position. Instead, they showed that the subset of trials in which the animals produced rapid eye movements contributed a large amount of across-trial response variability by altering the spatiotemporal statistics of the stimulus on the retina. Such results are complementary to those presented here, where the temporal dynamics of FEMs have virtually no effect on the ensemble of stimuli on the retina because of our use of a rapidly flashed white noise stimulus (McFarland et al., 2015). Therefore, together with these earlier studies, our results show that FEMs can contribute to trial-to-trial response variability in multiple ways, and the relative contribution of each will depend on the nature of the visual stimulus used. We also note that the FEM-corrected FF estimates reported by Snodderly and colleagues (~ 0.3) were substantially smaller than those reported here (~ 0.85). This difference is likely due to their use of a transient, “optimally tuned” stimulus given that the “onset transients” of visual cortical neurons have been shown to exhibit lower amounts of response variability (Müller et al., 2001; Churchland et al., 2010).

FEMs necessarily produce a source of trial-to-trial variability that is shared among neurons and thus could greatly affect measures of stimulus-independent “noise” correlation. Given that noise correlations between V1 neurons are typically found to be small in magnitude (Cohen and Kohn, 2011), precise measurement of noise correlations could be particularly sensitive to FEM-induced biases. For example, we found that FEMs inflate the magnitude of apparent stimulus-independent correlations by $\sim 40\%$ on average. In many cases, however, the effects were much

more dramatic, with FEMs producing large, apparently stimulus-independent correlations on the timescale of tens of milliseconds that disappeared entirely after correcting for FEM-induced variability (Fig. 9A). However, most studies measure noise correlations over the timescale of entire trials (hundreds of milliseconds to seconds; Cohen and Kohn, 2011) and previous work suggests that different mechanisms may generate noise correlations at such slower timescales (Smith and Kohn, 2008). Consistently, we found that significant positive stimulus-independent correlations remained after correcting for FEMs (Fig. 9C) and that these remaining correlations tended to have a slower timescale than stimulus-driven correlations (Fig. 9D). An interesting possibility is that these slower, stimulus-independent correlated fluctuations reflect the dominant patterns of spontaneous dynamics in the cortical network (Lakatos et al., 2005; Nauhaus et al., 2012; Ecker et al., 2014; Pachitariu et al., 2015; Cui et al., 2016).

Perhaps more important than FEM-induced effects on the overall magnitude of stimulus-independent correlations is the fact that FEM-induced biases will alter their measured relationship to stimulus-driven correlations. Indeed, we found that FEMs tended to produce larger apparent stimulus-independent correlations between neurons with more similar stimulus tuning (Fig. 9E). This raises the question of whether FEMs could alter the apparent relationship between so-called “signal” and “noise” correlations (Cohen and Kohn, 2011). There are several important differences, however, in how we measured these components of correlations compared with previous work. Most notably, although we quantified stimulus-driven correlations by measuring the similarity of firing rates driven by a one-dimensional noise stimulus, “signal” correlations are typically defined based on the similarity of tuning curves with respect to a set of parametric stimuli such as gratings of different orientations (Cohen and Kohn, 2011). The effects of FEMs on the measured relationship between stimulus-independent noise correlations and these more standard measures of “signal” correlation are expected to be different in general (e.g., neurons tuned to similar grating orientations need not have overlapping RFs). Furthermore, our dataset contained mostly neuron pairs with strongly overlapping RFs, resulting from our use of linear electrode arrays that sampled neurons largely from within a cortical column. The effects of FEMs on noise correlation measurements could be different when recording from neurons with more spatially distributed RFs. Nevertheless, our results clearly demonstrate that care must be taken when making comparisons between measures of “signal” and “noise” correlations in awake animals.

Although studies of V1 neurons in awake animals typically ignore the effects of FEMs, in many cases, control analyses are used to argue that FEMs are unlikely to affect the results significantly. Such analyses generally involve excluding a subset of the data based on recorded eye positions (e.g. time bins or trials with rapid eye movements; Roberts et al., 2007; Hansen et al., 2012; Herrero et al., 2013) and determining the impact of such exclusions on the results. Although excluding data with rapid eye movements can reduce trial-to-trial variability in the spatiotemporal statistics of the retinal stimulus (Gur et al., 1997; Gur and Snodderly, 2006), it will largely leave the overall distribution of eye positions unchanged. As a result, the biases shown here, which are generated by trial-to-trial variability in absolute eye position, would be largely unaffected. In principle, excluding data in which the measured eye position is far from the fixation target (Roberts et al., 2007) could reduce the effects of FEM-induced biases. However, the inability of standard eye-tracking methods to measure absolute eye position accurately at the scale

of FEMs (Read and Cumming, 2003; Kimmel et al., 2012; McFarland et al., 2014) may provide a fundamental limitation to resolving FEM-induced biases with such analyses.

References

- Abbott LF, Dayan P (1999) The effect of correlated variability on the accuracy of a population code. *Neural Comput* 11:91–101. [CrossRef Medline](#)
- Amarasingham A, Geman S, Harrison MT (2015) Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proc Natl Acad Sci U S A* 112:6455–6460. [CrossRef Medline](#)
- Arieli A, Sterkin A, Grinvald A, Aertsen A (1996) Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273:1868–1871. [CrossRef Medline](#)
- Averbeck BB, Latham PE, Pouget A (2006) Neural correlations, population coding and computation. *Nat Rev Neurosci* 7:358–366. [CrossRef Medline](#)
- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10:1157–1169. [CrossRef Medline](#)
- Brody CD (1999) Correlations without synchrony. *Neural Comput* 11:1537–1551. [CrossRef Medline](#)
- Buracas GT, Zador AM, DeWeese MR, Albright TD (1998) Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* 20:959–969. [CrossRef Medline](#)
- Butts DA, Weng C, Jin J, Alonso JM, Paninski L (2011) Temporal precision in the visual pathway through the interplay of excitation and stimulus-driven suppression. *J Neurosci* 31:11313–11327. [CrossRef Medline](#)
- Carandini M (2004) Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol* 2:E264. [CrossRef Medline](#)
- Cherici C, Kuang X, Poletti M, Rucci M (2012) Precision of sustained fixation in trained and untrained observers. *J Vis* 12:pii: 31. [CrossRef Medline](#)
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN (2011) Variance as a signature of neural computations during decision making. *Neuron* 69:818–831. [CrossRef Medline](#)
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378. [CrossRef Medline](#)
- Cohen MR, Kohn A (2011) Measuring and interpreting neuronal correlations. *Nat Neurosci* 14:811–819. [CrossRef Medline](#)
- Cohen MR, Maunsell JH (2009) Attention improves performance primarily by reducing interneuronal correlations. *Nat Neurosci* 12:1594–1600. [CrossRef Medline](#)
- Constantinople CM, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical networks. *Neuron* 69:1061–1068. [CrossRef Medline](#)
- Cornsweet TN, Crane HD (1973) Accurate two-dimensional eye tracker using first and fourth Purkinje images. *J Opt Soc Am* 63:921–928. [CrossRef Medline](#)
- Cui Y, Liu LD, McFarland JM, Pack CC, Butts DA (2016) Inferring cortical variability from local field potentials. *J Neurosci* 36:4121–4135. [CrossRef Medline](#)
- David SV, Vinje WE, Gallant JL (2004) Natural stimulus statistics alter the receptive field structure of v1 neurons. *J Neurosci* 24:6991–7006. [CrossRef Medline](#)
- Ecker AS, Berens P, Cotton RJ, Subramanian M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS (2014) State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82:235–248. [CrossRef Medline](#)
- Goris RL, Movshon JA, Simoncelli EP (2014) Partitioning neuronal variability. *Nat Neurosci* 17:858–865. [CrossRef Medline](#)
- Gur M, Snodderly DM (1987) Studying striate cortex neurons in behaving monkeys: benefits of image stabilization. *Vision Res* 27:2081–2087. [CrossRef Medline](#)
- Gur M, Snodderly DM (1997) Visual receptive fields of neurons in primary visual cortex (V1) move in space with the eye movements of fixation. *Vision Res* 37:257–265. [CrossRef Medline](#)
- Gur M, Snodderly DM (2006) High response reliability of neurons in primary visual cortex (V1) of alert, trained monkeys. *Cereb Cortex* 16:888–895. [Medline](#)
- Gur M, Beylin A, Snodderly DM (1997) Response variability of neurons in

- primary visual cortex (V1) of alert monkeys. *J Neurosci* 17:2914–2920. [Medline](#)
- Hafed ZM, Clark JJ (2002) Microsaccades as an overt measure of covert attention shifts. *Vision Res* 42:2533–2545. [CrossRef Medline](#)
- Hansen BJ, Chelaru MI, Dragoi V (2012) Correlated variability in laminar cortical circuits. *Neuron* 76:590–602. [CrossRef Medline](#)
- Herrero JL, Gieselmann MA, Sanayei M, Thiele A (2013) Attention-induced variance and noise correlation reduction in macaque V1 is mediated by NMDA receptors. *Neuron* 78:729–739. [CrossRef Medline](#)
- Honda S, Kohama T, Tanaka T, Yoshida H (2013) Quantitative evaluation of arousal level based on the analyses of microsaccade rates and pupil fluctuations. *Conf Proc IEEE Eng Med Biol Soc* 2013:2108–2111. [CrossRef Medline](#)
- Kara P, Reinagel P, Reid RC (2000) Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* 27:635–646. [CrossRef Medline](#)
- Kimmel DL, Mammo D, Newsome WT (2012) Tracking the eye non-invasively: simultaneous comparison of the scleral search coil and optical tracking techniques in the macaque monkey. *Front Behav Neurosci* 6:49. [CrossRef Medline](#)
- Ko HK, Poletti M, Rucci M (2010) Microsaccades precisely relocate gaze in a high visual acuity task. *Nat Neurosci* 13:1549–1553. [CrossRef Medline](#)
- Lakatos P, Shah AS, Knuth KH, Ulbert I, Karmos G, Schroeder CE (2005) An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *J Neurophysiol* 94:1904–1911. [CrossRef Medline](#)
- Martinez-Conde S, Macknik SL, Hubel DH (2004) The role of fixational eye movements in visual perception. *Nat Rev Neurosci* 5:229–240. [Medline](#)
- McAdams CJ, Maunsell JH (1999) Effects of attention on the reliability of individual neurons in monkey visual cortex. *Neuron* 23:765–773. [CrossRef Medline](#)
- McFarland JM, Cui Y, Butts DA (2013) Inferring nonlinear neuronal computation based on physiologically plausible inputs. *PLoS Comput Biol* 9:e1003143. [CrossRef Medline](#)
- McFarland JM, Bondy AG, Cumming BG, Butts DA (2014) High-resolution eye tracking using V1 neuron activity. *Nat Commun* 5:4605. [CrossRef Medline](#)
- McFarland JM, Bondy AG, Saunders RC, Cumming BG, Butts DA (2015) Saccadic modulation of stimulus processing in primary visual cortex. *Nat Commun* 6:8110. [CrossRef Medline](#)
- Mechler F, Ringach DL (2002) On the classification of simple and complex cells. *Vision Res* 42:1017–1033. [CrossRef Medline](#)
- Moreno-Bote R, Beck J, Kanitscheider I, Pitkow X, Latham P, Pouget A (2014) Information-limiting correlations. *Nat Neurosci* 17:1410–1417. [CrossRef Medline](#)
- Müller JR, Metha AB, Krauskopf J, Lennie P (2001) Information conveyed by onset transients in responses of striate cortical neurons. *J Neurosci* 21:6978–6990. [Medline](#)
- Nauhaus I, Busse L, Ringach DL, Carandini M (2012) Robustness of traveling waves in ongoing activity of visual cortex. *J Neurosci* 32:3088–3094. [CrossRef Medline](#)
- Nawrot MP, Boucsein C, Rodriguez Molina V, Riehle A, Aertsen A, Rotter S (2008) Measurement of variability dynamics in cortical spike trains. *J Neurosci Methods* 169:374–390. [CrossRef Medline](#)
- Niell CM, Stryker MP (2010) Modulation of visual responses by behavioral state in mouse visual cortex. *Neuron* 65:472–479. [CrossRef Medline](#)
- Pachitariu M, Lyamzin DR, Sahani M, Lesica NA (2015) State-dependent population coding in primary auditory cortex. *J Neurosci* 35:2058–2073. [CrossRef Medline](#)
- Pack CC, Conway BR, Born RT, Livingstone MS (2006) Spatiotemporal structure of nonlinear subunits in macaque visual cortex. *J Neurosci* 26:893–907. [CrossRef Medline](#)
- Park IM, Pillow JW (2011) Bayesian spike-triggered covariance analysis. *Adv Neural Inf Process Syst (NIPS)* 24:1692–1700.
- Park M, Horwitz G, Pillow JW (2011) Active learning of neural response functions with Gaussian processes. *Adv in Neural Inf Process Syst (NIPS)* 24:2043–2051.
- Perkel DH, Gerstein GL, Moore GP (1967) Neuronal spike trains and stochastic point processes. II. Simultaneous spike trains. *Biophys J* 7:419–440. [CrossRef Medline](#)
- Rad KR, Paninski L (2010) Efficient, adaptive estimation of two-dimensional firing rate surfaces via Gaussian process methods. *Network* 21:142–168. [CrossRef Medline](#)
- Read JC, Cumming BG (2003) Measuring V1 receptive fields despite eye movements in awake monkeys. *J Neurophysiol* 90:946–960. [CrossRef Medline](#)
- Roberts M, Delicato LS, Herrero J, Gieselmann MA, Thiele A (2007) Attention alters spatial integration in macaque V1 in an eccentricity-dependent manner. *Nat Neurosci* 10:1483–1491. [CrossRef Medline](#)
- Roorda A, Romero-Borja F, Donnelly Iii W, Queener H, Hebert T, Campbell M (2002) Adaptive optics scanning laser ophthalmoscopy. *Opt Express* 10:405–412. [CrossRef Medline](#)
- Saalmann YB, Pigarev IN, Vidyasagar TR (2007) Neural mechanisms of visual attention: how top-down feedback highlights relevant locations. *Science* 316:1612–1615. [CrossRef Medline](#)
- Sahani M, Linden J (2003) How linear are auditory cortical responses? In: *Advances in neural information processing systems* (Becker S, Thrun S, Obermayer K, eds). Cambridge: MIT.
- Schmitzer-Torbert N, Jackson J, Henze D, Harris K, Redish AD (2005) Quantitative measures of cluster quality for use in extracellular recordings. *Neuroscience* 131:1–11. [CrossRef Medline](#)
- Schölvinc ML, Saleem AB, Benucci A, Harris KD, Carandini M (2015) Cortical state determines global variability and correlations in visual cortex. *J Neurosci* 35:170–178. [CrossRef Medline](#)
- Shadlen MN, Newsome WT (1998) The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J Neurosci* 18:3870–3896. [Medline](#)
- Smith MA, Kohn A (2008) Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28:12591–12603. [CrossRef Medline](#)
- Snodderly DM (1987) Effects of light and dark environments on macaque and human fixational eye movements. *Vision Res* 27:401–415. [CrossRef Medline](#)
- Softky WR, Koch C (1993) The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13:334–350. [Medline](#)
- Teich M, Turcott R, Siegel R (1996) Temporal correlation in cat striate-cortex neural spike trains. *IEEE Eng Med Biol Mag* 15:79–87.
- Tolhurst DJ, Movshon JA, Dean AF (1983) The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23:775–785. [CrossRef Medline](#)
- van Hateren JH, van der Schaaf A (1998) Independent component filters of natural images compared with simple cells in primary visual cortex. *Proc Biol Sci* 265:359–366. [CrossRef Medline](#)
- Vintch B, Zaharia A, Movshon JA, Simoncelli EP (2012) Efficient and direct estimation of a neural subunit model for sensory coding. *Adv in Neural Inf Process Syst (NIPS)* 25:3113–3121.