

A Functional Model of Neuronal Response Variability in Primary Visual Cortex

Dylan Festa¹ (dylan.festa@einstein.yu.edu), Amir Aschner², Adam Kohn², Ruben Coen-Cagli^{1,2}

Albert Einstein College of Medicine, ¹Department of System and Computational Biology

²Department of Neuroscience

1300 Morris Park Ave, Bronx, New York, United States

Abstract

Responses of sensory neurons to repeated presentations of identical stimuli are variable. Despite extensive studies on the structure and mechanisms of this variability, its functional role remains debated. Here we propose and test a functional account of both response selectivity and variability, based on two recent hypotheses about neural coding: first, that probabilistic inference about localized visual features explains how primary visual cortex (V1) neurons integrate information inside and outside their receptive fields (RFs). Second, that the inferred probability distribution is reflected in the across-trial distribution of neuronal responses (termed sampling hypothesis), and therefore higher uncertainty in the inference implies higher variability. The resulting model predicts that stimuli surrounding the RF should reduce response variability, reflecting that surround information reduces uncertainty about stimuli inside the RF. We test the predictions on macaque V1 responses to compound gratings and natural images. We find that variability is generally suppressed by stimuli extending beyond the RF; that the suppression is weaker for uninformative surrounds (ie. with mismatched orientation); and that the modulation of variability and average firing rate can be dissociated. Our results offer strong evidence for a functional role of cortical variability in probabilistic inference.

Keywords: primary visual cortex; neural coding; visual coding; neuronal variability; neural variability; sampling hypothesis

Introduction

Cortical response variability has been studied extensively, because it can limit the information encoded by neuronal populations and therefore reduce behavioral accuracy (Zohary et al., 1994; Averbeck et al., 2006). However, it has also been proposed that variability could play functional roles (Stein et al., 2005), in particular in the representation of probabilities. Theory suggests that perception can be modeled as probabilistic inference (Knill & Pouget, 2004), and that these inferences rely on internal models adapted to the statistics of the natural environment (Barlow et al., 1961; Berkes et al., 2011). Understanding how neurons maintain internal representations of probability distributions is therefore key to understanding sensory processing, and for these reason different schemes have been proposed (Pouget et al., 2013). A recent influential proposal, termed the sampling hypothesis, is that the instantaneous neuronal activity represents samples from the target probability distribution (Hoyer & Hyvärinen, 2003; Fiser et al., 2010). Therefore, in this framework, response variability may

be indicative of the uncertainty associated to the variables represented by the firing activity.

Here we combine modeling and electrophysiology in macaques to test this functional hypothesis in primary visual cortex (V1), with a focus on how it is modulated by contextual stimuli. We focus on spatial context, i.e. stimuli in the surround of the neurons receptive field (RF), because it is known to strongly modulate firing rate (Cavanaugh et al., 2002), it has been linked to image statistics and probabilistic inference (Schwartz & Simoncelli, 2001; Coen-Cagli et al., 2012), and it has recently been shown to affect variability (Snyder et al., 2014). Importantly, spatial context provides additional information that could reduce uncertainty about the stimulus inside the RF, without modifying the RF stimulus itself. We relate neuronal responses to probabilistic inference in Gaussian scale mixture (GSMs) models. GSM models capture well the statistics of natural images (Wainwright et al., 2001) and successfully predict V1 responses including how spatial contextual stimuli modulate firing rate (Coen-Cagli et al., 2015), and how variability is affected by stimulus onset and contrast (Orbán et al., 2016). We use a simple formulation of the GSM model with spatially separate input regions representing the RF of a neuron and its surround, where the features encoded by the neuron – oriented edges – are subject to additive observation noise, and to the global influence of a single scalar multiplier, the mixer (representing e.g. the image contrast level). We assume that the goal of the neuron is, given a visual input, to invert the generative model and represent the posterior distribution of the feature inside the RF, while marginalizing out the nuisance variables (i.e. the mixer and the observation noise). Those unknowns therefore are a source of uncertainty that depends on the stimulus, and are reflected in the width of the probability distribution. Following the sampling hypothesis, we then assume that neuronal responses are samples from this posterior distribution.

We show that in this framework, model neurons exhibit supra-Poisson variability consistent with existing data (Goris et al., 2014). The model predicts that surround stimulation should reduce response variability beyond the known reduction in firing rate, reflecting that spatial context information reduces uncertainty, and that this effect on variability should be stronger when surround stimuli are more informative. We tested these predictions with recordings from V1 of anesthetized and awake fixating macaques, viewing compound gratings and natural images. Our results offer strong support for these predictions, providing new evidence for the theory that cortical variability has a precise functional role in probabilistic inference.



Methods

Generative model

To generate predictions about contextual modulation of response variability, we extended a Gaussian Scale Mixture (GSM) model. This model and related extensions capture well the spatial statistics of natural images (Wainwright et al., 2001) and predict the average firing rate of V1 neurons to compound gratings (Schwartz & Simoncelli, 2001) and natural images (Coen-Cagli et al., 2015).

We denote the observable quantity as the vector $\mathbf{x} \in \mathbb{R}^n$, which in our case corresponds to the output of a bank of n linear, Gabor-like filters (Simoncelli & Freeman, 1995) applied to an image patch. The filters are arranged into a center and a surround region (Fig. 1). Every position and orientation has two filters with complementary phases (in a quadrature pair), for convenience we denote the center vertical filters as x_c^+ and x_c^- . We now consider a generative process of the form:

$$\mathbf{x} = \mathbf{v} \mathbf{g} + \boldsymbol{\eta} \quad , \quad \text{with} \quad (1)$$

$$\mathbf{v} \sim \text{Rayleigh}(1) \quad ; \quad \mathbf{g} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{g}}) \quad ; \quad \boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_{\text{noise}}) \quad (2)$$

Here the vector $\mathbf{g} \in \mathbb{R}^n$ represents the image features, i.e. oriented edges corresponding to the filter position and orientation. The variable $\mathbf{v} \in \mathbb{R}^+$ is instead the mixer, a positive global multiplier, that can be intended as a global contrast level. $\boldsymbol{\eta} \in \mathbb{R}^n$ is an additional source of noise.

The task of the model neuron is to invert the generative process described above: given a visual stimulus and a corresponding observation \mathbf{x} , infer the latent image features inside the receptive field, denoted as g_c^+ and g_c^- . This requires the computation of the posterior distribution over two factors of uncertainty: the mixer values and the observation noise. We then model the phase-invariant neuronal response at each trial as:

$$r = r_0 + \alpha \sqrt{(g_c^+)^2 + (g_c^-)^2} \quad , \quad \text{for } \mathbf{g} \sim P(\mathbf{g}|\mathbf{x}) \quad (3)$$

where $P(\mathbf{g}|\mathbf{x})$ is the posterior probability of the features of interest, obtained by inference over the generative model. The parameters r_0 and α are not part of the GSM, and could be used for quantitative fits at the single-neuron level. Here, for a qualitative comparison, we set them heuristically to achieve realistic Fano factors.

This results in a direct dependency between the trial-to-trial variability of the spike responses and the shape of the posterior belief $P(\mathbf{g}|\mathbf{x})$. We derived analytical expressions for this relationships in a reduced model (not shown) with no observation noise (i.e. $\boldsymbol{\eta} = \mathbf{0}$ in Eq.1). In the model contextual stimuli (that is, the surround elements of \mathbf{x}) have a divisive scaling effect on both mean and variance of the inferred \mathbf{g} . This analytic result therefore predicts a precise link between divisive normalization (Carandini & Heeger, 2012) and contextual modulation of response variability, as has been recently suggested (Verhoef & Maunsell, 2017; Coen-Cagli & Solomon, 2018).

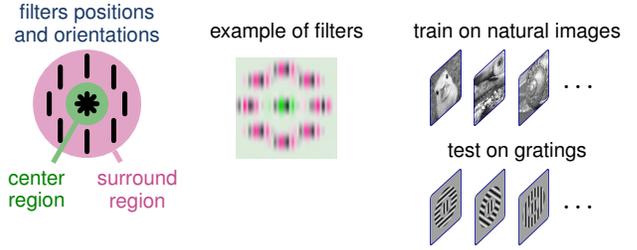


Figure 1: Model implementation

Model fitting and implementation

The $\boldsymbol{\Sigma}_{\mathbf{g}}$ is computed by moment-matching, based on an empirical estimate of the covariance of \mathbf{x} , computed applying the filter bank to a large sample of natural image patches ($N \approx 10^4$). The observation noise $\boldsymbol{\Sigma}_{\text{noise}}$ is found by applying white noise to the filters, and extracting the resulting correlation structure. The global scaling of the noise is chosen heuristically.

To test the model, we apply the same filter bank to patches of interest (gratings equivalent to those used experimentally), and sample from the posterior distribution $P(\mathbf{g}|\mathbf{x})$ numerically by a Hamiltonian Monte Carlo procedure (Stan Development Team, 2018). We compute spike-counts using Eq. 3.

Electrophysiology

Data were collected with microelectrode arrays implanted in V1 of anesthetized and awake macaques. In the anesthetized experiments, stimuli were static natural images and gratings of varying size, phase, and orientation (details below) presented for 100 msec followed by 200 msec blank screen, 20 times in pseudo-random order (see Coen-Cagli et al., 2015, for details). In the awake experiments animals performed a fixation task and stimuli were similar, except they were presented for 200 msec with a blank of 150 or 100 msec, and repeated 60 to 120 times.

We analyzed only neurons whose RFs were well-centered on the images, and for each neuron only stimuli that evoked robust responses above spontaneous activity.

Results

A general prediction of this framework is that the presence of contextual information should reduce response variability, as a consequence of the expected reduction of uncertainty about the features within the receptive field (RF). We therefore compared the Fano factors (FF) of V1 neurons in the presence of a large image (size \gg RF) or a small image (size \approx RF). The results consistently show that variability is reduced when a surround stimulus is present, for gratings (Fig. 2 a) and for natural image patches (Fig. 2 b).

This effect could reflect a generalized network mechanism, whereby providing a stronger input (i.e. a larger stimulus) to the population induces a more deterministic dynamical regime (Rajan et al., 2010). However our model also predicts that contextual stimuli have a stronger impact on variability when they are more informative about the stimulus inside the RF.

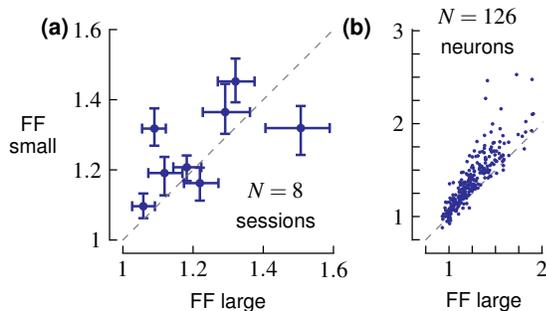


Figure 2: Fano factors (FF) measured in macaque V1. Ordinate: FF for a stimulus matched to the average RF size (1 deg), abscissa: larger stimulus (6.7 deg) that covers the RF and its surround. **(a)** Static gratings, average across neurons in each recording sessions; error bars: 68% c.i. ; **(b)** natural images, points represent single neurons.

Specifically, the model predicts that both mean response and FF decrease when center and surround are parallel, but less so when they have different orientations (Fig. 3 a).

We therefore presented compound gratings to both awake and anesthetized monkeys. The center region was kept fixed at the neuron’s preferred orientation, whereas the surround had a fixed size and contrast, but varying orientations. Consistent with the prediction, we found that the suppression of average response and FF were similarly tuned for the surround orientation (Fig. 4 a).

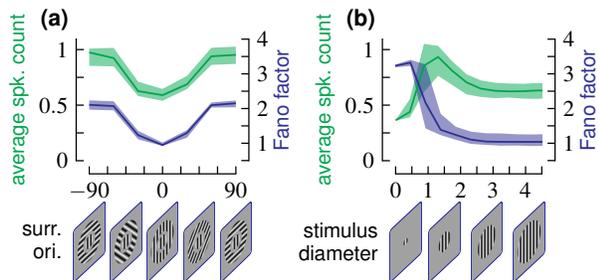


Figure 3: neuronal responses predicted by the GSM model. Left abscissa (green): trial-averaged normalized spike count; right abscissa (blue): Fano factor. Stimuli: **(a)** compound gratings, fixed center orientation (preferred angle) and surround of varying orientation; **(b)** circular patches of oriented gratings at preferred orientation of varying size. Shaded areas: 68% c.i. obtained from different noise realizations on the visual stimuli.

To exclude the possibility that the reduction in FF was simply a byproduct of the reduction in mean response, we performed a mean-matching analysis. We grouped neurons across conditions by mean response, using bins corresponding to 0.33 spike counts. We then consider mean-matched pairs, where one of the neuron is in the parallel surround condition, and the other in the orthogonal surround condition. We

then compute a Fano factor reduction score as

$$\text{score} = \frac{FF_{\perp} - FF_{\parallel}}{1/2 (FF_{\perp} + FF_{\parallel})} \quad (4)$$

Despite having very similar spike counts, the matched pairs show a consistently higher FF in the orthogonal condition compared to the parallel one (Fig. 4 b).

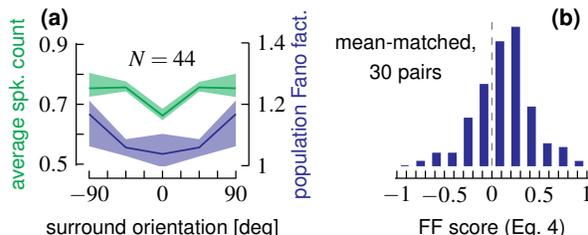


Figure 4: awake V1 responses to compound gratings with center and surround of varying orientation (see cartoons in Fig. 3 a). **(a)** Left abscissa (green): average spike counts, normalized to the response to RF-sized stimulus; right abscissa (blue): FF. Thick lines: average across neurons; shaded areas: 68% c.i. ; **(b)** FF reduction score computed as in Eq. 4 for mean-matched neuron pairs.

To further de-couple mean response modulation from effects on variability, we performed a second series of experiments. We considered circular patches of gratings and natural images, of varying size. We expect that, as size increases, the average response would change non-monotonically (Cavanaugh et al., 2002). However the GSM model makes a strong distinction between low firing due to a weak stimulus or due to surround suppression. In the former case the response is small due to a small input (the x in Eq. 1), and there is still high uncertainty in the model latent variables. In the latter case, the reduction is due to a higher estimate of the global mixer (v in Eq. 1 has a divisive effect on g), and there is more overall information about the hidden variables and the nature of the stimulus. The net result is that variability decreases monotonically with stimulus size (Fig. 3 b). We tested this model prediction in V1 of both anesthetized and awake macaque monkeys, finding a good qualitative correspondence between data and model (compare Fig. 3 b to Fig. 5).

Discussion

In this work, we combined computational modeling and experimental electrophysiology to investigate the functional role of response variability in visual coding. We extended the normative model of Coen-Cagli et al. (2012), including explicit predictions on contextual modulation of variability, based on the hypothesis of neural sampling (Orbán et al., 2016).

Our results indicate that models of neural coding based on the GSM statistics can explain the nonlinear effects of contextual modulation not only on average neuronal responses

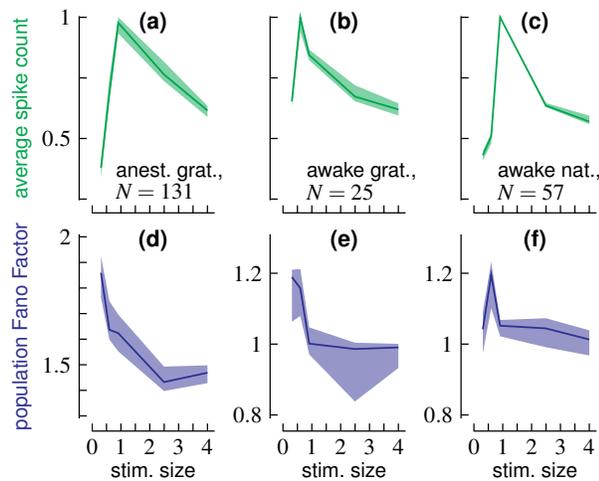


Figure 5: V1 responses in awake and anesthetized monkey to circular patches of varying size. The patch is either an oriented grating, at best orientation (as represented in Fig. 3 b), or a natural image. Spike counts of each neuron are normalized by their peak response. Same conventions as in Fig. 4 a. **(a,d)** Anesthetized, gratings; **(b,e)** awake, gratings; **(c,f)** awake, natural images.

(Coen-Cagli et al., 2015), but also on modulation of variability. Our data supports the model prediction that surround modulation of variability is stronger for more informative stimuli, and can be dissociated from the modulation of average responses. Furthermore, we observed similar effects across natural images; future work will address image-specific model predictions.

This work offers new evidence for a functional role of cortical variability in probabilistic inference, and underlines the importance of unified models of neural coding that account for both mean response and variability.

References

Averbeck, B. B., Latham, P. E., & Pouget, A. (2006, May). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5), 358–366.

Barlow, H. B., et al. (1961). Possible principles underlying the transformation of sensory messages. *Sensory communication*, 1, 217–234.

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011, January). Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment. *Science*, 331(6013), 83–87.

Carandini, M., & Heeger, D. J. (2012). Normalization as a canonical neural computation. *Nat Rev Neurosci*, 13(1).

Cavanaugh, J. R., Bair, W., & Movshon, J. A. (2002, November). Nature and Interaction of Signals From the Receptive Field Center and Surround in Macaque V1 Neurons. *Journal of Neurophysiology*, 88(5).

Coen-Cagli, R., Dayan, P., & Schwartz, O. (2012, March). Cortical Surround Interactions and Perceptual Saliency via Natural Scene Statistics. *PLOS Comput Biol*, 8(3).

Coen-Cagli, R., Kohn, A., & Schwartz, O. (2015, November). Flexible gating of contextual influences in natural vision. *Nature Neuroscience*, 18(11).

Coen-Cagli, R., & Solomon, S. S. (2018, December). Neural response variability and divisive normalization. *bioRxiv*.

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010, March). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, 14(3), 119–130.

Goris, R. L. T., Movshon, A. J., & Simoncelli, E. P. (2014). Partitioning neuronal variability. *Nat Neurosci*, 17(6).

Hoyer, P. O., & Hyvärinen, A. (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior. In *Advances in neural information processing systems* (pp. 293–300).

Knill, D. C., & Pouget, A. (2004). The bayesian brain: the role of uncertainty in neural coding and computation. *TRENDS in Neurosciences*, 27(12), 712–719.

Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016, October). Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex. *Neuron*, 92(2).

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013, September). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, 16(9), 1170–1178.

Rajan, K., Abbott, L., & Sompolinsky, H. (2010, July). Stimulus-dependent suppression of chaos in recurrent neural networks. *Physical Review E*, 82(1).

Schwartz, O., & Simoncelli, E. P. (2001). Natural signal statistics and sensory gain control. *Nat Neurosci*, 4(8), 819–825.

Simoncelli, E. P., & Freeman, W. T. (1995, October). The steerable pyramid: a flexible architecture for multi-scale derivative computation. In *Proceedings., International Conference on Image Processing (Vol. 3, pp. 444–447 vol.3)*.

Snyder, A. C., Morais, M. J., Kohn, A., & Smith, M. A. (2014, August). Correlations in V1 Are Reduced by Stimulation Outside the Receptive Field. *Journal of Neuroscience*, 34(34), 11222–11227.

Stan Development Team. (2018). *Stan modeling language users guide and reference*. (Version 2.18)

Stein, R. B., Gossen, E. R., & Jones, K. E. (2005, May). Neuronal variability: noise or part of the signal? *Nat Rev Neurosci*, 6, 389.

Verhoef, B.-E., & Maunsell, J. H. R. (2017, July). Attention-related changes in correlated neuronal activity arise from normalization mechanisms. *Nature Neuroscience*, 20(7).

Wainwright, M. J., Simoncelli, E. P., & Willsky, A. S. (2001, July). Random Cascades of Wavelet Trees and Their Use in Analyzing and Modeling Natural Images. *Applied and Computational Harmonic Analysis*, 11(1).

Zohary, E., Shadlen, M. N., & Newsome, W. T. (1994, July). Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370(6485), 140.