Neuronal Variability Reflects Probabilistic Inference Tuned to Natural Image Statistics

Dylan Festa¹ Amir Aschner² Aida Davila² Adam Kohn^{1,2,3}
Ruben Coen-Cagli^{1*}

¹*Department of Systems and Computational Biology*

²*Dominick Purpura Department of Neuroscience*

³*Department of Ophthalmology and Visual Sciences*

Albert Einstein College of Medicine, Bronx, NY, USA

* Corresponding author — ruben.coen-cagli@einsteinmed.org

Abstract

Neuronal activity in sensory cortex fluctuates over time and across repetitions of the same input. This variability is often considered detrimental to neural coding. The theory of neural sampling proposes instead that variability encodes the uncertainty of perceptual inferences. In primary visual cortex (V1), modulation of variability by sensory and non-sensory factors supports this view. However, it is unknown whether V1 variability reflects the statistical structure of visual inputs, as would be required for inferences correctly tuned to the statistics of the natural environment. Here we combine analysis of image statistics and recordings in macaque V1 to show that probabilistic inference tuned to natural image statistics explains the widely observed dependence between spike-count variance and mean, and the modulation of V1 activity and variability by spatial context in images. Our results show that the properties of a basic aspect of cortical responses — their variability — can be explained by a probabilistic representation tuned to naturalistic inputs.

1 Introduction

In sensory cortex, neuronal activity is typically variable, both in the absence of sensory input and for repeated presentations of a stimulus (Tolhurst, Movshon, et al. 1983; Tomko and Crapper 1974). This variability is modulated by several sensory (M. Chen, Yan, et al. 2014; Churchland, Yu, et al. 2010; Coen-Cagli and Solomon 2019; Goris, Movshon, et al. 2014; Kohn and Smith 2005; Orbán, Berkes, et al. 2016; Ponce-Alvarez, Thiele, et al. 2013; Rabinowitz, Goris, et al. 2015; Solomon, S. C. Chen, et al. 2014; Verhoef and Maunsell 2017) and non-sensory (Cohen and Maunsell 2009; Dadarlat and Stryker 2017; Mitchell, Sundberg, et al. 2009;

White, Abbott, et al. 2012) factors, suggesting it may play a functional role rather than simply reflecting noise. Understanding the functional role of variability is at the core of the inquiry of neural coding (Beck, Ma, et al. 2012; Dinstein, Heeger, et al. 2015; Fiser, Berkes, et al. 2010; Pouget, Dayan, et al. 2003; Shadlen and Newsome 1998; Sompolsky, Yoon, et al. 2001).

Parametric descriptive models can capture how stimuli modulate neuronal variability (Charles, Park, et al. 2018; Coen-Cagli and Solomon 2019; Goris, Movshon, et al. 2014; Stevenson 2016), but they do not address why modulation of variability occurs and what functional role it might play. Here we develop and test a normative model, based on efficient coding (Bell and Sejnowski 1997; Karklin and Lewicki 2009; Olshausen and Field 1996; Schwartz and Simoncelli 2001) and probabilistic inference (Berkes, Orbán, et al. 2011; Coen-Cagli, Kohn, et al. 2015; Lochmann, Ernst, et al. 2012; Ma, Beck, et al. 2006; Orbán, Berkes, et al. 2016), to explain the properties of response variability in sensory cortex. In this approach, we hypothesize about functional and computational principles of cortical processing, to generate predictions about cortical activity. Specifically, we propose that probabilistic inference tuned to the statistics of natural images can explain the properties of response variability in visual cortex.

Although normative models have typically been used to explain trial-averaged responses, they can also be used to explain response variability (Boerlin, Machens, et al. 2013; Fiser, Berkes, et al. 2010; Hoyer and Hyvärinen 2003; Hunsberger, Scott, et al. 2014). In particular, some aspects of variability in primary visual cortex (V1) can be explained by the theory of neural sampling. This theory builds on the broader idea that the brain approximates operations of probabilistic inference (Knill and Pouget 2004; Pouget, Beck, et al. 2013), and hypothesizes that instantaneous neuronal activity represents samples from a probability distribution (Fiser, Berkes, et al. 2010; Haefner, Berkes, et al. 2016; Hoyer and Hyvärinen 2003). According to this view, variability of neuronal activity reflects uncertainty about the visual input (i.e. the width of the inferred probability distribution). As a result, variability is reduced by stimulus onset (Churchland, Yu, et al. 2010) and stimulus contrast (Finn, Priebe, et al. 2007; Kohn and Smith 2005), because of a reduction in uncertainty (Orbán, Berkes, et al. 2016).

Here we hypothesize that modulation of uncertainty by visual input should reflect inferences tuned to the statistics of natural images, and thus that the properties of response variability should reflect the statistical structure of images. To test this prediction, we consider a successful modeling framework, the Gaussian Scale Mixture (GSM; Theis, Hosseini, et al. 2012; Wainwright, Simoncelli, et al. 2000). This model assumes that images are composed by local features (e.g. oriented edges; Fig. 1A) and global features (e.g. image contrast), and that V1 neurons aim to represent the local features while discarding the global features (Coen-Cagli, Dayan, et al. 2012; Coen-cagli, Dayan, et al. 2009; Coen-Cagli, Kohn, et al. 2015; Orbán, Berkes, et al. 2016; Schwartz and Simoncelli 2001). GSMs can explain the modulation of trial-averaged V1 responses by stimuli in the surround of the receptive field (RF; Angelucci, Levitt, et al. 2002; Cavanaugh, Bair, et al. 2002b; Sceniak, Ringach, et al. 1999; Sillito and Jones 1996; Walker, Ohzawa, et al. 1999). However, it is unclear whether this framework can also explain the surround modulation of variability (Haider, Krause, et al. 2010; Snyder, Morais, et al. 2014)

A

$$\text{Image} = \left(g_c \cdot \text{Feature}_c + g_{s1} \cdot \text{Feature}_{s1} + g_{s2} \cdot \text{Feature}_{s2} + \dots \right) \cdot v + \text{noise}$$

B

C

D

E

Figure 1: See next page for caption.

Figure 1: Sampling-based inference in the GSM model explains the dependence between spike-count variance and mean. **A.** Representation of the generative process of the Gaussian scale mixture (GSM) model, (Methods Eq. 1). The image (left) is described as the combination of local oriented features weighted by Gaussian coefficients, further multiplied by a global modulator and corrupted by additive, Gaussian noise. **B.** Encoding of sensory information according to the sampling hypothesis: the goal of a neuron is to represent the posterior distribution (orange, middle) of the feature it encodes. The activity of the neuron corresponds to samples from that distribution, therefore the histogram of spike counts over time or repetitions (green, right) reflects the distribution. **C.** Tuning of mean and variance in a 1-dimensional version of the GSM with no noise. For fixed input x , the visual feature g and the modulator ν are bound to lie on the hyperbole $\nu = x/g$ (black line). Therefore, a larger estimate of ν implies reduced mean and variance of g (blue versus brown curves). **D.** Mean versus variance of a GSM model neuron in response to 1,000 patches of natural images. Patches were selected randomly, with the requirement of sufficient signal strength inside the RF, i.e. above the median of the full distribution of $(x_{1+}^2 + x_{1-}^2)$ on natural scenes, where x_{1+} and x_{1-} are the odd and even phases of the center vertical filter (see Methods). **E.** The Fano factor (FF; ratio between mean and variance) as a function of the mean for the same GSM simulation reported in **D.** Red dashed line represents the best linear fit. Pearson corr. 0.214, $p < 10^{-4}$.

and whether this modulation reflects the statistical properties of natural inputs.

Here we combine modeling and electrophysiology in macaques to test our hypothesis that V1 variability is tuned to natural image statistics. First, we show analytically that the dependence between spike-count variance and mean observed empirically (Goris, Movshon, et al. 2014; Shadlen and Newsome 1998; Tolhurst, Movshon, et al. 1983) emerges in the GSM from the multiplicative interactions between local and global image features. Second, we show that stimuli in the RF surround modulate these interactions, and thus also response variability. Finally, we test predictions about surround modulation of firing rate and variability with recordings in V1 of awake and anesthetized macaques viewing natural images and gratings.

Our results show that visual context modulates neuronal response strength and variability independently, suggesting these modulations reflect probabilistic inference about local visual features. Our work thus provides evidence that the tuning of cortical variability can be explained assuming the brain performs operations of probabilistic inference of natural image statistics.

2 Results

2.1 The dependence between spike-count variance and mean reflects multiplicative interactions between latent variables

To study the relation between natural image statistics and V1 cortical variability, we considered the GSM because it captures the most prominent aspects of low-level image statistics, namely the sparseness of V1-like, oriented visual features and their nonlinear statistical dependence (Schwartz and Simoncelli 2001; Wainwright, Simoncelli, et al. 2000). We assumed that the instantaneous firing of V1 neurons (Methods Eq. 4) represents samples from the inferred probability distribution (termed posterior distribution; Berkes, Orbán, et al. 2011; Orbán, Berkes, et al. 2016) of oriented visual features encoded by the neurons. The inference of the posterior distribution requires inverting the so-called generative model of stimuli: that is, how features — small patches with different orientations and positions — are combined to produce images (Fig. 1A). Given an input image, model neurons then encode the inferred probability distribution of the coefficients of those features in the image. This is illustrated schematically for a vertical feature in Fig. 1B-top. The posterior distribution (middle column) in this case was broad with a large mean, indicating that the vertical feature was strongly present in the input image, though its precise coefficient was uncertain. Conversely, the image in Fig. 1B-bottom contains little evidence for the vertical feature, leading to a narrow posterior centered on zero. In the sampling framework, neuronal responses represent samples from this posterior distribution (Fig. 1B, right column). Thus, the variance of the spike count distribution (i.e. the neuronal variability) reflects the variance or width of the posterior, corresponding to the uncertainty about the coefficient of the encoded feature.

We studied whether, in the GSM, response variance depends on response mean, as observed in V1 (Goris, Movshon, et al. 2014; Shadlen and Newsome 1998; Tolhurst, Movshon, et al. 1983). The GSM assumes $\mathbf{x} = \nu \mathbf{g}$ where the sensory input $\mathbf{x}$ is the result of local features $\mathbf{g}$ (the variables encoded by the neurons) multiplied by a global modulator ν (e.g. image contrast). To gain intuition about the mean–variance relationship of the model, we first considered the simplest formulation of a GSM, where $\mathbf{x}$ and $\mathbf{g}$ are 1-dimensional. Although the expression relating these quantities — $x = \nu g$ — is deterministic, knowledge of x is insufficient to determine g , due to the unknown ν . Computing the probability distribution of g by accounting for the possible values of ν is a fundamental operation of probabilistic inference, called marginalization (Beck, Latham, et al. 2011; Pouget, Beck, et al. 2013). Crucially, because of the multiplication, both the inferred value of g and its uncertainty (i.e. the mean and standard deviation of the posterior over g) are divisively related to ν . For instance, assume we observed $x = 10$ and we inferred that ν is likely to be between 1 and 2 (Fig. 1C, dark blue), then by marginalization we would infer that g is with high probability between 5 and 10 (Fig. 1C, light blue). If instead ν was inferred to be in the interval 4–5 (Fig. 1C, dark brown), then g could only take values between 2 and 2.5, thus shrinking both in mean and variance (Fig. 1C,

light brown). This example illustrates why a neuron whose responses reflect samples from the inferred distribution of g should display a dependence between mean and variance in its response statistics. Note that this dependency is not linear, nor do mean and variance strictly follow each other as they would in a Poisson process. In general, the relative scaling depends on model choices, such as the uncertainty on the priors and, for high dimensional inputs, the stimulus structure (as explained in the next section). Notice too that if the mixer term were additive instead of multiplicative, then changes in its inferred value would only change the inferred mean of $\mathbf{g}$, not its variance, leading to different predictions (Supplementary Fig. S1).

To validate this intuition more rigorously, we considered GSM inference on real images. As in past normative models (Coen-Cagli, Dayan, et al. 2012; Coen-cagli, Dayan, et al. 2009; Coen-Cagli, Kohn, et al. 2015; Schwartz and Simoncelli 2001), we implemented a GSM with oriented filters (Simoncelli and Freeman 1995) spatially arranged to define both the RF of the model neuron and its surround (Fig. 2A; details in Methods). The model was trained on a large ensemble ($N=10,000$) of natural image patches extracted from the BSDS500 database (Arbelaez, Maire, et al. 2011, <https://github.com/BIDS/BSDS500>).

Given an input image, the visual inputs $\mathbf{x}$ (a vector) were determined by the activations of those filters applied to the image. We denoted by $\mathbf{g}$ the corresponding local visual features. First, we verified that the multiplicative effect of the modulator allows the GSM to capture the statistics of natural images (Wainwright, Simoncelli, et al. 2000) better than an additive modulator (Supplementary Fig. S1). We found through analytical derivations and simulations that the variance of the inferred $\mathbf{g}$ grows with the mean, and both are divisively scaled by the estimate of the global modulator ν , leading to a general reduction of uncertainty when the estimate of ν increases (Methods Eq. 2, 3; see Supplementary Text for derivation). We then simulated model responses to a wide range of natural images (Fig. 1D), and characterized the mean–variance relation. The response variance of the model neuron scaled proportionally with its mean. Furthermore the ratio of variance to mean, termed Fano factor (FF), increased on average for stimuli that elicited stronger mean responses (Fig. 1E), in qualitative agreement with the statistics of V1 neurons (Goris, Movshon, et al. 2014). Importantly, training a GSM on different image sets, such as white noise, led to different parameter values but qualitatively similar predictions for neural responses (Supplementary Fig. S2), indicating that the mean–variance dependence arises from matching the generative model’s structure to image statistics (i.e. multiplicative latent interactions) rather than fine-tuning its parameters.

These analyses confirm the intuition that the dependence between posterior variance and mean observed in the GSM emerges from the multiplicative interactions between the global modulator and the local variables. Because this partition between local and global variables in the GSM is known to capture well the statistics of natural images (Schwartz and Simoncelli 2001; Wainwright, Simoncelli, et al. 2000), our result establishes a precise link between image statistics and cortical variability.

Figure 2: Surround stimulation reduces GSM uncertainty and V1 variability for natural images.

A. In the GSM, the inputs to the model neuron are provided by the activity of quadrature pairs of oriented filters, corresponding to the spatial RF (green) and its surround (magenta). **B.** Noise-free GSM model applied to an image without (blue) and with (brown) surround. The surround stimulus does not change the constraint between g_c (the local feature associated with the RF center) and ν , but it influences the estimate of the modulator and therefore also the estimate of g_c . **C.** FF averaged across small (1 deg) and large (3.1 or 6.7 deg) natural image patches. Black and gray circles: average FF across images for each V1 neuron; black denotes a significant difference ($p < 0.05$) across the two conditions. Orange circle: average FF of the GSM response for the same set of images. For the conversion to spike counts (see Methods, Eq. 4) we used the scaling factor $c = 2$. **D-F.** Tuning of the mean spike count (green) and FF (blue), for natural image patches of varying size. **D.** GSM model, scaling factor $c = 15$. This constant was different than in **C**, because the experiments of **C** used images with a broader range of orientation and frequency content than **D**. **E,F.** Data from one awake fixating macaque V1, for two example neurons **E** and the population average **F**. Error bars represent the 68% c.i.

2.2 Surround stimulation reduces uncertainty and V1 variability

The previous analysis shows that variability in the GSM is influenced by the inferred values of the global modulator. Therefore, the framework predicts that variability is sensitive to stimulus manipulations that affect the inferred global modulator. Specifically, stimuli that lead to a higher estimate of the modulator present less uncertainty over the hidden feature, and thus should reduce response variability. To test this prediction, we considered the modulation of V1 activity induced by spatial context — by stimuli in the surround of a neuron’s RF — because spatial context can reduce stimulus uncertainty without modifying the stimulus drive inside the RF (Albright and Stoner 2002).

First, we verified for the GSM that surround stimuli (i.e. image regions that activate the surround filters) reduce uncertainty. The activity of the model neuron is associated with the oriented feature in the center. However, the surround input contributes to the estimate of the global modulator, and therefore influences the neuronal response. Specifically, our analytical results show that, for a fixed RF input, surround stimulation increased the estimated modulator and therefore had a suppressive influence both on the mean and variance of the neuronal response (Fig. 2B; Methods), validating our intuition that surround stimuli reduce uncertainty by resulting in a higher estimate of the global modulator.

Next, we tested whether surround stimulation reduces V1 variability, relative to RF stimulation alone, by analyzing previously published data on V1 surround modulation in anesthetized macaques (Coen-Cagli, Kohn, et al. 2015). In these experiments, natural image patches were presented at two different sizes, either masked to fit within the average RF (1 degree), or extending well beyond into the surround (3.1-6.7 degrees). Among the neurons with a significant change in FF across conditions (127/261 neurons, $p < 0.05$), the vast majority (91.3%) had a lower FF for large images than small ones, consistent with model predictions. The average FF, across all neurons, was also lower for large images than small ones (1.15 versus 1.22, $p < 10^{-6}$, $N = 261$ neurons). We verified with a mean-matching analysis that this difference in FF could not be explained by differences in spike-count mean (Supplementary Fig. S3). This result agrees qualitatively with the model (Fig. 2C, orange symbol), although surround suppression of FF was stronger in the model, possibly because surround modulation in the GSM is recruited by all images, whereas in V1 it is weak or absent for many images (Coen-Cagli, Kohn, et al. 2015). Consistent with this possibility, the strength of surround suppression of responsivity and of FF were positively correlated (Supplementary Fig. S4).

2.3 Distinct effects of RF and surround stimulation on variability suppression

Suppression of response variability by large stimuli might not be due solely to surround stimulation. Visual stimuli reduce the variability seen in spontaneous activity (Churchland, Yu, et al. 2010). Therefore large images might reduce variability by providing stronger drive

to the RF, in those cases where small images did not completely cover the RF. To test whether stimuli larger than the RF induced further reduction of the FF, beyond the reduction caused by the stronger RF drive, we considered responses to circular patches of natural images, with sizes ranging from much smaller to much larger than the typical RF.

We first studied the effects of stimulus size in the GSM. We found that the mean response peaked for images matched in size to the RF and decreased for larger stimuli, consistent with past work (Coen-Cagli, Dayan, et al. 2012). The FF, on the other hand, decreased monotonically with stimulus size, well after the stimulus filled the RF (Fig. 2D), because large stimuli lead to a larger estimate of the global modulator (Supplementary Fig. S5A). The difference between the behavior of the FF and the mean indicates that it should be possible to dissociate the effects of variability reduction from the modulation of spike-count mean: stimuli smaller than the RF and larger than the RF can elicit similar average responses but with different variability.

We tested these predictions in V1 responses to natural images of different sizes in one awake fixating macaque. For the two example neurons of Fig. 2E, the mean spike count displayed the typical non-monotonic size dependence (green), whereas the FF decreased monotonically (blue). Similar effects were evident across all recorded neurons for stimuli ranging from approximately half the RF size up to several times larger ($N = 86$; Fig. 2F). The FF decreased by 18.7% as stimuli increased from approximately $1/2$ RF size to RF size, and an additional 5.7% as stimuli increased from RF size to approximately twice that size (Table 1, left), which is the average extent of the suppressive surround in V1 (Angelucci, Levitt, et al. 2002; Cavanaugh, Bair, et al. 2002a; Sceniak, Ringach, et al. 1999). Furthermore, the FF decreased for stimuli larger than the RF compared to stimuli smaller than the RF, even when both stimuli evoked approximately the same number of spikes (Table 1, right). To be sure that our results were not affected by inaccurate estimates of RF size, due to variations in local contrast across natural images, we measured responses to static gratings in the same animal, and obtained similar results (Table 1, experiment 2; Supplementary Fig. S6A). New analyses of previously published data from anesthetized animals (Coen-Cagli, Kohn, et al. 2015) also confirmed these results (Table 1, experiment 3; Supplementary Fig. S6), ruling out the possibility that microsaccades in the awake animals might have introduced biases.

Note that the FF was lower on average for stimuli smaller than $1/2$ RF size (Fig. 2F, leftmost point). This was true for the subset of neurons with large RF ($N = 65/86$), whereas the FF decreased strictly monotonically for neurons with smaller RFs (Supplementary Fig. S7). Both the large apparent RF size and the non-monotonicity of the FF would be expected if stimuli were not perfectly centered on the RF (Supplementary Fig. S8). Furthermore, the FF decreased monotonically with stimulus size in the anesthetized dataset, for which stimulus centering could be controlled more tightly (Supplementary Fig S7).

These analyses show that stimulation of the RF surround reduces response variability, beyond the known reduction from spontaneous to stimulus-driven activity (Churchland, Yu, et al.

Experiment	FF decrease ($\frac{1}{2}$ RF) – (RF)	FF decrease (RF) – ($2 \times$ RF)	p -value	Mean-matched FF decrease (size > RF) – (size < RF)	p -value
1. Natural, awake ($N = 86$; Fig. 2F)	18.7%	5.7%	0.082	25.7%	$< 10^{-5}$
2. Gratings, awake ($N = 19$; Supp. Fig. S6)	31.7%	9.0%	0.05	47.7%	$< 10^{-3}$
3. Gratings, anesthetized ($N = 229$; Supp. Fig S6)	14.2%	7.05%	$< 10^{-3}$	22.6%	$< 10^{-5}$

Table 1: Rows. Separate experiments, with number of neurons selected in each experiment (selection criteria in Methods). **Columns.** Columns 1-3, changes in FF with stimulus size; columns 4 and 5, mean-matched (see Methods) change in FF with stimulus size. In all cases, a positive change denotes a reduction in FF for larger stimuli. First column: change in FF (Methods, Eq. 5) from the stimulus closest to $\frac{1}{2}$ of the RF size (out of all tested sizes) to the RF-sized stimulus. Second column: change in FF from the RF-sized stimulus to the large stimulus (closest to 2 RF size). Third column: the p -value for the second column. Fourth column: FF change from stimuli smaller to larger than RF size. Sizes are selected to match the mean spike count across neurons (spike count change $< 3\%$, $p > 0.05$, for all experiments). Fifth column: p -value for the fourth column.

2010).

2.4 Surround suppression of variability is orientation selective

Surround suppression of mean firing rate is known to be stronger for image patches with matched orientation inside and outside the RF, and weaker when the surround orientation is orthogonal to the center (Angelucci, Bijnanzadeh, et al. 2017; Cavanaugh, Bair, et al. 2002b; Gardner, Anzai, et al. 1999; Sillito and Jones 1996; Walker, Ohzawa, et al. 1999; Webb, Dhruv, et al. 2005). It is not known whether variability is similarly tuned. In our GSM model, surround tuning of mean responses (Fig. 3A, green) was obtained by using surround filters with the same orientation as the feature of interest inside the RF (details in Methods), as in past implementations (Coen-cagli, Dayan, et al. 2009; Schwartz and Simoncelli 2001).

Because the GSM predicts that surround suppression of both mean spike-counts and variability is controlled by the inferred strength of the global modulator, we found that surround suppression of model response variability and of mean spike-counts were similarly tuned (Fig. 3A). We verified that this corresponded to a smaller estimate of the global modulator for orthogonal surround stimuli (Supplementary Fig. S5B), which in turn resulted in weaker surround suppression of variability.

To test these model predictions, we measured V1 responses to compound static gratings

Figure 3: Surround reduction of variability is orientation tuned. **A.** Surround orientation tuning of the mean spike count (green line) and FF (blue line), relative to the center stimulus alone (dashed lines) in the GSM model. Error bars: 68% c.i. scaling factor (Methods Eq. 5). **B,C.** Percent change in the mean spike count **B** and Fano factor **C** from orthogonal to matched surround orientation, in V1 of two awake fixating macaques. Yellow bars denote neurons with a significant change across conditions. The difference is considered significant when the 68% c.i.'s of the two conditions do not overlap.

in two awake, fixating macaques ($N = 71$ neurons). Consistent with past literature, the mean response was suppressed (relative to no surround) more when the surround and center orientations were matched (Fig. 3B; average suppression matched 0.844, orthogonal 0.885; average reduction 6.28%, $p = 0.0043$). In agreement with model predictions, the FF was smaller for the matched surround (Fig. 3C; average FF matched 0.973, orthogonal 1.02; average reduction 4.73%, $p = 0.032$), and this was true in the majority ($N = 9/14$) of neurons with a significant change ($p < 0.05$). However, although consistent with the GSM prediction, the magnitude of the effect was small (see also Discussion). One reason might be that, in our data, 26/71 neurons responded more strongly to parallel than orthogonal surrounds (i.e. opposite to the surround tuning of our GSM implementation), which may be due both to imperfect stimulus centering and to the known heterogeneity in the orientation tuning of surround suppression of firing rate (Cavanaugh, Bair, et al. 2002b). Consistent with this explanation, we verified that if we restricted our analysis to neurons that responded more weakly to parallel than orthogonal surround ($N = 45/71$; average reduction 17.3%, $p < 10^{-5}$), the surround tuning of FF was also stronger (average reduction 7.37%, $p = 0.013$) than for the entire population (Supplementary Fig. S9).

Our analysis shows that surround suppression of variability in V1 is tuned to the orientation of surround stimuli, in a manner similar to the tuning of firing rate suppression, suggesting partly shared mechanisms. In the GSM framework, this tuning arises because only matched surround stimuli provide information about the global modulator and thus reduce uncertainty.

3 Discussion

We have presented a theoretical framework that explains V1 variability and its modulation by spatial context in natural images, as reflecting probabilistic inference about local features in visual inputs. Our work builds on the theory of neural sampling (Fiser, Berkes, et al. 2010; Haefner, Berkes, et al. 2016; Hoyer and Hyvärinen 2003; Orbán, Berkes, et al. 2016), in which neuronal variability encodes uncertainty of the inferences, and offers two main contributions. First, we established a precise link between V1 response variability and the statistics of natural images. We showed that the dependence between spike-count variance and mean, and the modulation of variability by spatial context are general consequences of probabilistic inference when there are multiplicative interactions between latent variables, which is a widely-adopted description of natural image statistics (Coen-Cagli, Dayan, et al. 2012; Dayan and Abbott 2005; Gao and Vasconcelos 2007; Geisler 2008; Schwartz and Simoncelli 2001; Wainwright, Simoncelli, et al. 2000). Second, we validated our model with measurements of V1 activity. Consistent with model predictions, spatial context in images modulated V1 variability beyond the known reduction of variability from spontaneous to stimulus-driven activity (Churchland, Yu, et al. 2010). Furthermore, the tuning of contextual modulation of variability was similar to (although weaker than) that of mean spike counts, suggesting shared mechanisms.

3.1 Natural image statistics and contextual modulation of response variability

Normative models of visual processing have explained properties of V1 representations from optimization and efficiency principles related to the statistics of the natural environment (Bell and Sejnowski 1997; Hyvärinen and Hoyer 2000; Karklin and Lewicki 2009; Olshausen and Field 1996; Rao and Ballard 1999; Schwartz and Simoncelli 2001). This work has typically addressed only the trial-averaged spike counts. However, across-trial variability is substantial in cortex and can strongly influence perception (Kohn, Coen-Cagli, et al. 2016; Moreno-Bote, Beck, et al. 2014; Pouget, Beck, et al. 2013; Shadlen and Newsome 1998). Understanding cortical processing requires addressing this variability, which we have done via the neural sampling theory.

The hypothesis that neuronal variability reflects sampling from a distribution (Hoyer and Hyvärinen 2003) is rooted in machine learning research focused on efficient inference schemes (Bishop 2006). Past work in neural network modeling has shown how samples might be generated dynamically, and in a manner that is fast enough for accurate inference within short, biologically relevant timescales (Echeveste, L. Aitchison, et al. 2020; Hennequin, Laurence Aitchison, et al. 2014; Legenstein and Maass 2014; Savin and Denève 2014).

While past work has addressed the plausibility of neural sampling, we have focused instead on contextual effects, for two important reasons. First, contextual effects disambiguate between

two key aspects of neural coding: the strength of the stimulus feature represented by the neuron, and the uncertainty about that feature. This is because stimuli in the RF surround do not directly affect the inputs to the RF, but they can modulate uncertainty. This is different from contrast modulation (Orbán, Berkes, et al. 2016) and other common experimental manipulations (e.g. adding stimulus noise; Britten, Shadlen, et al. 1992; Nienborg and Cumming 2014), that modulate both the strength of a visual feature and its uncertainty. Second, natural visual inputs have rich statistical structure that extends across the visual field. There is abundant evidence suggesting a relation between spatial structure in images and spatial contextual effects in cortex (Coen-Cagli, Kohn, et al. 2015; Rao and Ballard 1999; Schwartz and Simoncelli 2001; Spratling 2010; Zhu and Rozell 2013). Contextual modulation of V1 trial-averaged responses has been characterized extensively with artificial stimuli (Angelucci, Levitt, et al. 2002; Cavanaugh, Bair, et al. 2002a,b; Sceniak, Ringach, et al. 1999; Sillito and Jones 1996; Walker, Ohzawa, et al. 1999), and is also prominent for natural inputs (Haider, Krause, et al. 2010; Vinje and Gallant 2000). Past work using the GSM and its extensions has explained a wide range of those phenomena, as reflecting a computation optimized to the statistics of natural images (Coen-Cagli, Dayan, et al. 2012; Coen-cagli, Dayan, et al. 2009; Schwartz and Simoncelli 2001). The modeling and experimental results presented here are consistent with this prior work, as we report strong and tuned surround suppression of mean spike-counts (Fig. 2D-F, Fig. 3). But our findings go beyond this previous work, by establishing a general relation between response variability and natural image statistics (Fig. 1C) and relating surround influences on mean spike counts and on variability (Fig. 2D-F, Fig. 3).

Our model could be further extended to account for the fact that contextual modulation is weak or absent for some stimuli, such as when contextual inputs are not informative (Coen-Cagli, Kohn, et al. 2015). Variability reduction by stimulus context should be weaker or absent for such uninformative contextual stimuli, which would be consistent with our observations that, when we used natural images, the level of surround suppression of FF varied substantially across images (Fig. 2C and Supplementary Fig. S4), and that suppression was also weaker for orthogonal grating surrounds (Fig. 3C). Although V1 responses agreed well with model predictions, we observed a quantitative discrepancy between the two: contextual modulation of FF and its tuning were much stronger in the model. This could reflect that, in the model, the main source of uncertainty (particularly for the high-contrast stimuli we used), and therefore variability, is the unknown value of the global modulator. Model response variability is therefore extremely sensitive to contextual stimuli. In V1, there are likely multiple latent sources of uncertainty that could partly mask the effects of our experimental manipulation of spatial context. Addressing this discrepancy may require considering non-sensory contextual factors such as attention and behavioral state (Haefner, Berkes, et al. 2016; Rabinowitz, Goris, et al. 2015).

3.2 Influences of divisive normalization on variability and other response statistics

Our mathematical analysis of the GSM inference shows that, in the model, response strength and variability are jointly modulated by divisive normalization (Carandini and Heeger 2012; Heeger 1992). This is because the mean and variance of the inferred distribution of the local features depend divisively on the inferred value of the global modulator (Methods, Eq. 3), which in turn is obtained by combining the inputs corresponding to all features (Methods Eq. 2; Coen-Cagli, Dayan, et al. 2012). Therefore, our model points to divisive normalization as the key operation for surround modulation of rate and variability. There is abundant indirect evidence that normalization modulates responses beyond firing rate. For instance, stimulus manipulations that engage normalization, such as varying contrast and size (Cavanaugh, Bair, et al. 2002b; Heeger 1992), also modulate variability (Haider, Krause, et al. 2010; Kohn and Smith 2005; Snyder, Morais, et al. 2014). In addition, although the mechanisms of normalization are debated, network models based on inhibitory stabilization (Y. Ahmadian, D. B. Rubin, et al. 2013) reproduce many of those stimulus-induced effects, indicating a common mechanism that could control both firing rate (Hennequin, Yashar Ahmadian, et al. 2018; Daniel B. Rubin, Hooser, et al. 2015) and variability (Hennequin, Yashar Ahmadian, et al. 2018) consistently with normalization.

Other work has established the connection between normalization and variability more directly. A descriptive model of stochastic normalization has been shown to fit changes in variability with stimulus contrast (Coen-Cagli and Solomon 2019) and orientation noise (Henaff, Boudry-Singer, et al. 2020), and revealed that, even for fixed stimuli, variability is reduced during epochs of strong normalization (Coen-Cagli and Solomon 2019). Our analytical results on normalization and variability bridge the gap between this literature and a theory of the computational role of variability.

3.3 Relation to other descriptive models and functional explanations of cortical variability

Previous work used a GSM to demonstrate stimulus dependent changes in response statistics (Orbán, Berkes, et al. 2016). In particular, Orbán, Berkes, et al. (2016) suggested that a GSM could unify effects of response mean and variability. Our work extends this study in two important aspects. First, Orbán and colleagues used approximate inference in their GSM, based on a maximum a posteriori estimate for the global scaling variable. Consequently, posterior variance was exclusively due to observation noise, while variance resulting from uncertainty in the global scaling variable was ignored. This required tuning a nonlinear conversion from membrane potential to spike counts to account for realistic response variability (Carandini 2004). Here, we include both sources of uncertainty — input noise and the unknown global latent variable — and we show that the GSM framework is sufficient to capture the dependence

between response mean and variance for a wide range of inputs (Fig. 1C,D), without further tuning the conversion from membrane potential to spikes. Second, the treatment of Orbán et al was sufficient for a coarse grained account of contextual effects (such as changes in sparseness and reliability), but our analysis unveils a more complex repertoire of contextual effects for natural images, leading to detailed predictions that related statistical dependencies across visual space to contextual modulation of V1 variability.

Another recent model (Henaff, Boundy-Singer, et al. 2020) proposes that uncertainty is represented in the response variability, and is thus related to sampling and to our work. However, they propose that variability is partitioned into two terms, Poisson variability and fluctuations in response gain (Goris, Movshon, et al. 2014). Uncertainty is encoded specifically by the amplitude of the gain fluctuations. Different from our work, the Poisson term in that framework does not have a functional role and is left unexplained, and there is no precise relation between V1 variability and the statistics of natural images. In addition, whereas sampling-based representations can approximate the full posterior distribution, the model of Henaff, Boundy-Singer, et al. (2020) focuses only on the mean and variance (uncertainty) of the posterior. Therefore, future experimental work could further distinguish between these theories by comparing higher-order statistics of V1 responses to the corresponding statistics in the visual inputs.

4 Methods

4.1 Model of V1 responses

The Gaussian Scale Mixture (GSM) generative model

The observable variables are given by the outputs of linear, oriented filters (Simoncelli and Freeman 1995) applied to grayscale input images. We assume oriented filters because they approximate well those optimized to natural images, and also represent a canonical choice for V1 models that used the GSM framework (Coen-Cagli, Dayan, et al. 2012; Coen-cagli, Dayan, et al. 2009; Coen-Cagli, Kohn, et al. 2015; Echeveste, L. Aitchison, et al. 2020; Orbán, Berkes, et al. 2016). One pair of filters (even and odd phase, forming a quadrature pair) represents the RF of the model neuron, and another 8 pairs are uniformly distributed on a circle around the RF, all with the same orientation (represented in Fig. 2A as vertical). The surround filters slightly overlap with the RF filters, to reflect that suppressive surround mechanisms in V1 partly overlap with the RF (Cavanaugh, Bair, et al. 2002a) (see Fig. 2A). The responses of the 18 filters form a 18-dimensional input vector, denoted as $\mathbf{x}$.

The generative model uses latent variables to capture the statistics of $\mathbf{x}$, as follows:

$$\mathbf{x} = \nu \mathbf{g} + \boldsymbol{\eta} \quad (1)$$

$$\mathbf{g} \sim \mathcal{N}(0, C_g) \quad ; \quad \nu \sim \text{Rayleigh}(1) \quad ; \quad \boldsymbol{\eta} \sim \mathcal{N}(0, C_{\text{noise}})$$

The observable $\mathbf{x}$ results from the product of the feature vector $\mathbf{g}$, which has the same dimensionality of $\mathbf{x}$, and a positive scalar ν , that acts as global modulator. The additive noise $\boldsymbol{\eta}$ plays the role of observation noise in the generative model. That is, it accounts for the fact that the GSM is not a perfect model of the statistics of the observable $\mathbf{x}$ on natural images. As we explain below, this additive noise is also helpful to account for realistic response variability with weak stimuli (Supplementary Fig. S10). We assume that $\mathbf{g}$ and $\boldsymbol{\eta}$ are generated from multivariate normal distributions, with mean 0 and covariances C_g and C_{noise} , respectively; ν follows a Rayleigh distribution with mean 1. Note that changing the Rayleigh parameter is equivalent to rescaling C_g .

Model optimization

The covariance of the noise term, denoted as C_{noise} in Eq. 1, is found numerically, by applying the filters to 10,000 white-noise patches. We take the empirical covariance of the resulting outputs and scale it by a free parameter, set heuristically at 0.1 to ensure a realistic response variability for weak inputs (Supplementary Fig. S10). The covariance matrix C_g is computed by moment-matching (Doulgeris and Eltoft 2009), based on the empirical covariance of filter outputs over 10,000 natural image patches, scaled by a term that accounts for the mixer. This ensures that the model is adapted to natural image statistics, as in previous work (Coen-Cagli, Dayan, et al. 2012). The image patches used for training are considered noise-free, and the noise level in the trained model is tuned heuristically. This choice was motivated by convenience, and by noticing that pixel noise tended to be small, reflecting the digital quality of images and not indicative of sensory noise.

Probabilistic inference and sampling

Having defined the generative process, we can express the posterior distribution of the latent feature of interest, for example the center-vertical feature with odd spatial phase, g_{1+} , given the filters response $\tilde{\mathbf{x}}$ to a test image. This quantity is denoted $P(g_{1+}|\tilde{\mathbf{x}})$, and results from an operation of Bayesian inference and marginalization over the other latent variables (Supplemental Text, Section 1). In particular, the global modulator plays a key role in the inference of g_{1+} . To gain further insight, we first derived analytical solutions for the regime in which input noise is negligible, i.e. $\boldsymbol{\eta} = 0$. First, can be expressed analytically and

approximated (Supplemental Text, Section 4) as:

$$E[\nu|\tilde{\mathbf{x}}] = \sqrt{\lambda} (1 + \mathcal{O}(\lambda^{-1})) \quad , \quad \text{with} \quad \lambda = \sqrt{\sum_{i,j} (C_g^{-1})_{i,j} \tilde{x}_i \tilde{x}_j} \quad (2)$$

where $\mathcal{O}(\lambda^{-1})$ represents a generic function that drops to zero asymptotically with λ^{-1} . This shows that the estimate of the mixer depends on the outputs of all filters. Second, the distribution of the feature of interest, $P(g_{1+}|\tilde{\mathbf{x}})$, can also be expressed in closed-form in the low-noise limit (Supplemental Text, Section 4). Its mean and Fano factor can be approximated as:

$$E[g_{1+}|\tilde{\mathbf{x}}] = \frac{\tilde{x}_{1+}}{\sqrt{\lambda}} (1 + \mathcal{O}(\lambda^{-1})) \quad , \quad \text{and} \quad \text{FF}[g_{1+}|\tilde{\mathbf{x}}] = \frac{\tilde{x}_{1+}}{4\lambda\sqrt{\lambda}} (1 + \mathcal{O}(\lambda^{-1})) \quad (3)$$

In the approximation above (derived in Supplemental Text, Section 4), the expected value of the feature of interest depends linearly on the input inside the RF, $\tilde{x}_{1+}$. However it is scaled by $\sqrt{\lambda}$, a quantity approximately equal to the expected value of the global modulator (Eq. 2), which includes the influence of the surround. The variance instead scales divisively with the square of λ , which in turn determines the reduction of variability (the FF in Eq. 3) by surround stimulation. This analysis thus shows that, in the GSM inference, divisive normalization influences both the mean and the variance of the posterior distribution, thus providing a normative explanation for the dependence between spike-count variance and mean observed in sensory neurons. Notice also that the expected value and the FF are not always monotonically related, because λ depends both on inputs inside and outside the RF, and appears with different exponents in the FF and expected value. For instance, surround stimulation affects only λ and thus changes the FF and expected value in the same direction, whereas changing contrast affects both numerator and denominator resulting in opposite scaling of the expected value and FF.

The analytical results in Eq. 3 refer to the reduced model without additive noise. In this formulation, for very small inputs $\tilde{\mathbf{x}} \approx 0$ the inferred mean and variance converge to zero, resulting in model neurons with an unrealistically silent and very stable baseline activity. We therefore extended the generative model to non-zero additive noise, and determined the model neuron responses numerically, by Monte Carlo sampling, implemented through the Stan programming language (<https://mc-stan.org/>). When comparing the analytical solution for the noiseless model with the simulation results for the full model, we found that, as expected, they differ predominantly in the regime of small inputs, where the model with noise still preserves a non-zero response and variability (Supplementary Fig. S10).

Our choice of a fixed Rayleigh prior for the mixer (in line with past work; Coen-Cagli, Dayan, et al. 2012; Coen-cagli, Dayan, et al. 2009; Coen-Cagli, Kohn, et al. 2015; Schwartz, Sejnowski, et al. 2006) is mainly due to mathematical convenience, as it allows us to obtain analytical insights on the scaling of mean and variance with $\tilde{x}_{1+}$ and λ . Although we focused here on qualitative predictions, for quantitative fits of GSM models to neural data one could leverage the flexibility afforded by modifying the mixer prior and introducing additional free parameters.

529 Conversion to spike counts

530 For the purpose of our analysis, in Eq. 3 $\tilde{x}_{1+}$ is assumed greater than 0 (e.g. a grating stimulus
531 in-phase with the filter). To cover the general case, and appropriately express neural response
532 and FF in terms of spike counts, we performed the following transformation:

$$533 \quad r = c \sqrt{g_{1+}^2 + g_{1-}^2} \quad (4)$$

534 where c is a fixed parameter set heuristically so that mean responses and FF are in a realistic
535 range (values are reported in the figure captions), and the $\pm$ represent the two spatial phases
536 at the RF position. One strength of this framework (following Orbán, Berkes, et al. 2016)
537 is that it is a fully normative model of response variability, and does not need to assume
538 additional noise in the spiking process. We can therefore directly consider the instantaneous
539 response r as a spike count, with a rounding error that is small for sufficiently high c . In
540 the no-noise approximation, the mean and variance of r can be expressed analytically, and
541 preserve the behavior of Eq. 3 (see Supplemental Text, Section 5). For the full model, we
542 compute a single-trial response r for each sample of g_{1+} , g_{1-} . The mean, variance and FF of
543 the model neuron are then found numerically, using 400 samples.

544 The simple form of Eq. 4 allows for analytical results that provide useful intuitions. However,
545 when testing the GSM response to stimuli of fixed size, we found that an increase in contrast led
546 to a decrease in variance, in conflict with V1 data (Supplementary Fig. S11A,B). This behavior
547 can be easily corrected (Supplementary Fig. S11E,F) by using a different transformation
548 between the latent variable $\mathbf{g}$ and the neural response r , in the form of a rectified expansive
549 nonlinearity (Orbán, Berkes, et al. 2016). Note however that the GSM predictions for size
550 tuning and surround-orientation tuning stimuli are qualitatively robust to the specific choice
551 of transformation (Supplementary Fig. S12).

552 4.2 Neurophysiology

553 Animal preparation and data collection

554 We recorded data from male adult macaque monkeys (*Macaca fascicularis*), either anes-
555 thetized (3 animals) or awake (2 animals). The protocol and general methods employed for
556 the anesthetized experiments have been described previously (Smith and Kohn 2008). In
557 short, anesthesia was induced with ketamine (10 mg/kg of body weight) and maintained
558 during surgery with isoflurane (1.5–2.5% in 95% O₂), switching to sufentanil (6–18 μg /kg per
559 h, adjusted as needed) during recordings. Eye movements were reduced using vecuronium
560 bromide (0.15 mg/kg per h). Temperature was maintained in the 36–37 C° range, and relevant
561 vital signs (EEG, ECG, blood pressure, end-tidal PCO₂, temperature, and airway pressure)

were monitored continuously to ensure sufficient level of anesthesia and well-being. We implanted a 10×10 multielectrode array (400 spacing, 1-mm length) in V1.

For awake experiments the animal was first familiarized with a restraining chair (Crist Instruments). Then a titanium headpost was implanted under full isoflurane anesthesia in an aseptic environment. Postoperative analgesic (buprenorphine) and antibiotic (enrofloxacin) were provided. After a six week recovery period, the animal was trained to fixate in a $1 \text{ deg} \times 1 \text{ deg}$ window. Eye position was monitored with a high-speed infrared camera (Eyelink, 1000 Hz). Once sufficient performance was reached, a second surgery was performed in which a craniotomy and durotomy were performed over the occipital cortex. A 96-channel and a 48-channel microelectrode array were implanted in V1 (and a third, 48-channel array in V4, not considered here). The dura was sutured over the arrays and covered with a gelatin film (Duragen). The craniotomy was covered with titanium mesh, held in place with titanium screws. On the first day of recording we mapped the spatial receptive fields of the sampled neurons by presenting small patches of drifting full contrast gratings (0.5 deg diameter; 4 orientations, 1 cycle/deg , 3 Hz drift rate, 250 ms presentation) at 25 distinct positions spanning a $3 \text{ deg} \times 3 \text{ deg}$ region of visual space. Subsequent stimuli were centered in the aggregate RF of the recorded units.

All procedures were approved by the Albert Einstein College of Medicine and followed the guidelines in the United States Public Health Service Guide for the Care and Use of Laboratory Animals.

4.2.1 Visual stimuli

Visual stimuli were generated with custom software (EXPO) and displayed on a cathode ray tube monitor (Hewlett Packard p1230; 1024×768 pixels, with cd/mean luminance and 100 Hz frame rate) viewed at a distance of 110 cm (for anesthetized) or 60 cm (for awake). In each session, stimuli were randomly interleaved, separated by a uniform gray screen (blank stimulus). All grating stimuli were presented at 100% contrast.

Surround modulation experiments. We measured surround modulation in anesthetized animals with grayscale natural images (as described in Coen-Cagli, Kohn, et al. 2015). Briefly, we presented 270 images in total, each at two sizes (1 degree and $3.1\text{-}6.7 \text{ degrees}$). These included 90 distinct images. For images with a dominant orientation, we presented four variants rotated in steps of 45 degrees , to increase the probability that each variant would drive at least some of the recorded neurons. Images were presented for 200 ms followed by 100 ms blank screen in pseudo-random order, each repeated 20 times.

Size tuning experiments. We measured size tuning with grayscale natural images, and both static and drifting gratings (Table 1 and Supplement Fig. S6). In each session of the awake experiments, we presented 10 natural images (a subset of the 270 described above) masked by

a circular window with diameters of 0.34, 0.55, 0.90, 2.4 and 3.8 deg, with stimulus duration 200 ms and inter-stimulus interval of 100 ms. Images were presented 60-74 times each. We chose images that evoked strong average responses in a majority of the neurons reported in Coen-Cagli, Kohn, et al. (2015). In separate sessions, we measured size tuning with static circular gratings, with diameters of 0.34, 0.55, 0.90, 2.4 and 3.8 deg; orientations of 0, 45, 90 and 135 deg; duration of 250 ms, and inter-stimulus duration of 100 ms. We set the spatial frequency (1 cycle/deg) to be appropriate for V1 neurons at the recorded eccentricity. Each stimulus was repeated 114-124 times. In the anesthetized experiments, we measured size tuning with static circular gratings, testing a larger range of conditions (diameters of 0.34, 0.55, 0.90, 1.5, 2.4, 3.8 and 6.2 deg; orientations of 0, 45, 90 and 135 deg), and repeated each stimulus 20 times.

Surround orientation tuning experiments. We measured orientation tuning of surround modulation in two awake monkeys, using static compound gratings with a spatial frequency of 1 cyc/deg presented for 200 ms (100 ms inter-stimulus interval). For monkey M we used a central grating of diameter 1 deg, orientations of 0 and 90 deg, and an annular surround with inner diameter of 1 deg and outer diameter of 6 deg, with orientation either matched or orthogonal to the center. For monkey C, the central grating was 0.5 deg in diameter; orientations were 0, 45, 90 and 135 deg; and a surround with inner diameter of 1.5 deg and outer diameter of 5 deg, with orientation either matched or orthogonal to the center. We introduced this gap between center and surround stimulus, to reduce the extent to which the surround stimulus encroached on the neurons' RFs. The results for both monkeys were qualitatively similar. Therefore, we combined them in our analyses.

4.3 Data Analysis

For each electrode, we extracted waveform signals (sampled at 30 kHz) whenever the extracellular voltage exceeded a threshold of 5 times the square root of the mean square signal on each channel. We then sorted waveforms manually using the Plexon Offline Sorter, and isolated both single and multi-unit clusters, here both referred to as neurons. Data analysis was then performed in Julia (<https://julialang.org>).

Characterization of neuronal responses and inclusion criteria. We computed spike counts in a fixed window with length equal to the stimulus duration, shifted by 50 ms after stimulus onset. We also computed baseline activity in the 50 ms window from 20 ms before to 30 ms after stimulus onset. We excluded from further analyses all neurons that were not driven by any stimulus above baseline + 1 std. We also excluded all natural images and grating orientations that, when presented at a size closest to 1 degree (out of those presented), did not drive the neurons above the baseline + 1std. Next, we defined the response latency of each neuron as the first time at which the peristimulus time histograms (regularized using a smoothing cubic spline with parameter $2 \cdot 10^{-6}$) at the preferred stimulus size (for size-tuning

experiments) or at the smallest size presented (0.5 or 1 degree, for the other experiments) crossed a threshold of baseline + 1 std. All further analyses were performed on spike counts in windows shifted by the latency of each individual neuron. In the surround modulation experiments on anesthetized monkeys with natural image patches (Fig. 2C) we selected only neurons that responded significantly to at least 10 distinct images.

We computed the mean spike-count by averaging across trials, and characterized variability by the Fano factor (FF), the ratio between across-trial variance and mean of the spike count. We focused on the FF because, when compared across conditions, it quantifies changes in variability beyond the changes in mean activity. We excluded neurons whose average FF across all stimulus conditions was larger than 2.

Because we were interested in surround modulation of variability, we excluded neurons with RFs not well centered on the stimuli. In experiments with anesthetized animals, we measured multi-unit spatial RFs using small circular oriented gratings (size 0.5 deg, 4 orientations, 250 ms presentation), fitting the spike-counts with a two-dimensional Gaussian. We only kept for further analysis those neurons whose RF center was within 0.4 deg of the stimulus center. Due to the limited duration of the awake sessions, we could not measure spatial RFs prior to each session. We therefore relied only on the responsivity to small stimuli (described above), and on the following additional criteria (for size-tuning experiments, Fig. 2F and Table 1), as a proxy for appropriate stimulus centering. First, we excluded the neurons that had maximum response for very small (0.3 deg) or very large (> 4 deg) stimuli, because this was indicative of poor centering. Second, we excluded natural images that elicited weak surround suppression of the mean spike count (below 15%). We verified that our results did not change qualitatively when we changed this threshold (Supplementary Fig. S9).

Lastly, we excluded the neurons whose mean spike count was zero for any given stimulus size (for size tuning experiments) or surround condition (for surround orientation tuning experiments), because the FF is not defined in those cases. For the surround orientation tuning experiments (Fig. 3) we analyzed only the preferred orientation out of those presented, to ensure responses were robust enough that we could measure surround suppression effects reliably.

Statistical analysis. In the size tuning experiments (Fig. 2E-F and Table 1) we first computed mean spike-count and FF for each neuron, each stimulus size and condition (natural image identity or grating orientation). We then averaged across conditions, using mean for spike-counts and geometric mean for FFs, obtaining an area summation curve for both spike count and FF for each neuron (e.g. Fig. 2E). The differences in FF across sizes were measured as:

$$\% \text{difference in FF} = 100 \cdot \frac{(\text{FF}_{\alpha} - \text{FF}_{\beta})}{(\text{FF}_{\alpha} + \text{FF}_{\beta}) / 2} \quad (5)$$

Where α refers to the stimulus size closer to RF and β to the stimulus size approximately

twice the RF. To visualize population averages in Fig. 2F and Supplementary Fig. S6-S7, we expressed stimulus size relative to RF size, and then averaged across neurons for each relative size. Note that some of the relative sizes were available only in a subpopulation with a specific RF size. In those points, averages refer to the available neurons. Supplementary Fig. S7 shows instead the groups as separate. For the surround orientation tuning experiments, we quantified differences in FF also by Eq. 5, but with α representing the stimulus with orthogonal surround, and β the stimulus with matching surround. Confidence intervals in the population plots were estimated by bootstrapping.

For the mean-matching tests in the area-summation experiment (Table 1 and Supplement, Fig. S3, we compared the FF between stimuli that were smaller versus larger than the RF, and elicited a similar trial-averaged spike count. Specifically, we pooled the mean spike counts of all neurons and stimuli smaller than the RF in one group, and all neurons and stimuli larger than the RF in a second group. We then subsampled the same number of cases from each group, so as to obtain identical histograms of mean spike counts. Lastly, we compared the FF distributions of the two groups. The p -values in Table 1 were computed using a paired sample t-test of the null hypothesis that differences between samples from the two conditions (i.e. RF size versus $2 \times$ RF size) had mean ≤ 0 .

Acknowledgments

We thank the Coen-Cagli and Kohn labs for helpful discussion, the Kohn lab for help with experiments, and Gergo Orbán for discussion and comments on a previous version of the manuscript. This work was supported by NIH grants EY030578 and EY021371.

Author Contributions

DF, AK and RCC designed the project; DF and RCC developed the theory and models; AA and AK performed the experiments; DF and AD analyzed the data; DF, AK and RCC wrote the paper.

Data Availability

Data for Fig. 2C and Supplementary Fig. S6B S6B are publicly available on the CRCNS data sharing site <https://crcns.org>. Data for the other figures will be made available by the authors upon request.

Code Availability

Code for model simulations and data analysis will be available without restrictions on GitHub (<https://github.com/rubencoencagli/festa-et-al-2020>).

Materials & Correspondence

Correspondence should be addressed to Ruben Coen-Cagli at ruben.coen-cagli@einsteinmed.org.

Competing Interests

The authors have no competing interests.

References

- Ahmadian, Y., D. B. Rubin, and K. D. Miller (2013). “Analysis of the stabilized supralinear network”. In: *Neural Comput* 25.8, pp. 1994–2037. DOI: [10.1162/NECO_a_00472](https://doi.org/10.1162/NECO_a_00472).
- Albright, Thomas D. and Gene R. Stoner (2002). “Contextual Influences on Visual Processing”. In: *Annual Review of Neuroscience* 25.1, pp. 339–379. DOI: [10.1146/annurev.neuro.25.112701.142900](https://doi.org/10.1146/annurev.neuro.25.112701.142900).
- Angelucci, Alessandra et al. (2002). “Circuits for Local and Global Signal Integration in Primary Visual Cortex”. In: *Journal of Neuroscience* 22.19, pp. 8633–8646. DOI: [10.1523/JNEUROSCI.22-19-08633.2002](https://doi.org/10.1523/JNEUROSCI.22-19-08633.2002).
- Angelucci, Alessandra et al. (2017). “Circuits and Mechanisms for Surround Modulation in Visual Cortex”. In: *Annu Rev Neurosci* 40, pp. 425–451. DOI: [10.1146/annurev-neuro-072116-031418](https://doi.org/10.1146/annurev-neuro-072116-031418).
- Arbelaez, P., M. Maire, C. Fowlkes, and J. Malik (2011). “Contour detection and hierarchical image segmentation”. In: *IEEE Trans Pattern Anal Mach Intell* 33.5, pp. 898–916. DOI: [10.1109/TPAMI.2010.161](https://doi.org/10.1109/TPAMI.2010.161).
- Beck, Jeffrey M., Peter E. Latham, and Alexandre Pouget (2011). “Marginalization in neural circuits with divisive normalization”. In: *Journal of Neuroscience* 31.43, pp. 15310–15319.
- Beck, Jeffrey M., Wei Ji Ma, Xaq Pitkow, Peter E. Latham, and Alexandre Pouget (2012). “Not noisy, just wrong: the role of suboptimal inference in behavioral variability”. In: *Neuron* 74.1, pp. 30–9. DOI: [10.1016/j.neuron.2012.03.016](https://doi.org/10.1016/j.neuron.2012.03.016).
- Bell, A. J. and T. J. Sejnowski (1997). “The ”independent components” of natural scenes are edge filters”. In: *Vision Res* 37.23, pp. 3327–38. DOI: [10.1016/s0042-6989\(97\)00121-1](https://doi.org/10.1016/s0042-6989(97)00121-1).

- 729 Berkes, Pietro, Gergő Orbán, Máté Lengyel, and József Fiser (2011). “Spontaneous Cortical
730 Activity Reveals Hallmarks of an Optimal Internal Model of the Environment”. In: *Science*
731 331.6013, pp. 83–87.
- 732 Bishop, Christopher M. (2006). *Pattern recognition and machine learning*. Information science
733 and statistics. New York: Springer, p. 738.
- 734 Boerlin, Martin, Christian K. Machens, and Sophie Denève (2013). “Predictive Coding of
735 Dynamical Variables in Balanced Spiking Networks”. In: *PLoS Comput Biol* 9.11, e1003258.
736 DOI: [10.1371/journal.pcbi.1003258](https://doi.org/10.1371/journal.pcbi.1003258).
- 737 Britten, K. H., Michael N. Shadlen, William T. Newsome, and J. A. Movshon (1992). “The
738 analysis of visual motion: a comparison of neuronal and psychophysical performance”. In:
739 *J Neurosci* 12.12, pp. 4745–65.
- 740 Carandini, Matteo (2004). “Amplification of Trial-to-Trial Response Variability by Neurons in
741 Visual Cortex”. In: *PLOS Biology* 2.9, e264. DOI: [10.1371/journal.pbio.0020264](https://doi.org/10.1371/journal.pbio.0020264).
- 742 Carandini, Matteo and David J. Heeger (2012). “Normalization as a canonical neural computa-
743 tion”. In: *Nat Rev Neurosci* 13.1.
- 744 Cavanaugh, James R., Wyeth Bair, and J. Anthony Movshon (2002a). “Nature and Interaction
745 of Signals From the Receptive Field Center and Surround in Macaque V1 Neurons”. In:
746 *Journal of Neurophysiology* 88.5.
- 747 – (2002b). “Selectivity and Spatial Distribution of Signals From the Receptive Field Surround
748 in Macaque V1 Neurons”. In: *Journal of Neurophysiology* 88.5, pp. 2547–2556. DOI: [10.1152/
749 jn.00693.2001](https://doi.org/10.1152/jn.00693.2001).
- 750 Charles, Adam S., Mijung Park, J. Patrick Weller, Gregory D. Horwitz, and Jonathan W. Pillow
751 (2018). “Dethroning the Fano Factor: A Flexible, Model-Based Approach to Partitioning
752 Neural Variability”. In: *Neural Computation* 30.4, pp. 1012–1045. DOI: [10.1162/neco_a_
753 01062](https://doi.org/10.1162/neco_a_01062).
- 754 Chen, Minggui et al. (2014). “Incremental Integration of Global Contours through Interplay
755 between Visual Cortical Areas”. In: *Neuron* 82.3, pp. 682–694. DOI: [10.1016/j.neuron.
756 2014.03.023](https://doi.org/10.1016/j.neuron.2014.03.023).
- 757 Churchland, Mark M. et al. (2010). “Stimulus onset quenches neural variability: a widespread
758 cortical phenomenon”. In: *Nat Neurosci* 13.3, pp. 369–378. DOI: [10.1038/nn.2501](https://doi.org/10.1038/nn.2501).
- 759 Coen-Cagli, Ruben, Peter Dayan, and Odelia Schwartz (2012). “Cortical Surround Interactions
760 and Perceptual Salience via Natural Scene Statistics”. In: *PLOS Comput Biol* 8.3.
- 761 Coen-cagli, Ruben, Peter Dayan, and Odelia Schwartz (2009). “Statistical Models of Linear
762 and Nonlinear Contextual Interactions in Early Visual Processing”. In: *Advances in Neural
763 Information Processing Systems* 22.
- 764 Coen-Cagli, Ruben, Adam Kohn, and Odelia Schwartz (Nov. 2015). “Flexible gating of con-
765 textual influences in natural vision”. en. In: *Nature Neuroscience* 18.11, pp. 1648–1655. DOI:
766 [10.1038/nn.4128](https://doi.org/10.1038/nn.4128).
- 767 Coen-Cagli, Ruben and Selina S. Solomon (2019). “Relating Divisive Normalization to Neuronal
768 Response Variability”. In: *J Neurosci* 39.37, pp. 7344–7356. DOI: [10.1523/JNEUROSCI.0126-
769 19.2019](https://doi.org/10.1523/JNEUROSCI.0126-19.2019).

- Cohen, Marlene R. and John H. R. Maunsell (2009). "Attention improves performance primarily by reducing interneuronal correlations". In: *Nature Neuroscience* 12.12, pp. 1594–1600. DOI: [10.1038/nn.2439](https://doi.org/10.1038/nn.2439).
- Dadarlat, Maria C. and Michael P. Stryker (2017). "Locomotion Enhances Neural Encoding of Visual Stimuli in Mouse V1". In: *Journal of Neuroscience* 37.14, pp. 3764–3775. DOI: [10.1523/JNEUROSCI.2728-16.2017](https://doi.org/10.1523/JNEUROSCI.2728-16.2017).
- Dayan, Peter and Larry Abbott (2005). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*. The MIT Press.
- Dinstein, Ilan, David J. Heeger, and Marlene Behrmann (2015). "Neural variability: friend or foe?" In: *Trends in Cognitive Sciences* 19.6, pp. 322–328.
- Doulgeris, Anthony P and Torbjørn Eltoft (2009). "Scale mixture of Gaussian modelling of polarimetric SAR data". In: *EURASIP Journal on Advances in Signal Processing* 2010.1, p. 874592.
- Echeveste, Rodrigo, L. Aitchison, Guillaume Hennequin, and Máté Lengyel (2020). "Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference". In: *Nat Neurosci* 23.9, pp. 1138–1149. DOI: [10.1038/s41593-020-0671-1](https://doi.org/10.1038/s41593-020-0671-1).
- Finn, I. M., N. J. Priebe, and D. Ferster (2007). "The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex". In: *Neuron* 54.1, pp. 137–52. DOI: [10.1016/j.neuron.2007.02.029](https://doi.org/10.1016/j.neuron.2007.02.029).
- Fiser, József, Pietro Berkes, Gergő Orbán, and Máté Lengyel (2010). "Statistically optimal perception and learning: from behavior to neural representations". In: *Trends in Cognitive Sciences* 14.3, pp. 119–130.
- Gao, Dashan and Nuno Vasconcelos (2007). *Bottom-up saliency is a discriminant process*, pp. 1–6. DOI: [10.1109/ICCV.2007.4408851](https://doi.org/10.1109/ICCV.2007.4408851).
- Gardner, Justin L., Akiyuki Anzai, Izumi Ohzawa, and Ralph D. Freeman (1999). "Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex". In: *Visual Neuroscience* 16.6, pp. 1115–1121. DOI: [10.1017/S0952523899166112](https://doi.org/10.1017/S0952523899166112).
- Geisler, Wilson S. (2008). "Visual Perception and the Statistical Properties of Natural Scenes". In: 59.1, pp. 167–192. DOI: [10.1146/annurev.psych.58.110405.085632](https://doi.org/10.1146/annurev.psych.58.110405.085632).
- Goris, Robbe L. T., J. Anthony Movshon, and Eero P. Simoncelli (2014). "Partitioning neuronal variability". In: *Nature neuroscience* 17.6, p. 858.
- Haefner, Ralf M., Pietro Berkes, and József Fiser (2016). "Perceptual Decision-Making as Probabilistic Inference by Neural Sampling". In: *Neuron* 90.3, pp. 649–660. DOI: [10.1016/j.neuron.2016.03.020](https://doi.org/10.1016/j.neuron.2016.03.020).
- Haider, Bilal et al. (2010). "Synaptic and Network Mechanisms of Sparse and Reliable Visual Cortical Activity during Nonclassical Receptive Field Stimulation". In: *Neuron* 65.1, pp. 107–121. DOI: [10.1016/j.neuron.2009.12.005](https://doi.org/10.1016/j.neuron.2009.12.005).
- Heeger, David J. (1992). "Normalization of cell responses in cat striate cortex". In: *Visual Neuroscience* 9.2, pp. 181–197. DOI: [10.1017/s0952523800009640](https://doi.org/10.1017/s0952523800009640).
- Henaff, O. J., Z. M. Boundy-Singer, K. Meding, C. M. Ziemba, and Robbe L. T. Goris (2020). "Representation of visual uncertainty through neural gain variability". In: *Nat Commun* 11.1, p. 2513. DOI: [10.1038/s41467-020-15533-0](https://doi.org/10.1038/s41467-020-15533-0).

- 811 Hennequin, Guillaume, Yashar Ahmadian, Daniel B. Rubin, Máté Lengyel, and Kenneth D.
812 Miller (2018). “The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single
813 Stimulus-Tuned Attractor Account for Patterns of Noise Variability”. In: *Neuron* 98.4, 846–
814 860.e5. DOI: [10.1016/j.neuron.2018.04.017](https://doi.org/10.1016/j.neuron.2018.04.017).
- 815 Hennequin, Guillaume, Laurence Aitchison, and Máté Lengyel (2014). “Fast Sampling-Based
816 Inference in Balanced Neuronal Networks”. In: *Advances in Neural Information Processing*
817 *Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger.
818 Vol. 27. Curran Associates, Inc.
- 819 Hoyer, P. O. and Aapo Hyvärinen (2003). “Interpreting neural response variability as Monte
820 Carlo sampling of the posterior”. In: pp. 293–300.
- 821 Hunsberger, E., M. Scott, and C. Eliasmith (2014). “The competing benefits of noise and
822 heterogeneity in neural coding”. In: *Neural Comput* 26.8, pp. 1600–23. DOI: [10.1162/NECO_a_00621](https://doi.org/10.1162/NECO_a_00621).
- 823
- 824 Hyvärinen, Aapo and P. O. Hoyer (2000). “Emergence of phase- and shift-invariant features by
825 decomposition of natural images into independent feature subspaces”. In: *Neural Comput*
826 12.7, pp. 1705–20. DOI: [10.1162/089976600300015312](https://doi.org/10.1162/089976600300015312).
- 827 Karklin, Y. and M. S. Lewicki (2009). “Emergence of complex cell properties by learning to
828 generalize in natural scenes”. In: *Nature* 457.7225, pp. 83–6. DOI: [10.1038/nature07481](https://doi.org/10.1038/nature07481).
- 829 Knill, D. C. and Alexandre Pouget (2004). “The Bayesian brain: the role of uncertainty in
830 neural coding and computation”. In: *Trends Neurosci* 27.12, pp. 712–9. DOI: [10.1016/j.tins.2004.10.007](https://doi.org/10.1016/j.tins.2004.10.007).
- 831
- 832 Kohn, Adam, Ruben Coen-Cagli, Ingmar Kanitscheider, and Alexandre Pouget (2016). “Cor-
833 relations and Neuronal Population Information”. In: *Annual Review of Neuroscience* 39.1,
834 pp. 237–256. DOI: [10.1146/annurev-neuro-070815-013851](https://doi.org/10.1146/annurev-neuro-070815-013851).
- 835 Kohn, Adam and Matthew A. Smith (2005). “Stimulus dependence of neuronal correlation in
836 primary visual cortex of the macaque”. In: *Journal of Neuroscience* 25.14, pp. 3661–3673.
- 837 Legenstein, R. and W. Maass (2014). “Ensembles of spiking neurons with noise support optimal
838 probabilistic inference in a dynamically changing environment”. In: *PLoS Comput Biol* 10.10,
839 e1003859. DOI: [10.1371/journal.pcbi.1003859](https://doi.org/10.1371/journal.pcbi.1003859).
- 840 Lochmann, T., U. A. Ernst, and S. Deneve (2012). “Perceptual inference predicts contex-
841 tual modulations of sensory responses”. In: *J Neurosci* 32.12, pp. 4179–95. DOI: [10.1523/JNEUROSCI.0817-11.2012](https://doi.org/10.1523/JNEUROSCI.0817-11.2012).
- 842
- 843 Ma, Wei Ji, Jeffrey M. Beck, Peter E. Latham, and Alexandre Pouget (2006). “Bayesian inference
844 with probabilistic population codes”. In: *Nature Neuroscience* 9.11, pp. 1432–1438. DOI:
845 [10.1038/nn1790](https://doi.org/10.1038/nn1790).
- 846 Mitchell, Jude F., Kristy A. Sundberg, and John H. Reynolds (2009). “Spatial Attention Decor-
847 relates Intrinsic Activity Fluctuations in Macaque Area V4”. In: *Neuron* 63.6, pp. 879–888.
848 DOI: [10.1016/j.neuron.2009.09.013](https://doi.org/10.1016/j.neuron.2009.09.013).
- 849 Moreno-Bote, Ruben et al. (2014). “Information-limiting correlations”. In: *Nat Neurosci* 17.10,
850 pp. 1410–1417.

- Nienborg, H. and B. G. Cumming (2014). “Decision-Related Activity in Sensory Neurons May Depend on the Columnar Architecture of Cerebral Cortex”. In: *Journal of Neuroscience* 34.10, pp. 3579–3585. DOI: [10.1523/jneurosci.2340-13.2014](https://doi.org/10.1523/jneurosci.2340-13.2014).
- Olshausen, B. A. and D. J. Field (1996). “Emergence of simple-cell receptive field properties by learning a sparse code for natural images”. In: *Nature* 381.6583, pp. 607–9. DOI: [10.1038/381607a0](https://doi.org/10.1038/381607a0).
- Orbán, Gergő, Pietro Berkes, József Fiser, and Mátè Lengyel (2016). “Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex”. In: *Neuron* 92.2, pp. 530–543. DOI: <https://doi.org/10.1016/j.neuron.2016.09.038>.
- Ponce-Alvarez, A., A. Thiele, T. D. Albright, G. R. Stoner, and G. Deco (2013). “Stimulus-dependent variability and noise correlations in cortical MT neurons”. In: *Proceedings of the National Academy of Sciences* 110.32, pp. 13162–13167. DOI: [10.1073/pnas.1300098110](https://doi.org/10.1073/pnas.1300098110).
- Pouget, Alexandre, Jeffrey M. Beck, Wei Ji Ma, and Peter E. Latham (2013). “Probabilistic brains: knowns and unknowns”. In: *Nature Neuroscience* 16.9, pp. 1170–1178.
- Pouget, Alexandre, Peter Dayan, and R. S. Zemel (2003). “Inference and computation with population codes”. In: *Annu Rev Neurosci* 26, pp. 381–410. DOI: [10.1146/annurev.neuro.26.041002.131112](https://doi.org/10.1146/annurev.neuro.26.041002.131112).
- Rabinowitz, N. C., Robbe L. T. Goris, M. Cohen, and E. P. Simoncelli (2015). “Attention stabilizes the shared gain of V4 populations”. In: *Elife* 4, e08998. DOI: [10.7554/eLife.08998](https://doi.org/10.7554/eLife.08998).
- Rao, R. P. and D. H. Ballard (1999). “Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects”. In: *Nat Neurosci* 2.1, pp. 79–87. DOI: [10.1038/4580](https://doi.org/10.1038/4580).
- Rubin, Daniel B., Stephen D. Van Hooser, and Kenneth D. Miller (2015). “The Stabilized Supralinear Network: A Unifying Circuit Motif Underlying Multi-Input Integration in Sensory Cortex”. In: *Neuron* 85.2, pp. 402–417. DOI: [10.1016/j.neuron.2014.12.026](https://doi.org/10.1016/j.neuron.2014.12.026).
- Savin, Cristina and Sophie Denève (2014). “Spatio-temporal Representations of Uncertainty in Spiking Neural Networks”. In: *Advances in Neural Information Processing Systems*. Ed. by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger. Vol. 27. Curran Associates, Inc.
- Sceniak, Michael P., Dario L. Ringach, Michael J. Hawken, and Robert Shapley (1999). “Contrast’s effect on spatial summation by macaque V1 neurons”. In: *Nature Neuroscience* 2, p. 733.
- Schwartz, O., T. J. Sejnowski, and Peter Dayan (2006). “Soft mixer assignment in a hierarchical generative model of natural scene statistics”. In: *Neural Comput* 18.11, pp. 2680–718. DOI: [10.1162/neco.2006.18.11.2680](https://doi.org/10.1162/neco.2006.18.11.2680).
- Schwartz, O. and E. P. Simoncelli (2001). “Natural signal statistics and sensory gain control”. In: *Nat Neurosci* 4.8, pp. 819–825.
- Shadlen, Michael N. and William T. Newsome (1998). “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding”. In: *The Journal of Neuroscience* 18.10, pp. 3870–3896. DOI: [10.1523/jneurosci.18-10-03870.1998](https://doi.org/10.1523/jneurosci.18-10-03870.1998).

- 891 Sillito, A. M. and H. E. Jones (1996). "Context-dependent interactions and visual processing in
892 V1". In: *Journal of Physiology-Paris* 90.3, pp. 205–209. DOI: [10.1016/S0928-4257\(97\)81424-](https://doi.org/10.1016/S0928-4257(97)81424-6)
893 [6](https://doi.org/10.1016/S0928-4257(97)81424-6).
- 894 Simoncelli, E. P. and W. T. Freeman (1995). "The steerable pyramid: a flexible architecture
895 for multi-scale derivative computation". In: *Proceedings., International Conference on Image*
896 *Processing*. Vol. 3, 444–447 vol.3. DOI: [10.1109/ICIP.1995.537667](https://doi.org/10.1109/ICIP.1995.537667).
- 897 Smith, Matthew A. and Adam Kohn (2008). "Spatial and temporal scales of neuronal correlation
898 in primary visual cortex". In: *J Neurosci* 28.48, pp. 12591–603. DOI: [10.1523/JNEUROSCI.](https://doi.org/10.1523/JNEUROSCI.2929-08.2008)
899 [2929-08.2008](https://doi.org/10.1523/JNEUROSCI.2929-08.2008).
- 900 Snyder, A. C., M. J. Morais, A. Kohn, and Matthew A. Smith (2014). "Correlations in V1 Are
901 Reduced by Stimulation Outside the Receptive Field". In: *Journal of Neuroscience* 34.34,
902 pp. 11222–11227.
- 903 Solomon, Selina S., Spencer C. Chen, John W. Morley, and Samuel G. Solomon (2014). "Local
904 and Global Correlations between Neurons in the Middle Temporal Area of Primate Visual
905 Cortex". In: *Cerebral Cortex* 25.9, pp. 3182–3196. DOI: [10.1093/cercor/bhu111](https://doi.org/10.1093/cercor/bhu111).
- 906 Sompolinsky, Haim, Hyungsoo Yoon, Kukjin Kang, and Maoz Shamir (2001). "Population
907 coding in neuronal systems with correlated noise". In: *Physical Review E* 64.5, p. 051904.
- 908 Spratling, M. W. (2010). "Predictive coding as a model of response properties in cortical area
909 V1". In: *J Neurosci* 30.9, pp. 3531–43. DOI: [10.1523/JNEUROSCI.4911-09.2010](https://doi.org/10.1523/JNEUROSCI.4911-09.2010).
- 910 Stevenson, Ian H. (2016). "Flexible models for spike count data with both over- and under-
911 dispersion". In: *Journal of Computational Neuroscience* 41.1, pp. 29–43. DOI: [10.1007/s10827-](https://doi.org/10.1007/s10827-016-0603-y)
912 [016-0603-y](https://doi.org/10.1007/s10827-016-0603-y).
- 913 Theis, L., R. Hosseini, and M. Bethge (2012). "Mixtures of conditional Gaussian scale mixtures
914 applied to multiscale image representations". In: *PLoS One* 7.7, e39857. DOI: [10.1371/](https://doi.org/10.1371/journal.pone.0039857)
915 [journal.pone.0039857](https://doi.org/10.1371/journal.pone.0039857).
- 916 Tolhurst, D. J., J. Anthony Movshon, and A. F. Dean (1983). "The statistical reliability of signals
917 in single neurons in cat and monkey visual cortex". In: *Vision research* 23 8, pp. 775–85.
- 918 Tomko, George J. and Donald R. Crapper (1974). "Neuronal variability: non-stationary re-
919 sponses to identical visual stimuli". In: *Brain Research* 79.3, pp. 405–418.
- 920 Verhoef, Bram-Ernst and John H. R. Maunsell (2017). "Attention-related changes in correlated
921 neuronal activity arise from normalization mechanisms". In: *Nature Neuroscience* 20.7.
- 922 Vinje, William E. and Jack L. Gallant (2000). "Sparse Coding and Decorrelation in Primary
923 Visual Cortex During Natural Vision". In: *Science* 287.5456, p. 1273. DOI: [10.1126/science.](https://doi.org/10.1126/science.287.5456.1273)
924 [287.5456.1273](https://doi.org/10.1126/science.287.5456.1273).
- 925 Wainwright, M. J., E. P. Simoncelli, and A. S. Willsky (2000). "Random cascades of Gaussian
926 scale mixtures and their use in modeling natural images with application to denoising". In:
927 vol. 1, 260–263 vol.1. DOI: [10.1109/ICIP.2000.900944](https://doi.org/10.1109/ICIP.2000.900944).
- 928 Walker, Gary A., Izumi Ohzawa, and Ralph D. Freeman (1999). "Asymmetric Suppression
929 Outside the Classical Receptive Field of the Visual Cortex". In: *The Journal of Neuroscience*
930 19.23, pp. 10536–10553. DOI: [10.1523/jneurosci.19-23-10536.1999](https://doi.org/10.1523/jneurosci.19-23-10536.1999).

- 931 Webb, Ben S., Neel T. Dhruv, Samuel G. Solomon, Chris Tailby, and Peter Lennie (2005). “Early
932 and Late Mechanisms of Surround Suppression in Striate Cortex of Macaque”. In: *Journal*
933 *of Neuroscience* 25.50, pp. 11666–11675. DOI: [10.1523/JNEUROSCI.3414-05.2005](https://doi.org/10.1523/JNEUROSCI.3414-05.2005).
- 934 White, B., Larry Abbott, and J. Fiser (2012). “Suppression of cortical neural variability is
935 stimulus- and state-dependent”. In: *J Neurophysiol* 108.9, pp. 2383–92. DOI: [10.1152/jn.](https://doi.org/10.1152/jn.00723.2011)
936 [00723.2011](https://doi.org/10.1152/jn.00723.2011).
- 937 Zhu, M. and C. J. Rozell (2013). “Visual nonclassical receptive field effects emerge from sparse
938 coding in a dynamical system”. In: *PLoS Comput Biol* 9.8, e1003191. DOI: [10.1371/journal.](https://doi.org/10.1371/journal.pcbi.1003191)
939 [pcbi.1003191](https://doi.org/10.1371/journal.pcbi.1003191).