Neuronal variability reflects probabilistic inference tuned to natural image statistics

Dylan Festa[1], Amir Aschner[2], Aida Davila[2], 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 10461, USA

* Corresponding author: Ruben Coen-Cagli

1300 Morris Park Avenue, Price Building Rm 353B

Bronx NY 10461

Email: ruben.coen-cagli@einsteinmed.org

Tel: 718-678-1150

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 Poisson-like variability, 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.

Introduction

In sensory cortex, neuronal activity is typically variable, both in the absence of sensory input and for repeated presentations of a stimulus (Tomko and Crapper 1974, Tolhurst, Movshon et al. 1983). This variability is modulated by several sensory (Kohn and Smith 2005, Churchland, Yu et al. 2010, Ponce-Alvarez, Thiele et al. 2013, Chen, Yan et al. 2014, Goris, Movshon et al. 2014, Solomon, Chen et al. 2014, Rabinowitz, Goris et al. 2015, Orbán, Berkes et al. 2016, Verhoef and Maunsell 2017, Coen-Cagli and Solomon 2019) and non-sensory (Cohen and Maunsell 2009, Mitchell, Sundberg et al. 2009, White, Abbott et al. 2012, Dadarlat and Stryker 2017) 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 (Shadlen and Newsome 1998, Sompolinsky, Yoon et al. 2001, Pouget, Dayan et al. 2003, Fiser, Berkes et al. 2010, Beck, Ma et al. 2012, Dinstein, Heeger et al. 2015).

Parametric descriptive models can capture how stimuli modulate neuronal variability (Goris, Movshon et al. 2014, Stevenson 2016, Charles, Park et al. 2018, Coen-Cagli and Solomon 2019), 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 (Olshausen and Field 1996, Bell and Sejnowski 1997, Schwartz and Simoncelli 2001, Karklin and Lewicki 2009) and probabilistic inference (Ma, Beck et al. 2006, Berkes, Orbán et al. 2011, Lochmann, Ernst et al. 2012, Coen-Cagli, Kohn et al. 2015), 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 (Hoyer and Hyvärinen 2003, Fiser, Berkes et al. 2010, Boerlin, Machens et al. 2013, 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 (Hoyer and Hyvärinen 2003, Fiser, Berkes et al. 2010, Haefner, Berkes et al. 2016). 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), in a manner that depends on stimulus contrast (Kohn and Smith 2005, Finn, Priebe et al. 2007), 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, Wainwright, Simoncelli et al. 2000, Theis, Hosseini et al. 2012). 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 (Schwartz and Simoncelli 2001, Coen-cagli, Dayan et al. 2009, Coen-Cagli, Dayan et al. 2012, Coen-Cagli, Kohn et al. 2015). GSMs can explain the modulation of trial-averaged V1 responses by stimuli in the surround of the receptive field (RF, Sillito and Jones 1996, Sceniak, Ringach et al. 1999, Walker, Ohzawa et al. 1999, Angelucci, Levitt et al. 2002, Cavanaugh, Bair et al. 2002). 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) 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 Poisson-like variability (Tolhurst, Movshon et al. 1983, Shadlen and Newsome 1998, Goris, Movshon et al. 2014) 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.

Results

Poisson-like variability 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 (Wainwright, Simoncelli et al. 2000, Schwartz and Simoncelli 2001). 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 a model of how features—small patches with different orientations and positions (Fig. 1A)—contribute to images (termed a generative model). Given an input image, model neurons then code for 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 (known as Poisson-like response statistics; known as Poisson-like response statistics; Tolhurst, Movshon et al. 1983, Shadlen and Newsome 1998, Goris, Movshon et al. 2014). The GSM assumes $\mathbf{x} = \nu\,\mathbf{g}$ where the sensory input *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 **x** and **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 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 (Figure 1C, dark blue), then by marginalization we would infer that g is with high probability between 5 and 10 (Figure 1C, light blue). If instead ν was inferred to be in the interval 4–5 (Figure 1C, dark brown), then g could only take values between 2 and 2.5, thus shrinking both in mean and variance (Figure 1C, light brown). This example illustrates why a neuron whose responses reflect samples from the inferred distribution of g should display Poisson-like statistics. Notice that, if the relation between *x* and *g* were linear instead of multiplicative, then changes in the inferred value of ν would only change the inferred mean of g, not its variance.

To validate this intuition more rigorously, we considered GSM inference on real images. As in past normative models (Schwartz and Simoncelli 2001, Coen-cagli, Dayan et al. 2009, Coen-Cagli, Dayan et al. 2012, Coen-Cagli, Kohn et al. 2015), we implemented a GSM with oriented filters (Simoncelli and Freeman 1995) spatially arranged to cover the RF of the model neuron and its surround (Fig. 2A; details in Methods). Given an input image, the visual inputs **x** (a vector) were determined by the activations of those filters applied to the image. We denoted by **g** the corresponding local visual features. 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 for stimuli that elicited stronger mean responses, in qualitative agreement with the statistics of V1 neurons (Goris, Movshon et al. 2014).

These analyses confirm the intuition that Poisson-like variability 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 (Wainwright, Simoncelli et al. 2000, Schwartz and Simoncelli 2001), our result establishes a precise link between image statistics and cortical variability.

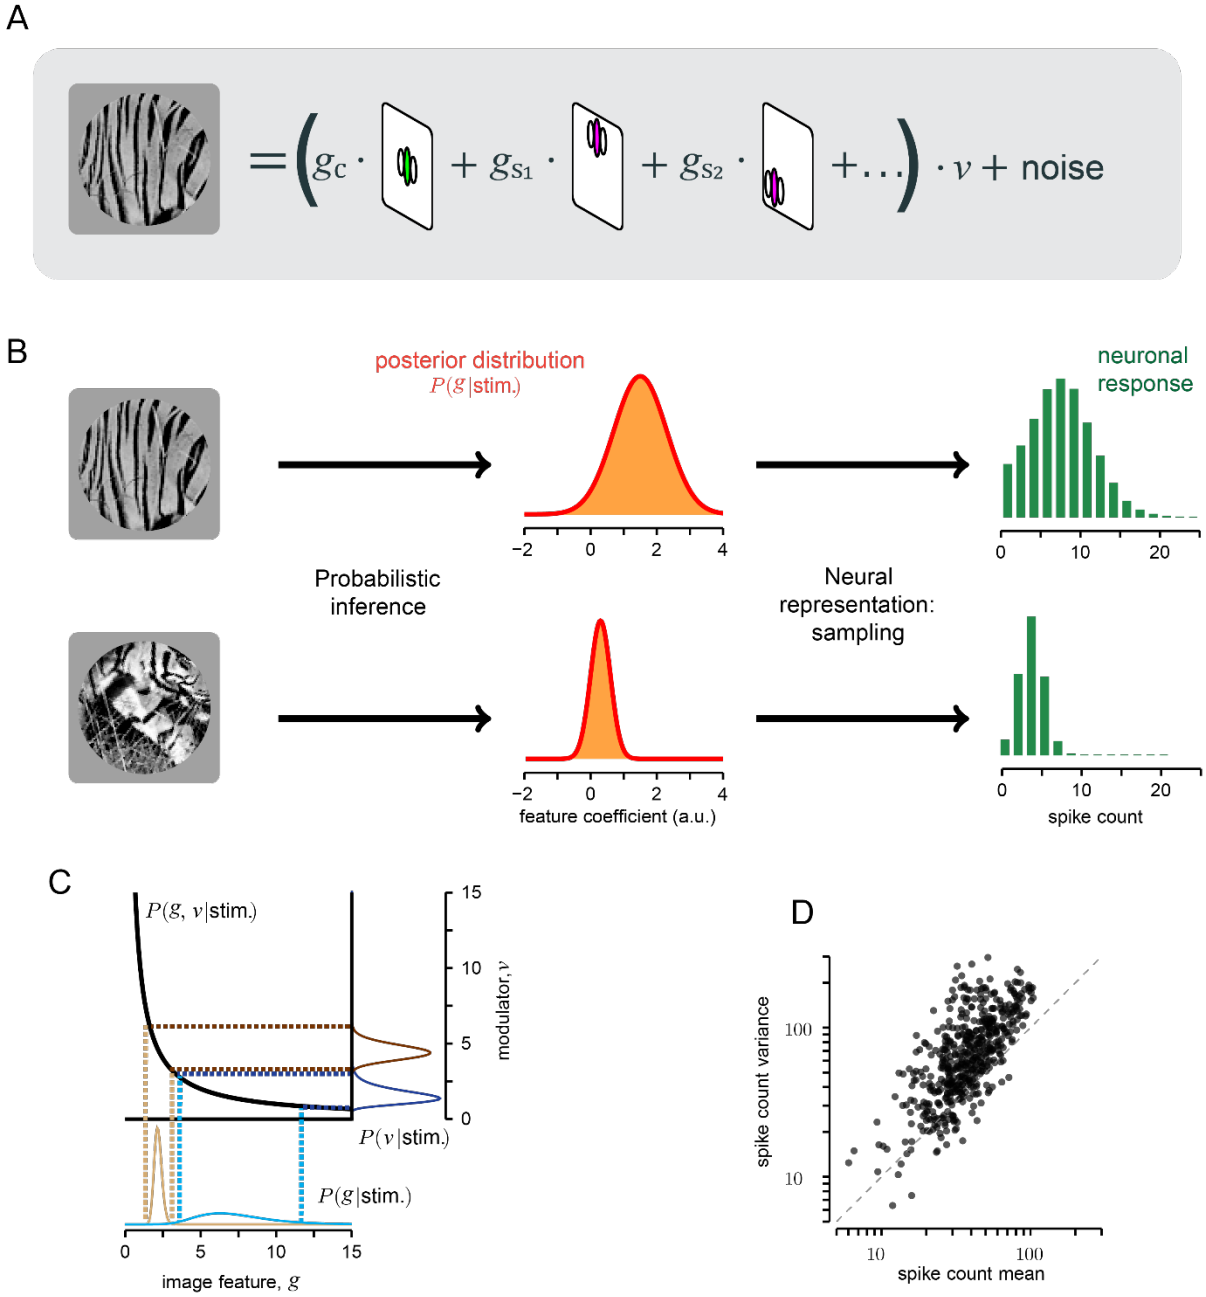

Figure 1 - Poisson-like variability explained by sampling-based inference in the GSM model. (A) Representation of the generative process of the Gaussian scale mixture (GSM) model, (Methods Eq. 1). The image (left) is described as the combination of local oriented features weighted by Gaussian coefficients, 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 v are

bound to lie on the hyperbole $v = x / g$ (black line). Therefore, a larger estimate of v implies reduced mean and variance of g (blue versus brown curves). (D) Mean versus variance of a GSM model neuron in response to 500 patches of natural images.

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 provide more information about the modulator reduce uncertainty 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 Eq. 3), validating our intuition that surround stimuli reduce uncertainty by providing information about 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 firing rate (Supplementary Fig S1). 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. S2).

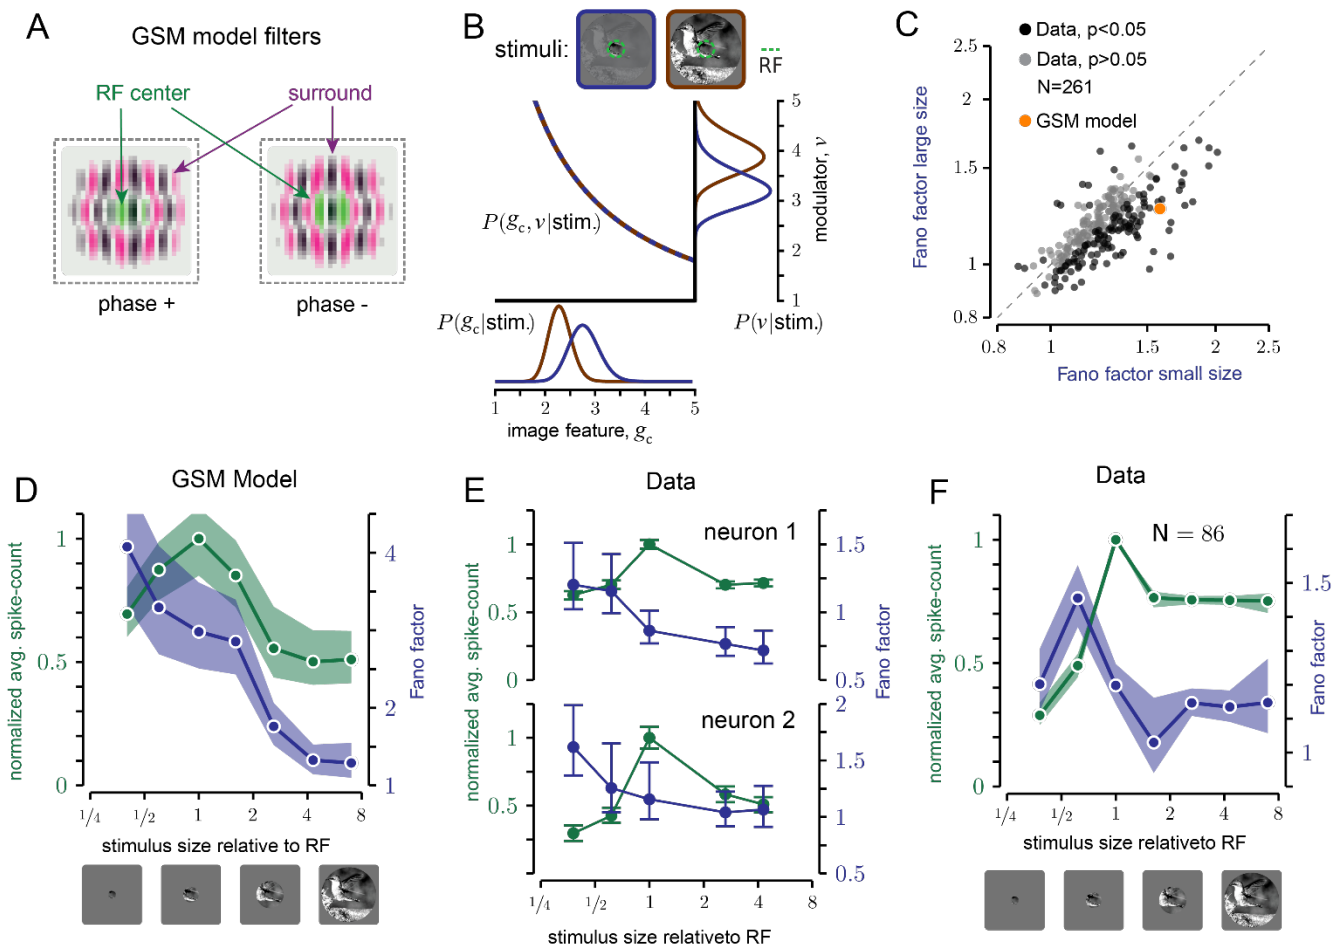

Figure 2 : Surround stimulation reduces GSM uncertainty and V1 variability for natural images. **(A)** In the GSM, the inputs to the model neuron are provided by the activity of quadrature pairs of oriented filters, corresponding to the spatial RF (green) and its surround (magenta). **(B)** Noise-free GSM model applied to an image with (blue) and without (brown) surround**.** The surround stimulus does not change the joint distribution of g_c (the local feature associated with the RF center) and ν, but it influences the estimate of the modulator ν and therefore also the estimate of g_c. **(C)** Fano factor (FF, ratio between variance and mean) 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 neuron in anesthetized macaque V1; 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.

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 variability relative to the 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. S3). 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 firing rate: 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 trends 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 ½ RF size to RF size, and an additional 5.74% 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 (Sceniak, Ringach et al. 1999, Angelucci, Levitt et al. 2002, Cavanaugh, Bair et al. 2002). 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 firing rate (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 both in the same animal and in a second awake monkey, and obtained similar results (Table 1, experiment 2; Supplementary Fig. S4A). We also observed the same trends (Table 1, experiment 3; Supplementary Fig. S4B) in data from anesthetized macaques (Coen-Cagli, Kohn et al. 2015), ruling out the possibility that microsaccades in the awake animals might have biased the result.

Experiment	FF change ½ vs. 1 RF	FF change 1 vs. 2 RF	p-value	Mean-matched FF change Smaller vs. larger than RF	p-value
1. Natural, awake (N=86; Fig 2F)	18.7%	5.74%	0.0082	25.7%	<1E-5
2. Gratings, awake (N=19; Supp Fig S4)	31.7 %	9.0%	0.05	47.7%	<1E-3
3. Gratings, anesthetized (N=229; Supp Fig S4)	14.2 %	7.05%	<1E-3	22.6%	<1E-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 ½ 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 2xRF 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.

Note that the FF was lower on average for stimuli smaller than ½ 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. S5). 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. S6). Furthermore, the FF decreased monotonically with stimulus size in the anesthetized dataset, for which stimulus centering could be controlled more tightly (Supplementary Fig. S5).

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. 2010).

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 (Sillito and Jones 1996, Gardner, Anzai et al. 1999, Walker, Ohzawa et al. 1999, Cavanaugh, Bair et al. 2002, Webb, Dhruv et al. 2005, Angelucci, Bijanzadeh et al. 2017). 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 (Schwartz and Simoncelli 2001, Coen-cagli, Dayan et al. 2009). Because the GSM predicts that surround suppression of both firing rate and variability is controlled by the inferred strength of the global modulator, we found that surround suppression of model response variability and firing rate were similarly tuned (Fig. 3A). We verified that this corresponded to a smaller estimate of the global modulator for orthogonal surround stimuli (Supplementary Fig. S3), which in turn resulted in weaker surround suppression of variability.

To test these model predictions, we measured V1 responses to compound static gratings 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. 2002). 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<1E-5), the surround tuning of FF was also stronger (average reduction 7.37%, p=0.013) than for the entire population (Supplementary Fig. S7).

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.

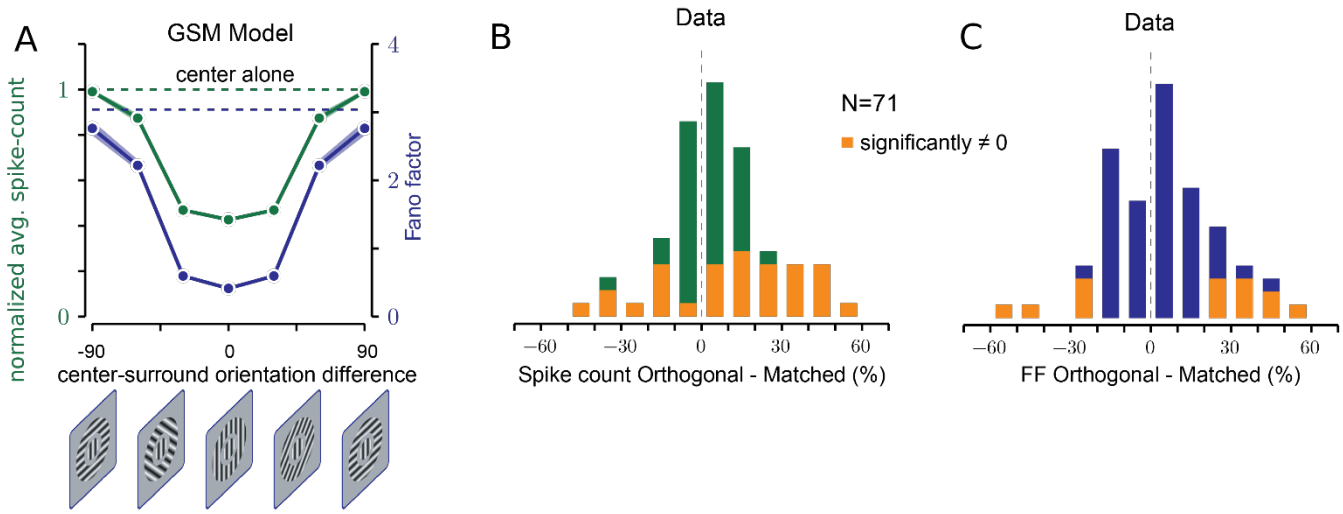

Figure 3: Surround reduction of variability is orientation tuned. **(A)** Surround orientation tuning of the mean spike count (green line) and FF (blue line), relative to the center stimulus alone (dashed lines) in the GSM model. Error bars: 68% c.i. Scaling factor (Methods Eq.4) $c = 40$ **(B,C)** Percent 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.

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 (Hoyer and Hyvärinen 2003, Fiser, Berkes et al. 2010, Haefner, Berkes et al. 2016, 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 Poisson-like variability and its modulation by spatial context is a general consequence of probabilistic inference when there are multiplicative interactions between latent variables, which is a widely-adopted description of natural image statistics (Wainwright, Simoncelli et al. 2000, Schwartz and Simoncelli 2001, Dayan and Abbott 2005, Gao and Vasconcelos 2007, Geisler 2008, Coen-Cagli, Dayan et al. 2012). 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 firing rate, suggesting shared mechanisms.

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 (Olshausen and Field 1996, Bell and Sejnowski 1997, Rao and Ballard 1999, Hyvarinen and Hoyer 2000, Schwartz and Simoncelli 2001, Karklin and Lewicki 2009). This work has typically addressed only the trial-

averaged firing rate. However, across-trial variability is substantial in cortex and can strongly influence perception (Shadlen and Newsome 1998, Pouget, Beck et al. 2013, Moreno-Bote, Beck et al. 2014, Kohn, Coen-Cagli et al. 2016). 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 has shown how samples might be generated dynamically with networks of spiking neurons, and in a manner that is fast enough for accurate inference within short, biologically relevant timescales (Hennequin, Aitchison et al. 2014, Legenstein and Maass 2014, Savin and Deneve 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 (Rao and Ballard 1999, Schwartz and Simoncelli 2001, Spratling 2010, Zhu and Rozell 2013, Coen-Cagli, Kohn et al. 2015). Contextual modulation of V1 trial-averaged responses has been characterized extensively with artificial stimuli (Sillito and Jones 1996, Sceniak, Ringach et al. 1999, Walker, Ohzawa et al. 1999, Angelucci, Levitt et al. 2002, Cavanaugh, Bair et al. 2002, Cavanaugh, Bair et al. 2002), and is also prominent for natural inputs (Vinje and Gallant 2000, Haider, Krause et al. 2010). 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 (Schwartz and Simoncelli 2001, Coen-cagli, Dayan et al. 2009, Coen-Cagli, Dayan et al. 2012). The modeling and experimental results presented here are consistent with this prior work, as we report strong and tuned surround suppression of firing rate (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 firing rate 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. S2), 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 (Rabinowitz, Goris et al. 2015, Haefner, Berkes et al. 2016).

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 (Heeger 1992, Carandini and Heeger 2012). 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 (Heeger 1992, Cavanaugh, Bair et al. 2002), also modulate variability (Kohn and Smith 2005, Haider, Krause et al. 2010, Snyder, Morais et al. 2014). 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, Boundy-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.

Relation to other descriptive models and functional explanations of cortical variability

Our study builds on previous work that used the GSM to demonstrate stimulus dependent changes in response statistics (Orbán, Berkes et al. 2016). In particular, Orbán et al. (2016) suggested that the GSM could unify effects of firing rate and variability. Our work extends this study in two important aspects. First, Orbán and colleagues utilized an approximate solution for the GSM model response, whereby uncertainty is due to input noise, not the unknown state of the global latent variable. This required tuning the nonlinear conversion from membrane potential to firing rate to account for Poisson-like 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 Poisson-like responses for a wide range of inputs (Fig. 1C,D), without tuning the spiking nonlinearity (although tuning may be necessary to quantitatively match the data). Second, our analytical results addressed more directly the relation between V1 responses and natural image statistics, and tested detailed predictions relating 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.

Methods

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. Here, rather than learning this linear transform as in other applications of the GSM (e.g. Orbán, Berkes et al. 2016), we assume oriented filters because they approximate well those learned on natural images, and also represent a canonical choice for V1 models. 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. 2002) (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:

$$\begin{gathered}\mathbf{x} = \nu\,\mathbf{g} + \boldsymbol{\eta} \\ \mathbf{g} \sim \mathcal{N}\left(0, \mathbf{C}_g\right) \quad ; \quad \nu \sim \text{Rayleigh}(1) \quad ; \quad \boldsymbol{\eta} \sim \mathcal{N}\left(0, \mathbf{C}_{\text{noise}}\right)\end{gathered} \qquad \text{Eq. 1}$$

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 $\boldsymbol{\eta}$ is necessary to account for realistic response variability with weak stimuli (Supplementary Fig. S8). We assume that $\mathbf{g}$ and $\boldsymbol{\eta}$ are generated from multivariate normal distributions, with mean 0 and covariances $\mathbf{C}_g$ and $\mathbf{C}_{\text{noise}}$, respectively ; ν follows a Rayleigh distribution with mean 1. Note that changing the Rayleigh parameter is equivalent to rescaling $\mathbf{C}_g$.

Model optimization. The covariance of the noise term, denoted as $\mathbf{C}_{\text{noise}}$ 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. S8). The covariance matrix $\mathbf{C}_g$ is computed by moment-matching (Doulgeris and Eltoft 2009), using the empirical covariance of filter outputs over 10,000 natural image patches. This ensures that the model is adapted to natural image statistics, as in previous work (Coen-Cagli, Dayan et al. 2012).

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 a test image $\tilde{\mathbf{x}}$. This quantity is denoted $P(g_{1+} \mid \tilde{\mathbf{x}})$, and results from an operation of Bayesian inference and marginalization over the other latent variables (Supplemental Text, Section 1). In particular, the global modulator ν 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. $\boldsymbol{\eta} = \mathbf{0}$. First, ν can be expressed analytically and approximated (Supplemental Text, Section 1) as:

$$\mathrm{E}\left[\nu \mid \tilde{\mathbf{x}}\right] = \sqrt{\lambda}\ \left[1 + \mathcal{O}(\lambda^{-1})\right]\ , \ \text{with} \quad \lambda = \sqrt{\sum_{i,j}\left(\mathbf{C}_g^{-1}\right)_{ij}\tilde{\mathbf{x}}_i^T\tilde{\mathbf{x}}_j^T} \qquad \text{Eq. 2}$$

where $\mathcal{O}(\lambda^{-1})$ represents a generic function that drops to zero asymptotically with λ^{-1} for $\lambda \gg 0$. This shows that the estimate of the mixer depends on the outputs of all filters. Second, the distribution of the feature of interest, $P(g_{1+} \mid \tilde{\mathbf{x}})$, 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:

$$\mathrm{E}\left[g_{1+} \mid \tilde{\mathbf{x}}\right] = \frac{\tilde{x}_{1+}}{\sqrt{\lambda}}\left[1+\mathcal{O}(\lambda^{-1})\right] \quad \text{and} \quad \mathrm{FF}\left[g_{1+} \mid \tilde{\mathbf{x}}\right] = \frac{\tilde{x}_{1+}}{4\lambda\sqrt{\lambda}}\left[1+\mathcal{O}(\lambda^{-1})\right] \qquad \text{Eq. 3}$$

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 Poisson-like variability.

The analytical results in Eq. 3 refer to the reduced model without additive noise. In this formulation, for very small inputs $\tilde{\mathbf{x}} \approx \mathbf{0}$ 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. S8).

Conversion to spike counts. In the GSM, the elements of $\mathbf{g}$ are not bounded to positive values. In order to simulate a neuronal response (the instantaneous firing rate) we therefore apply the following transformation:

$$r = c\sqrt{g_{1+}^2 + g_{1-}^2} \qquad \text{Eq. 4}$$

where c is a fixed parameter set heuristically so that mean responses and FF fall into a realistic range (values are reported in the figure captions), and the $\pm$ represent the two spatial phases at the RF position. In the no-noise approximation, the mean and variance of r can be expressed analytically, and preserve the behavior of Eq. 4 (see Supplemental Text, Section 5). For the full model, we compute a single-trial response r for each sample of g_{1+}, g_{1-}. The mean, variance and FF of the model neuron are then found numerically, using 400 samples.

Neurophysiology

Animal preparation and data collection. We recorded data from male adult macaque monkeys (*Macaca fascicularis*), either anesthetized (3 animals) or awake (2 animals). The protocol and general methods employed for the anesthetized experiments have been described previously (Smith and Kohn 2008). In short, anesthesia was induced with ketamine (10 mg/kg of body weight) and maintained during surgery with isoflurane (1.5–2.5% in 95% O2), switching to sufentanil (6–18 μg /kg per h, adjusted as needed) during recordings. Eye movements were reduced using vecuronium bromide (0.15 mg/kg per h). Temperature was maintained in the 36 –37 C° range, and relevant 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 were implanted a 10x10 multielectrode array (400 μm spacing, 1-mm length) into 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 deg x 1 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 deg x 3 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.

Visual stimuli. Visual stimuli were generated with custom software (EXPO) and displayed on a cathode ray tube monitor (Hewlett Packard p1230; 1024x768 pixels, with ~ 40 cd/m^2 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 3.1-6.7 degrees). These included 90 distinct images, and, for images with a well-defined 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. S4). 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 diameter [0.34, 0.55, 0.90, 2.4, 3.8] deg , with stimulus duration 200 ms and blank duration 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 diameter [0.34, 0.55, 0.90, 2.4, 3.8] deg, orientation [0, 45, 90, 135] deg, duration 250 ms and blank duration 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 (diameter [0.34, 0.55, 0.90, 1.5, 2.4, 3.8, 6.2] deg, orientation [0,45,90,135] 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 blanks). For monkey M we used a central grating of diameter 1 deg, orientations [0, 90], 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 [0, 45, 90, 135] 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 together in our analyses.

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 50ms 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 + 1std. 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. S7).

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:

$$\%\text{change in FF} = 100 \cdot (\mathrm{FF}_\alpha - \mathrm{FF}_\beta) / \left((\mathrm{FF}_\alpha + \mathrm{FF}_\beta)/2\right) \qquad \text{Eq. 5}$$

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. S4-S5, 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. S5 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. S1, we compared the FF between stimuli that were smaller versus larger than the RF, and elicited similar firing rates. 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 2xRF size) had mean ≤ 0.

Acknowledgments

We thank the Coen-Cagli and Kohn labs for helpful discussion, and the Kohn lab for help with experiments. 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.

Competing Interests

The authors have no competing interests.

Data Availability

Data for Figure 2C and Supplementary Figure S4B are publicly available on the CRCNS data sharing site 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 (www.github.com).

Materials & Correspondence

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

Bibliography

Albright, T. D. and G. R. Stoner (2002). "Contextual Influences on Visual Processing." Annual Review of Neuroscience **25**(1): 339-379.

Angelucci, A., M. Bijanzadeh, L. Nurminen, F. Federer, S. Merlin and P. C. Bressloff (2017). "Circuits and Mechanisms for Surround Modulation in Visual Cortex." Annu Rev Neurosci **40**: 425-451.

Angelucci, A., J. B. Levitt, E. J. S. Walton, J.-M. Hupé, J. Bullier and J. S. Lund (2002). "Circuits for Local and Global Signal Integration in Primary Visual Cortex." Journal of Neuroscience **22**(19): 8633-8646.

Beck, J. M., P. E. Latham and A. Pouget (2011). "Marginalization in neural circuits with divisive normalization." Journal of Neuroscience **31**(43): 15310-15319.

Beck, J. M., W. J. Ma, X. Pitkow, P. E. Latham and A. Pouget (2012). "Not noisy, just wrong: the role of suboptimal inference in behavioral variability." Neuron **74**(1): 30-39.

Bell, A. J. and T. J. Sejnowski (1997). "The "independent components" of natural scenes are edge filters." Vision Res **37**(23): 3327-3338.

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

Bishop, C. M. (2006). Pattern recognition and machine learning. New York, Springer.

Boerlin, M., C. K. Machens and S. Denève (2013). "Predictive Coding of Dynamical Variables in Balanced Spiking Networks." PLoS Comput Biol **9**(11): e1003258.

Britten, K. H., M. N. Shadlen, W. T. Newsome and J. A. Movshon (1992). "The analysis of visual motion: a comparison of neuronal and psychophysical performance." J Neurosci **12**(12): 4745-4765.

Carandini, M. (2004). "Amplification of Trial-to-Trial Response Variability by Neurons in Visual Cortex." PLOS Biology **2**(9): e264.

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

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

Cavanaugh, J. R., W. Bair and J. A. Movshon (2002). "Selectivity and Spatial Distribution of Signals From the Receptive Field Surround in Macaque V1 Neurons." Journal of Neurophysiology **88**(5): 2547-2556.

Charles, A. S., M. Park, J. P. Weller, G. D. Horwitz and J. W. Pillow (2018). "Dethroning the Fano Factor: A Flexible, Model-Based Approach to Partitioning Neural Variability." Neural Computation **30**(4): 1012-1045.

Chen, M., Y. Yan, X. Gong, C. D. Gilbert, H. Liang and W. Li (2014). "Incremental Integration of Global Contours through Interplay between Visual Cortical Areas." Neuron **82**(3): 682-694.

Churchland, M. M., B. M. Yu, J. P. Cunningham, L. P. Sugrue, M. R. Cohen, G. S. Corrado, W. T. Newsome, A. M. Clark, P. Hosseini, B. B. Scott, D. C. Bradley, M. A. Smith, A. Kohn, J. A. Movshon, K. M. Armstrong, T. Moore, S. W. Chang, L. H. Snyder, S. G. Lisberger, N. J. Priebe, I. M. Finn, D. Ferster, S. I. Ryu, G. Santhanam, M. Sahani and K. V. Shenoy (2010). "Stimulus onset quenches neural variability: a widespread cortical phenomenon." Nat Neurosci **13**(3): 369-378.

Coen-cagli, R., P. Dayan and O. Schwartz (2009). Statistical Models of Linear and Nonlinear Contextual Interactions in Early Visual Processing. Advances in Neural Information Processing Systems 22.

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

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

Coen-Cagli, R. and S. S. Solomon (2019). "Relating Divisive Normalization to Neuronal Response Variability." J Neurosci **39**(37): 7344-7356.

Cohen, M. R. and J. H. R. Maunsell (2009). "Attention improves performance primarily by reducing interneuronal correlations." Nature Neuroscience **12**(12): 1594-1600.

Dadarlat, M. C. and M. P. Stryker (2017). "Locomotion Enhances Neural Encoding of Visual Stimuli in Mouse V1." Journal of Neuroscience **37**(14): 3764-3775.

Dayan, P. and L. Abbott (2005). Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems, The MIT Press.

Dinstein, I., D. J. Heeger and M. Behrmann (2015). "Neural variability: friend or foe?" Trends in Cognitive Sciences **19**(6): 322-328.

Doulgeris, A. P. and T. Eltoft (2009). "Scale mixture of Gaussian modelling of polarimetric SAR data." EURASIP Journal on Advances in Signal Processing **2010**(1): 874592.

Finn, I. M., N. J. Priebe and D. Ferster (2007). "The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex." Neuron **54**(1): 137-152.

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

Gao, D. and N. Vasconcelos (2007). Bottom-up saliency is a discriminant process.

Gardner, J. L., A. Anzai, I. Ohzawa and R. D. Freeman (1999). "Linear and nonlinear contributions to orientation tuning of simple cells in the cat's striate cortex." Visual Neuroscience **16**(6): 1115-1121.

Geisler, W. S. (2008). "Visual Perception and the Statistical Properties of Natural Scenes." **59**(1): 167-192.

Goris, R. L. T., J. A. Movshon and E. P. Simoncelli (2014). "Partitioning neuronal variability." Nature neuroscience **17**(6): 858.

Haefner, R. M., P. Berkes and J. Fiser (2016). "Perceptual Decision-Making as Probabilistic Inference by Neural Sampling." Neuron **90**(3): 649-660.

Haider, B., M. R. Krause, A. Duque, Y. Yu, J. Touryan, J. A. Mazer and D. A. McCormick (2010). "Synaptic and Network Mechanisms of Sparse and Reliable Visual Cortical Activity during Nonclassical Receptive Field Stimulation." Neuron **65**(1): 107-121.

Heeger, D. J. (1992). "Normalization of cell responses in cat striate cortex." Visual Neuroscience **9**(2): 181-197.

Henaff, O. J., Z. M. Boundy-Singer, K. Meding, C. M. Ziemba and R. L. T. Goris (2020). "Representation of visual uncertainty through neural gain variability." Nat Commun **11**(1): 2513.

Hennequin, G., L. Aitchison and M. Lengyel (2014). Fast sampling-based inference in balanced neuronal networks. Advances in neural information processing systems.

Hoyer, P. O. and A. Hyvärinen (2003). Interpreting neural response variability as Monte Carlo sampling of the posterior.

Hunsberger, E., M. Scott and C. Eliasmith (2014). "The competing benefits of noise and heterogeneity in neural coding." Neural Comput **26**(8): 1600-1623.

Hyvarinen, A. and P. Hoyer (2000). "Emergence of phase- and shift-invariant features by decomposition of natural images into independent feature subspaces." Neural Comput **12**(7): 1705-1720.

Karklin, Y. and M. S. Lewicki (2009). "Emergence of complex cell properties by learning to generalize in natural scenes." Nature **457**(7225): 83-86.

Knill, D. C. and A. Pouget (2004). "The Bayesian brain: the role of uncertainty in neural coding and computation." Trends Neurosci **27**(12): 712-719.

Kohn, A., R. Coen-Cagli, I. Kanitscheider and A. Pouget (2016). "Correlations and Neuronal Population Information." Annual Review of Neuroscience **39**(1): 237-256.

Kohn, A. and M. A. Smith (2005). "Stimulus dependence of neuronal correlation in primary visual cortex of the macaque." Journal of Neuroscience **25**(14): 3661-3673.

Legenstein, R. and W. Maass (2014). "Ensembles of spiking neurons with noise support optimal probabilistic inference in a dynamically changing environment." PLoS Comput Biol **10**(10): e1003859.

Lochmann, T., U. A. Ernst and S. Deneve (2012). "Perceptual inference predicts contextual modulations of sensory responses." J Neurosci **32**(12): 4179-4195.

Ma, W. J., J. M. Beck, P. E. Latham and A. Pouget (2006). "Bayesian inference with probabilistic population codes." Nature Neuroscience **9**(11): 1432-1438.

Mitchell, J. F., K. A. Sundberg and J. H. Reynolds (2009). "Spatial Attention Decorrelates Intrinsic Activity Fluctuations in Macaque Area V4." Neuron **63**(6): 879-888.

Moreno-Bote, R., J. Beck, I. Kanitscheider, X. Pitkow, P. Latham and A. Pouget (2014). "Information-limiting correlations." Nat Neurosci **17**(10): 1410-1417.

Nienborg, H. and B. G. Cumming (2014). "Decision-Related Activity in Sensory Neurons May Depend on the Columnar Architecture of Cerebral Cortex." Journal of Neuroscience **34**(10): 3579-3585.

Olshausen, B. A. and D. J. Field (1996). "Emergence of simple-cell receptive field properties by learning a sparse code for natural images." Nature **381**(6583): 607-609.

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

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." Proceedings of the National Academy of Sciences **110**(32): 13162-13167.

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

Pouget, A., P. Dayan and R. S. Zemel (2003). "Inference and computation with population codes." Annu Rev Neurosci **26**: 381-410.

Rabinowitz, N. C., R. L. Goris, M. Cohen and E. P. Simoncelli (2015). "Attention stabilizes the shared gain of V4 populations." Elife **4**: e08998.

Rao, R. P. and D. H. Ballard (1999). "Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects." Nat Neurosci **2**(1): 79-87.

Savin, C. and S. Deneve (2014). Spatio-temporal representations of uncertainty in spiking neural networks. Advances in Neural Information Processing Systems.

Sceniak, M. P., D. L. Ringach, M. J. Hawken and R. Shapley (1999). "Contrast's effect on spatial summation by macaque V1 neurons." Nature Neuroscience **2**: 733.

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

Shadlen, M. N. and W. T. Newsome (1998). "The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding." The Journal of Neuroscience **18**(10): 3870-3896.

Sillito, A. M. and H. E. Jones (1996). "Context-dependent interactions and visual processing in V1." Journal of Physiology-Paris **90**(3): 205-209.

Simoncelli, E. P. and W. T. Freeman (1995). The steerable pyramid: a flexible architecture for multi-scale derivative computation.

Smith, M. A. and A. Kohn (2008). "Spatial and temporal scales of neuronal correlation in primary visual cortex." J Neurosci **28**(48): 12591-12603.

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

Solomon, S. S., S. C. Chen, J. W. Morley and S. G. Solomon (2014). "Local and Global Correlations between Neurons in the Middle Temporal Area of Primate Visual Cortex." Cerebral Cortex **25**(9): 3182-3196.

Sompolinsky, H., H. Yoon, K. Kang and M. Shamir (2001). "Population coding in neuronal systems with correlated noise." Physical Review E **64**(5): 051904.

Spratling, M. W. (2010). "Predictive coding as a model of response properties in cortical area V1." J Neurosci **30**(9): 3531-3543.

Stevenson, I. H. (2016). "Flexible models for spike count data with both over- and under- dispersion." Journal of Computational Neuroscience **41**(1): 29-43.

Theis, L., R. Hosseini and M. Bethge (2012). "Mixtures of conditional Gaussian scale mixtures applied to multiscale image representations." PLoS One **7**(7): e39857.

Tolhurst, D. J., J. A. Movshon and A. F. Dean (1983). "The statistical reliability of signals in single neurons in cat and monkey visual cortex." Vision research **23 8**: 775-785.

Tomko, G. J. and D. R. Crapper (1974). "Neuronal variability: non-stationary responses to identical visual stimuli." Brain Research **79**(3): 405-418.

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

Vinje, W. E. and J. L. Gallant (2000). "Sparse Coding and Decorrelation in Primary Visual Cortex During Natural Vision." Science **287**(5456): 1273.

Wainwright, M. J., E. P. Simoncelli and A. S. Willsky (2000). Random cascades of Gaussian scale mixtures and their use in modeling natural images with application to denoising.

Walker, G. A., I. Ohzawa and R. D. Freeman (1999). "Asymmetric Suppression Outside the Classical Receptive Field of the Visual Cortex." The Journal of Neuroscience **19**(23): 10536-10553.

Webb, B. S., N. T. Dhruv, S. G. Solomon, C. Tailby and P. Lennie (2005). "Early and Late Mechanisms of Surround Suppression in Striate Cortex of Macaque." Journal of Neuroscience **25**(50): 11666-11675.

White, B., L. F. Abbott and J. Fiser (2012). "Suppression of cortical neural variability is stimulus- and state-dependent." J Neurophysiol **108**(9): 2383-2392.

Zhu, M. and C. J. Rozell (2013). "Visual nonclassical receptive field effects emerge from sparse coding in a dynamical system." PLoS Comput Biol **9**(8): e1003191.