Modeling the Diverse Effects of Divisive Normalization on Noise Correlations

Oren Weiss[1], Hayley A. Bounds[2], Hillel Adesnik[3], and Ruben Coen-Cagli[1,4,5]

[1] *Department of Systems and Computational Biology, Albert Einstein College of Medicine*

[2] *Helen Wills Neuroscience Institute, University of California, Berkeley*

[3] *Department of Molecular and Cell Biology, University of California, Berkeley*

[4] *Dominick P. Purpura Department of Neuroscience, Albert Einstein College of Medicine*

[5] *Department of Ophthalmology and Visual Sciences, Albert Einstein College of Medicine*

July 7, 2023

Abstract

Divisive normalization, a prominent descriptive model of neural activity, is employed by theories of neural coding across many different brain areas. Yet, the relationship between normalization and the statistics of neural responses beyond single neurons remains largely unexplored. Here we focus on noise correlations, a widely studied pairwise statistic, because its stimulus and state dependence plays a central role in neural coding. Existing models of covariability typically ignore normalization despite empirical evidence suggesting it affects correlation structure in neural populations. We therefore propose a pairwise stochastic divisive normalization model that accounts for the effects of normalization and other factors on covariability. We first show that normalization modulates noise correlations in qualitatively different ways depending on whether normalization is shared between neurons, and we discuss how to infer when normalization signals are shared. We then apply our model to calcium imaging data from mouse primary visual cortex (V1), and find that it accurately fits the data, often outperforming a popular alternative model of correlations. Our analysis indicates that normalization signals are often shared between V1 neurons in this dataset. Our model will enable quantifying the relation between normalization and covariability in a broad range of neural systems, which could provide new constraints on circuit mechanisms of normalization and their role in information transmission and representation.

Author Summary

Cortical responses are often variable across identical experimental conditions, and this variability is shared between neurons (noise correlations). These noise correlations have been studied extensively to understand how they impact neural coding and what mechanisms determine their properties. Here we show how correlations relate to divisive normalization, a mathematical operation widely adopted to describe how the activity of a neuron is modulated by other neurons via divisive gain control. We introduce the first statistical model of this

relation. We extensively validate the model and investigate parameter inference in synthetic data. We find that our model, when applied to data from mouse visual cortex, outperforms a popular model of noise correlations that does not include normalization, and it reveals diverse influences of normalization on correlations. Our work demonstrates a framework to measure the relation between noise correlations and the parameters of the normalization model, which could become an indispensable tool for quantitative investigations of noise correlations in the wide range of neural systems that exhibit normalization.

1 Introduction

Neurons in the sensory cortices of the brain exhibit substantial response variability across identical experimental trials [Shadlen and Newsome, 1998; Tolhurst, Movshon, and Dean, 1983]. These fluctuations in activity are often shared between pairs of simultaneously recorded neurons, called noise correlations [Cohen and Kohn, 2011]. Because the presence of these correlations can constrain the amount of information encoded by neural populations and impact behavior [Rumyantsev et al., 2020; Kafashan et al., 2021; Bartolo et al., 2020; Abbott and Dayan, 1999; Averbeck and Lee, 2006; Kohn et al., 2016; Hu, Zylberberg, and Shea-Brown, 2014; Moreno-Bote et al., 2014; Zohary, Shadlen, and Newsome, 1994; Kanitscheider, Coen-Cagli, and Pouget, 2015], noise correlations have been widely studied. This work has also revealed that correlations are often modulated by stimulus and state variables [Cohen and Kohn, 2011; Panzeri et al., 2022], and therefore can play an important role in computational theories of sensory coding. For instance, noise correlations could emerge from neurons performing Bayesian inference and reflect the statistics of sensory inputs [Bányai et al., 2019; Berkes et al., 2011; Orbán et al., 2016; Orbán and Bányai, 2019] and prior expectations [Haefner, Berkes, and Fiser, 2016; Lange and Haefner, 2017; 2022]. From a mechanistic point of view, such a statistical structure of noise correlations poses strong constraints on circuit models of cortical activity [Bondy, Haefner, and Cum-

ming, 2018; Doiron et al., 2016; Litwin-Kumar and Doiron, 2012; Hennequin et al., 2018; Heeger and Zemlianova, 2020]. To better understand the functional impact and underlying mechanisms of noise correlations on neural coding and behavior, we need to be able to quantitatively characterize and interpret how noise correlations in neural populations are affected by experimental variables.

For this reason, successful descriptive models of neural activity have been developed to capture noise correlations [Pillow et al., 2008; Archer et al., 2014; Gardella, Marre, and Mora, 2019; Schneidman et al., 2006; Granot-Atedgi et al., 2013; Zhao and Park, 2017; Sokoloski, Aschner, and Coen-Cagli, 2021; Josić et al., 2009]. However, none of those models considers divisive normalization [Carandini and Heeger, 2012; Heeger, 1992; Albrecht and Geisler, 1991], an operation observed in a wide range of neural systems [Louie, Khaw, and Glimcher, 2013; Ohshiro, Angelaki, and DeAngelis, 2011; Olsen, Bhandawat, and Wilson, 2010] which has also been implicated in modulating the structure of noise correlations. Experimental phenomena that are accompanied by changes in noise correlations, including contrast saturation [Kohn and Smith, 2005], surround suppression [Liu, Haefner, and Pack, 2016; Snyder et al., 2014; Henry and Kohn, 2020], and attentional modulations of neural activity [Cohen and Maunsell, 2009; Ruff and Cohen, 2014; Mitchell, Sundberg, and Reynolds, 2009] have been successfully modeled using divisive normalization [Cavanaugh, Bair, and Movshon, 2002; Heeger, 1992; Reynolds and Heeger, 2009; Coen-Cagli, Dayan, and Schwartz, 2012], although those models only captured average firing rates of individual neurons. Additionally, some numerical simulation studies have shown how normalization can affect noise correlations [Tripp, 2012; Verhoef and Maunsell, 2017]. These results indicate that it is important to quantify the relative contribution of normalization and other factors to modulation of noise correlations in experimental data.

We propose a stochastic normalization model, the pairwise Ratio of Gaussians (RoG), to capture the across-trial joint response statistics for pairs of simultaneously recorded neurons.

This builds on our previous method that considered the relationship between normalization and single-neuron response variability (hence we refer to it as the independent RoG; [Coen-Cagli and Solomon, 2019]). In these RoG models, neural responses are described as the ratio of two random variables: the numerator, which represents excitatory input to the neuron, and the denominator (termed normalization signal), which represents the suppressive effect of summed input from a pool of neurons [Carandini and Heeger, 2012; Sawada and Petrov, 2017]. Our pairwise RoG allows for the numerators and denominators that describe the individual responses to be correlated across pairs; in turn, these correlations induce correlations between the ratio variables (i.e. the model neurons' activity). In this paper, we derive and validate a bivariate Gaussian approximation to the joint distribution of pairwise responses, which greatly simplifies the problem of fitting the model and interpreting its behavior. The model provides a mathematical relationship between noise correlations and normalization, which predicts qualitatively different effects of normalization on noise correlations, depending on the relative strength and sign of the correlation between numerators and between denominators. This could explain the diversity of modulations of noise correlations observed in past work [Ruff and Cohen, 2014; 2016; Verhoef and Maunsell, 2017]. To provide practical guidance for data-analytic applications of our model, we investigate the accuracy and stability of parameter inference, and illustrate the conditions under which our pairwise RoG affords better estimates of single-trial normalization signals compared to the independent RoG. We then demonstrate that the model fits accurately pairwise responses recorded in the mouse primary visual cortex (V1), and often outperforms a popular alternative that ignores normalization. In our dataset, we find that when the correlation parameter between denominators is significantly different from zero, it is positive, indicating that those pairs share their normalization signals.

Our results highlight the importance of modeling the relation between normalization and covariability to interpret the rich phenomenology of noise correlations. Our model and code provide a data-analytic tool that will allow researchers to further investigate such a

relationship, and to quantitatively evaluate predictions made by normative and mechanistic models regarding the role of correlated variability and normalization in neural coding and behavior.

2 Methods

2.1 Generative Model – Pairwise Ratio of Gaussians (RoG)

Here we describe in detail the RoG model and derive the Gaussian approximation to the ratio variable. Note that our RoG is entirely different from that used in [Pillow and Simoncelli, 2006] and in [Cavanaugh, Bair, and Movshon, 2002], despite the same acronym: [Pillow and Simoncelli, 2006] refers to the distribution obtained from the ratio between two Gaussian distributions, whereas we refer to the random variable that results from the ratio of two Gaussian variables. Similarly, [Cavanaugh, Bair, and Movshon, 2002] does not refer to probability distributions at all, but rather to a surround suppression model in which the center and surround mechanisms are characterized by Gaussian integration over their respective receptive fields, while the RoG considered here is a model of neural covariability in general.

We build from the standard normalization model [Carandini and Heeger, 2012; Sawada and Petrov, 2017] which computes the across-trial average neural response (e.g. firing rate) as the ratio between a driving input to a neuron (N) and the summed input from a pool of nearby neurons (D):

$$R = f_{\div}(N, D) = \frac{N}{D} \quad \text{where } D > 0 \tag{1}$$

Where $f_{\div}$ is the division operator; this functional notation is convenient for later derivations in which we consider the derivative of division. Our goal is to model the joint activity of pairs of neurons, so we extend the normalization model by considering two model neurons R_1, R_2.

Since we are interested in trial-to-trial variability, we assume that a pair of neural responses $\boldsymbol{R}_t = (R_1, R_2)_t$ on a single trial t can be written as the element-wise ratio of two Gaussian random vectors, $\boldsymbol{N}_t = (N_1, N_2)_t$ and $\boldsymbol{D}_t = (D_1, D_2)_t$, with additive Gaussian random noise $\boldsymbol{\eta}_t = (\eta_1, \eta_2)_t$ to capture the residual (i.e. stimulus-independent) variability.

As detailed further below, the numerators of two neurons can be correlated, and similarly for the denominators. Moreover, in principle there can be correlations between the numerators and denominators (e.g. (N_1, D_2) may be correlated), requiring us to consider the joint, four-dimensional Gaussian distribution for the vector $(\boldsymbol{N}_t, \boldsymbol{D}_t)$. However, in this paper we consider the simpler model in which $\boldsymbol{N}_t$ and $\boldsymbol{D}_t$ are independent and are each distributed according to their respective two-dimensional Gaussian distributions. This assumption allows for simplified mathematical derivations, and is supported by our previous work which found that including a parameter for the correlation between N and D caused over-fitting to single-neuron data [Coen-Cagli and Solomon, 2019]. However, we have also derived the equations for the case that numerators and denominators are correlated (see Appendix section 4), and implemented them in the associated code toolbox, so that interested researchers can test if their data warrant the inclusion of those additional free parameters.

We therefore write the generative model for the pairwise RoG as:

$$\begin{aligned} \boldsymbol{N}_t &\sim \mathcal{N}(\boldsymbol{\mu_N}, \boldsymbol{\Sigma_N}) \\ \boldsymbol{D}_t &\sim \mathcal{N}(\boldsymbol{\mu_D}, \boldsymbol{\Sigma_D}) \\ \boldsymbol{\eta}_t &\sim \mathcal{N}(\boldsymbol{\mu_\eta}, \boldsymbol{\Sigma_\eta}) \end{aligned} \tag{2}$$

$$\boldsymbol{R}_t = (R_1, R_2)_t = f_{\div}(\boldsymbol{N}_t, \boldsymbol{D}_t) + \boldsymbol{\eta}_t = \frac{\boldsymbol{N}_t}{\boldsymbol{D}_t} + \boldsymbol{\eta}_t$$

Where $f_{\div}$ is applied element-wise, $\boldsymbol{\mu_N}, \boldsymbol{\mu_D}$ are the two-dimensional vectors of means of the numerator and denominator, respectively, $\boldsymbol{\Sigma_N}, \boldsymbol{\Sigma_D}$ are the respective 2×2 covariance matrices, and $(\boldsymbol{\mu_\eta}, \boldsymbol{\Sigma_\eta})$ is the mean and covariance matrix for the residual component of the

model.

For the independent RoG, the ratio variable in general follows a Cauchy distribution whose moments are not well defined, but we used the result that when the denominator has negligible probability mass at values less than or equal to zero, the ratio distribution can be approximated by a Gaussian distribution with mean and variance that can be derived from a Taylor expansion [Díaz-Francés and Rubio, 2013; Hayya, Armstrong, and Gressis, 1975; Marsaglia et al., 2006; Pham-Gia, Turkkan, and Marchand, 2006]. This assumption is justified since the denominator is the sum of the non-negative responses from a pool of neurons [Heeger, 1992], and is therefore unlikely to attain values less than or equal to zero. A similar result applies for the pairwise extension, in which we can use the multivariate delta method (an application of a Taylor expansion) to compute the mean and covariance for the joint distribution of ratio variables [Ver Hoef, 2012] under the assumption that $\boldsymbol{\mu_D} > 0$.

To derive equations for the mean and covariance of the pairwise RoG, we use a Taylor expansion around the point $(\boldsymbol{\mu_N}, \boldsymbol{\mu_D})$:

$$\begin{aligned} \boldsymbol{R} &= f_{\div}(\boldsymbol{N}, \boldsymbol{D}) + \boldsymbol{\eta} \\ &= f_{\div}(\boldsymbol{\mu_N}, \boldsymbol{\mu_D}) + \nabla f_{\div}|_{\boldsymbol{\mu_N}, \boldsymbol{\mu_D}} \left(\begin{pmatrix} \boldsymbol{N} \\ \boldsymbol{D} \end{pmatrix} - \begin{pmatrix} \boldsymbol{\mu_N} \\ \boldsymbol{\mu_D} \end{pmatrix} \right) + \text{higher order terms} + \boldsymbol{\eta} \end{aligned} \tag{3}$$

Using only the first order terms, we derive expressions for the mean and covariance matrix

of the RoG:

$$\boldsymbol{\mu_R} = \mathrm{E}[\boldsymbol{R}] \approx \frac{\boldsymbol{\mu_N}}{\boldsymbol{\mu_D}} + \boldsymbol{\mu_\eta} \tag{4}$$

$$\begin{aligned}
\boldsymbol{\Sigma_R} &= \mathrm{E}[(\boldsymbol{R} - \boldsymbol{\mu_R})(\boldsymbol{R} - \boldsymbol{\mu_R})^\top] \\
&\approx \nabla f_\div|_{\boldsymbol{\mu_N},\boldsymbol{\mu_D}} \mathrm{E}\left[\left(\begin{pmatrix}\boldsymbol{N}\\ \boldsymbol{D}\end{pmatrix} - \begin{pmatrix}\boldsymbol{\mu_N}\\ \boldsymbol{\mu_D}\end{pmatrix}\right)\left(\begin{pmatrix}\boldsymbol{N}\\ \boldsymbol{D}\end{pmatrix} - \begin{pmatrix}\boldsymbol{\mu_N}\\ \boldsymbol{\mu_D}\end{pmatrix}\right)^\top\right] \nabla f_\div|^\top_{\boldsymbol{\mu_N},\boldsymbol{\mu_D}} + \boldsymbol{\Sigma_\eta} \\
&\approx \nabla f_\div|_{\boldsymbol{\mu_N},\boldsymbol{\mu_D}} \begin{pmatrix}\Sigma_N & 0\\ 0 & \Sigma_D\end{pmatrix} \nabla f_\div|^\top_{\boldsymbol{\mu_N},\boldsymbol{\mu_D}} + \boldsymbol{\Sigma_\eta}
\end{aligned} \tag{5}$$

Note that the variance of the denominator influences the mean of the ratio variable through a second order term, hence it does not appear in Eq. 4 (see [Hayya, Armstrong, and Gressis, 1975] for the second order Taylor expansion for the mean of a ratio variable). From Eq. 5, we can obtain expressions for the variance of each neuron in the pair and their covariance and correlation. First, we adopt the following notation to simplify the equations: let $\widehat{\sigma_{\eta_i}} = \frac{\mu_{D_i}}{\mu_{N_i}}\sigma_{\eta_i}$ and let $\delta_{X_i} = \frac{\sigma_{X_i}}{\mu_{X_i}}$. Then:

$$\sigma^2_{R_i} \approx \frac{\mu^2_{N_i}}{\mu^2_{D_i}}\left(\delta^2_{N_i} + \delta^2_{D_i} + \widehat{\sigma_{\eta_i}}^2\right) \quad \text{for } i = 1, 2$$

$$\mathrm{Cov}\,(R_1, R_2) \approx \frac{\mu_{N_1}\mu_{N_2}}{\mu_{D_1}\mu_{D_2}}\left(\rho_N\delta_{N_1}\delta_{N_2} + \rho_D\delta_{D_1}\delta_{D_2} + \rho_\eta\widehat{\sigma_{\eta_1}}\widehat{\sigma_{\eta_2}}\right) \tag{6}$$

$$\mathrm{Corr}(R_1, R_2) \approx \frac{\rho_N\delta_{N_1}\delta_{N_2} + \rho_D\delta_{D_1}\delta_{D_2} + \rho_\eta\widehat{\sigma_{\eta_1}}\widehat{\sigma_{\eta_2}}}{\sqrt{\delta^2_{N_1} + \delta^2_{D_1} + \widehat{\sigma_{\eta_1}}^2}\sqrt{\delta^2_{N_2} + \delta^2_{D_2} + \widehat{\sigma_{\eta_2}}^2}} \tag{7}$$

Equation 7 is commonly referred to as the formula for "spurious" correlation of ratios found when comparing ratios of experimental variables [Kronmal, 1993], and we further generalize this in Appendix section 4. To the extent that tuning similarity between neurons reflects similarity in the driving inputs, and that those driving inputs are variable, neurons with more similar tuning would have larger ρ_N, which in turn implies larger noise correlations

according to Eq. 7. This is consistent with the widespread empirical observation that signal correlations and noise correlations are correlated [Cohen and Kohn, 2011].

2.2 Parametrization of the Pairwise RoG for Contrast Responses

In the form described so far, the pairwise RoG has 10 stimulus dependent parameters $\left(\mu_{N_1}, \mu_{N_2}, \mu_{D_1}, \mu_{D_2}, \sigma_{N_1}, \sigma_{N_2}, \sigma_{D_1}, \sigma_{D_2}, \rho_N, \rho_D\right)$ and 5 stimulus independent parameters $\left(\mu_{\eta_1}, \mu_{\eta_2}, \sigma_{\eta_1}, \sigma_{\eta_2}, \rho_\eta\right)$ for the additive noise. For any stimulus condition, there are only five relevant measurements that can be derived from the neural data (the response means and variances for each neuron in a pair, and their correlation), so the model is over-parametrized. Therefore, to apply the RoG to neural data, we need to reduce the total number of parameters.

The generality of this model provides a procedure for converting a standard normalization model (i.e. a model for the mean response) into a RoG model that specifies both mean and (co)-variance. In this paper, we use the example of contrast-gain control, which has been widely used to model the response of single neurons and neural populations to visual input with varying contrast [Albrecht and Hamilton, 1982; Clatworthy et al., 2003; Geisler and Albrecht, 1992; Heeger, 1992]. By adapting such a model, we can reduce the stimulus dependence of the means of the numerator and denominator $(\mu_{N_1}, \mu_{N_2}, \mu_{D_1}, \mu_{D_2})$. In the contrast-gain control model, the neural response as a function of contrast c $(0 - 100\%)$ is computed as a "hyperbolic ratio" [Albrecht and Hamilton, 1982; Heeger, 1992]:

$$R(c) = \frac{R^{\max} c^2}{\epsilon^2 + c^2} + R_0 \tag{8}$$

Where $R^{\max}$ is the maximum response rate, ϵ is the semi-saturation constant (the contrast at which $R(\epsilon) = R^{\max}/2$) to prevent division by 0, and R_0 is the spontaneous activity of the neuron (the response at 0% contrast). We can convert this standard model into an RoG by setting the mean of the numerator and denominator in the RoG to the numerator and

denominator in this equation:

$$\mu_N = R^{\max} c^2, \;\; \mu_D = \epsilon^2 + c^2 \;\; , \mu_\eta = R_0 \tag{9}$$

By using this functional form, we can substitute the stimulus dependent parameters of the RoG $(\mu_{N_1}, \mu_{N_2}, \mu_{D_1}, \mu_{D_2})$ with the stimulus independent parameters $R_1^{\max}, R_2^{\max}, \epsilon_1, \epsilon_2$. Another model simplification is to assume that individual neural variability and mean neural response are related by a power function as has been observed in the visual cortex [Carandini, Heeger, and Movshon, 1997; Goris, Movshon, and Simoncelli, 2014; Gur, Beylin, and Snodderly, 1997]:

$$\sigma^2 = \alpha \cdot \mu^\beta \tag{10}$$

This parametrization allows the Fano Factor (the ratio of the variance to the mean) to vary with stimulus input (as long as $\beta \neq 1$) and for both over-dispersion (Fano factor > 1) and under-dispersion (Fano factor < 1). Moreover, as with the mean, the four stimulus dependent variance parameters of the model $(\sigma_{N_1}, \sigma_{N_2}, \sigma_{D_1}, \sigma_{D_2})$ can be replaced with four pairs of stimulus independent parameters $(\boldsymbol{\alpha_N}, \boldsymbol{\beta_N}, \boldsymbol{\alpha_D}, \boldsymbol{\beta_D})$. Lastly, in principle, the parameters controlling correlation (ρ_N, ρ_D) can vary with stimulus conditions but for computational simplicity we assume that (ρ_N, ρ_D) are stimulus independent. However, even with this assumption, our model can capture stimulus dependent noise correlations (Fig. 5) as often observed *in vivo* [Kohn and Smith, 2005; Ponce-Alvarez et al., 2013; Sadagopan and Ferster, 2012].

2.3 Fitting the RoG to Data

We optimize the values of the parameters, given a dataset, by maximum likelihood estimation. In this paper, we validate various properties of the pairwise RoG using synthetic data produced from the generative model. We will demonstrate the applicability of this

model to neural data analysis by fitting the pairwise RoG to calcium imaging data (see subsection 2.7).

Based on our previous discussion, we assume that the model parameters (collectively denoted Θ) are stimulus independent. We consider our dataset $\{\boldsymbol{R}_t(s)\}$ where s is the stimulus and t indexes the trial. We assume that, for each stimulus, our data is independent and identically distributed according to $\mathcal{N}\left(\boldsymbol{\mu_R}(s), \boldsymbol{\Sigma_R}(s); \Theta\right)$, and that data is independent across stimuli. We can therefore compute the negative log-likelihood of the data using the following equation (see Appendix 4 for derivation):

$$\begin{aligned} -\log p(\boldsymbol{R}_t(s);\Theta) \approx \frac{T}{2}\sum_{s\in S}\Big\{ &\log|\boldsymbol{\Sigma_R}(s,\Theta)| + \mathrm{Tr}(\boldsymbol{\Sigma_R}^{-1}(s,\Theta)\widehat{\boldsymbol{\Sigma_R}}(s)) \\ &+ (\widehat{\boldsymbol{\mu_R}}(s) - \boldsymbol{\mu_R}(s,\Theta))^\top \boldsymbol{\Sigma_R}^{-1}(s,\Theta)(\widehat{\boldsymbol{\mu_R}}(s) - \boldsymbol{\mu_R}(s,\Theta))\Big\} \end{aligned} \tag{11}$$

Where $\widehat{\boldsymbol{\mu_R}}(s)$ and $\widehat{\boldsymbol{\Sigma_R}}(s)$ are the empirical mean and covariance across trials computed from the data.

In practice, we have found that it is computationally faster to first optimize the parameters for each neuron in the pair separately (which is equivalent to fitting the independent RoG model), and then optimize the correlation parameters (i.e. the ρ parameters) with the single-neuron model parameters fixed. This two-step optimization process is referred to as the inference functions for marginals method in the copula literature, and is known to be mathematically equivalent to maximum likelihood estimation for Gaussian copulas [Joe and Xu, 1996], which is the case we consider here. This points to an extension or alternative to the pairwise RoG presented here that considers the bivariate distribution to be some non-Gaussian copula with Gaussian marginals, which we leave for future work. We assumed that the pairwise distribution is Gaussian for computational simplicity, but others have used non-Gaussian copulas to model neural populations [Berkes, Wood, and Pillow, 2008].

2.3.1 Cross-validated Goodness of Fit

To measure the quality of model fit, we used a cross-validated pseudo-R^2 measure [Coen-Cagli, Kohn, and Schwartz, 2015], as follows. During fitting, we divided the recording trials for each pair and for each stimulus into training and test sets (for simulation studies, we used two or ten-fold cross validation; for the calcium analysis we used leave-one-out cross-validation). We then fit the parameters of the model for each training set and used the following equation to assess the model prediction on the held out data:

$$\text{goodness of fit} = \frac{LL_{\text{fit}} - LL_{\text{null}}}{LL_{\text{oracle}} - LL_{\text{null}}} \tag{12}$$

Where LL_{fit} is the (negative) log-likelihood (using Eq. 11) for the test data using the optimized parameters, LL_{null} is the log-likelihood of the data assuming that there is no modulation of the responses by stimulus contrast, and LL_{oracle} is the log-likelihood of the data using the empirical mean, variance, and covariance of the training data. The reported goodness of fit score is the median across all training and test splits of the computed score (Eq. 12). Because of this cross-validation, goodness of fit values can be < 0 (the fit model is worse than the null model) or > 1 (the fit model performance is better than the oracle).

2.3.2 Quantifying the Accuracy of the Estimated Correlation Parameters

As we are interested in interpreting the correlation model parameters (ρ_N, ρ_D), we need to assess the accuracy of the maximum likelihood estimator. For simulations, we directly compare the estimated ρ values to the true values used to generate the data, measured by the mean squared error and its decomposition into bias and variance. For real neural data, however, we do not have access to the true values: instead, we compute confidence intervals. To do so, we perform a bootstrap fit procedure: given a set of pairwise neural responses $\{\boldsymbol{R}_t(s)\}$ with T simultaneously recorded trials, we sample these trials with replacement T times and then fit the pairwise RoG using the resampled set of neural responses as our

observations. Repeating this procedure for a large number of samples (in the analysis in subsection 3.2 and subsection 3.4 we used 1000 bootstrap samples) gives us sets of fit ρ_N, ρ_D, which we use to compute a 95% confidence interval. We then demonstrate one possible use of the confidence intervals, with an application focused specifically on the sign of the ρ parameters. We consider our estimate of the sign accurate if the ρ parameters are significantly different from zero, which in turn corresponds to cases in which the confidence interval does not include zero. Note however that when the ρ parameters are not significantly different from zero, that does not imply that the estimate is inaccurate.

2.4 Model Comparison

We compared the pairwise RoG to a modified version of the modulated Poisson model [Goris, Movshon, and Simoncelli, 2014], using Gaussian noise instead of Poisson [Aljadeff et al., 2016]. We call this model the *Modulated Gaussian* (MG). The original model is a compound Poisson-Gamma distribution, in which the Poisson rate parameter is the product of the mean tuning curve and a random gain variable that is Gamma distributed. The parameters of the Gamma distribution depend on the mean tuning curve ($f(s)$) and the variance of the gain variable (σ_G). Additionally, there are two sources contributing to (tuned) covariability: the correlation between the multiplicative gains (ρ_G), and the correlation between the Poisson processes (ρ_P). For the modulated Gaussian model, we use a bivariate Gaussian distribution whose moments (i.e. mean, variance, and covariance) are parametrized according to the moments of the modulated Poisson model. We made this modification to the modulated Poisson model for two main reasons. First, because we are examining continuously valued fluorescence traces as opposed to discrete spike count data, a continuous distribution is more appropriate for analysis. Second, the original modulated Poisson, while including a parametrization of the noise covariance between neurons, has no simple closed form for the bivariate distribution, which complicates the comparison of goodness of fit between the two models. By using a bivariate Gaussian distribution, we can more directly compare this model

to our proposed pairwise RoG.

More explicitly, the pairwise neural responses are distributed as:

$$
\begin{aligned}
\boldsymbol{R}_t &= (R_1, R_2)_t \sim \mathcal{N}(\boldsymbol{\mu}_{\boldsymbol{R}}^G, \boldsymbol{\Sigma}_{\boldsymbol{R}}^G) \\
\mu_{R_i} &= f(s) + \mu_{\eta_i} = R_i^{\max} \cdot \frac{c^2}{\epsilon_i^2 + c^2} + \mu_{\eta_i} \quad \text{for } i = 1, 2 \\
\sigma_{R_i}^2 &= \mu_{R_i} + \sigma_{G_i}^2 \mu_{R_i}^2 + \sigma_{\eta_i}^2 \\
\mathrm{Cov}(R_1, R_2) &= \rho_P \sqrt{\mu_{R_1} \mu_{R_2}} + \rho_G \sigma_{G_1} \sigma_{G_2} \mu_{R_1} \mu_{R_2} + \rho_\eta \sigma_{\eta_1} \sigma_{\eta_2}
\end{aligned}
\tag{13}
$$

where we assume the mean tuning cure is the contrast-response curve (Eq. 8), σ_{G_i} is the standard deviation of the multiplicative gain for neuron i, and ρ_G is the correlation between the multiplicative gains. ρ_P is no longer interpreted as the point process correlation; instead, ρ_P controls the portion of the tuned covariability that is independent of the shared gain. Similarly to the RoG, we also model the untuned variability η as additive bivariate Gaussian noise. We then fit the model parameters to data by minimizing the negative log-likelihood (Eq. 11, with $\boldsymbol{\mu}_{\boldsymbol{R}}^G$ and $\boldsymbol{\Sigma}_{\boldsymbol{R}}^G$ defined in Eq. 13). As with the pairwise RoG, we use cross-validation to account for model complexity and compute the goodness of fit scores using Eq. 12. An extension to this model was recently proposed that incorporates normalization by assuming the rate parameter is a ratio term in which the denominator is a Gaussian random variable, then deriving moments of the distribution similar to our work [Hénaff et al., 2020]. However, this model does not currently account for noise correlations, so we chose to instead adapt the Poisson-Gamma model.

The most relevant difference between RoG and MG is in how each model accounts for the effect of normalization on (co)variability. In the RoG, normalization directly influences variability by division operating on random variables. This creates flexible dependencies between the mean firing rate, individual neuron variability and shared covariability. In the MG, normalization influences the gain of neurons through the interaction between the mean

firing rate (i.e. the standard normalization model) and the gain parameter σ_G, which is assumed be a slowly fluctuating source of variability that scales how the mean firing rate effects variability. In this way, the normalization signal for the MG is a deterministic factor. The MG is therefore a simpler model that can only account for overdispersion, whereas the RoG allows for both overdispersion and underdispersion, and for diverse patterns of covariability, as we show in Fig. 5, albeit at the cost of additional parameters.

2.5 Inference of Single Trial Normalization from Measured Neural Activity

Because of the probabilistic formulation of the RoG, we can use Bayes theorem to compute the posterior probability of the normalization variable $\boldsymbol{D}_t$ in a single trial, given the observed neural responses $\boldsymbol{R}_t$:

$$\begin{aligned} p(\boldsymbol{D}_t|\boldsymbol{R}_t) &\propto p(\boldsymbol{R}_t|\boldsymbol{D}_t)p(\boldsymbol{D}_t) = p((R_1, R_2), (D_1, D_2)) \\ &\propto p((N_1, N_2), (D_1, D_2))|D_1 D_2| = p(N_1, N_2)p(D_1, D_2)|D_1 D_2| \end{aligned} \tag{14}$$

Where multiplication by $|D_1 D_2|$ occurs due to the change of variables formula for probability density functions. From this distribution, we can find the maximum a posteriori (MAP) estimate of the normalization strength in a single trial by differentiating the posterior distribution with respect the denominator variables and finding the maxima by setting the partial derivatives to 0. For ease of computation, we solve the equivalent problem of finding the zeros of the partial derivatives of the negative logarithm of the posterior distribution. Additionally, we assume that there is zero additive noise. In our previous work [Coen-Cagli and Solomon, 2019], we found that, when subtracting the mean additive noise from the simulated activity, the MAP estimate remained unbiased. Thus, for simplicity, we assume that we can subtract off the mean spontaneous activity and consider instead the posterior $p\left(\boldsymbol{D}|\boldsymbol{R} - \boldsymbol{\mu}_{\boldsymbol{\eta}}\right)$. To obtain an estimate for the denominator strength $\boldsymbol{D}_t$ we look at the partial derivatives

of the negative log posterior with respect to D_1, D_2 and solve to obtain the MAP estimate. This procedure leads to a two dimensional system of bivariate quadratic equations:

$$\begin{aligned} 2A_1{D_1}^2 + B_1D_1 + CD_1D_2 - 2 = 0 \\ 2A_2{D_2}^2 + B_2D_2 + CD_1D_2 - 2 = 0 \end{aligned} \tag{15}$$

Where the coefficients A_1, A_2, B_1, B_2, C are functions of the parameters of the model (see section 4 for the derivation of Eq. 15 and the full expressions of these coefficients).

A basic result from the algebra of polynomial systems (Bézout's theorem) tells us that this system has four pairs of possibly complex valued solutions [Sturmfels, 2002]. In fact, as solving this system amounts to solving a quartic equation in one variable, there exists an algebraic formula (with radicals) for solutions to this system as a function of the coefficients. This solution is too long to include here and uninformative, but was found using the Symbolic Math toolbox from MATLAB and is included in our toolbox subsection 2.8.

Because all of the variables involved are real valued, we are only interested in the existence of real solutions to this two-dimensional system. However, there is no theoretical guarantee that there will be any real solutions. In practice we take the real part of the algebraic solution to this system and find which pair of solutions $\left(D_1^{MAP}, D_2^{MAP}\right)$ minimize the negative log posterior. Alternatively, we can consider finding the MAP by directly minimizing Eq. 29 (see section 4) using numerical optimization. We have verified that, when real-valued solutions exist to Eq. 15, these coincide with numerically minimizing Eq. 29. However, as optimization of Eq. 29 must be computed on a per-trial basis, it is far too time consuming to perform when there are many experimental trials, so we utilize the algebraic solution to Eq. 15.

2.6 Generating Realistic Pairwise Neural Activity from the Model

To constrain our simulations to realistic parameter values for the contrast response function (Eq. 8), we took the single-neuron best fit parameters to macaque V1 data analyzed in our

previous work (for details see [Coen-Cagli and Solomon, 2019]) and created parameter pairs by considering all combinations ($N = 11628$ pairs) of these parameters. Using the generative model for the pairwise RoG (Eq. 2) and the contrast-response parametrization (Eq. 8), we can simulate single-trial neural activity from these parameter pairs and specific values for (ρ_N, ρ_D). These synthetic data allow us to explore properties of the pairwise model without having to exhaustively explore the full parameter space.

2.7 Data Collection and Processing

2.7.1 Animal Preparation

Data were collected from CaMKII-tTA;tetO-GCaMP6s mice [Wekselblatt et al., 2016], expressing GCaMP6s in cortical excitatory neurons. Mice were implanted with headplates and cranial windows over V1 [Sridharan et al., 2022]. Briefly, mice were anesthetized with 2% isoflurane and administered 2 mg/kg of dexamethasone and 0.5 mg/kg of buprenorphine. Animals were secured in a stereotaxic frame (Kopf) and warmed with a heating pad. The scalp was removed and the skull was lightly etched. A craniotomy was made over V1 using a 3.5 mm skin biopsy bunch. A cranial window, consisting of two 3 mm diameter circular coverslips glued to a 5 mm diameter circular coverslip, was placed onto the craniotomy, and secured into place with Metabond (C&B). Then a custom made titanium headplate was secured via Metabond (C&B) and the animals were allowed to recover in a heated cage.

2.7.2 Behavioral Task and Visual Stimuli

During imaging, mice were head-fixed in a tube, and were performing an operant visual detection task [Bounds et al., 2021]. Briefly, mice were trained to withhold licking when no stimulus was present, and lick within a response window after stimulus presentation. Mice were water-restricted and given a water reward for correct detection. Visual stimuli were drifting sinusoidal gratings (2 Hz, 0.08 cycles/degree) presented for 500 ms followed by a 1000 ms response window. Stimuli were generated and presented using PsychoPy2 [Peirce

et al., 2019]. Visual stimuli were presented using a gamma corrected LCD monitor (Podofo, 25 cm, 1024x600 pixels, 60 Hz refresh rate) located 10 cm from the right eye. Contrast of gratings were varied between 7 different levels: {2, 8, 16, 32, 64, 80, 100}, except for 2 recording sessions in which contrast level 80% was omitted. This did not alter any of the analysis, allowing sessions to be combined into a single dataset.

2.7.3 Calcium Imaging

Once they learned the task, mice started performing under the 2p microscope, and V1 was imaged via cranial window. Imaging was performed using a 2-photon microscope (Sutter MOM, Sutter inc.), with a 20X magnification (1.0 NA) water-immersion objective (Olympus Corporation). Recordings were done in L2/3 in a 800 x 800 µm field of view, with 75-100 mW of 920 nm laser light (Chameleon; Coherent inc). An electrically-tunable lens (Optotune) was used to acquire 3 plane volumetric images at 6.36 Hz. Planes were 30 µm apart. Acquisition was controlled with ScanImage (Vidrio Technologies).

Calcium imaging data was motion-corrected and ROI extracted using suite2p [Pachitariu et al., 2017], and all data was neuropil subtracted with a coefficient of 0.7 (we also analyzed data using neuorpil coefficients of 0.4 and 1, not shown, and see section 4 for additional analysis with deconvolved data; all the results presented in the main text were qualitatively similar across preprocessing methods).

2.7.4 Data Processing

Processing of calcium imaging data was performed using custom MATLAB code. Fluorescence traces for individual trials and cells (average of the neuropil subtracted fluorescence across a ROI) consisted of 24 frames: 4 frames of pre-stimulus blank, followed by 3 frames of stimulus presentation and 17 frames of post-stimulus blanks corresponding to the response window for the behavioral task. In our analyses, we considered one extra frame to account for onset delays and calcium dynamics. Baseline fluorescence (F_0) was computed as the

median across pre-stimulus frames (1-5), and the stimulus evoked fluorescence ($\Delta F/F$) was computed as the mean of the normalized fluorescence per frame ($(F(i) - F_0)/F_0$ for $F(i)$ the fluorescence of frame i) across frames corresponding to stimulus response. Spontaneous $\Delta F/F$ was computed as above during blank trials. Cells were included in further analysis if the evoked response at the highest contrast was at least 2 standard deviations above the spontaneous mean fluorescence. Across 9 recording sessions, 295/8810 neurons met this inclusion criterion. This small percentage is due to sessions recorded using gratings with fixed spatial frequency and orientation; thus, included neurons are visually responsive and selective for this combination.

2.8 Code and Data Availability

Code for running model simulations and fitting model to experimental data was written in MATLAB and will be made publicly available on Github (https://github.com/oren-weiss/pairwiseRatioGaussians). Calcium imaging data will be made publicly available on Zenodo.

3 Results

We developed the pairwise Ratio of Gaussians model (RoG, Fig. 1a) to quantify the relationship between normalization and response covariability (i.e. noise correlations) that has been suggested in some empirical studies of neural activity in visual cortex [Lin et al., 2015; Rikhye and Sur, 2015; Ruff and Cohen, 2016; Verhoef and Maunsell, 2017]. In the standard divisive normalization model (Eq. 1), the mean response is computed as the ratio of the excitatory drive (numerator $\boldsymbol{N}$) to a normalization signal summarizing inhibitory drive (denominator $\boldsymbol{D}$). Our pairwise RoG considers a pair of neurons where each individual neural response is well characterized by the standard normalization equation with corresponding numerators (N_1, N_2) and denominators (D_1, D_2). We then assume that the numerators and

denominators are bivariate Gaussian random vectors - which allows the possibility for correlations to exist among the numerators (denoted ρ_N) and among the denominators (ρ_D). From this, we derived equations for the mean responses and covariance matrix of the pair as a function of the numerators and denominators (Eqs. 4, 5). These equations depend on the Gaussian approximation to the ratio of two Gaussian random variables. We verified the validity of this approximation, by simulating the activity of pairs of neurons, and comparing the covariance (Fig. 1b,c) and correlation (Fig. 1d,e) of the true ratio distribution and of the approximate distribution (Eqs. 6, 7). The mean and variance are identical to the independent RoG model (a special case of the pairwise RoG, with numerators and denominators independent between neurons) which we validated previously [Coen-Cagli and Solomon, 2019].

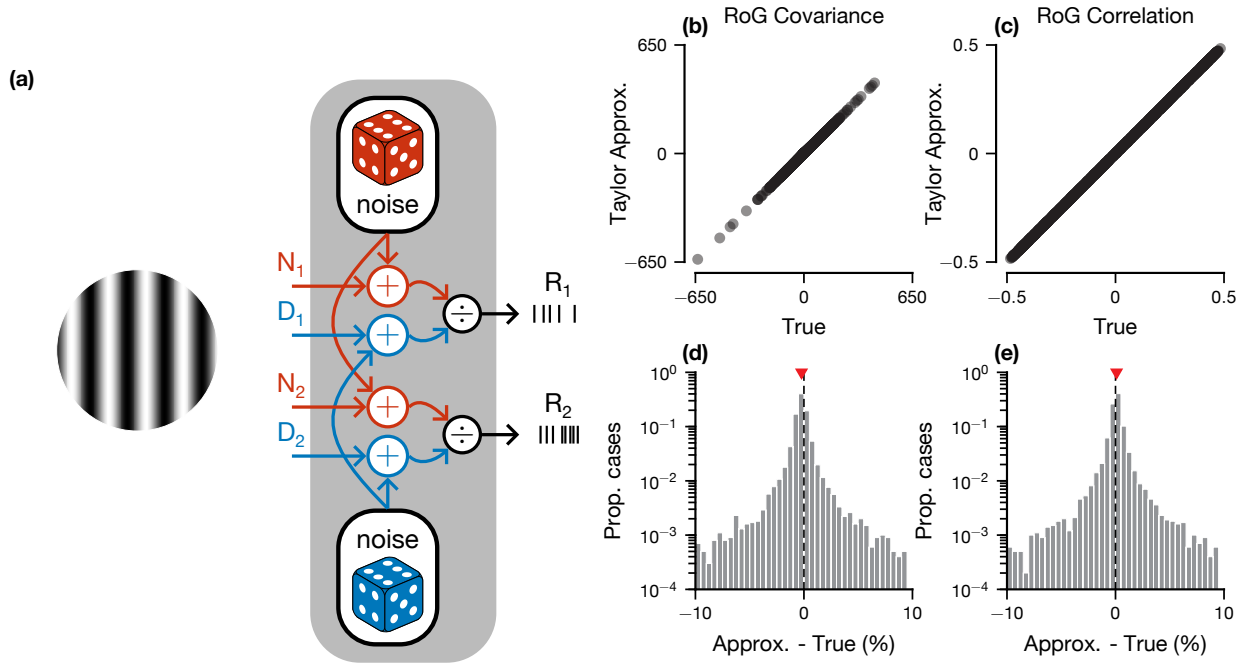

Figure 1: Definition and Validation of the Pairwise Ratio of Gaussians Model
(a) The pairwise RoG model describes pairs of neural responses (R_1, R_2), where each response is computed as the ratio of two stimulus-driven signals: numerators (N_1, N_2) and denominators (D_1, D_2). These variables are corrupted by additive bivariate Gaussian noise with a power-law variance-mean relationship.
(b-e) Comparison of the normal approximation we derived for the pairwise RoG noise covariance (b,c) and noise correlation (d,e) and the true values (estimated across $1e6$ simulated trials) for $1e4$ experiments (i.e. simulated pairs of neural responses). Each experiment used different model parameters and each trial was randomly drawn from the corresponding distribution. (b,d) scatter plot; (c,e) histogram of the percent difference between the Taylor approximation and the true value. The red marker indicates the median percent difference.
Model parameters were drawn uniformly from the following intervals:
$\boldsymbol{\mu_N} \in [0, 100]$, $\boldsymbol{\mu_D} \in [0.5, 1.5]$, $\boldsymbol{\mu_\eta} = 0$, $\boldsymbol{\alpha_N} = 1$, $\boldsymbol{\alpha_D} = 0.001$, $\boldsymbol{\beta_N}, \boldsymbol{\beta_D} \in [1, 1.5]$, $\rho_N, \rho_D \in [-0.5, 0.5]$, $\boldsymbol{\sigma_\eta^2} = 0.1 \cdot \boldsymbol{\sigma_N^2}/\boldsymbol{\sigma_D^2}$, $\rho_\eta = 0$. The ranges of the mean parameters were chosen to reproduce realistic firing rates of V1 cortical neurons, while the $\boldsymbol{\alpha}, \boldsymbol{\beta}$ parameters were chosen such that the variances of the $\boldsymbol{N}$ and $\boldsymbol{D}$ are relatively small and the probability that $\boldsymbol{D} \leq 0$ is negligible [2013].

3.1 Modulations of Correlated Variability Depend on Sharing of Normalization

Within the RoG modeling framework, there are two main sources of response (co)-variability: the numerator (excitatory drive) and the denominator (normalization). Depending on the

value of the corresponding ρ parameters, each of these sources of variability can be independent ($\rho = 0$) or shared ($\rho \neq 0$), and therefore contribute differently to noise correlations. Consequently, understanding modulations of noise correlations in the pairwise RoG requires understanding how normalization and external stimuli affect the relative amplitude of these sources, and how the effects depend on whether those sources are shared.

First, we studied the relationship between normalization and noise correlations for the lowest contrast stimuli (Fig. 2, yellow symbols). We define normalization strength for a given neuron as the mean of the denominator; for a fixed stimulus contrast, this is determined by the semi-saturation constant ϵ in Eq. 8. When the normalization signals are positively correlated (shared normalization, $\rho_D = 0.5$) but the excitatory drive is independent ($\rho_N = 0$), increasing normalization strength tends to decrease the magnitude of noise correlations (Fig. 2a). Conversely, when the normalization signals are independent ($\rho_D = 0$) but there is shared driving input ($\rho_N = 0.5$), the magnitude of noise correlations tends to increase with increasing normalization (Fig. 2b). Intuitively, this reflects that increasing the normalization strength reduces the contribution of fluctuations in normalization signal to the overall response fluctuations [Coen-Cagli and Solomon, 2019]. Therefore, increasing normalization strength reduces a source of noise correlations when normalization is shared, whereas it reduces a source of independent noise when normalization is not shared.

We found similar effects of normalization on noise correlations at high contrast levels (Fig. 2, orange and red symbols) although the slope of the relationship became shallower as contrast increased. This is due to the saturating effect of the contrast response function (Eq. 8): at high contrast, (co)fluctuations of the normalization signal across trials have relatively little effect on the responses of a neural pair, so the correlation in neural responses will be relatively unaffected by normalization strength. We also observed that the magnitude of noise correlations generally increased with contrast when the denominators were correlated and the numerators independent (Fig. 2a and S3), whereas it decreased when the numerators

were correlated (Fig. 2b,c and S3). Importantly, for the analysis of normalization strength at fixed contrast, we used the contrast semi-saturation constants (Eq. 8; i.e. a pure change in the denominator) as a measure of normalization strength. Conversely, increasing stimulus contrast increases both the numerator and denominator (Eq. 8). This explains our observation that, even though normalization is stronger at higher contrast, noise correlations can be modulated in different ways by contrast and by normalization, because changing stimulus contrast also affects the numerator term.

Indeed, these results can be derived from examining the equation for the correlation in the pairwise RoG (Eq. 7) rearranged as follows (see Appendix section 4 for more details):

$$\mathrm{corr}\left(R_1, R_2\right) = \frac{\rho_N}{\sqrt{\left(\frac{\delta_{D_1}^2}{\delta_{N_1}^2}+1\right)\left(\frac{\delta_{D_2}^2}{\delta_{N_2}^2}+1\right)}} + \frac{\rho_D}{\sqrt{\left(\frac{\delta_{N_1}^2}{\delta_{D_1}^2}+1\right)\left(\frac{\delta_{N_2}^2}{\delta_{D_2}^2}+1\right)}} \tag{16}$$

and by recognizing that the coefficient of variation of D_i (δ_{D_i}) is a decreasing function of normalization strength (μ_D or ϵ), whereas the ratio $\frac{\delta_{D_i}}{\delta_{N_i}}$ is *often* an increasing function of contrast.

We further analyze this equation, first in the case of changing normalization strength while keeping contrast fixed. The term on the right-hand side proportional to ρ_N is an increasing function of the mean normalization strength since the denominator is a decreasing function of μ_D. Conversely, the term proportional to ρ_D decreases with normalization since the denominator increases with μ_D. In these two cases, the monotonic dependence of noise correlation on normalization strength is guaranteed by Eq. 16 regardless of the specific parameter values. Similar patterns emerge when $\rho_N, \rho_D < 0$, except the signs of the noise correlations are reversed (Fig. S3). When the correlations of input and normalization signals are both different from zero, the relationship between noise correlation and normalization strength resembles a combination of the two previously described scenarios, and the specific parameter values determine which of the two terms in Eq. 16 dominates. For instance, in

our simulations with ($\rho_N = 0.5, \rho_D = 0.5$) (Fig. 2c), the magnitude of noise correlations increased with normalization strength on average similar to (Fig. 2b), indicating that the magnitude of the term proportional to ρ_N is usually larger than the magnitude of the ρ_D term, but this trend is not consistent, as evidenced by the increased spread of the scatter. When the input strength and normalization signal have opposite correlations (e.g. $\rho_N = 0.5, \rho_D = -0.5$), we obtained similar results but the magnitude of noise correlations was on average closer to 0 due to cancellation between the two sources of covariability (Fig. S3). A similar analysis of Eq. 16 shows that the effects of stimulus contrast are opposite to those of a pure change of normalization strength, because $\frac{\delta_{D_i}}{\delta_{N_i}}$ is often an increasing function of contrast but a decreasing function of normalization strength. Notice that we assumed there is no residual noise component ($\eta \equiv 0$), but all the analyses above remain valid when the amplitude of noise variance σ_η is relatively small compared to (δ_N, δ_D) (Appendix, Eq. 21).

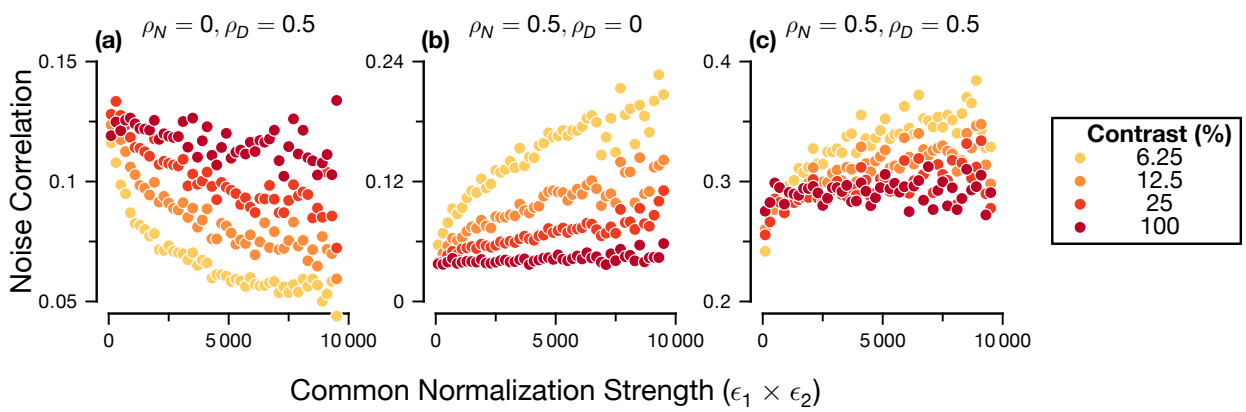

Figure 2: Relationship between Noise Correlations, Normalization Strength and Contrast Depends on the Source of Variability

Each panel shows, for a combination of ρ_N, ρ_D specified in the panel title, the median noise correlation of all generated neural pairs binned according to $\epsilon_1\epsilon_2$, a contrast independent measure of the common normalization strength. Bins with less than 100 pairs were discarded. Neural responses were generated from the contrast-response parametrization (Eq. 8). Noise correlation strength was computed across 1e3 simulated trials drawn from the pairwise RoG model. For each contrast level and combination of ρ_N, ρ_D, 1e5 simulated experiments were created. See Fig. Figure S4 for a more systematic exploration of the factors influencing modulation of noise correlations for two simulated pairs that matches the large scale experiment considered here. Model parameters were drawn uniformly from the following intervals:
$\boldsymbol{R^{max}} \in [5, 50],\ \boldsymbol{\epsilon} \in [10, 100],\ \boldsymbol{\alpha_N}, \boldsymbol{\alpha_D} \in [0.1, 1],\ \boldsymbol{\beta_N}, \boldsymbol{\beta_D} \in [1, 2],\ \boldsymbol{\eta} \equiv 0$

In summary, our analysis shows that normalization and stimulus contrast can have diverse effects on noise correlations, depending on whether neurons share their normalization signals and on the interplay between multiple sources of variability.

3.2 Inference of Correlation Parameters

The above analysis demonstrates that the relationship between noise correlations and normalization depends on how this correlated variability arises: either through cofluctuations in the excitatory drive or normalization signal (determined by ρ_N, ρ_D respectively). To employ these insights when fitting to data, we need to know how well we can infer these parameters from data. To do so, we generated synthetic neural data using realistic values for the single-neuron parameters (see subsection 2.6) and uniformly randomly sampled ρ_N, ρ_D parameters in the range [-0.9,0.9]. We then assessed the quality of the maximum likelihood estimate of the parameters by measuring the mean squared error with the true values, and we also calculated bootstrapped confidence intervals (see subsubsection 2.3.2).

We see in Figure 3a the comparison between the true and inferred ρ values. In general, the maximum likelihood estimators are largely unable to recover the true generating values (the overall Pearson correlation between the fit and true $(\rho_N, \rho_D) = (0.48, 0.18)$, the mean (across pairs) squared error is $(0.32, 0.57)$), indicating that these parameters are not identifiable with the parametrization we considered (contrast tuning). Further, this analysis indicates that inference of ρ_D is in general more difficult than it is for ρ_N. The lack of identifiability of these parameters is likely due to the multiple multiplicative interactions between the parameters. By looking closer at Eq. 7, we can see that the contribution of the ρ parameters to the noise correlation is multiplied by the respective standard deviations for the numerator or denominator variables. Because of our power law parametrization of the variance of these Gaussian variables (Eq. 10), we have an additional multiplicative parameter α that potentially makes it difficult to infer the exact value of the ρ parameters (see section 4 for

further discussion).

However, rather than the exact magnitude of the ρ parameters, we are often only interested in the *sign* of these parameters, as in deriving the relationship between normalization strength and noise correlations (subsection 3.1). Moreover, because all of the parameters in Eq. 7 are *a priori* positive besides the ρ parameters, we reasoned that the accuracy of sign inference might higher. Therefore, we considered the subset of cases where the parameters are significantly different from 0, which we define as cases where the 95% confidence interval does not include 0 and the pairwise goodness of fit is greater than the independent goodness of fit (see subsubsection 2.3.1). First, for those pairs (for ρ_N, 4306/11628 pairs were significant, for ρ_D, 1227/11628), we see a much stronger relationship between the fit and true ρ values (Fig 3a, darker colors, Pearson correlation between significant fit and true $(\rho_N, \rho_D) = (0.89, 0.61)$ mean (across pairs) squared error = (0.09,0.40)). Importantly, the plots also show that, for pairs with ρ parameters significantly different from 0, the sign of the inferred ρ parameter is very frequently equivalent to the sign of the true ρ parameter, in a similar proportion of cases for ρ_N and ρ_D. From this, we conclude that, for these significant ρ parameters, the *sign* of the inferred ρ parameter is accurate.

To better understand the quality of the maximum likelihood estimators, from the bootstrap samples used for the confidence intervals (subsubsection 2.3.2), we can compute the mean squared error on a per-pair basis (Fig. 3b-d). We reasoned that, as the strength of noise correlations increased, the quality of the estimator would also increase as this would provide more of a constraint on the optimization space. For ρ_N, across all synthetic pairs, we found this seems to be the case, as the mean squared error seems to decrease with increasing magnitude of noise correlations (Fig. 3b, left), while for ρ_D this did not seem to be the case (Fig. 3b, right). When we only consider the pairs in which the fit estimator was significantly different than 0 (Fig. 3b, darker colors), the overall bootstrapped mean squared error was lower than the nonsignifcant cases (with most lying below 0.5), confirming that the accuracy

for these pairs is generally better than for nonsignificant pairs, but with no clear dependence on the magnitude of the noise correlations. By decomposing the mean squared error into contributions from the bias and variance of the estimator (Fig. 3c,d), we can gain further understanding of the quality of the ρ estimators. For both ρ_N and ρ_D, it appears that the bias of the estimator is generally a larger contributor to the mean squared error than the variance of the estimator, i.e. that the estimators themselves are biased but the fits, at least for significant pairs, are largely stable under resampling of synthetic trials. Again, it appears there is a weak decrease in the bias and variance of the estimator with increasing noise correlation magnitude.

In summary, our analysis indicates that it is difficult to estimate the precise value of the ρ parameters in general, for the stimulus parametrization considered here (i.e. the classical normalization model for contrast tuning). The estimated values of ρ_N are more accurate than ρ_D, particularly when noise correlations are large. Lastly, we have provided a method to calculate confidence intervals on the estimates of ρ, and demonstrated how to use them to accurately recover the sign of the ρ parameters.

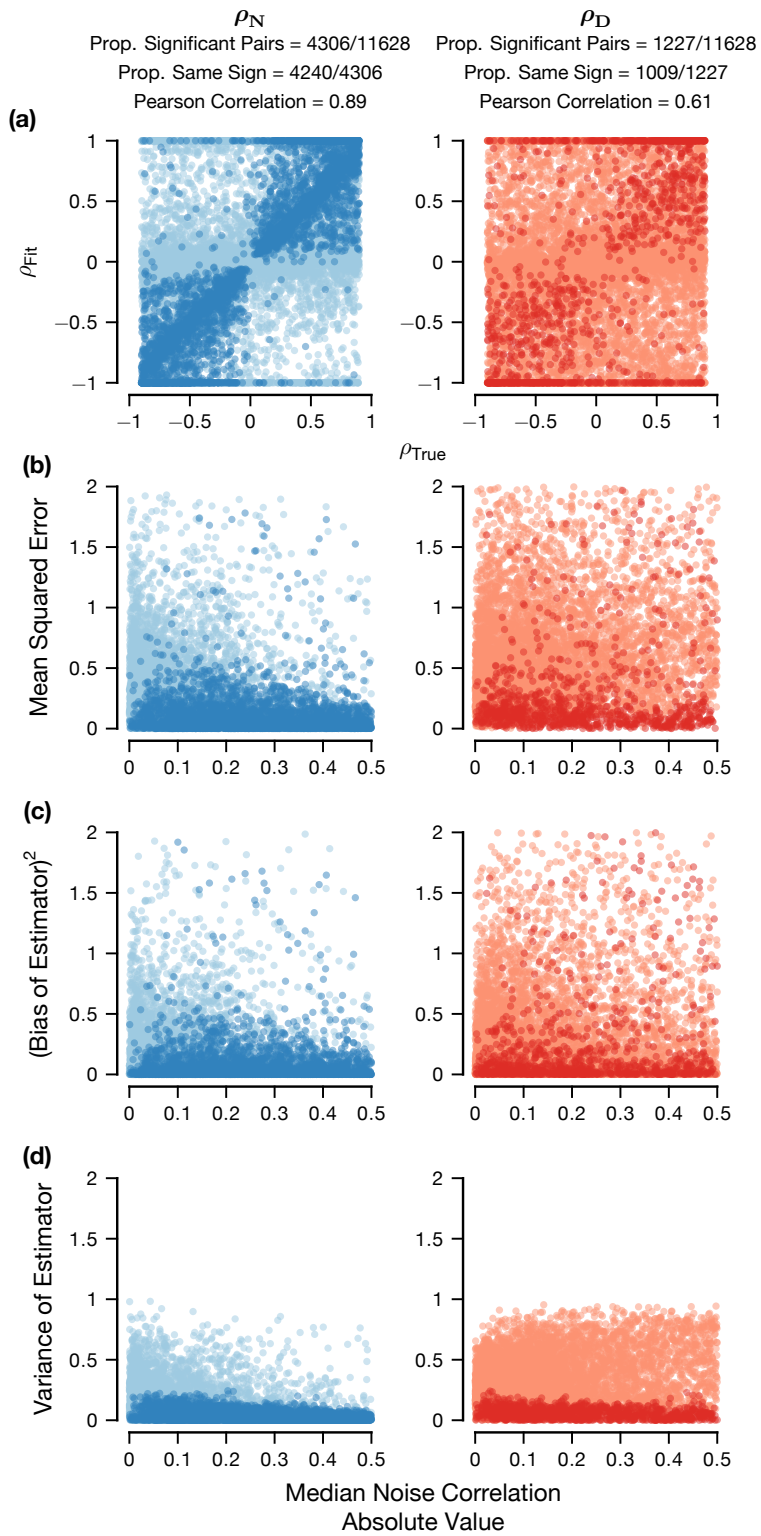
ρ_N
Prop. Significant Pairs = 4306/11628
Prop. Same Sign = 4240/4306
Pearson Correlation = 0.89
ρ_D
Prop. Significant Pairs = 1227/11628
Prop. Same Sign = 1009/1227
Pearson Correlation = 0.61
(a)
ρ_{Fit}
ρ_{True}
(b)
Mean Squared Error
(c)
(Bias of Estimator)^2
(d)
Variance of Estimator
Median Noise Correlation
Absolute Value

Figure 3: Accuracy of Inference of ρ Parameters
Plots were generated with 11628 synthetic parameter pairs with uniformly randomly generated $\rho_N, \rho_D \in [-0.9, 0.9]$, contrast levels $\{6.25, 12.5, 25, 50, 100\}$, 1000 synthetic trials and 1000 bootstrap resamples. The left column are the results of the analysis for ρ_N, the right column for ρ_D. The darker colors are the pairs where the 95% confidence interval for the respective ρ parameter excluded 0, whereas paler colors indicate where the c.i. included 0.
(a) Direct comparison between the true ρ value and the maximum likelihood estimator for the ρ value. (b) Mean squared error between the true ρ value and the fitρ value. The mean is computed over bootstrap samples (see subsubsection 2.3.2) as a function of the magnitude of the median noise correlations. We chose the magnitude as the plots are symmetric around 0. (c) The squared bias of maximum likelihood estimator as a function of the magnitude of the median noise correlations. The bias here is the difference between the mean across bootstrap samples and the true ρ value. (d) The variance of bootstrapped fit ρ values as a function of the magnitude of the median noise correlations.

3.3 Pairwise Model Improves Single-Trial Inference of Normalization Strength Even When Noise Correlations are Small

In past work that connected normalization to modulation of noise correlations, stimulus and experimental manipulations (e.g. contrast, attention) are used as proxies for normalization strength [Verhoef and Maunsell, 2017; Tripp, 2012; Ruff and Cohen, 2016] because normalization strength cannot be measured directly. However, these manipulations also affect other factors that drive neural responses (as we have illustrated in subsection 3.2, Fig. 3), which could confound these as measures of normalization signal. Therefore, quantitatively testing the relationship between noise correlations and normalization requires estimating the single-trial normalization strength for a pair of neurons. One of the advantages of our probabilistic, generative formulation of the pairwise RoG model (Eq. 2) is that it allows us to infer the single-trial normalization strength from measured neural activity (see subsection 2.5). The independent RoG model also provides an estimate for the single-trial normalization, which is known to be a valid estimator for the ground-truth normalization strength for data generated from the independent RoG [Coen-Cagli and Solomon, 2019]. We found similar results for the pairwise RoG (not shown), so we examined how the pairwise estimates for the single-trial normalization strength compares to estimates based on the independent model.

One possibility is that the estimate derived from the pairwise model would outperform the independent model as the magnitude of noise correlations increase. However, this is not necessarily the case. Figure 4a,b demonstrates this with two example synthetic neuron pairs. Because the single-trial normalization inference depends on the single-trial neural activity, correlations between neurons will induce correlations between the inferred normalization signals even for the independent RoG (Fig. 4a). However, when noise correlations are small due to cancellation between ρ_N and ρ_D, the independent model will infer minimal correlation between normalization signals while the pairwise model will correctly infer that the single-trial normalization is correlated (Fig. 4b).

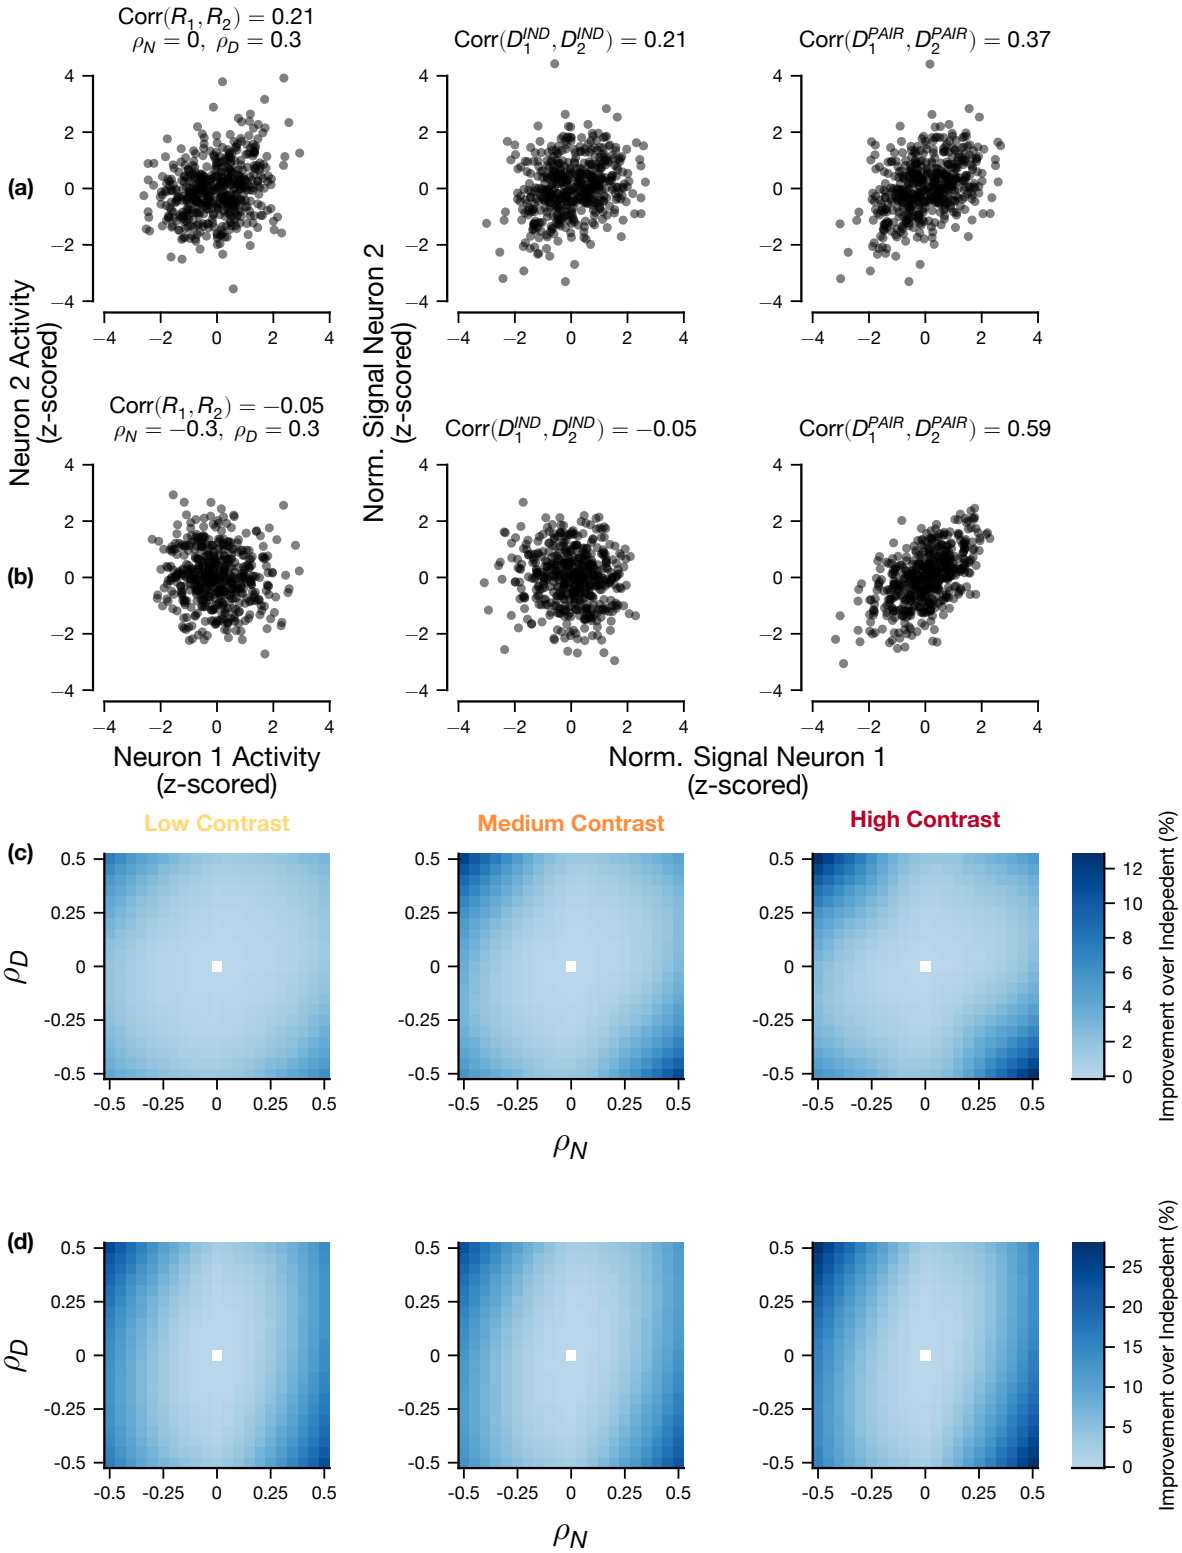

Figure 4 *(caption on next page)*

Figure 4: Inference of Single-Trial Normalization Depends on ρ Parameters and Contrast Level
(a,b) Two simulated experiments drawn from the pairwise RoG with different noise correlations arising from different underlying ρ_N, ρ_D values to compare the pairwise and independent estimators of single-trial normalization with the ground truth normalization signal. (a) has overall noise correlation of 0.21 across contrasts generated with $\rho_N = 0, \rho_D = 0.3$, (b) has overall noise correlation of -0.05 across contrasts generated with $\rho_N = -0.3, \rho_D = 0.3$. Z-scoring performed across trials. Random parameters drawn from $R^{max} \in [10, 100], \boldsymbol{\epsilon} \in [15, 25], \boldsymbol{\alpha_N}, \boldsymbol{\alpha_D} \in [0.1, 1], \boldsymbol{\beta_N}, \boldsymbol{\beta_D} \in [1, 1.5], \boldsymbol{\eta} \equiv 0$. Contrasts levels were {6.25, 12.5, 25, 50, 100}.
(c,d) Comparison of the single-trial normalization inference in the pairwise and independent RoG models for the mean-squared error (c) and the correlation between the estimate and the true value (d), as it depends on ρ_N, ρ_D and the contrast levels. Left: lowest contrast level (6.25); middle: intermediate contrast level (25); right: full contrast (100). Each bin corresponds to the median difference between the pairwise and independent models across simulated experiments.
(c,d) used 11628 synthetic pairs (see subsection 2.6).

To demonstrate this principle, we computed the difference of the mean squared error (Fig. 4c) and correlations (Fig. 4d) between the pairwise RoG and independent RoG normalization inference (normalized with respect to the ground truth values) for a large number of simulated pairs (11628) with systematically varying ρ_N, ρ_D values. First, we found that the distinction between the two models depended on the contrast level: as contrast levels increase, the quality of the pairwise and independent estimators of normalization signal become more distinct. Second, the improvement of the pairwise RoG estimate over the independent estimate increased when the magnitude of the ρ_D parameter increased. Consistent with the intuition provided by Fig. 4b, the largest improvement occurred when the ρ_N parameter had a large value that was the opposite sign of ρ_D. The dependence on the ρ_D parameter reflects that the estimator for the pairwise model incorporates knowledge about correlation between the normalization signals.

In summary, this analysis shows that the pairwise model produces accurate estimates of the single-trial normalization strength of neural pairs and outlines the conditions in which those estimates are preferable to those obtained with the independent model.

3.4 Pairwise Ratio of Gaussians Model Captures Correlated Variability in Mouse V1

To test how well the model captures experimental data, we applied it to calcium imaging data recorded in V1 of mice responding to sinusoidal gratings of varying contrast levels (see subsection 2.7). We analyzed neurons that were strongly driven by the visual stimuli (N=295 neurons, 5528 simultaneously recorded pairs; see subsubsection 2.7.4 for inclusion criteria). We focused on stimulus contrast tuning (Eq. 8) because the formulation of the corresponding standard normalization model captures firing rate data well [Heeger, 1992], and visual contrast affects both normalization strength and the strength of noise correlations [Kohn and Smith, 2005].

Because the RoG framework has not been validated before on mouse V1 fluorescence data, we first applied the independent RoG and found that it provided excellent fits (average cross-validated goodness of fit 0.85, 95% c.i [0.846,0.858], 5163/5528 pairs with goodness of fit > 0.5) on par with that found in macaque V1 data recorded with electrode arrays [Coen-Cagli and Solomon, 2019], thus demonstrating that the RoG framework is flexible enough to capture datasets with different statistics. In both cases, the analysis was performed on visually responsive neurons, that therefore exhibited strong contrast tuning of the firing rate, partly explaining the high performance. However, the independent RoG could not capture correlated variability, which was prominent in the data (median noise correlation across all pairs and contrasts 0.117, c.i. [0.115, 0.120], 2781/5528 pairs had noise correlations significantly different from 0).

Therefore, we tested if the pairwise RoG could capture correlated variability in the data. Figure 5a,b demonstrates that the model is able to capture contrast-dependent noise correlations, both for pairs with positive (example in Fig. 5a; 4327/5528 pairs) and negative median noise correlations (example in Fig. 5b; 1201/5528 pairs). Importantly, even though the ρ_N, ρ_D parameters were stimulus-independent, the pairwise model captured substan-

tial changes in noise correlations with contrast for many of the pairs analyzed (2991/5528 had greater than 0.5 correlation between the noise correlation values of data and model fit, across contrast levels). However, this ability to capture correlations comes at the cost of larger model complexity. To account for this, we next compared quantitatively the pairwise and independent models, using a cross-validated goodness of fit score (Eq. 12). The pairwise model slightly outperformed the independent model on average and for the majority of pairs (median difference in goodness of fit $= 0.0121, p < 0.001$, 4123/5528 pairs with pairwise goodness of fit greater than independent), denoting that the additional free parameters are warranted. Furthermore, because the independent model is a special case of the pairwise with noise correlations fixed at zero, we found as expected that the performance difference between the two models increased for pairs of neurons with larger noise correlations (see Appendix, Fig. S1).

To benchmark the RoG against a widely adopted alternative model, we considered the modulated Poisson model that was previously shown to capture noise correlations in macaque V1 [Goris, Movshon, and Simoncelli, 2014]. For application to our imaging dataset and for a fair comparison with the RoG, we used Gaussian noise instead of Poisson [Aljadeff et al., 2016], and termed this the Modulated Gaussian (MG) model (see subsection 2.4). The example pairs demonstrate that, while in some pairs the MG can capture the modulation of noise correlations with contrast as well as the RoG (Fig. 5a), it is not able to capture it in other pairs while the RoG can (Fig. 5b). Across the dataset, for the majority of pairs (5238/5528), the pairwise RoG had a higher goodness of fit score (Eq. 12) than the MG (Fig. 5c,d, median difference between goodness of fit for RoG and MG =0.238, 95% c.i. [0.232,0.245]). These results were also largely independent of the specific preprocessing method applied to the calcium imaging data (see Appendix section 4 and Appendix Fig. S2). Moreover, although both models capture the tuning of noise correlations with contrast level by using stimulus-independent correlation parameters, the RoG model better predicts the trend in noise correlations with contrast than the MG (median difference in correlations

between model fit noise correlations and empirical noise correlations = 0.0771, 95% c.i. [0.071, 0.084]; 1450/5528 pairs had statistically significant correlation between the pairwise RoG predictions and empirical noise correlations compared to 1008/5528 for the MG). In principle the MG model's ability to capture the modulation of noise correlations with contrast could be improved by including contrast dependence in the correlation parameters explicitly, although this would increase model complexity.

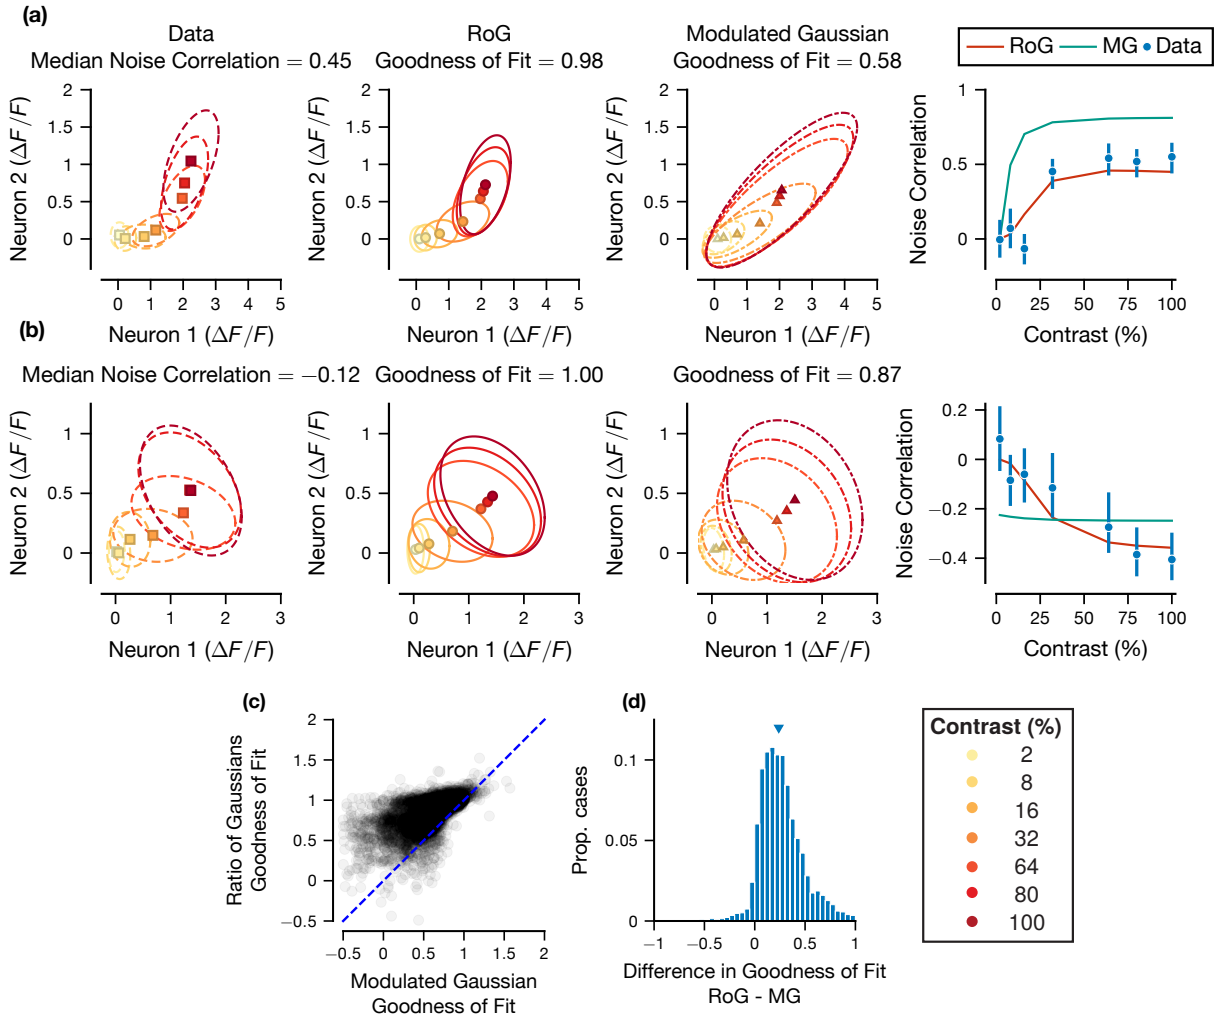

Figure 5: Pairwise RoG Captures Contrast-Dependent Noise Correlations in Mouse V1
(a,b) Pairwise neural responses for two example pairs of neurons in mouse V1 with positive (a) and negative (b) median noise correlations. From left to right: 1) empirical mean and covariance ellipses ($\sim$ 1 standard deviation from the empirical mean) for pairwise responses at each contrast level; 2) the RoG model predicted means and covariance ellipses, the panel title includes the cross-validated goodness of fit score; 3) the Modulated Gaussian (MG) predicted means and covariance ellipses; 4) compares the two model fit noise correlation values (continuous lines), with the empirical values as a function of contrast (error bars are 68% bootstrapped confidence interval). Neuronal pair in (a) had 93 repeats of each stimulus contrast, pair in (b) had 68 repeats.
(c) Scatter plot across all pairs of the goodness of fit score for Modulated Gaussian vs. the goodness of fit for the RoG. (d) Histogram of the difference between the scores in (c). Contrast levels are {2,8,16,32,64,80,100}.

These results demonstrate that the pairwise RoG captures a range of effects of stimulus contrast on noise correlations observed in experimental data, and performs competitively

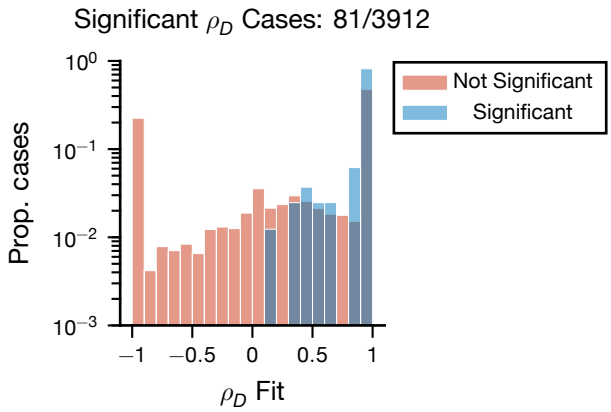

Figure 6: Inferred ρ_D are Positive in Mouse V1
Histogram comparing the inferred ρ_D values for those significantly different from 0 (81/3912 pairs) to the ρ_D values for pairs not significantly different from 0 but still meeting the goodness of fit criterion (3831 pairs).

against a popular alternative model that does not account for normalization explicitly.

Next, we analyzed the correlation parameters (ρ_N, ρ_D) in the model fit. We first only selected those pairs of neurons whose pairwise goodness of fit exceeded 0.5 and the independent goodness of fit measure (3912/5528 total), and we computed the bootstrapped 95% confidence interval for the (ρ_N, ρ_D) parameters (see subsubsection 2.3.2). As discussed in subsection 3.2, we then focused on the cases where the parameters were assessed to be significantly different from zero (855/3912 pairs for ρ_N, 81/3912 pairs for ρ_D; 76/3912 pairs for both). This is roughly in line with our theoretical results (see 3.2), although we had proportionally less pairs with significantly different from 0 ρ parameters. For these pairs that were significant, we found that nearly all of ρ_N parameters (844) were positive, and all of the inferred ρ_D parameters were positive (see Fig. 6), suggesting that normalization signals are generally shared for these pairs of neurons.

In summary, our result demonstrate a new approach to quantify how strongly normalization signals are shared between neurons, and to explain the diverse effects of normalization on noise correlations.

4 Discussion

We introduced a stochastic model of divisive normalization, the pairwise RoG, to characterize the trial-to-trial covariability between cortical neurons (i.e. noise correlations). The model provides excellent fits to calcium imaging recordings from mouse V1, capturing diverse effects of stimulus contrast and normalization strength on noise correlations (Figs. 5,6). We demonstrated that the effect of normalization on noise correlations differs depending on the sources of the variability, and that the model can accommodate both increases and decreases in noise correlations with normalization (Fig. 2) as past experiments had suggested. We then investigated the accuracy of inference of a key model parameter, which determines whether normalization is shared between neurons and provided a procedure for quantifying the significance of this inference using bootstrapping (Fig. 3). Lastly, we derived a Bayesian estimator for the single-trial normalization signals of simultaneously recorded pairs. Surprisingly, this estimator can be more accurate than the estimator based on the model that ignores noise correlations (the independent RoG) even when noise correlations are negligible (Fig 4).

As a descriptive, data analytic tool, our modeling framework complements normative and mechanistic theories of neural population variability. For instance, normative probabilistic accounts of sensory processing have suggested that divisive normalization may play a role in the inference of perceptual variables by modulating neural representations of uncertainty [Beck, Latham, and Pouget, 2011; Dehaene, Coen-Cagli, and Pouget, 2021; Hénaff et al., 2020; Bányai et al., 2019; Coen-Cagli, Kohn, and Schwartz, 2015; Festa et al., 2021; Orbán and Bányai, 2019]. Similarly, normalization could play a key role in multisensory cue combination [Hayashi, Kato, and Nozaki, 2020; Ohshiro, Angelaki, and DeAngelis, 2017; 2011]. However, the posited effect of normalization on covariability has not been tested quantitatively, as normalization signals are often not measurable. The pairwise RoG will allow researchers to test these hypotheses by providing a means with which to estimate

normalization signals from neural data and relate these to measures of neural covariability. In circuit-based models of neural dynamics such as the stabilized supralinear network [Hennequin et al., 2018] and the ORGaNICs architecture [Heeger and Mackey, 2019], the normalization computation emerges naturally from the network dynamics [Heeger and Zemlianova, 2020] and shapes the structure of stimulus-dependent noise correlations [Echeveste et al., 2019]. By quantifying the parametric relation between normalization and covariability, our descriptive tool will enable mapping those parameters onto the different circuit motifs and cell types posited by these network models.

When comparing the RoG to the Modulated Gaussian model (see subsection 2.4), we found that the RoG had better performance for the majority of pairs (Fig. 5c). We chose to adapt the modulated Poisson model [Goris, Movshon, and Simoncelli, 2014] as a comparison to the RoG because it was shown to successfully capture noise correlations in recordings from macaque V1. Moreover, this model belongs to the class of Generalized Linear Models, which are among the most widely used encoding models for neural activity [Pillow et al., 2008; Paninski, Pillow, and Lewi, 2007]. There are numerous alternative descriptive models of correlated neural population activity, among the most popular of these being latent variable models (LVMs), in which population-wide activity arises from interactions between a small set of unobserved variables [Yu et al., 2008; Ecker et al., 2014; Whiteway and Butts, 2019; Archer et al., 2014; Whiteway et al., 2019; Sokoloski, Aschner, and Coen-Cagli, 2021]. This effectively partitions the population noise covariance into underlying causes (i.e. latents) that are responsible for coordinating neural responses, which resembles our attribution of noise correlations to either shared input drive or shared normalization pools (Fig. 2). The RoG, on the other hand, is a pairwise model that seeks to explicitly characterize neural interactions through divisive normalization, which cannot be done with any existing LVMs; integrating normalization into the LVM framework is an important future extension of our model. One benefit of our current approach is that the RoG can be applied to any scenario in which two or more neurons are simultaneously recorded. LVMs can only be applied to relatively

large populations of simultaneously recorded neurons in order to estimate the globally shared latent factors. This is not always feasible for regions of the brain that are difficult to record from or using techniques such as intracellular voltage recordings [Kodandaramaiah et al., 2018; Hurwitz et al., 2021]. The downside of a method such as the RoG is scalability to large populations, as the model parameters have to be optimized for each recorded pair, which can be computationally expensive for modern datasets with thousands of neurons [Stevenson and Kording, 2011]. For instance, we were able to fit the RoG to data across multiple different preprocessing methods ($\sim$ 90000 pairs total) in a reasonable time ($\sim$ 27 hours running in parallel on a 28 core server without GPU acceleration), suggesting that it is not entirely impractical to use the RoG on a large dataset.

Three models have directly studied the relationship between normalization and across trial covariability [Ruff and Cohen, 2017; Tripp, 2012; Verhoef and Maunsell, 2017]. Tripp's [Tripp, 2012] simulation work on velocity tuning in the medial temporal cortex (MT) consistently predicted that normalization would decorrelate neural responses. However, we found that noise correlations could also increase with normalization. This is because Tripp modeled correlations to solely arise from tuning similarity between neurons. Conversely, in the RoG framework, noise correlations originate from input correlations ρ_N and correlations between normalization signals ρ_D. Our model then offers more flexibility than Tripp's by allowing relationships between normalization and correlation to depend on the sources of correlations. Verhoef and Maunsell [Verhoef and Maunsell, 2017] investigated the effect of attention on noise correlations by using a recurrent network implementation of the normalization model of attention [Reynolds and Heeger, 2009]. They describe multiple different patterns of the effect of normalization on noise correlations depending on tuning similarity between a pair of neurons and where attention is directed. Our model does not currently account for the effect of attention, but this would be possible by adapting the standard normalization model of attention which would require an additional parameter for the attentional gain. These prior two models are also primarily simulation based, while our model is meant to be data

analytic. Lastly, Ruff and Cohen [Ruff and Cohen, 2017] proposed a normalization model to explain how attention increases correlations between V1 and MT neurons [2016]. They modeled the trial-averaged MT neural responses as a function of trial-averaged responses of pools of V1 neurons. After fitting the parameters, single-trial MT responses were predicted by feeding the pooled single-trial V1 responses into the equation. By construction, variability in predicted MT neural responses only arises from variability in the V1 neural responses, which only occur in the numerator of their normalization model. Our model also allows for variability in the denominator of the normalization equation and therefore their model can be seen as a special case of the pairwise RoG.

An important limitation of our model is that the correlation parameters (ρ_N, ρ_D) are not identifiable (Fig. 3), meaning the model parametrization is such that multiple different parameter sets result in equivalent models (e.g. equivalent likelihoods and moments). This is a common issue when using complex nonlinear models as proposed here [Audoly et al., 2001]: in our model, this is due to multiplicative interactions between model parameters. Improving parameter estimation will require better constraints on model parameters, alternative optimization algorithms, or different objective functions (see Appendix section 4 for further discussion). Future extensions of the model to population level interactions through latent variable models offer another avenue to improve parameter estimation: the variability of population activity is often low-dimensional, which could naturally impose parameter constraints. Nonetheless, we developed a method to calculate confidence intervals and use them to determine when the sign of those parameters, which is an important factor in controlling noise correlations (Fig. 2), can be recovered accurately. We showed that we were accurately able to recover the sign of the correlation in synthetic datasets when the bootstrap confidence interval for the parameter of interest excluded 0. In the V1 dataset analyzed, we found that (855,81)/3912 pairs meet this criterion for ρ_N, ρ_D respectively. Although this is a minority of cases, it demonstrates that typical datasets with existing recording technologies could nonetheless provide sufficient power for studies that focus on the ρ parameters values. It will

be important in future work to understand which experimental conditions would maximize the yield of pairs with accurate estimates of the ρ parameters.

We chose in this study to primarily analyze the normalized fluorescence traces ($\Delta F/F$) rather than using deconvolution or spike inference methods (see [Pnevmatikakis, 2019; Stringer and Pachitariu, 2019; Evans, Petersen, and Humphries, 2020] for a review). Deconvolution methods were developed in part due to the slow temporal dynamics of the calcium indicators relative to membrane potentials generating spiking activity [Yaksi and Friedrich, 2006; Benisty et al., 2022]. Deconvolution and other spike inference techniques attempt to mitigate this limitation for analyses that depend on more exact measures of spike timing, and developers note these methods should be avoided when temporal information is not relevant and the raw calcium traces provide "sufficient information" [Stringer and Pachitariu, 2019]. Because of the construction of the contrast detection task (see subsection 2.7) and the temporally invariant nature of contrast responses in V1 [Albrecht et al., 2002], the analysis of the dataset presented here does not require precise temporal information so the use of normalized fluorescence traces was sufficient. Additionally, deconvolution changes the statistics of the data greatly, such as altering the distribution of noise correlations and increasing the sparsity of the fluorescence signal [Rupasinghe et al., 2021]. One recent work attempted to account for these differences by using more appropriate probabilistic models [Wei et al., 2020], but does not currently model noise correlations. On the other hand, calcium fluorescence is an indirect measure of neuronal communication and coding, being related to the underlying action potentials through a complex generative model [Vogelstein et al., 2010; Triplett and Goodhill, 2019]. As such, it might be inappropriate or insufficient to apply an encoding model directly to the $\Delta F/F$ traces, as we have done here. As additional analysis, we analyzed deconvolved traces using two variants of the OASIS method [Friedrich, Zhou, and Paninski, 2017]: unconstrained OASIS as found in suite2p [Pachitariu et al., 2017]), or OASIS with an ℓ_1 sparsity constraint as in [Lyall et al., 2021]. As expected, the deconvolution techniques significantly altered the distribution of noise correlations, but the results of

our analysis of these deconvolved data was qualitatively in-line with the results obtained on the raw calcium traces (see section 4).

The generality of the modeling framework presented here leaves room for future expansion. One such direction would be to increase the dimensionality to model correlations among a neural population. This would require more correlation parameters, which could make the model more difficult to fit to data. However, reasoning that population variability is low-dimensional [Cunningham and Yu, 2014; Huang et al., 2019; Lin et al., 2015; Rabinowitz et al., 2015; Schölvinck et al., 2015; Umakantha et al., 2021], it is likely this issue could be circumvented by applying dimensionality reduction techniques within the model or by allowing the sharing of correlation parameters across a neural population. Another interesting application of this model would look directly at the effects of normalization on information transmission and representation. The relationship between noise correlations and the amount of information that can be represented by a neural population has been widely discussed [Abbott and Dayan, 1999; Averbeck and Lee, 2006; Averbeck, Latham, and Pouget, 2006; Kohn et al., 2016; Moreno-Bote et al., 2014; Zohary, Shadlen, and Newsome, 1994]. Moreover, some experimental and theoretical work has connected modulations of information in neural populations with computations that have been modeled with normalization models, such as surround suppression and attention [Henry and Kohn, 2020; Kanashiro et al., 2017; Ringach, 2010; Tripp, 2012]. Our model could be modified to investigate this connection and further illuminate the effects of normalization on information transmission.

Acknowledgement

We thank members of the Kohn and Coen-Cagli laboratories for feedback on the manuscript. We also thank Daniel Quintana for help with animal training.

This work was supported by the National Institutes of Health (NIH) Grant R01- EY030578

(R.C.C.), NIH Grant RF1- DA056400 (R.C.C.), NIH Grant UF1- NS107574 (H.A.), NIH Grant U19- NS107613 (H.A.), NIH Grant RF1- RF1- MH120680 (H.A.), NIH Grant RO1-EY023756 (H.A.), New York Stem Cell Foundation (H.A.), National Science Foundation Graduate Research Fellowship (DGE 1752814 (H.A.B.)).

References

Abbott, LF and P Dayan (1999). "The effect of correlated variability on the accuracy of a population code." In: *Neural Computation* 11.1, pp. 91–101.

Albrecht, DG and WS Geisler (1991). "Motion selectivity and the contrast-response function of simple cells in the visual cortex." In: *Visual neuroscience* 7.6, pp. 531–546.

Albrecht, DG and DB Hamilton (1982). "Striate Cortex of Monkey and Cat: Contrast Response Function." In: *Journal of neurophysiology* 48.1, pp. 217–237.

Albrecht, DG et al. (2002). "Visual cortex neurons of monkeys and cats: temporal dynamics of the contrast response function." In: *Journal of neurophysiology* 88.2, pp. 888–913.

Aljadeff, J et al. (2016). "Analysis of neuronal spike trains, deconstructed." In: *Neuron* 91.2, pp. 221–259.

Archer, EW et al. (2014). "Low-Dimensional Models of Neural Population Activity in Sensory Cortical Circuits." In: vol. 27. Advances in Neural Information Processing Systems. Curran Associates, Inc.

Audoly, S et al. (2001). "Global identifiability of nonlinear models of biological systems." In: *IEEE Transactions on biomedical engineering* 48.1, pp. 55–65.

Averbeck, BB, PE Latham, and A Pouget (2006). "Neural Correlations, Population Coding and Computation." In: *Nature Reviews Neuroscience* 7.5, p. 358.

Averbeck, BB and D Lee (2006). "Effects of Noise Correlations on Information Encoding and Decoding." In: *Journal of Neurophysiology* 95.6, pp. 3633–3644.

Bányai, M et al. (2019). “Stimulus complexity shapes response correlations in primary visual cortex.” In: *Proceedings of the National Academy of Sciences* 116.7, pp. 2723–2732.

Bartolo, R et al. (2020). “Information-limiting correlations in large neural populations.” In: *Journal of Neuroscience* 40.8, pp. 1668–1678.

Beck, JM, PE Latham, and A Pouget (2011). “Marginalization in Neural Circuits with Divisive Normalization.” In: *Journal of Neuroscience* 31.43, pp. 15310–15319.

Benisty, H et al. (2022). “Data processing of functional optical microscopy for neuroscience.” In: *arXiv preprint arXiv:2201.03537*.

Berkes, P, F Wood, and J Pillow (2008). “Characterizing neural dependencies with copula models.” In: *Advances in neural information processing systems* 21.

Berkes, P et al. (2011). “Spontaneous Cortical Activity Reveals Hallmarks of an Optimal Internal Model of the Environment.” In: *Science (New York, N.Y.)* 331.6013, pp. 83–87.

Bondy, AG, RM Haefner, and BG Cumming (2018). “Feedback Determines the Structure of Correlated Variability in Primary Visual Cortex.” In: *Nature Neuroscience* 21.4, pp. 598–606.

Bounds, HA et al. (2021). “Multifunctional Cre-dependent transgenic mice for high-precision all-optical interrogation of neural circuits.” In: *bioRxiv*.

Carandini, M, DJ Heeger, and JA Movshon (1997). “Linearity and Normalization in Simple Cells of the Macaque Primary Visual Cortex.” In: *Journal of Neuroscience* 17.21, pp. 8621–8644.

Carandini, M and DJ Heeger (2012). “Normalization as a Canonical Neural Computation.” In: *Nature Reviews Neuroscience* 13.1, pp. 51–62.

Cavanaugh, JR, W Bair, and JA Movshon (2002). “Nature and Interaction of Signals from the Receptive Field Center and Surround in Macaque V1 Neurons.” In: *Journal of Neurophysiology* 88.5, pp. 2530–2546.

Clatworthy, PL et al. (2003). “Coding of the Contrasts in Natural Images by Populations of Neurons in Primary Visual Cortex (V1).” In: *Vision research* 43.18, pp. 1983–2001.

Coen-Cagli, R, P Dayan, and O Schwartz (2012). "Cortical surround interactions and perceptual salience via natural scene statistics." In: *PLoS computational biology* 8.3, e1002405.

Coen-Cagli, R, A Kohn, and O Schwartz (2015). "Flexible Gating of Contextual Influences in Natural Vision." In: *Nature Neuroscience* 18.11, pp. 1648–1655.

Coen-Cagli, R and SS Solomon (2019). "Relating Divisive Normalization to Neuronal Response Variability." In: *Journal of Neuroscience* 39.37, pp. 7344–7356.

Cohen, MR and A Kohn (2011). "Measuring and interpreting neuronal correlations." In: *Nature neuroscience* 14.7, pp. 811–819.

Cohen, MR and JHR Maunsell (2009). "Attention Improves Performance Primarily by Reducing Interneuronal Correlations." In: *Nature Neuroscience* 12.12, pp. 1594–1600.

Cunningham, JP and BM Yu (2014). "Dimensionality Reduction for Large-Scale Neural Recordings." In: *Nature Neuroscience* 17.11, pp. 1500–1509.

Dehaene, GP, R Coen-Cagli, and A Pouget (2021). "Investigating the Representation of Uncertainty in Neuronal Circuits." In: *PLOS Computational Biology* 17.2, e1008138.

Díaz-Francés, E and FJ Rubio (2013). "On the Existence of a Normal Approximation to the Distribution of the Ratio of Two Independent Normal Random Variables." In: *Statistical Papers* 54.2, pp. 309–323.

Doiron, B et al. (2016). "The Mechanics of State-Dependent Neural Correlations." In: *Nature Neuroscience* 19.3, pp. 383–393.

Echeveste, R et al. (2019). *Cortical-like Dynamics in Recurrent Circuits Optimized for Sampling-Based Probabilistic Inference*. 9. Neuroscience, pp. 1138–1149.

Ecker, AS et al. (2014). "State dependence of noise correlations in macaque primary visual cortex." In: *Neuron* 82.1, pp. 235–248.

Evans, MH, RS Petersen, and MD Humphries (2020). "On the use of calcium deconvolution algorithms in practical contexts." In: *bioRxiv*, p. 871137.

Festa, D et al. (2021). "Neuronal Variability Reflects Probabilistic Inference Tuned to Natural Image Statistics." In: *Nature Communications* 12.1, p. 3635.

Friedrich, J, P Zhou, and L Paninski (2017). “Fast online deconvolution of calcium imaging data.” In: *PLoS computational biology* 13.3, e1005423.

Gao, Y et al. (2015). “High-dimensional neural spike train analysis with generalized count linear dynamical systems.” In: *Advances in neural information processing systems* 28.

Gardella, C, O Marre, and T Mora (2019). “Modeling the correlated activity of neural populations: A review.” In: *Neural Computation* 31.2, pp. 233–269.

Geisler, WS and DG Albrecht (1992). “Cortical Neurons: Isolation of Contrast Gain Control.” In: *Vision research* 32.8, pp. 1409–1410.

Goris, RL, JA Movshon, and EP Simoncelli (2014). “Partitioning Neuronal Variability.” In: *Nature neuroscience* 17.6, pp. 858–865.

Granot-Atedgi, E et al. (2013). “Stimulus-Dependent Maximum Entropy Models of Neural Population Codes.” In: *PLOS Computational Biology* 9.3, e1002922.

Gur, M, A Beylin, and DM Snodderly (1997). “Response Variability of Neurons in Primary Visual Cortex (V1) of Alert Monkeys.” In: *Journal of Neuroscience* 17.8, pp. 2914–2920.

Haefner, RM, P Berkes, and J Fiser (2016). “Perceptual decision-making as probabilistic inference by neural sampling.” In: *Neuron* 90.3, pp. 649–660.

Hayashi, T, Y Kato, and D Nozaki (2020). “Divisively Normalized Integration of Multisensory Error Information Develops Motor Memories Specific to Vision and Proprioception.” In: *Journal of Neuroscience* 40.7, pp. 1560–1570.

Hayya, J, D Armstrong, and N Gressis (1975). “A note on the ratio of two normally distributed variables.” In: *Management Science* 21.11, pp. 1338–1341.

Heeger, DJ and KO Zemlianova (2020). “A Recurrent Circuit Implements Normalization, Simulating the Dynamics of V1 Activity.” In: *Proceedings of the National Academy of Sciences* 117.36, pp. 22494–22505.

Heeger, DJ (1992). “Normalization of Cell Responses in Cat Striate Cortex.” In: *Visual neuroscience* 9.2, pp. 181–197.

Heeger, DJ and WE Mackey (2019). "Oscillatory Recurrent Gated Neural Integrator Circuits (ORGaNICs), a Unifying Theoretical Framework for Neural Dynamics." In: *Proceedings of the National Academy of Sciences* 116.45, pp. 22783–22794.

Hénaff, OJ et al. (2020). "Representation of Visual Uncertainty through Neural Gain Variability." In: *Nature Communications* 11.1, p. 2513.

Hennequin, G et al. (2018). "The Dynamical Regime of Sensory Cortex: Stable Dynamics around a Single Stimulus-Tuned Attractor Account for Patterns of Noise Variability." In: *Neuron* 98.4, 846–860.e5.

Henry, CA and A Kohn (2020). "Spatial Contextual Effects in Primary Visual Cortex Limit Feature Representation under Crowding." In: *Nature Communications* 11.1, p. 1687.

Hu, Y, J Zylberberg, and E Shea-Brown (2014). "The Sign Rule and beyond: Boundary Effects, Flexibility, and Noise Correlations in Neural Population Codes." In: *PLOS Computational Biology* 10.2, e1003469.

Huang, C et al. (2019). "Circuit Models of Low-Dimensional Shared Variability in Cortical Networks." In: *Neuron* 101.2, 337–348.e4.

Hurwitz, C et al. (2021). "Building population models for large-scale neural recordings: Opportunities and pitfalls." In: *Current opinion in neurobiology* 70, pp. 64–73.

Joe, H and JJ Xu (1996). "The Estimation Method of Inference Functions for Margins for Multivariate Models." In.

Josić, K et al. (2009). "Stimulus-Dependent Correlations and Population Codes." In: *Neural Computation* 21.10, pp. 2774–2804.

Kafashan, M et al. (2021). "Scaling of sensory information in large neural populations shows signatures of information-limiting correlations." In: *Nature communications* 12.1, pp. 1–16.

Kanashiro, T et al. (2017). "Attentional Modulation of Neuronal Variability in Circuit Models of Cortex." In: *eLife* 6. Ed. by P Latham, e23978.

Kanitscheider, I, R Coen-Cagli, and A Pouget (2015). "Origin of Information-Limiting Noise Correlations." In: *Proceedings of the National Academy of Sciences* 112.50, E6973–E6982.

Kodandaramaiah, SB et al. (2018). "Multi-neuron intracellular recording in vivo via interacting autopatching robots." In: *Elife* 7, e24656.

Kohn, A and MA Smith (2005). "Stimulus Dependence of Neuronal Correlation in Primary Visual Cortex of the Macaque." In: *Journal of Neuroscience* 25.14, pp. 3661–3673.

Kohn, A et al. (2016). "Correlations and Neuronal Population Information." In: *Annual Review of Neuroscience* 39.1, pp. 237–256.

Kronmal, RA (1993). "Spurious correlation and the fallacy of the ratio standard revisited." In: *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 156.3, pp. 379–392.

Lange, RD and RM Haefner (2017). "Characterizing and interpreting the influence of internal variables on sensory activity." In: *Current opinion in neurobiology* 46, pp. 84–89.

Lange, RD and RM Haefner (2022). "Task-induced neural covariability as a signature of approximate Bayesian learning and inference." In: *PLoS computational biology* 18.3, e1009557.

Lin, IC et al. (2015). "The nature of shared cortical variability." In: *Neuron* 87.3, pp. 644–656.

Litwin-Kumar, A and B Doiron (2012). "Slow Dynamics and High Variability in Balanced Cortical Networks with Clustered Connections." In: *Nature Neuroscience* 15.11, pp. 1498–1505.

Liu, LD, RM Haefner, and CC Pack (2016). "A Neural Basis for the Spatial Suppression of Visual Motion Perception." In: *eLife* 5, e16167.

Louie, K, MW Khaw, and PW Glimcher (2013). "Normalization is a general neural mechanism for context-dependent decision making." In: *Proceedings of the National Academy of Sciences* 110.15, pp. 6139–6144.

Lyall, EH et al. (2021). "Synthesis of a comprehensive population code for contextual features in the awake sensory cortex." In: *ELife* 10, e62687.

Marsaglia, G et al. (2006). "Ratios of Normal Variables." In: *Journal of Statistical Software* 16.4, pp. 1–10.

Mitchell, JF, KA Sundberg, and JH Reynolds (2009). “Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4.” In: *Neuron* 63.6, pp. 879–888.

Moreno-Bote, R et al. (2014). “Information-limiting correlations.” In: *Nature Neuroscience* 17.10, pp. 1410–1417.

Ohshiro, T, DE Angelaki, and GC DeAngelis (2017). “A Neural Signature of Divisive Normalization at the Level of Multisensory Integration in Primate Cortex.” In: *Neuron* 95.2, 399–411.e8.

Ohshiro, T, DE Angelaki, and GC DeAngelis (2011). “A Normalization Model of Multisensory Integration.” In: *Nature Neuroscience* 14.6, pp. 775–782.

Olsen, SR, V Bhandawat, and RI Wilson (2010). “Divisive normalization in olfactory population codes.” In: *Neuron* 66.2, pp. 287–299.

Orbán, G and M Bányai (2019). “Noise Correlations and Perceptual Inference.” In: *Current Opinion in Neurobiology* 58.

Orbán, G et al. (2016). “Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex.” In: *Neuron* 92.2, pp. 530–543.

Pachitariu, M et al. (2017). “Suite2p: beyond 10,000 neurons with standard two-photon microscopy.” In: *BioRxiv*.

Paninski, L, J Pillow, and J Lewi (2007). “Statistical Models for Neural Encoding, Decoding, and Optimal Stimulus Design.” In: *Progress in Brain Research*. Ed. by P Cisek, T Drew, and JF Kalaska. Vol. 165. Computational Neuroscience: Theoretical Insights into Brain Function. Elsevier, pp. 493–507.

Panzeri, S et al. (2022). “The structures and functions of correlations in neural population codes.” In: *Nature Reviews Neuroscience* 23.9, pp. 551–567.

Peirce, J et al. (2019). “PsychoPy2: Experiments in behavior made easy.” In: *Behavior research methods* 51.1, pp. 195–203.

Pham-Gia, T, N Turkkan, and E Marchand (2006). "Density of the Ratio of Two Normal Random Variables and Applications." In: *Communications in Statistics-Theory and Methods* 35.9, pp. 1569–1591.

Pillow, JW and EP Simoncelli (2006). "Dimensionality reduction in neural models: an information-theoretic generalization of spike-triggered average and covariance analysis." In: *Journal of vision* 6.4, pp. 9–9.

Pillow, JW et al. (2008). "Spatio-Temporal Correlations and Visual Signalling in a Complete Neuronal Population." In: *Nature* 454.7207, pp. 995–999.

Pnevmatikakis, EA (2019). "Analysis pipelines for calcium imaging data." In: *Current opinion in neurobiology* 55, pp. 15–21.

Ponce-Alvarez, A et al. (2013). "Stimulus-Dependent Variability and Noise Correlations in Cortical MT Neurons." In: *Proceedings of the National Academy of Sciences* 110.32, pp. 13162–13167.

Rabinowitz, NC et al. (2015). "Attention Stabilizes the Shared Gain of V4 Populations." In: *eLife* 4, e08998.

Reynolds, JH and DJ Heeger (2009). "The normalization model of attention." In: *Neuron* 61.2, pp. 168–185.

Rikhye, RV and M Sur (2015). "Spatial Correlations in Natural Scenes Modulate Response Reliability in Mouse Visual Cortex." In: *Journal of Neuroscience* 35.43, pp. 14661–14680.

Ringach, DL (2010). "Population Coding under Normalization." In: *Vision research* 50.22, pp. 2223–2232.

Ruff, DA and MR Cohen (2014). "Attention Can Either Increase or Decrease Spike Count Correlations in Visual Cortex." In: *Nature Neuroscience* 17.11, pp. 1591–1597.

Ruff, DA and MR Cohen (2017). "A Normalization Model Suggests That Attention Changes the Weighting of Inputs between Visual Areas." In: *Proceedings of the National Academy of Sciences* 114.20, E4085–E4094.

Ruff, DA and MR Cohen (2016). “Stimulus dependence of correlated variability across cortical areas.” In: *Journal of Neuroscience* 36.28, pp. 7546–7556.

Rumyantsev, OI et al. (2020). “Fundamental bounds on the fidelity of sensory cortical coding.” In: *Nature* 580.7801, pp. 100–105.

Rupasinghe, A et al. (2021). “Direct extraction of signal and noise correlations from two-photon calcium imaging of ensemble neuronal activity.” In: *Elife* 10, e68046.

Sadagopan, S and D Ferster (2012). “Feedforward Origins of Response Variability Underlying Contrast Invariant Orientation Tuning in Cat Visual Cortex.” In: *Neuron* 74.5, pp. 911–923.

Sawada, T and AA Petrov (2017). “The Divisive Normalization Model of V1 Neurons: A Comprehensive Comparison of Physiological Data and Model Predictions.” In: *Journal of Neurophysiology* 118.6, pp. 3051–3091.

Schneidman, E et al. (2006). “Weak Pairwise Correlations Imply Strongly Correlated Network States in a Neural Population.” In: *Nature* 440.7087, pp. 1007–1012.

Schölvinck, ML et al. (2015). “Cortical State Determines Global Variability and Correlations in Visual Cortex.” In: *Journal of Neuroscience* 35.1, pp. 170–178.

Shadlen, MN and WT Newsome (1998). “The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding.” In: *Journal of Neuroscience* 18.10, pp. 3870–3896.

Snyder, AC et al. (2014). “Correlations in V1 Are Reduced by Stimulation Outside the Receptive Field.” In: *Journal of Neuroscience* 34.34, pp. 11222–11227.

Sokoloski, S, A Aschner, and R Coen-Cagli (2021). “Modelling the neural code in large populations of correlated neurons.” In: *eLife* 10, e64615.

Sridharan, S et al. (2022). “High-performance microbial opsins for spatially and temporally precise perturbations of large neuronal networks.” In: *Neuron*.

Stevenson, IH (2016). “Flexible Models for Spike Count Data with Both Over-and under-Dispersion.” In: *Journal of computational neuroscience* 41.1, pp. 29–43.

Stevenson, IH and KP Kording (2011). “How advances in neural recording affect data analysis.” In: *Nature neuroscience* 14.2, pp. 139–142.

Stringer, C and M Pachitariu (2019). “Computational processing of neural recordings from calcium imaging data.” In: *Current opinion in neurobiology* 55, pp. 22–31.

Sturmfels, B (2002). *Solving Systems of Polynomial Equations*. 97. American Mathematical Soc.

Tolhurst, DJ, JA Movshon, and AF Dean (1983). “The Statistical Reliability of Signals in Single Neurons in Cat and Monkey Visual Cortex.” In: *Vision research* 23.8, pp. 775–785.

Triplett, MA and GJ Goodhill (2019). “Probabilistic encoding models for multivariate neural data.” In: *Frontiers in neural circuits*, p. 1.

Tripp, BP (2012). “Decorrelation of spiking variability and improved information transfer through feedforward divisive normalization.” In: *Neural Computation* 24.4, pp. 867–894.

Umakantha, A et al. (2021). “Bridging neuronal correlations and dimensionality reduction.” In: *Neuron* 109.17, pp. 2740–2754.

Ver Hoef, JM (2012). “Who Invented the Delta Method?” In: *The American Statistician* 66.2, pp. 124–127.

Verhoef, BE and JHR Maunsell (2017). “Attention-related changes in correlated neuronal activity arise from normalization mechanisms.” In: *Nature Neuroscience* 20.7, pp. 969–977.

Vogelstein, JT et al. (2010). “Fast nonnegative deconvolution for spike train inference from population calcium imaging.” In: *Journal of neurophysiology* 104.6, pp. 3691–3704.

Wei, XX et al. (2020). “A zero-inflated gamma model for post-deconvolved calcium imaging traces.” In: *Neurons, Behavior, Data analysis, and Theory* 3.2, pp. 1–21.

Wekselblatt, JB et al. (2016). “Large-scale imaging of cortical dynamics during sensory perception and behavior.” In: *Journal of neurophysiology* 115.6, pp. 2852–2866.

Whiteway, MR and DA Butts (2019). “The Quest for Interpretable Models of Neural Population Activity.” In: *Current Opinion in Neurobiology* 58, pp. 86–93.

Whiteway, MR et al. (2019). “Characterizing the nonlinear structure of shared variability in cortical neuron populations using latent variable models.” In: *Neurons, behavior, data analysis and theory* 3.1.

Yaksi, E and RW Friedrich (2006). “Reconstruction of firing rate changes across neuronal populations by temporally deconvolved Ca2+ imaging.” In: *Nature methods* 3.5, pp. 377–383.

Yu, BM et al. (2008). “Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity.” In: *Advances in neural information processing systems* 21.

Zhao, Y and IM Park (2017). “Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population Spike Trains.” In: *Neural Computation* 29.5, pp. 1293–1316.

Zohary, E, MN Shadlen, and WT Newsome (1994). “Correlated neuronal discharge rate and its implications for psychophysical performance.” In: *Nature* 370.6485, pp. 140–143.