J Neurophysiol 118: 29–46, 2017.
First published March 15, 2017; doi:10.1152/jn.00931.2016.

RESEARCH ARTICLE | *Neural Circuits*

Population activity statistics dissect subthreshold and spiking variability in V1

Mihály Bányai,[1] Zsombor Koman,[1] and Gergő Orbán[1,2]
[1]*Computational Systems Neuroscience Lab, MTA Wigner Research Centre for Physics, Budapest, Hungary; and* [2]*NAP B-Population Activity Research Unit, MTA Wigner Research Centre for Physics, Budapest, Hungary*

Submitted 9 December 2016; accepted in final form 13 March 2017

Bányai M, Koman Z, Orbán G. Population activity statistics dissect subthreshold and spiking variability in V1. *J Neurophysiol* 118: 29–46, 2017. First published March 15, 2017; doi:10.1152/jn.00931.2016.—Response variability, as measured by fluctuating responses upon repeated performance of trials, is a major component of neural responses, and its characterization is key to interpret high dimensional population recordings. Response variability and covariability display predictable changes upon changes in stimulus and cognitive or behavioral state, providing an opportunity to test the predictive power of models of neural variability. Still, there is little agreement on which model to use as a building block for population-level analyses, and models of variability are often treated as a subject of choice. We investigate two competing models, the doubly stochastic Poisson (DSP) model assuming stochasticity at spike generation, and the rectified Gaussian (RG) model tracing variability back to membrane potential variance, to analyze stimulus-dependent modulation of both single-neuron and pairwise response statistics. Using a pair of model neurons, we demonstrate that the two models predict similar single-cell statistics. However, DSP and RG models have contradicting predictions on the joint statistics of spiking responses. To test the models against data, we build a population model to simulate stimulus change-related modulations in pairwise response statistics. We use single-unit data from the primary visual cortex (V1) of monkeys to show that while model predictions for variance are qualitatively similar to experimental data, only the RG model's predictions are compatible with joint statistics. These results suggest that models using Poisson-like variability might fail to capture important properties of response statistics. We argue that membrane potential-level modeling of stochasticity provides an efficient strategy to model correlations.

NEW & NOTEWORTHY Neural variability and covariability are puzzling aspects of cortical computations. For efficient decoding and prediction, models of information encoding in neural populations hinge on an appropriate model of variability. Our work shows that stimulus-dependent changes in pairwise but not in single-cell statistics can differentiate between two widely used models of neuronal variability. Contrasting model predictions with neuronal data provides hints on the noise sources in spiking and provides constraints on statistical models of population activity.

spiking variability; phenomenological models; population activity; noise correlations; visual cortex

MULTINEURON ACTIVITY PATTERNS are critical features of the neural code. Discovering the organizational principles of these activity patterns requires the characterization of the response statistics of neuronal populations. Traditionally, neuronal responses have solely been characterized by the mean activity (Adrian 1926). However, prediction of future states of the population activity, prediction of activity of left-out neurons, obtaining low-variance estimates of physical/cognitive variables, or revealing hidden processes underlying the observed neural dynamics requires the characterization of higher order moments of response statistics. Indeed, stimulus and cognitive state dependence of variability (Churchland et al. 2010; Finn et al. 2007; Goris et al. 2015) and correlations (Berkes et al. 2011; Cohen and Kohn 2011; Ecker et al. 2010, 2014; Fiser et al. 2004; Rabinowitz et al. 2015) has recently drawn considerable attention and has been the subject of theoretical analysis (Haefner et al. 2013; Moreno-Bote et al. 2008).

Instead of a complete characterization of the multivariate activity distribution, approximations are used, which characterize the distribution with a limited set of parameters. Maximum entropy models (Schneidman et al. 2006) or their variants (Lin et al. 2015; Macke et al. 2011; Tkačik et al. 2010; Yu et al. 2011) directly parametrize some low-dimensional features (the marginals) of the response statistics, while latent variable models trace the covariability back to changes in a lower-dimensional space (Yu et al. 2009). A critical difference between these models lies in the assumed source and form of variability. The source of variability in neural responses is diverse and can be traced back to bottom-up influences (Anderson et al. 2000), effects of unrecorded neurons of the population, top-down influences (Goris et al. 2015; Renart and Machens 2014), and mechanisms intrinsic to individual neurons (Faisal et al. 2008). These different sources of variability imply different combinations of private and shared variabilities and different forms of dependencies on stimulus attributes; therefore interpretation of observed patterns strongly depends on the assumptions on the properties of variability.

Detailed single-cell models for the building blocks of population models require a prohibitive amount of data, and therefore phenomenological models are used instead, which can be directly constrained by the details of single-cell statistics. Spike count measurements are characterized by a distinctive linear relationship between mean firing rate and spike count variability (Tolhurst et al. 1981). The resemblance of this statistic to the characteristics of the Poisson process motivated the use of Poisson process for modeling spike count variability, where variability arises as a consequence of a renewal process with independent spikes sampled in finite time windows with a given expected value but independent across time windows.

Address for reprint requests and other correspondence: M. Bányai, MTA Wigner RCP, 29-33 Konkoly-Thege út, 1121 Budapest, Hungary (e-mail: banyai.mihaly@wigner.mta.hu).

This form of single-cell statistics is central to many theories of coding (Churchland et al. 2011; Ecker et al. 2016; Froudarakis et al. 2014; Jazayeri and Movshon 2006; Ma and Jazayeri 2014; Pillow 2007; Simoncelli et al. 2004). Modeling neural variability by directly assuming Poisson variability has proven efficient in identifying the effect of attention both at the level of single cells (Goris et al. 2015) and at the level of pairwise correlations (Ecker et al. 2010; Goris et al. 2015; Rabinowitz et al. 2015).

The Poisson model assumes that variability is introduced beyond the firing rate: once the instantaneous firing rate has been established through the interactions of the neuron, spikes are generated in a stochastic manner. However, neurons are known to exhibit variability at their membrane potential (Carandini and Ferster 2000; Finn et al. 2007; Haider et al. 2013; Tan et al. 2014), i.e., at a stage before the generation of spikes, and the spike generation process has been shown to be highly reliable (Mainen and Sejnowski 1995). The rectified Gaussian model (RG) formulates an alternative account of single cell response statistics, which is compatible with the Poisson-like aspect of variability but assumes that variability is dominated by processes before the stage of spike generation (practically at the level of membrane potential) but assumes spike generation to be close to deterministic (Carandini 2004; Dorn and Ringach 2003). This model has also proven successful in accounting for patterns in mean spiking responses of V1 neurons (Finn et al. 2007). The Poisson-like spiking model and the RG model are two extremes of a spectrum of phenomenological models that account for the basic single-cell statistics of cortical neurons while also being capable of describing pairwise phenomena. It is unclear, however, how well the alternative models of spike count variability can predict the patterns in response correlations as well as the patterns in mean responses and response variance.

In this study we contrast competing approaches proposed for describing spike count variability and use their predictions on pairwise response statistics. We define the doubly stochastic Poisson (DSP) model and the RG model to analyze the relationship between membrane potential statistics and spike count statistics. In particular, we focus on the joint spiking statistics of a pair of neurons to demonstrate a dissociation between the two models based on changes in spike count correlation resulting from changes in pairwise membrane potential statistics. Using these analyses as a starting point, we simulate stimulus-dependent modulation of spike count statistics for both of the models in populations of model neurons and compare the predictions to orientation-dependent and contrast-dependent changes in the activity statistics of extracellularly recorded V1 neurons in awake monkeys. We argue that changes in stimulus attributes give rise to distinctive patterns in response correlations that are compatible with the RG model but contradict the DSP model. This dissociation provides constraints to the construction of descriptive models of population activity when bottom-up effects modulate the firing statistics.

MATERIALS AND METHODS

Model of membrane potential responses. We modeled the responses of direction selective neurons of V1 at the level of membrane potentials. Characterization of the response statistics of neurons requires modeling not only the first-order statistics (mean) of membrane potential but also higher level statistics. Instead of aiming for a complete description of the response statistics, we focus on the second-order statistics of neural responses. This is motivated by two considerations: *1*) experimental designs are typically limited in terms of the number of trials, rendering higher-order joint statistics of neurons hard to estimate; and *2*) pairwise statistics are the simplest measure of population activity going beyond individual cell response properties and are already able to capture definitive signatures of population-level cortical computation (Haefner et al. 2016; Orbán et al. 2016; Singh et al. 2016).

The main determinant of neuronal responses is the receptive field that defines what stimulus features the neuron is sensitive to. We based the responses of model neurons on the most widely studied of these changes in V1, the orientation dependence of responses: the tuning curve was used to summarize systematic modulation of the mean response as the direction of an oriented stimulus is changing. Within-trial and across-trial fluctuations in membrane potential responses of individual cells are prevalent in direction-selective neurons (Tomko and Crapper 1974). Furthermore, variations of membrane potential responses are not independent: deviations from the mean responses tend to be correlated across neurons (Yu and Ferster 2010). We summarized individual and joint variations in membrane potential responses in terms of a probability distribution that defines the probability with which any given combination of membrane potentials is present. We represented variances and correlations through a covariance matrix, which together with the mean membrane potential response define a Gaussian distribution. We obtained the response of the population of neurons at any given time bin b by drawing an independent sample from this distribution:

$$\mathbf{u}_b \sim N(\mathbf{u}_b;\ \mu, \mathrm{C_{mp}}) \tag{1}$$

The mean of the distribution, $\mu \in \mathbb{R}^N$, determines the trial-averaged membrane potential level of each neuron, and the sequences of individual samples give rise to within-trial variability. The full second-order structure of membrane potential incorporates temporal dependencies beyond the spatial structure. A sample from the Gaussian distribution in our model specified a single membrane potential value in a 20-ms bin, assuming that there is no temporal dependence on longer timescales, and omitting variability on shorter ones. Our choice of independent samples was motivated by the typical support of the autocorrelation function of membrane potentials, which is on the scale of 20 ms (Azouz and Gray 1999), and therefore membrane potential values beyond this scale can be regarded as independent.

Similar to systematic changes in trial-averaged mean membrane potential responses, variances and covariances are also characterized by systematic changes (Finn et al. 2007; Tan et al. 2014). As a consequence, response variance and covariance were not taken to be fixed parameters of the model but were dependent on the stimulus.

Intensity of spike responses of neurons is determined by the instantaneous firing rate function. We obtained the mapping from membrane potential to firing rate by the firing rate nonlinearity, for which we take a parametric form from the literature (Carandini 2004), parametrized identically for each neuron:

$$\mathbf{r}_b = k[\mathbf{u}_b - V_{th}]_+^{\beta} \tag{2}$$

V_{th} denotes the membrane potential threshold under which the firing rate is zero. Above the threshold the firing rate is a power law function with exponent β, where $\beta = 1$ corresponds to a linear mapping. Gain parameter, k, indicates how many spikes are to be expected within a single time bin (which is always 20 ms in our study). Using this mapping, we obtain an instantaneous rate in every time bin, which serves as an intermediate quantity between membrane potentials and spikes. The parameters in the rate model are chosen to be typical to V1 based on Carandini 2004. The values used throughout this paper are $V_{th} = 0$ (the threshold only contributes to the rate through the difference with the membrane potential, thus we need to choose μ and V_{th} together to produce realistic membrane potential

dynamics), $\beta = 1.4$, and $k = 0.4$, which together with the 20-ms time bin corresponds to a rate of 20 Hz, typically observed in V1 in response to a high-contrast stimulus of parameters preferred by the tuning curve of the cell (Finn et al. 2007).

Since the focus of the study is to understand the implications of membrane potential fluctuations on spiking statistics, we do not seek to find a match between membrane potential recordings and the model. Rather, membrane potential statistics is assessed in terms of its consequences on spiking statistics. The two competing models, the DSP and RG models, differ in the way how spikes are generated. A detailed description of the models follows but differences in the way spikes are generated from firing rates are summarized in a cartoon depicting the main components of the models (Fig. 1).

Spike frequency adaptation. Phenomenological models of spiking are constrained to models with limited number of parameters which are feasible to fit and fitting requires limited data. These models, however, omit some dynamical properties of the neurons that can affect the response statistics. Since adaptation introduces a joint modulation of firing rates it potentially has an effect on the correlation structure of the neurons. We introduced firing rate adaptation by modulating the firing rate of the neurons by an additive term that decayed exponentially in a simulated trial. The relative modulation of the firing rate was 100% and a decay time constant of 40 ms was used (Ahmed et al. 1998).

Doubly stochastic Poisson spike-generation model. The firing rate nonlinearity establishes the link between average membrane potential and average firing rate. A widely used approach (Gur et al. 1997) assumes that the firing rate determines the probability with which a spike is generated in any particular time window. When these probabilities are independent across time we formally obtain the Poisson process. In this model, spike counts are sampled from a Poisson distribution, parametrized by the instantaneous rate, independently for each neuron n.

$$s_b^n \sim Poisson(s_b^n; r_b^n) \tag{3}$$

In summary, the Poisson model of spiking, the DSP model, relies on two sources of variability: *1*) membrane potentials are stochastically generated and these samples are correlated to represent covariability of neuronal responses; and *2*) a second source of stochasticity comes from the generation of spikes, which introduces noise that is independent across neurons.

Rectified Gaussian model. The DSP model assumes that the firing rate defines the probability with which spikes are generated in a given time window. Inspired by the relatively little stochasticity found in sensory neurons in vitro (Mainen and Sejnowski 1995), an alternative model can be formulated, which assumes a deterministic process for spike generation (Carandini 2004). Self-consistency requires that the average number of spikes with constant stimulus is equal to the firing rate. A model that assumes no further source of variability but fulfils the self-consistency criterion can be formulated by integrating the firing rate over time and generating spikes whenever the integral crosses integer values. If we consider the rate to be constant within a time bin b, scaled appropriately with the base rate parameter k, the integral becomes a finite sum, for neuron n:

$$s_t^n = \left\lfloor \sum_{b=1}^{\tau} r_b^n \right\rfloor \tag{4}$$

The spiking model defined this way, the RG model, is formally equivalent to an integrate-and-fire neuron model without refractory period with the addition of the firing rate nonlinearity. The main determinant of variability in the RG model is coming from the variability present in the membrane potentials. An additional, though minor, source of variability for the timing of spikes originates from an uncertainty of the state of the integrator at the beginning of a trial. This component is small relative to the stochasticity of the membrane

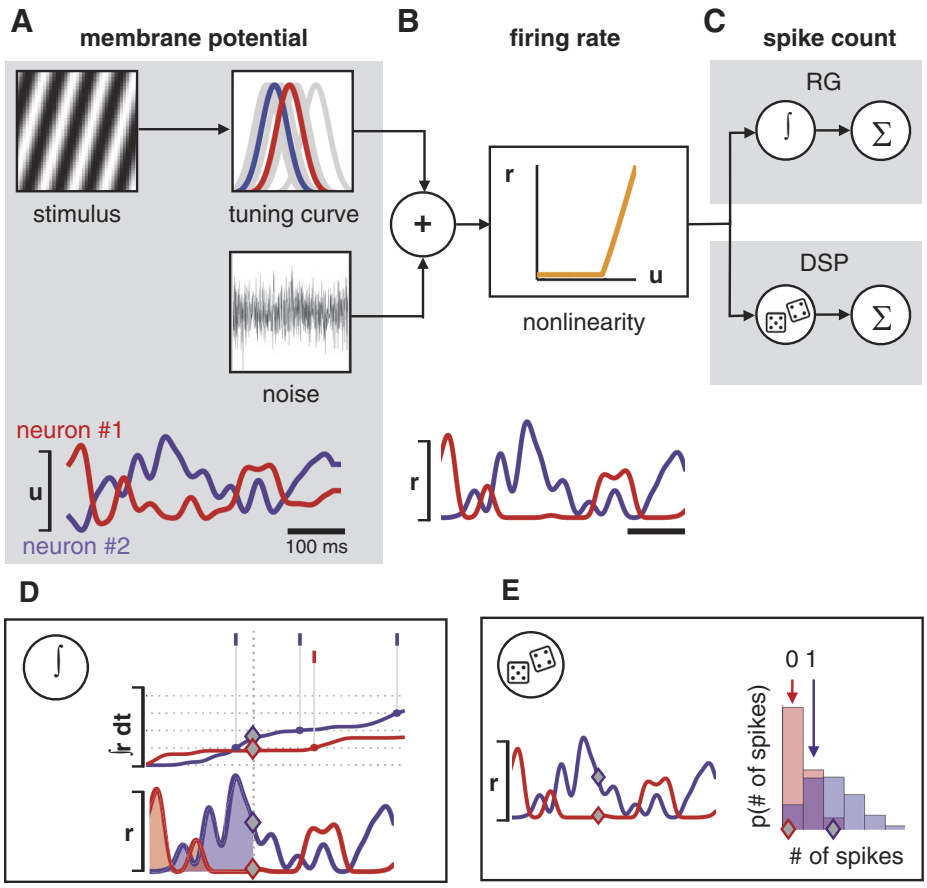

Fig. 1. Schematic illustration of the rectified Gaussian (RG) and doubly stochastic (DSP) models. *A*: membrane potential responses are characterized by a mean activation and a stochastic component that varies over time. The mean is assumed to be stimulus dependent and is determined by the orientation tuning curve. Importantly, the stochastic component is not independent but is correlated among neurons. *B*: firing rates are calculated from membrane potentials by transforming them using the firing rate nonlinearity. *C*: spike counts are specific to the spiking model used. Spike counts are obtained by summing action potentials generated throughout a trial. *D*: according to RG spiking model, spikes are generated by integrating the firing rate and deterministically registering a spike every time the integrated rate crosses an integer value. *E*: in the DSP spiking model, spike counts are stochastically generated: the time varying firing rate (normalized by the time window being considered, diamonds, *left*) determines the mean of the Poisson distribution (diamonds, *right*), which assigns probabilities to the number of spikes to be generated in the time window (*right*).

potentials at high firing rates but can affect not only spike timing but also spike counts at low firing rates.

To gain additional insight to the interdependence between membrane potential and spike count statistics in the two spiking models, moments of the firing rate can be calculated. For integer exponents (β) of firing rate nonlinearity and normal-distributed membrane potential Fano factor and pairwise correlations can be calculated for the firing rate of threshold-power-law neurons (Hennequin and Lengyel 2016). The rate moments are close approximations of the RG spike count moments, as the two distributions only differ in a quantization step. These calculations can be extended to assess the response statistics in a DSP model by using the law of total covariance when combining the rate covariance with the covariance matrix implied by Poisson variability.

Hodgkin-Huxley model neuron. To provide an insight how active membranes behave under the conditions investigated with the DSP and RG models, we set up a two-neuron simulation in which correlated current injection was provided to a pair of identical Hodgkin-Huxley neurons (Hodgkin and Huxley 1952). Similar to the DSP and RG models, samples were drawn from a bivariate normal distribution in every 20 ms. Both mean and variance of the normal distribution were identical for the two neurons; 500-ms trials were simulated and firing rate and spike count correlations were calculated from repeated trials.

Simulation of population measurements. To compare the predictions of the RG and DSP models with experimental recordings of the activity of a neuron population, population models were constructed and spike response statistics of the two models were contrasted with experimental data. Predictions of the models could be directly tested by recording pairs of neurons and assessing the changes in pairwise statistics upon stimulus manipulations. A direct mapping from stimulus to spiking responses requires a very precise receptive field characterization for the pairs and a precise control of the receptive field content. Since characterization of correlations requires tens of trials, ensuring invariant stimulus content across these trials is particularly challenging since a minor change in the fixation would introduce biases in the measured noise correlation. Therefore, instead of requiring neurons with well-characterized receptive fields, we chose an alternative approach in which variations in tuning characteristics can be accounted for. An additional benefit of the approach is that data is abundantly available and the analysis can be easily generalized for other sensory modalities.

Determining the second-order statistics of membrane potentials for a population of N neurons requires the specification of their means, variances, and correlations, as specified in *Eq. 1*. We seek to model population activity in response to oriented grating stimuli. Mean of the membrane potential response μ^n for cell n is determined by the orientation of grating stimuli through the tuning curve. Each neuron is characterized by a preferred orientation that corresponds to the stimulus orientation at which the tuning curve has a maximum. Preferred orientations are sampled independently for each neuron from a uniform distribution. The height of the tuning curve is varying from neuron to neuron and is sampled from a Gaussian distribution with a standard deviation of 0.1. The width of the tuning curve is fixed at $SD_{TC} = 0.2\pi$. The time-varying stochastic component of the membrane potential response is coming from a Gaussian distribution that is characterized by its variance. The level of variance is inhomogeneous across the population and is set randomly by sampling an inverse gamma distribution. Parameters of the inverse gamma distribution are the shape and scale parameters, with values three and four, respectively, chosen to reproduce the scale of spike count Fano factors observed in experiments. Membrane potential correlation matrices are generated algorithmically by specifying the width of the distribution with a scalar parameter (Lewandowski et al. 2009). Distribution of membrane potential correlations were tuned such that the distributions of spike count correlations were matched for the RG and DSP models.

Simulation of changes in stimulus orientation is straightforward since tuning curves define the changes in membrane potential mean, while other aspects of the statistics are assumed to be unchanged. Indeed, earlier studies on membrane potential variance have demonstrated that variance is mostly invariant to changes in stimulus orientation (Finn et al. 2007). The question of the orientation dependence of membrane potential correlations is still open but we took a conservative approach by assuming orientation independence at the level of membrane potential correlations. Modulation of stimulus contrast was simulated by modulating both the tuning curve and the level of variance. Lowered contrast was modeled by a contrast-invariant modulation of tuning curves that corresponded to a scaling of the gain of the tuning curve. Variance was increased in a homogeneous manner across the population: decreased contrast resulted in an increased variance. Direct effect of contrast change on membrane potential correlations is an open question; therefore we split the analysis into two parts. First, we analyzed contrast-dependent modulation of spiking statistics with contrast-independent membrane potential correlations. Second, inspired by indirect evidence of the effect of contrast on membrane potential synchrony (Tan et al. 2014), we simulated contrast-dependent change in membrane potential correlations by inhomogeneously modulating pairwise correlations resulting in a shifted population mean.

When the membrane potential statistics of the cell populations were fully defined, we simulated membrane potential responses by taking membrane potential samples from the population of 50 neurons. A single trial was 500 ms long; the number of trials was 1,000. The samples were then transformed to instantaneous firing rates by the rectifier nonlinearity (*Eq. 2*), and then spike counts were obtained by using the DSP and RG spiking models.

Electrophysiological data. To test the predictions of the models, we analyzed spiking data from electrophysiological recordings. Aiming for population-level characterization of responses puts constraints on the kind of measurements that are applicable for comparison. As anesthesia is known to introduce significant biases in neuronal response correlations (Ecker et al. 2014), we sought to test model predictions against data recorded from awake animals. We used publicly available data recorded in the laboratories of Matthias Bethge and Andreas Tolias (Ecker et al. 2010). Detailed description of the recording settings are available at the original publication. Briefly, single-unit recordings were obtained by extracellular electrode arrays from a population of direction selective cells (comprising both simple and complex cells) in V1 of awake monkeys. Stimuli consisted of static and moving full-field gratings. We constrained our analysis to static gratings because the static grating data set featured multiple contrast levels besides eight grating orientations. Using static gratings ensured that both simple and complex cells could be modeled solely by setting up tuning curves to define direction selectivity. Furthermore, because of the firing rate dependence of spike count variance, spike count correlations are prone to be underestimated when neural responses depend on stimulus phase, which effect can be prevented by using static stimuli. We used 400-ms segments extracted from the evoked activity period of the trials in which the spike counts were calculated.

To reliably estimate pairwise correlations, we needed to exclude some of the recordings. We only considered pairs in which both units had an average firing rate over 0.1 Hz to avoid biased correlation estimates due to the insufficient number of spiking events. Pilot analyses (data not shown) have demonstrated that low number of stimulus repetitions can lead to highly inconsistent estimates of the spike count correlations; therefore we only included recording sessions in which the number of repetitions was sufficiently high. Using a limit over trial repetitions between 30 and 40 ensures that the variance of the correlation estimates does not exceed 0.03 for small and 0.02 for large correlation levels (Kenney and Keeping 1951). Thus we included five sessions in the analysis, with repetition num-

bers (39, 40, 85, 72, 39). The filtering criteria for firing rates and trial numbers allowed us to use 41 units from the recordings.

To be able to compare spike count correlations for preferred and nonpreferred stimuli, we used a broader definition of preferred orientation. A stimulus orientation was taken as belonging to a specific neuron's preferred orientation if the firing rate of a unit averaged over the trials belonging to the given orientation was higher than the average firing rate over all trials. This binary classification scheme helped us to avoid errors in orientation preference estimation and in the same time lead to more reliable estimation of correlations in the nonpreferred condition, since including only the orientation perpendicular to the most preferred one would have produced very low firing rates. Preferred-orientation correlations were calculated for neurons where the orientation of the stimulus was preferred for both neurons of the pair.

Analysis of neural responses. We characterize the distributions of spiking responses of neurons up to second-order statistics, similarly to the descriptions of membrane potentials. However, due to specific properties of spike trains, the applied measures are slightly different. Spike count responses are characterized by variances that grow linearly with spike count means. Therefore, we are interested in changes in the variance that are independent of changes in the mean. By using Fano factor, we can control for this effect and can obtain a trial-by-trial measure of response variability for neuron n:

$$\mathrm{FF}[s_t^n]_t = \frac{\mathrm{Var}[s_t^n]_t}{\mathrm{E}[s_t^n]_t} \tag{5}$$

The choice is also supported by the fact that systematic changes in the membrane potential variance are similarly observed in the spike count Fano factor, as described by Churchland et al. 2010.

The second-order statistics of a population of N units with temporal dynamics is completely characterized by N autocorrelation functions and $N(N-1)/2$ cross-correlation functions, specifying the linear dependence between each pair at every timescale (Moreno-Bote et al. 2008). While correlations may occur at different timescales, it is a typical choice in experiments to use the correlation of spike counts over entire trials. Doing so has the advantage of taking interactions with different delays into account similarly, by sacrificing the finer temporal structure of coactivations (Smith and Kohn 2008). To account for irregularities in the firing rates and individual variances of experimentally recorded spike trains, correlations are calculated between z-scored spike counts, defined as follows:

$$\varrho^{ij} = \mathrm{Corr}\left[\frac{s_t^i - \mathrm{E}[s_t^i]_t}{\mathrm{SD}[s_t^i]_t}, \frac{s_t^j - \mathrm{E}[s_t^j]_t}{\mathrm{SD}[s_t^j]_t}\right]_t \tag{6}$$

RESULTS

We explored two models of neuronal spiking in V1 to test the effects of stimulus-change-related modulations of spiking statistics, focusing on the assessment of modulations of pairwise statistics. The DSP model and the RG model differ in the sources of the variability the models assume. Instantaneous firing rate of the models is defined in an identical manner but they differ in the way spikes are generated. The DSP model introduces substantial variability during the spike-generation process while the RG model introduces only a minimal amount of variability at this stage. Below we provide a brief overview of the main features of the models.

The DSP model builds on extensive data showing a linear relationship between spike count mean and spike count variance (Britten et al. 1993; Softky and Koch 1993; Tolhurst et al. 1983), which motivated a model of spiking activity that assumes a Poisson process at spike generation. The basic Poisson assumption constrains both the level of variability and the form of covariability. The noise introduced by the Poisson process is "private" to the neuron and therefore covariability is not accounted for by this source of variability.

Variability of spiking activity is characterized by the Fano factor. A Poisson process with a constant mean, i.e., constant firing rate, has a Fano factor of 1. Rate modulation, such as those induced by changing stimulus attributes like phase of a grating stimulus, can introduce additional variability (so called stimulus variability), thus the Fano factor can exceed 1. However, earlier studies have found changes in Fano factor that cannot be derived from the linear relationship between the mean and variance assumed by a Poisson account: in V1 Fano factor was shown to shrink when increased mean response was induced by stimulus contrast (Churchland et al. 2010; Orbán et al. 2016). A further motivation for extending the Poisson model is the form of covariability it assumes.

Covariability can be introduced in a network of Poisson-spiking neurons by joint rate modulation. Indeed, this modulation was shown to account for top-down modulation of spiking responses in V1 (Ecker et al. 2014; Goris et al. 2015). This form of joint modulation formally means a completely correlated modulation of firing rates. We relaxed the assumption of completely correlated firing rates to be able to tune both variability and covariability: a multivariate normal distribution was used to model an additional source of variability independent of the Poisson stochasticity. To increase the compatibility of this model with the formulation of the RG model, we assumed that this stochasticity was fed into the firing rate through a firing rate nonlinearity. Although the firing rate nonlinearity is not a necessary component of the Poisson model, this novel source of stochasticity has an intuitive interpretation: it can be regarded as a subthreshold variability. Based on different parametrizations of the multivariate normal distribution of the "membrane potential" dynamics, both the Fano factor and spike count correlations could be controlled largely independently of the firing rate of the model neurons. To obtain an instantaneous firing rate, the membrane potential was sampled from the multivariate normal distribution, and consecutive samples were assumed to be independent across 20-ms time bins. Independence of samples is a simplifying assumption and was motivated by the fast-decaying autocorrelation function of V1 neurons (Azouz and Gray 1999). The nonlinearity to obtain firing rates was reflecting the nonlinearity found in simple cells of V1 (Carandini 2004) (Fig. 1).

In the RG model a single source of stochasticity was assumed, which was a stochastic process at the level of membrane potential (Dorn and Ringach 2003). Intracellular measurements in awake animals have confirmed that stimulus-evoked activity of single neurons can be well characterized by a normal distribution (Haider et al. 2013; Tan et al. 2014) and using a multivariate form of this distribution is a natural extension to model covariability between neurons. Again, we approximated traces of membrane potentials as being independently sampled in discrete time bins (Fig. 1*A*). Membrane potentials were mapped through the same nonlinearity as in the DSP model, which ensured a similar evolution of firing rate with increased mean membrane potential (Fig. 1*B*). Spikes were obtained by a process that bears as little stochasticity as possible: firing rate was integrated over time and spikes were generated when the integral crossed integer values (Fig. 1*D*).

While this model does not rely on a Poisson process to ensure the scaling of spike count variance with the mean, it has been demonstrated to account for the approximately linear relationship between spike count mean and spike count variance under physiological conditions (Carandini 2004). This relationship is solely a result of the interaction of the subthreshold variability and the firing rate nonlinearity: as the mean membrane potential increases, the firing threshold together with the convex firing rate nonlinearity maps the range in which the membrane potential is varying onto an increasingly wider range of firing rates. Importantly, a consequence of this mechanism is that when the mean membrane potential is high, the effect of nonlinearity realized by the threshold is reduced since the variability results in less frequent threshold crossings. As a consequence, at high firing rates the rate of change in the variability of firing rate will decrease, ultimately resulting in reduced Fano factor, as also confirmed by recordings from V1 neurons (Carandini 2004).

Identical parametrization of the models means that the number of parameters characterizing the models are equal; thus the expressive power of the two models is similar. Parameters were established based on previously published data (see MATERIALS AND METHODS) (Carandini 2004). There are important differences, however, in the statistics of spiking activities the two models predict (Fig. 2, *A* and *B*. Assuming identical distributions for membrane potentials for a pair of model neurons, the firing rate nonlinearity and consistent spike generation processes ensure similar mean spike counts (Fig. 2*C*). The variances, however, differ for the two models: while the variance of responses of RG model neurons is dominantly determined by the appropriately scaled variance of the membrane potentials, the variance of the DSP model is the sum of the membrane potential variance and a term coming from the Poisson stochasticity. As a result, spike count variance, as well as Fano factor, of the DSP model exceeds that of the RG model (Fig. 2*C*). The correlation measured from the spike count distribution can differ from the membrane potential correlations as a result of multiple factors (Cohen and Kohn 2011; Ecker et al. 2010). Most importantly, the firing rate nonlinearity can truncate the subthreshold part of the membrane potential distribution causing a decrease in spike count correlations relative to the membrane potential correlations (de la Rocha et al. 2007), which is evident for the RG model (Fig. 2*B*). Another important consequence of the excess private variability introduced by the Poisson spike generation process is a drop in the spike count correlation (Fig. 2*C*). A simple intuition for this effect can be obtained by considering that the covariance matrix of the spike count correlation is the sum of the membrane potential covariance and the covariance of the spike generation. The latter, however, is a diagonal matrix, as this source of noise is independent among neurons; therefore it only increases the diagonal elements of the resulting covariance matrix, and, since correlation is the ratio of the covariance and the geometric mean of the variances, there is an overall decrease in the magnitude of correlations.

In the coming sections we analyze the consequences of these differences in simplified models of a pair of neurons before moving to the analysis of the pairwise response statistics of populations of neurons. When using the simplified models we do not simulate the tuning curve-mediated changes in membrane potentials; rather, we directly investigate the effects of changes in membrane potential statistics. These analyses provide predictions on changes in spike count statistics expected in response to changes in stimulus orientation and contrast.

Matching spike response statistics. To be able to contrast the effects of stimulus change on response statistics of the competing models, we first establish a method for matching the spiking statistics of the DSP and RG models in a pair of neurons. Since equal membrane potential statistics lead to different spiking statistics, it is clear that either membrane potential statistics or parameters of the firing rate nonlinearity need to be adjusted to have matching firing rates, Fano factors, and spike count correlations. To keep our arguments simple, we keep the firing rate nonlinearity unchanged. In fact, the scale of membrane potential (which is determined together by the distance of the membrane potential from the firing rate threshold and the variance of the membrane potential) and the scale of firing rate nonlinearity (parameter k) can be altered

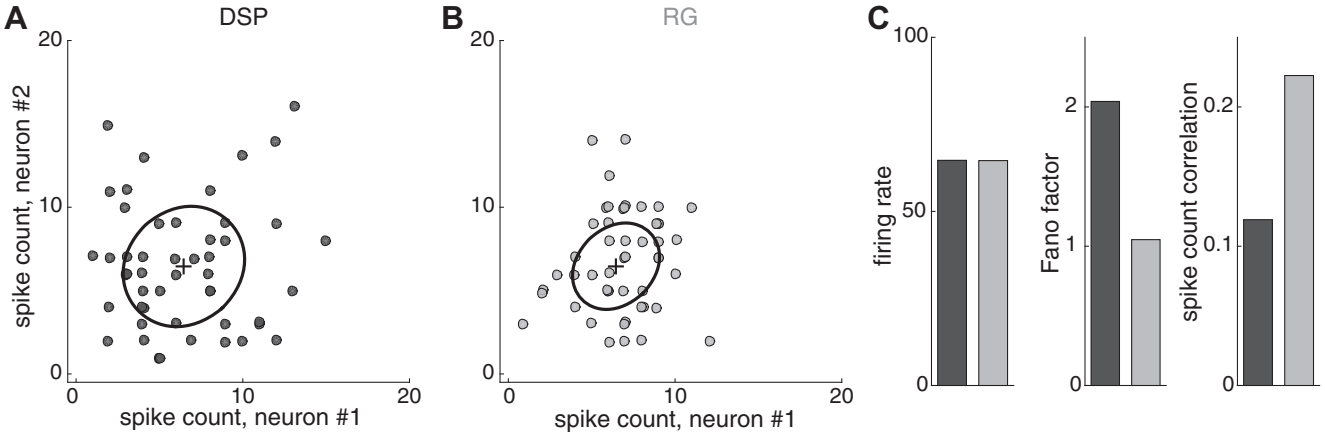

Fig. 2. Spike response statistics for the DSP and RG models. *A* and *B*: spike count distribution for a pair of model neurons with the same membrane potential statistics (mean, variance, correlation) for DSP (*A*) and RG models (*B*) in a simulated 100-ms time window. Membrane potential means and variances were identical for the 2 neurons and membrane potential correlation was set to 0.25. Circles indicate spike counts for individual trials. Cross shows the across-trial mean while ellipse represents the across-trial covariance ellipse of the joint spike count distribution. Small jitter was added to spike counts for illustration purposes. *C*: while the spiking models are consistent in predicting equal mean spike count responses for both DSP (dark bars) and RG models (light bars) at matching membrane potential distributions, Fano factors and spike count correlations show characteristic differences. Note that spike count correlations are systematically lower than membrane potential correlations at both models.

largely interchangeably (see Fig. 4, *G–I*); therefore the argument can be translated into changes in the scale parameter of the firing rate. The DSP and RG models reside at two extremes of the spectrum of doubly stochastic processes: DSP features high level of variance at spike generation while RG demonstrates minimal variability at that point. It has been established that spike count is not sufficient to uniquely determine the spiking model (Amarasingham et al. 2015) even if Fano factors and spike count correlations are also fixed. This finding ensures that matching the second-order spiking statistics is possible for a particular setting of spiking mean, variance, and covariance, but it does not ensure that modulating one aspect of the response statistics will result in equal changes in other aspects of response statistics at both the DSP and RG models.

By exploring membrane potential parameters for the two models, we can obtain a parameter setting where firing rates (Fig. 3*A*) and Fano factors (Fig. 3*B*) are matched. Both of these criteria constrain the parameter sets up to a linear combination of the tested parameters, and therefore the intersection of the lines allows matched firing rate and Fano factor. Difference between the spike count correlations along the explored parameter range is relatively untouched (Fig. 3*C*) and can be adjusted by tuning membrane potential correlations. The resulting statistics-matched models (Fig. 3, *D–F*) have markedly different membrane potential variances (0.2 and 2.75 mV^2 for the DSP and RG models, respectively) and different membrane potential correlations (0.95 and 0.13 for the DSP and RG models, respectively). Because of the extra variance and the decorrelation effect of the DSP model, these differences are expected and highlight that the Poisson process introduces a private variability that can easily wash out membrane potential correlations. While the level of membrane potential correlations required in the DSP model seems to be extreme, it only serves the purpose of matching the spiking statistics. At lower firing rates the excess variance added to the membrane potential covariance would be lower and therefore statistics matching would require considerably lower membrane potential correlations.

Dependence of second-order spiking statistics on the membrane potential mean. The membrane potential mean in direction selective V1 neurons is sensitive to stimulus orientation. To understand the effects incurred by orientation change on spiking statistics, we need to separate the effects on different aspects of the response statistics. Changes in membrane potential mean have obvious effects on the firing rate. Effects of mean membrane potential on other characteristics of the re-

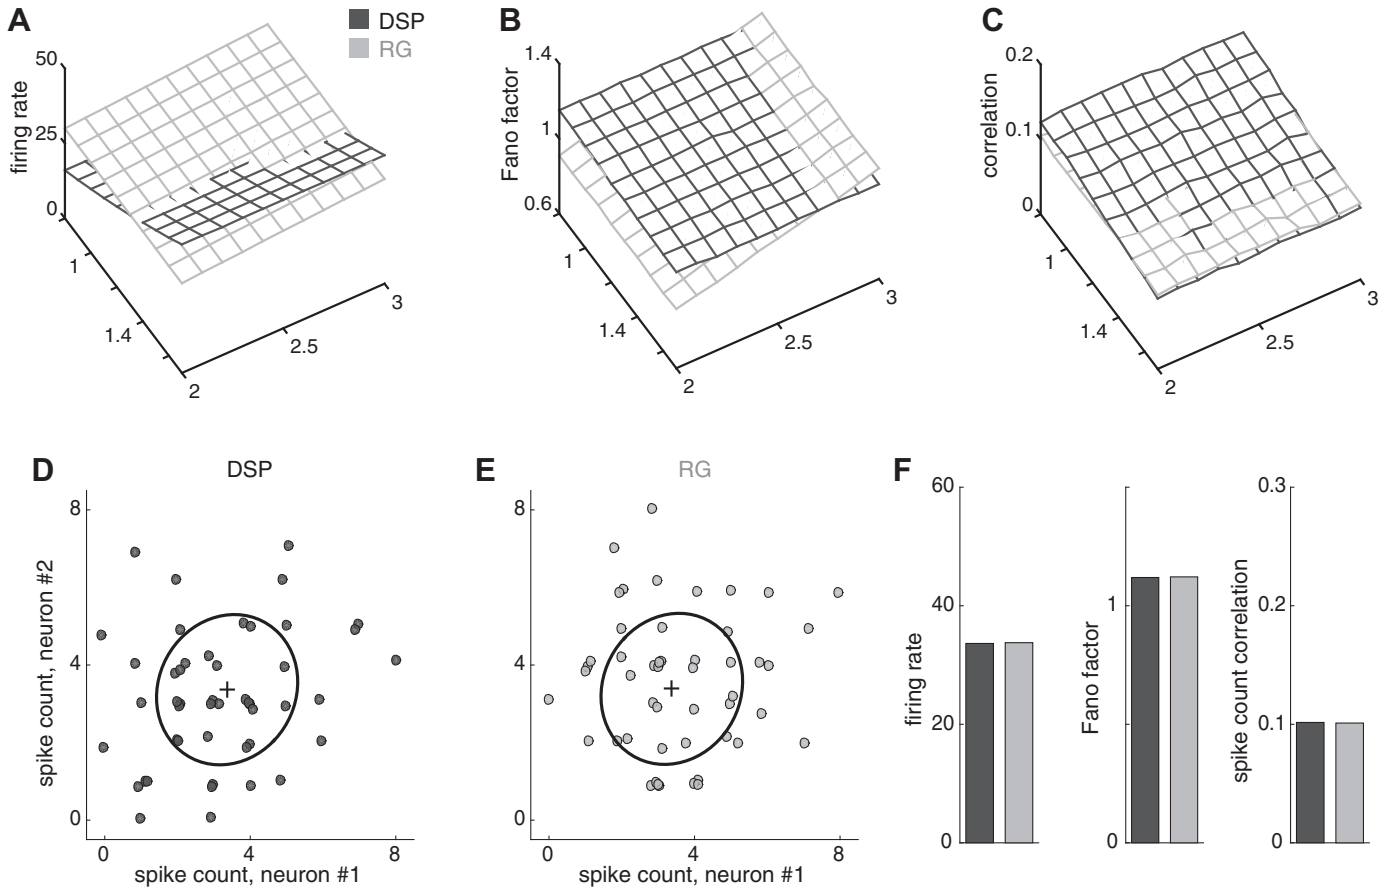

Fig. 3. Matching the output statistics of DSP and RG models. *A–C*: spike count (*A*), Fano factor (*B*) and spike count correlation (*C*) of the neurons as a function of membrane potential parameters for the DSP (dark plane) and RG model (light plane). Simulations show responses in a 100-ms time window. Intersections of the plane indicate equal statistical measure for the 2 models, which then help to determine the spike count statistics. Parameters explored for matching spiking statistics were the mean membrane potential of the DSP model and membrane potential variance of the RG model. *D–F*: by tuning membrane potential statistics for the DSP (*D*) and RG (*E*) models individually, the spike count statistics of the models can be matched up to the mean, variability (as measured by the Fano factor), and correlation. Spike count distribution representation is the same as that on Fig. 2. Matched means. *F*: spiking statistics for the matched models. Colors are identical to panels *A–C*.

sponse statistics are less straightforward. To be able to see how the membrane potential mean changes affect Fano factors and spike count correlations in DSP and RG models, we aim to see changes in these measures independent of changes in firing rates.

Membrane potential statistics are different in the statistics-matched DSP and RG models. Therefore, we first establish rates of change for the membrane potentials in the two models, which guarantee that even upon deviating from the statistics-matched levels of membrane potentials, the firing rates change at a similar rate (Fig. 4*A*). Using this firing rate-matched scenario, we can directly contrast mean-related changes in Fano factors (Fig. 4*B*) and spike count correlations (Fig. 4*C*). Fano factors are approximately equal (Fig. 4*B*) across the range of membrane potentials for the DSP and RG models, which resonates with the results of Amarasingham et al. (2015) that single-cell spiking statistics are not sufficient to unambiguously identify the spiking model even if responses are analyzed for a wide range of stimulus parameters. Independence of Fano factors from membrane potentials (Fig. 4*B*), and from firing rates as well, confirms the original results of the RG model (Carandini 2004) and is also expected for the DSP model. Growing firing rate resulting from increased mean membrane potential has a differential effect on spike count correlations in the two models (Fig. 4*C*). In the DSP, increased firing rate incurs increased private variability, which suppresses the contribution of the membrane potential covariance to the total covariance and therefore spike count correlation diminishes. In the RG model, however, a different mechanism dominates: increased mean activation results in a higher proportion of the membrane potential covariance to be above firing threshold, and, consequently, smaller truncation of this distribution boosts the magnitude of the measured correlation (de la Rocha et al. 2007). Firing rate frequency adaptation might introduce a form of collective change in firing rates. We briefly tested whether the qualitative behavior of the models is affected by the presence of firing rate adaptation (Ahmed et al. 1998). Our two-neuron simulations confirmed that the mean membrane potential dependence of spike count correlations is only marginally affected by firing rate adaptation. Taken together, under controlled change in mean activity, spiking variability is relatively insensitive to the choice of DSP or RG models. Analysis of spike count correlations, however, provides opposing predictions for DSP and RG models when mean responses are systematically changed.

Differential effects of the membrane potential mean on spike count correlations in the DSP and RG models highlight an opportunity to distinguish the spiking models when stimulus orientation is varied. A detailed exploration of the evolution of spike count correlation with changing membrane potential mean can confirm the robustness of the effect seen on Fig. 4*C*. We tested robustness by assessing spike count correlations at different levels of membrane potential correlations (Fig. 4, *D* and *E*). Analysis of the RG model reveals a monotonic rise of the magnitude of spike count correlations from zero toward the level of the membrane potential correlation as membrane potential mean increases (Fig. 4*E*). This dependency of spike count correlation on the level of excitation is reflecting the behavior of Hodgkin-Huxley neurons that exhibit increasing spike count correlation with elevated levels of current injections (Fig. 4*J*), In the DSP model, the effect of a less truncated membrane potential joint distribution, when a larger proportion of the distribution gets above threshold, is shown at a low membrane potential regime (Fig. 4*D*): at low levels of activations, a rise similar to the RG model can be observed. The range of similar evolution of spike count correlation in the DSP and RG models, however, is severely limited (Fig. 4*F*) and can only be observed at moderate firing rates but high Fano factors, which combination is not characteristic of cortical neurons. Beyond that point, a steady decline of spike count correlation takes place that converges to zero (Fig. 4, *D* and *F*). Thus patterns in the change of spike count correlations upon stimulus manipulation can reliably remove the ambiguity between the DSP and RG models.

The biphasic profile of membrane potential dependence of spike count correlation raises the possibility that an increased gain in the firing rate nonlinearity can simply scale the firing rate profile of the DSP model. Thus the regime where the spike count correlation is positively correlated with membrane potential mean could possibly overcome the limited range shown on Fig. 4*D* and could reach higher firing rates. We tested this question by scaling the firing rate gain together with inverse scaling of the membrane potential (Fig. 4*G*). To keep not only the firing rate but also the Fano factor constant, the variance of the membrane potential was also scaled together with the mean and rate gain parameters (Fig. 4*G*). In the resulting setting we could test a wide range of the gain parameter k, while keeping the mean, the Fano factor, and the correlation of spiking responses constant (data not shown). The firing rate profile was identical for the different parameter settings (Fig. 4*H*). Importantly, the evolution of correlations was very close at different settings of the gain parameter k, with no visible shift in the membrane potential (or alternatively firing rate) value maximizing the correlation (Fig. 4*I*). This analysis demonstrates that the regime where increasing correlations are present with increasing firing rates is constrained to low firing rate levels and high Fano factors.

Firing rate nonlinearity has a major contribution both to shaping the response variability (Fano factor) and the correlation between neuronal responses. To gain an insight into the effects of firing rate nonlinearity, we used analytical calculations to assess the evolution of response statistics at two forms of firing rate nonlinearity: at the threshold-linear model (corresponding to an exponent of 1) and at the threshold-quadratic model (corresponding to an exponent of 2). Analytical results can be obtained for normal-distributed subthreshold activity with threshold-power-law nonlinearity at integer exponents (Hennequin and Lengyel 2016). These calculations approximate our RG model well, and can be easily extended to the DSP model. Firing rate responses in the case of a threshold-linear firing rate transformation are characterized by a soft-threshold behavior (Fig. 5*A*), which is a consequence of response variability: while the membrane potential is transformed through a hard threshold, membrane potential variance affects the mean firing rate, which results in a gradual transformation. Membrane potential mean dependence of Fano factors (Fig. 5*B*) reveals opposing trends for the different firing rate exponents, indicating that the firing rate independence of Fano factor, characteristic of Poisson spiking, is sensitive to the choice of nonlinearity. Membrane potential dependence of correlations (Fig. 5*C*) in the RG model are consistent at the two firing rate exponents and confirm our findings using physio-

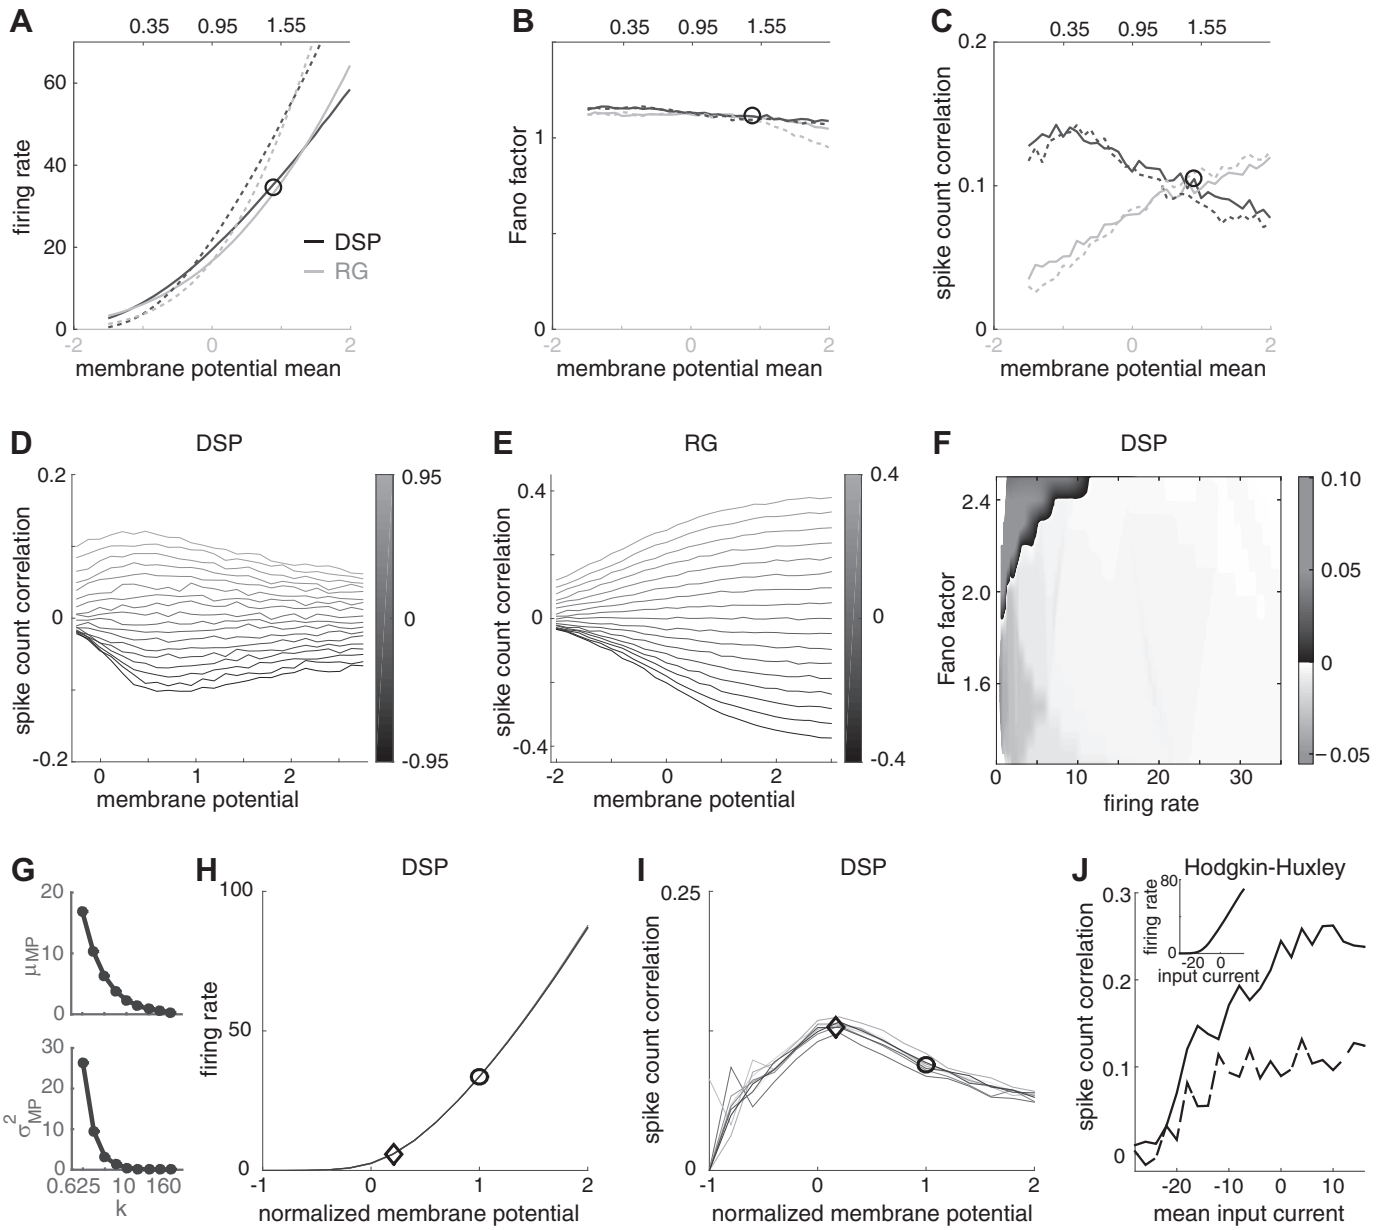

Fig. 4. Dependence of spike count statistics on membrane potential mean. *A–C*: predictions of the 2 spiking models (solid lines) for firing rate (*A*), Fano factor (*B*), and spike count correlation (*C*) upon changes in membrane potential mean. Matched-statistics DSP and RG are tested with changing membrane potential means but constant membrane potential variances. Spiking with adapting firing rates are included for comparison (dashed line). *A*: approximately equal firing rates are achieved by appropriate linear scaling of the membrane potentials. Note that to match the spiking statistics, the membrane potential statistics (including the mean, variance, and correlation of neurons) differ in the 2 models (see different horizontal axes at the bottom and the top of panels). *B*: Fano factors are close to equal and invariant across the whole range of membrane potentials. *C*: while firing rates and Fano factors cannot explicitly differentiate between the DSP and RG models, correlation shows systematic differences: pairwise correlation in the DSP model decreases with increasing membrane potential levels, but it increases in the RG model. Except for the changing mean membrane potential, parameters of the models are the same as those used on Fig. 3 (matched mean is denoted by black circle). *D* and *E*: evolution of spike count correlations with changing membrane potentials at different membrane correlation levels in the matched statistics DSP (*D*) and RG (*E*) models. The DSP model initially exhibits a brief increase of the correlation magnitude toward the level of membrane potential correlation but the tendency reverses and a decline toward 0 spike count correlation takes place. Color bar: range of membrane potential correlations tested. The RG model is characterized by a steadily increasing magnitude of spike count correlation, saturating at the correlation of the membrane potentials. Color bar: range of membrane potential correlations tested. *F*: exploration of spike count statistics in the DSP model to assess the regimes in which correlation magnitude increases with the firing rate. Positive slope regimes (dark hues) are constrained to low firing rates and high Fano factors, while the majority of mean firing rate Fano factor combinations are characterized by negative slope. Sharp edge between the negative and positive slope regimes correspond to the peak of curves on panel *D*. *G–I*: analysis of the dependence of spike count statistics on the base rate. Membrane potential mean (μ_{MP}) and variance (σ^2_{MP}) (*G*) are tuned for a range of values of the base rate parameter *k* to produce matching spike count statistics (firing rate, Fano factor, spike count correlation at one particular firing rate level, circle in *H* and *I*). *H*: membrane potential dependence of firing rates is identical for the different gain levels, *k* (lines are overlapping). Membrane potentials are normalized to the values established in *D*; diamond denotes the maximum on *I*. *I*: membrane potential dependence of correlations is similar across different levels of the gain, *k* (gray lines). *J*: spike count correlation between a pair of Hodgkin-Huxley neurons at different current injection levels. Standard deviation of the Gaussian-distributed current was 9 mS; correlation between the currents was 0.2 and 0.4 for the dashed and solid lines, respectively. *Inset* shows the firing rate of the neurons (curves of different current correlations overlap).

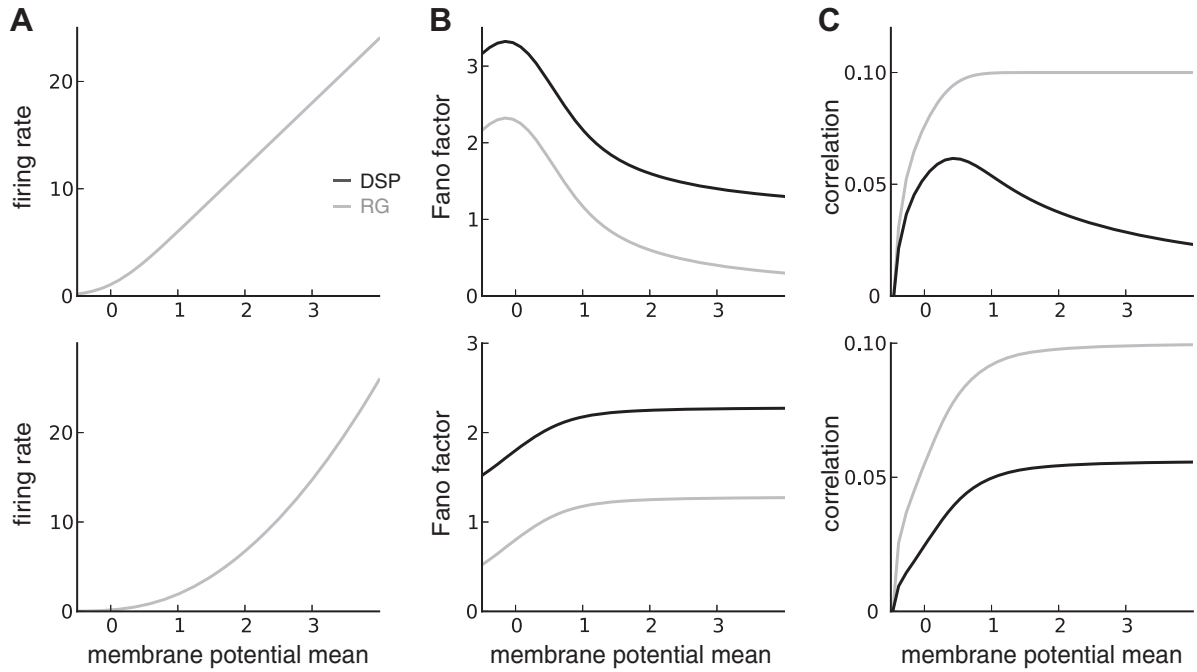

Fig. 5. Analytical evaluation of membrane potential mean dependence of activity statistics. Analysis of firing rate (*A*), Fano factor (*B*), and pairwise correlations (*C*) for the DSP (black lines) and RG (gray lines) models at different firing rate nonlinearities (firing rate exponent was 1 and 2 for *top* and *bottom*, respectively). Note that for the transparency of the analytical treatment Fano factor and correlations were not matched.

logical firing rate exponents. Correlations in the DSP model, however, show different patterns. At an exponent of one membrane potential mean dependence of correlation qualitatively matches that observed at the firing rate exponent characteristic of V1 neurons. At an exponent of 2 the DSP correlation shows a qualitative behavior matching that of the RG model: the correlation monotonously increases with increasing membrane potential mean. At high-firing-rate exponent the two interacting sources of variance at the DSP model cause a switch in qualitative behavior. Increased mean firing rate increases the variance of a Poisson process. Meanwhile the nonlinearity stretches the variance (and covariance) of the normal distribution. At higher values of the exponent, the stretching resulting from the firing rate exponent overcomes the amount of stretch in the variance resulting from the Poisson process, the covariance elements of the normal distribution will dominate the zero-covariance of the Poisson, and thus the correlation will grow. This analysis highlights that different characteristics of response statistics show joint changes with altered firing rate exponents.

Dependence of spiking statistics on membrane potential variance. Stimulus contrast was demonstrated to have a combined effect on membrane potential mean and variance (Finn et al. 2007): while the mean of the membrane potential response shrinks as contrast goes to zero, membrane potential variance grows. Thus changes in membrane potential response variance have relevant consequences on the spiking statistics. To get insights into the effects of joint changes in membrane potential mean and membrane potential variance, we first explored these characteristics separately and then turned to the combined effects of parallel changes.

After the mean-dependent changes discussed in the previous section, we set out to analyze the membrane potential variance dependence of spiking statistics (Fig. 6). Similar to the protocol followed at testing the effects of membrane potential mean, we started from the matched-statistics DSP and RG models and set the range of variance scaling such that the resulting firing rate changes in the two models are approximately equal (Fig. 6*A*). Again, this mean firing rate-matched approach ensures that changes seen in the Fano factors and spike count correlations are not related to differences in firing rates. Contrasting the Fano factors at firing rate-matched settings of the DSP and RG models revealed similar tendencies but slightly differing values for the Fano factors (Fig. 6*B*). Increased membrane potential variance translated into increased Fano factors in both of the spiking models, but the DSP model was characterized by systematically larger Fano factors at higher membrane potential variances (Fig. 6*B*). This difference is due to the excess variance of the DSP model coming from the increased spike count variance of the Poisson stochasticity at higher mean spike counts. Changing membrane potential variance showed conflicting effects in the two models on spike count correlations (Fig. 6*C*). While the spike count correlation in the RG model was relatively insensitive to changes in membrane potential variance and thus firing rate, the DSP model was shown to exhibit increased spike count correlation with increased membrane potential variance (Fig. 6*C*). This increase can be easily understood by recognizing that a scaled membrane potential covariance results in a larger relative contribution of the membrane potential covariance to the total covariance, thus the spike count correlation will be more dependent on the correlated membrane potential stochasticity than on uncorrelated spiking stochasticity. Thus analysis of the dependence of spike count statistics on membrane potential variance reveals that DSP and RG models can be distinguished based on spike count correlation of the responses when subthreshold variance is modulated.

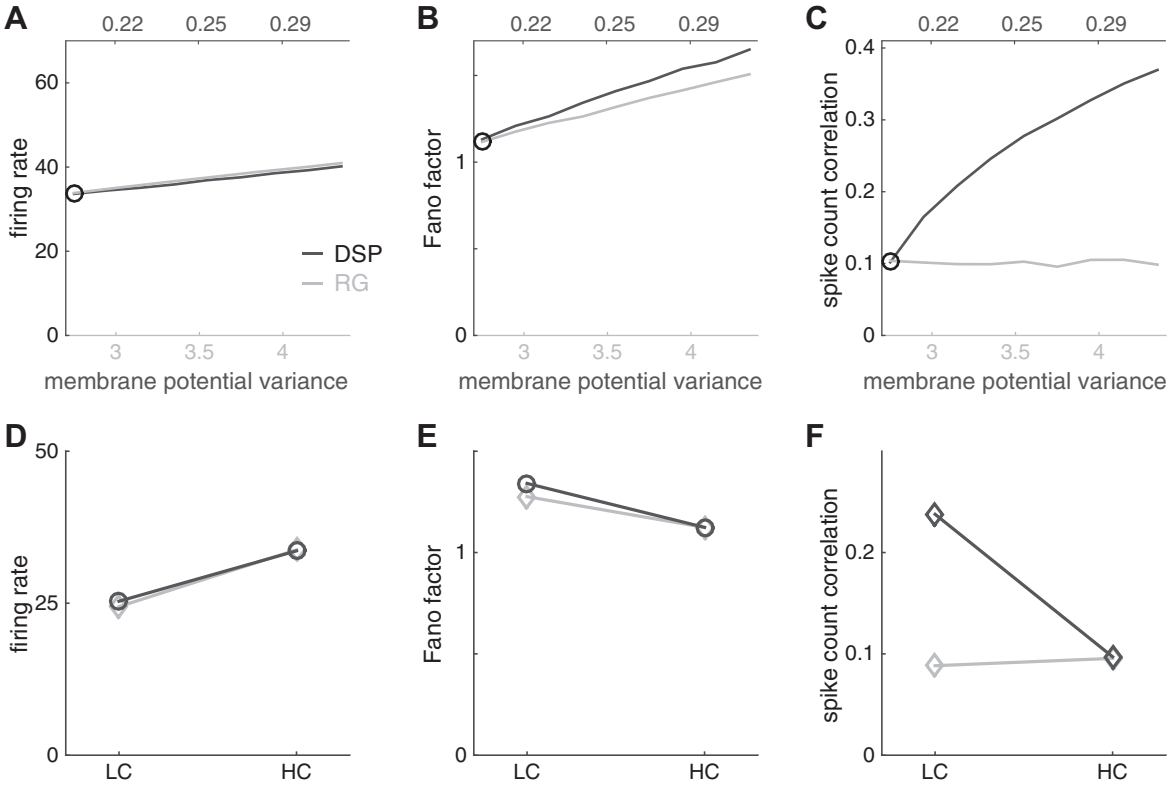

Fig. 6. Effect of membrane potential variance on spike count statistics. *A–C*: matched-statistics DSP and RG are tested with changing membrane potential variances but constant membrane potential means. Firing rates (*A*) and Fano factors (*B*) are approximately equal but display an increasing tendency with increased membrane potential variance in both models, but spike count correlations (*C*) show systematic differences. While spike count correlations are invariant across the range of variances for the RG model, the Poisson model shows a steady increase with increasing membrane potential variance. Dark lines, DSP model; light lines, RG model. Note that to match the spiking statistics, the membrane potential statistics (including the mean, variance, and correlation of neurons) differ in the 2 models (see different horizontal axes at the bottom and the top of panels). Except for the changing mean membrane potential, parameters of the models are the same as those used on Fig. 3 (matched variance is denoted by black circle). *D–F*: contrast dependence of spiking statistics. Spike counts (*D*), Fano factors (*E*), and spike count correlations (*F*) in the 2 models with concomitant changes in membrane potential mean and variance stimulating changes expected in V1 upon switch from high-contrast (HC) to low-contrast (LC) stimulus: decreased mean was followed by increased variance. Under such conditions the spike count correlation is expected to decrease for the DSP model but increase for the RG model when contrast is increased (*F*).

Contrast change incurs concomitant changes in membrane potential mean and variance. Based on the previous analyses we can conclude how these parallel changes interact when the spiking statistics are considered in response to a stimulus at lower contrast. In terms of firing rate, decreased mean membrane potential and increased variance due to reduced contrast have opposing effects, but the effect of decreased membrane potential dominates patterns in firing rate (Fig. 6*D*). In terms of Fano factor, it is the change in variance that causes increased Fano factors in both models (Fig. 6*E*). In terms of spike count correlations, the two models have distinct predictions (Fig. 6*F*). It is only the change in the mean membrane potential that contributes to a shrinking magnitude of correlations in the RG model (Fig. 4, *C* and *E*). In the DSP model, one component contributing to contrast-related changes is the increased variance, which causes larger spike count correlations (Fig. 6*C*). The effect of decreased mean seems to be more complex: it results in increased correlations in a wide range of parameters and decreased correlations in a specific subspace of the parameters (Fig. 4*F*). Remarkably, this subspace is characterized by low firing rates and relatively high Fano factors. In summary, contrast modulation-related changes in membrane potential variance introduce changes in Fano factors and spike count correlations in both models. The magnitude of the variance change determines the difference in Fano factors between high-contrast and low-contrast conditions but the direction of deviation is the same for both models. Contrast has opposing effects on spike count correlations in the two models. In the DSP model both increased membrane potential variance and decreased membrane potential mean incur higher spike count correlations at lower contrast levels. In the RG model, however, excess variance does not affect spike count correlations and therefore changes in spike count correlations are solely determined by changes in mean membrane potential which ultimately results in decreasing correlations with decreasing contrast. Thus, similar to stimulus orientation-related changes in the two-neuron model, changes in stimulus contrast induce changes in the spike count correlations that distinguish the DSP and RG models.

Contrast- and orientation-dependent modulation of spiking correlations in a population of direction selective neurons. To test the predictions of the two models against experimental data, we simulated the activity of a population of V1 direction selective neurons in response to changes in stimulus. As demonstrated by our analysis of the two-neuron model, stimulus change-related changes in variability do not distinguish between the DSP and RG models. Therefore it is expected that population distributions of Fano factor are indistinguishable too. Spike count Fano factors show little dependence on stimulus orientation and are increased when stimulus contrast

is lowered (Orbán et al. 2016), as predicted by both spiking models. As a consequence, we focus on the analysis of stimulus dependence of correlated variability in spiking responses of a population of V1 neurons.

We tested how decreasing stimulus contrast affects the population distribution of spike count correlations in the two models and compared the results to experimental data. Traditionally, studies focus on the population mean of the distribution of correlations; however, as positive and negative correlations may be affected differentially by stimulus manipulation (as in Fig. 4, *D* and *E*), it is important to take the shape of the distribution into account as well. In particular, we focused on the change in the standard deviation of the spike count correlation distribution from lower to higher contrast stimuli (ΔSD_{ctr}). First, we assumed that contrast affects the mean and variance of the membrane potential responses but leave correlations intact and investigated the effects of varying membrane potential correlations in a subsequent analysis. We set up two populations of 50 model neurons each, implementing different spiking profiles corresponding to the DSP and the RG models. We simulated membrane potential activity using tuning curve responses to a full-field grating stimulus (see MATERIALS AND METHODS). Single-cell spiking statistics in the model populations were matched (mean firing rates were 7.9 and 5.4 Hz; mean Fano factors were and 1.1 and 0.9 in the high-contrast condition in the DSP and RG populations, respectively). The membrane potential correlations in the two populations were chosen such that mean and width of the two spike count correlation distributions are matched at high-contrast stimulus presentation (Fig. 7, *A* and *B*, *insets*). Constructing membrane potential covariance matrices such that they reflect the correlation between spike count correlations and similarity in orientation preference (Lin et al. 2015) did not affect the results (data not shown). The distribution of spike count correlations was tuned to have width and mean comparable to physiological values observed in the data recorded (−0.02 and 0.16 for the mean and standard deviation in high-contrast conditions, respectively) (Ecker et al. 2010). The distribution was wider for membrane potentials than for spike counts both in the case of the DSP and RG models and, as expected from previous analyses, the membrane potential correlation distribution was wider for the DSP model (0.42 and 0.22 for the DSP and RG

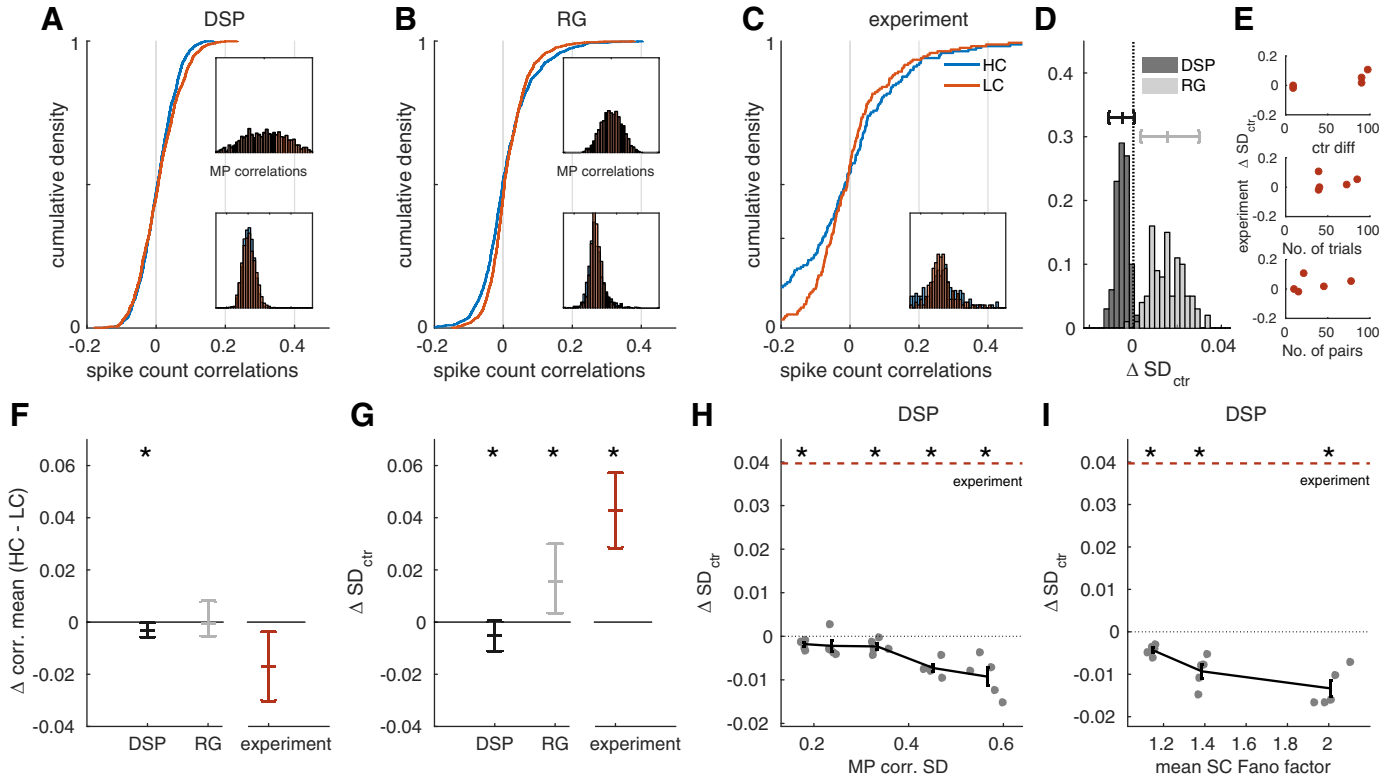

Fig. 7. Comparison of model predictions with experimental data: effect of stimulus contrast. *A*: Cumulative distributions (CDF) of pairwise spike count correlations in a simulated population of 50 DSP neurons in response to a low-contrast (LC) and high-contrast (HC) stimulus, using membrane potentials (MP) from tuning curves proportional to the contrast. *Insets* show the histogram of membrane potential correlations and the spike count correlation histograms corresponding to the CDFs. *B*: distributions of pairwise spike count correlations in a simulated population of 50 RG neurons in response to the same stimuli. *C*: distributions of measured spike count correlations in response to grating stimuli of 10 and 100% contrast (data from Ecker et al. 2010). *D*: distributions of ΔSD_{ctr}, in 100 simulated sessions using the 2 model populations, error bars representing 95% confidence intervals (CIs). *E*: ΔSD_{ctr} in the experimental sessions. Since recording parameters and stimulus properties vary across sessions, ΔSD_{ctr} are shown against 3 relevant characteristics of the sessions: contrast difference (ctr diff) between the 2 conditions (*top*), number of trials (*middle*), number of recorded pairs. All 3 measures show a positive correlation with ΔSD_{ctr}. *F*: differences in the means of the 2 correlation distributions in the 2 simulated and experimentally recorded populations. Error bars represent 95% CIs for simulated, SE over 5 bootstrapping samples from the experimental population. Asterisks represent significant difference of the mean from 0 at the $P = 0.05$ level. *G*: ΔSD_{ctr}. Mean of the differences are shown for the 2 models across different populations and for experimental recordings across different sets of left-out neurons. The mean of ΔSD_{ctr} is significantly different from 0 in all populations, with the RG agreeing with data regarding the sign of ΔSD_{ctr} and the DSP not. *H*: dependence of ΔSD_{ctr} in the DSP model on the width of MP correlation distribution (MP corr. SD). *I*: ΔSD_{ctr} in the DSP model at different values of MP mean and variance, shown as a function of the spike count (SC) Fano factor, testing whether low mean and high variance in MP can reverse contrast effects in spike count correlations.

models, respectively). Lowering the stimulus contrast made the spike count correlation distribution wider in the DSP population (Fig. 7*A*). The increase in the magnitude of the correlations at lower contrast levels is a direct consequence of the effect demonstrated in a pair of neurons (Fig. 6*F*), namely that lower firing rates and increased membrane potential variance tilts the balance between correlated membrane potential and independent spiking variability toward the former. Conversely, in the RG population, the correlation distribution becomes narrower with decreasing stimulus contrast (Fig. 7*B*). This effect is again in agreement with the results from the analysis of a pair of neurons (Fig. 6*F*).

Analyzing population responses from V1 reveals contrast-dependent changes in the distribution of spike count correlations. The distribution of spike count correlations is narrower in response to low-contrast stimuli than in response to high contrast (standard deviations of 0.16 for high and 0.12 for low contrast, respectively; Fig. 7*C*). The uncertainty of these estimates was assessed by bootstrapping (discarding 20% of all pairs 5 times), providing an estimate for the standard error of mean for the change in the standard deviation of the spike count correlation distribution. The same analysis was performed for the two model populations to obtain a comparable estimate of uncertainty using 100 simulations and 41 randomly selected units, similar to the number of single-units available from the experiment. Repeated simulations were necessary to assess the uncertainty of the statistical estimates arising due to the limited population size. ΔSD_{ctr} was significantly positive in the experiment (Fig. 7*G*, one-sample *t*-test $t(4) = 2.96$, $P = 0.042$), similarly to the RG model (standard deviation of 0.08 and 0.06 in the high- and low-contrast conditions, respectively, $P = 0.010$; Fig. 7*G*).This result is in contrast with the negative ΔSD_{ctr} in the DSP model (standard deviations of 0.051 and 0.060 in the high- and low-contrast conditions, respectively, $P = 0.040$; Fig. 7*G*). Changes in the mean of correlations are relatively small and are consistent across both models and experimental data (in the experimental population, −0.015 and 0.001 in the high- and low-contrast condition, respectively, one-sample *t*-test $t(4) = -1.27$, $P = 0.27$; in the DSP population, 0.004 and 0.009, $P = 0.010$; in the RG population, 0.009 and 0.012, $P = 0.38$; Fig. 7*F*).

The membrane potential statistics in both models were chosen to match the spiking statistics in the analyzed data set as precisely as possible. However, we do not have a direct assessment of the correlation distribution nor of the private response statistics of the membrane potential. Thus, to test the robustness of the dissociation of the two models by contrast-related changes in spike count correlations, we explored how much ΔSD_{ctr} depends on the specific settings we used in the model. In the case of the RG model, intuitions obtained from the two-neuron analysis confirm that the positive ΔSD_{ctr} is robust against changes in parameters. In the case of the DSP model, however, the specific settings might affect the sign of ΔSD_{ctr} (see Fig. 7, *B* and *C*). We varied the magnitude of membrane potential correlations in the DSP model to assess whether the sign of ΔSD_{ctr} could be reversed at specific settings. Simulations confirmed that this manipulation is not capable of reproducing the experimentally observed pattern (ΔSD_{ctr} being significantly or nonsignificantly negative at all tested values; Fig. 7*H*). The analysis described in Fig. 7 motivated us to investigate whether ΔSD_{ctr} could also be positive in the DSP model population when intensively decreasing means and increasing variances concomitantly in the membrane potential. Such changes did not produce a positive ΔSD_{ctr} within (nor substantially above) the Fano factor range plausibly observed in measurements (the changes being significantly or nonsignificantly negative at all tested values; Fig. 7*I*). These analyses highlight that second-order statistics of population responses to contrast-varied stimuli reliably discriminate between the DSP and RG models of neural spiking and it is the RG model that provides predictions compatible with contrast change-related data.

Experimental studies suggest that changes in stimulus contrast might be reflected in changes not only in mean and variance of membrane potentials but also in correlations (Tan et al. 2014). We investigated this possibility by assessing how much ΔSD_{ctr} changes in the models if the magnitude (average absolute value) of membrane potential correlations changes with contrast. Simultaneously, we tested different levels of changes in the membrane potential mean (which was multiplied by a factor of 2 from the low- to the high-contrast stimulus in the analyses described above), and the average absolute value of correlations (which did not change between the two contrast conditions in the analyses shown in Fig. 7). The larger the increase was in the membrane potential mean with contrast, the larger the reported effect was in ΔSD_{ctr} (Fig. 8). As the analysis of Tan et al. (2014) indicated a decrease in membrane potential correlations with increasing contrast, average membrane potential correlation ratios below 1 are of particular interest. At higher levels of increase in the mean, ΔSD_{ctr} in the RG population remains positive for 10 or 20% decreases in membrane potential correlations, while in the DSP population it is consistently negative.

When using grating stimuli, a well-controllable manipulation of stimuli besides stimulus contrast is altering the grating orientation. Thus we compared model predictions in response to such changes with experimental data as well, in particular regarding the change in the standard deviation of the spike count correlation distribution from nonpreferred to preferred contrast stimuli (ΔSD_{ori}). Population models are constructed in an analogous manner to the analysis of contrast dependence. Preferred orientations are uniformly distributed, and therefore a simple orientation change is not expected to cause changes in population response statistics. Therefore, to assess orientation-related changes in population responses, we needed to separate neurons that respond strongly to a specific orientation from those that respond weakly to the same stimulus. We classified any particular orientation as preferred or nonpreferred orientation based on whether the response of the neuron was above or below its average response intensity. To construct distribution of correlations for preferred and nonpreferred directions, pairs of cells were selected based on whether both of the cells had the actual stimulus among their preferred orientation or both had it among their nonpreferred orientations. Simulation results were contrasted with recordings of populations of V1 single-units (Ecker et al. 2010) in which we also separated pairs of units observing preferred and nonpreferred stimuli (see MATERIALS AND METHODS). In the experimental data set, we observed a positive ΔSD_{ori} (standard deviations of 0.15 for preferred and 0.10, for nonpreferred stimuli, one-paired *t*-test $t(4) = 6.33$, $P = 0.003$; Fig. 9, *C* and *E*). In the DSP population, ΔSD_{ori} was not significantly different from zero (standard

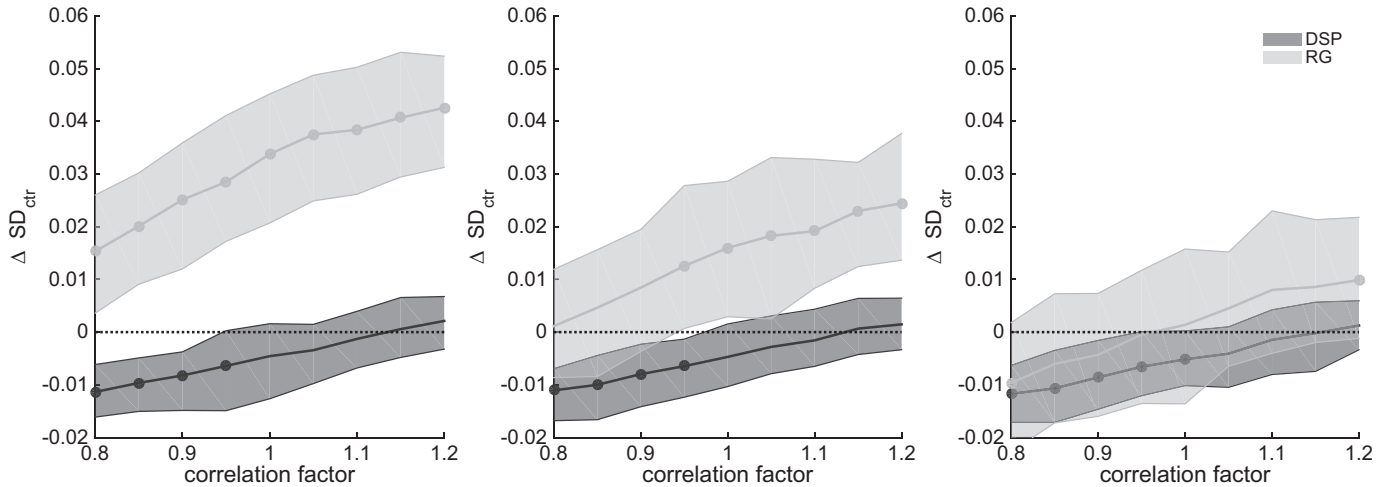

Fig. 8. Effect of the modulation of membrane potential correlations with stimulus contrast. ΔSD_{ctr} at different levels of scaling of membrane potential correlations assumed. A correlation factor of 1.1 means 10% larger correlations at high contrast than at low contrast. Different panels explore the relative scale of mean response at high- vs. low-contrast stimulation (mean scaling is 2.5, 2, and 1.5 for the *left*, *middle*, and *right*, respectively). Shaded areas indicate 95% CIs around the mean ΔSD_{ctr}. Beads at the mean values indicate P values lower than 0.05 for testing the hypothesis that ΔSD_{ctr} deviates from 0. Relative changes in the magnitude of correlations systematically change ΔSD_{ctr}. Moderate changes in membrane potential correlation do not cause qualitative changes in ΔSD_{ctr}, but larger correlation changes can even invert its sign. Reducing the effect of contrast on the mean response has strong effect on the significance of ΔSD_{ctr}.

deviations of 0.05 and 0.05, $P = 0.31$; Fig. 9, *A* and *E*). In the RG population, the ΔSD_{ori} was positive (standard deviations of 0.09 and 0.07, $P = 0.120$; Fig. 9, *B* and *E*). Similar to the simulations and to the experimental data of contrast dependence of responses, changes in the means of correlation distributions were small (in the experimental population, −0.001 and 0.001 in the preferred and nonpreferred condition, respectively, one-sample *t*-test $t(4) = -0.20$, $P = 0.85$; in the DSP population, −0.001 and 0.015, $P = 0.50$; in the RG population, 0.009 and 0.007, $P = 0.39$; Fig. 9*D*). Varying the membrane potential mean, variance, and correlation values in the DSP model revealed that such tuning is incapable to produce the significantly positive ΔSD_{ori} observed in the experiment (Fig. 9, *F* and *G*). While experimental results do not provide a direct dissociation between the two models in terms of predictions related to the dependence of correlations on orientation preference, they are reproduced without any particular tuning of model parameters in the case of the RG model, but the DSP model can only account for the patterns in experimental data with specific parameter tuning.

DISCUSSION

We analyzed a pair of widely used models of spiking to investigate their power to predict stimulus-dependent changes not only in single-cell statistics but in joint statistics of activity too. The DSP model (Gur et al. 1997) assumes stochastic spiking to account for linear scaling of spike count variance with spike count mean and relies on a separate stochastic process to model the covariance structure of spiking responses. The RG model (Dorn and Ringach 2003) assumes a single source of stochasticity at the level of membrane potentials but relies on a quasi-deterministic process of spike generation. We have demonstrated that, while in terms of single-cell statistical measures the models make similar predictions for stimulus change-related modulations in response statistics, predictions on pairwise correlations are distinct. In a model of a pair of neurons, simulated changes in stimulus orientation and stimulus contrast revealed opposing changes in spike count correlations. The key intuition behind this finding is that the level of spike count correlation with a given level of membrane potential correlation is determined by two phenomena: *1*) higher mean membrane potential implies higher magnitude spike count correlations (de la Rocha et al. 2007) and *2*) higher firing rate implies lower level of spike count correlations when private variability scales with spike count mean. The interaction of these two processes result in opposing changes in the correlation structure of the population. To assess which of the two models is compatible with neural recordings, response statistics of a population of neurons was simulated in both models and contrasted with response statistics of neurons recorded in V1 of awake monkeys. We have shown that the predictions of the RG model are in line with electrophysiological data while the DSP model is not capable of reproducing the patterns of noise correlations in V1 neurons. Our analyses highlight that joint responses can inform us on the possible building blocks of population models and also highlight that the width of the correlation distribution carries information beyond the mean of the distribution. These analyses provide important constraints on the models that can be used effectively to characterize computations in neural populations.

In our analysis we adopted an approach where we focused on matching the single-cell spike count statistics of the models without particular emphasis on the interpretation of the actual levels of membrane potentials. This approach is motivated by the aim to directly contrast changes in pairwise spiking statistics between the two models. It is obvious that tuning the membrane potential parameters of the two models separately could result in membrane potential values that are easier to interpret physiologically. Nonetheless, the assumptions of the DSP model necessitate choices that are hard to reconcile with neuronal data. For instance, in the analysis of the pair of neurons we used a membrane potential correlation of 0.95 in the DSP model to obtain a spike count correlation level matching that of the RG model (~0.1) when firing rate was ~35 Hz. In the population model of spiking statistics less extreme

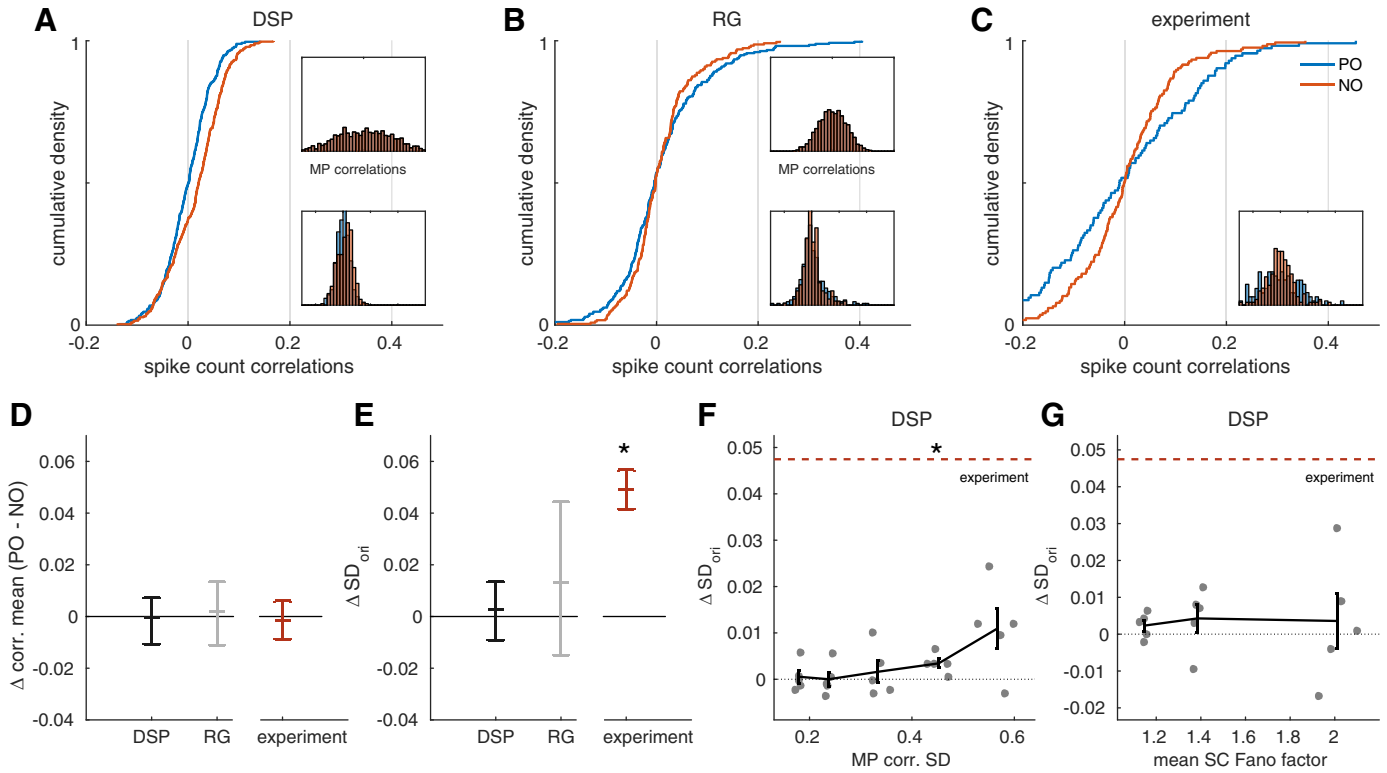

Fig. 9. Comparison of model predictions with experimental data: effect of stimulus orientation. *A*: cumulative distributions (CDF) of pairwise spike count correlations in a simulated population of 50 DSP neurons from pairs observing stimuli of preferred orientation (PO) and nonpreferred orientation (NO), using membrane potentials (MP) from tuning curves distributed uniformly over orientations. *Insets* show the histogram of membrane potential correlations and the spike count correlation histograms corresponding to the CDFs. *B*: distributions of pairwise spike count correlations in a simulated population of 50 RG neurons in response to the same stimuli. *C*: distributions of measured spike count correlations in response to grating stimuli of different orientations (data from Ecker et al. 2010). *D*: differences in the means of the 2 correlation distributions in the 2 simulated and the measured populations. Asterisks represent significant difference of the mean from 0 at the 0.05 level. *E*: ΔSD_{ori}. Mean of ΔSD_{ori} are shown for the 2 models across different populations and for experimental recordings across different sets of left-out neurons. The mean of ΔSD_{ori} is significantly different from 0 in the experimental and the RG populations, but not in the DSP population. *F*: dependence of ΔSD_{ori} in the DSP model on the width of MP correlation distribution. *G*: ΔSD_{ori} in the DSP model at different values of MP mean and variance, shown as a function of the spike count Fano factor, testing whether low mean and high variance in MP can reproduce experimentally measured orientation effects in spike count correlations.

values were used; still, the tendency of the DSP model to wash out correlations necessitated a wider correlation distribution for the DSP than for the RG model.

To make use of a limited amount of data without detailed characterization of receptive fields we adapted an approach for modeling population activity that required the fitting of population distributions of single-cell activity statistics to measurements, and not to individual cells. Using data sets containing a high number of stimulus repetitions and additional receptive field characterizations for recorded units would enable the fitting of response statistics of model units to measurement and a direct comparison of the models under the specific recording conditions used in the experiments.

We chose to use independent temporal evolution for membrane potentials, which meant independent membrane potential samples in fixed, short time intervals. This was motivated by the relatively fast decay of membrane potential correlations (Azouz and Gray 1999). This treatment explicitly incorporates within-trial variability to model the shared variability of neurons. This is in contrast with other models of response variability (Ecker et al. 2016), which consider processes that produce trial-to-trial variability since changes in attentional modulation occur on a slower timescale. We argued in the paper that spiking statistics provide constraints on the spiking models; therefore we believe that patterns in the auto- and cross-correlation functions (Smith and Kohn 2008) provide further constraints on neural models of spiking. Estimation of the amount of variability that is not caused by stimulus variation is contaminated by uncontrolled variability introduced by eye movements (Gur et al. 1997; Gur and Snodderly 2006). Eye movements affect both Fano factor and correlation estimations (McFarland et al. 2016) and their contribution cannot be ruled out in the case of the data analyzed here. Our critical conclusions, however, concerned changes in correlations rather than the magnitude of correlations. These measurements would only be affected by eye movements if the statistics of eye movements changed across the tested conditions. While changes in eye movement statistics can in principle be a consequence of altered contrast, it is not relevant when orientations are considered. Our analysis revealed qualitatively similar changes in correlations both when changes in orientation and contrast were considered. Therefore we expect our results to hold even if eye movement-related modulations of variability were considered (McFarland et al. 2016).

Earlier investigations have revealed that identifying the spiking model based on spiking statistics of neurons can be challenging and can even be prohibitive under a wide range of circumstances (Amarasingham et al. 2015). In our study we

capitalized on the stimulus dependence of pairwise response statistics to distinguish between the models. In their paper, Amarasingham et al. (2015) argued that the covariance of a pair of neurons is not sufficient to uniquely determine the spiking model without extra assumptions being made on the spiking covariance, which was formulated as a sum of firing rate covariability and spiking covariability. We argued that spiking is a process characteristic of an individual neuron, and therefore the variability associated with this process is private to the neuron. As a consequence, we assumed a spiking covariance that is diagonal. Indeed, the assumption that spiking variability is independent across neurons is the primary reason for the differential effect of both stimulus contrast and orientation on spike count correlations.

Recent studies using approaches close to our DSP model (Ecker et al. 2016; Rabinowitz et al. 2015) discuss top-down modulation of response correlations in contrast with the bottom-up modulations discussed here. The important contribution of these models is that different forms of correlations in spiking responses are explained based on simple computational principles. The differences between those approaches and ours are important and can be instructive for future work. First, correlations are introduced at the level of firing rates instead of membrane potentials. Since the variability is introduced before spike generation, the conclusions of those studies can be easily translated to our approach solely by assessing the effect of the firing rate nonlinearity. Second, these studies used a single scalar (or a low dimensional) stochastic variable to model gain modulation of visual cortical neurons. The collective gain modulation implies a firing rate correlation of 1 between those neurons that share the modulatory signal. In our case correlated activity was introduced by a multivariate stochastic process that implements a softer coupling between neurons. Nevertheless, the analysis of the DSP model reveals that realistic firing rates, Fano factors, and spike count correlations require a surprisingly high level of coordination between neurons. Third, the use of Poisson stochasticity in these models implies that response variability explained by the model is constrained to that beyond Poisson variability. Our analysis shows that the RG model can be a more effective model for variability than the DSP. Since the RG model only assumes a single source of variability at the level of membrane potentials, it provides an opportunity to account for a larger proportion of variability and therefore provides an opportunity for models with better predictive power. As demonstrated here and in earlier studies (Carandini and Ferster 2000), single-cell statistics is not distinctive for the two models at moderate firing rates, Fano factors only start to deviate from a constant value at high firing rates. Our results on distinctive patterns in pairwise correlations highlight that once these correlations become relevant for the analysis, special care is required to assess whether different predictions on correlations can affect conclusions. This is especially true for population-level analyses of neuronal responses where assessment of the joint statistics is quintessential for understanding processes underlying the collective behavior of neurons. The magnitude of the demonstrated effect can be different under different experimental conditions. In conditions where the stimulus drive is large strong signal correlations are induced. Therefore the relative contributions of noise correlations investigated here are expected to be smaller. In conditions where stimulus-driven changes are absent, however, noise correlations have larger effect. The RG and DSP models represent two extremes of a continuous scale of spiking models: while RG assumes a deterministic spike generation process, DSP relies on highly variable spike generation. Our aim was to demonstrate qualitative differences of the predictions of the two approaches. Models interpolating between the extremes, including the negative binomial model of spike generation (Lin et al. 2015), can be used to quantitatively match physiological data.

Poisson-like firing has been used extensively in the literature (Froudarakis et al. 2014; Goris et al. 2015; Jazayeri and Movshon 2006; Ma et al. 2006; Ma and Jazayeri 2014; Pillow 2007; Simoncelli et al. 2004). Besides its capability to provide a parsimonious explanation of the relationship between spiking intensity and spiking variability, Poisson neurons have much theoretical appeal too. First, Poisson-like spiking distribution ensures that fitting network parameters is a convex optimization problem, or, in other words, there is a single (global) maximum in optimization (Paninski 2004). Second, in theories of encoding information via populations of neurons, a Poisson-like likelihood function provides a representation in which the log likelihood contains linear terms, which enables simple, neurally plausible computations (Jazayeri and Movshon 2006; Ma 2006). In contrast with models based on the Poisson assumption, alternative approaches have applied models that have a closer relationship to the RG model (Brette and Gerstner 2005; Lin et al. 2015; Paninski 2004) and argued for the capability of such models to predict both for single-cell response properties (Brette and Gerstner 2005) and population statistics (Lin et al. 2015). Response variability in these two model classes is approached in two markedly different ways: in one, stochasticity is part of the spike generation process; in the other it originates before spike generation and can be related to membrane potential-level processes. In this context, our study contributes to the field by a direct and controlled comparison of the predictions of these approaches on spiking statistics and will help to identify the conditions under which they can effectively be used.

Recent advances in modeling data recorded from a large number of neurons have helped to assess neural responses in a trial-by-trial manner and to relate them to variances in behavior (Churchland et al. 2010; Yu et al. 2009), disentangle mixed sensitivities (Kobak et al. 2016), eliminate noise by tracing neural variability back to changes in latent factors (Machens et al. 2010), and implement closed-loop brain-computer interfacing (Sadtler et al. 2014). Modeling and predicting correlations is a critical factor in population-level analyses of neuronal data (Cunningham and Yu 2014). Those are precisely the correlated changes in neuronal activity that help to eliminate the effect of uncontrolled variables, to reduce apparent noise in the measurements, and to predict the activity of missing neurons. Furthermore, recent advances in functional interpretation of population response statistics have pointed out that correlations can be the hallmarks of computations taking place in the brain that reflect adaptation (Orbán et al. 2016), attention (Haefner et al. 2016), or task learning (Singh et al. 2016). Therefore, it is crucial to have an adequate model of response variability that can predict the effects of changes in stimulus on the correlation structure. The stochasticity assumed to underlie observed spiking variability can take on the form of Poisson variability (Archer et al. 2014; Macke et al. 2011) or Gaussian noise (Yu

et al. 2009; Sadtler et al. 2014). Our study aims to provide constraints on the forms of stochasticity in these models by emphasizing that bottom-up driven changes in the response statistics can differentiate between alternative models. As demonstrated in the paper, even when single-cell response statistics have limited power to distinguish between alternative models, joint statistics can reveal properties that are incompatible with the predictions of one or the other. We expect that proper understanding and characterization of stochasticity in the nervous system helps to better interpret joint statistics and especially correlations present in the activity of neural populations.

ACKNOWLEDGMENTS

We thank M. Lengyel for useful discussions, D. G. Nagy for comments on the manuscript, and especially A. Ecker, P. Berens, M. Bethge, and A. Tolias for making their data publicly available.

GRANTS

This work was supported by a Lendület Award of the Hungarian Academy of Sciences (G. Orbán, M. Bányai) and an award from the National Brain Research Program of Hungary (NAP-B, KTIA_NAP_12-2-201). The authors declare no competing financial interests.

DISCLOSURES

No conflicts of interest, financial or otherwise, are declared by the authors.

AUTHOR CONTRIBUTIONS

M.B. and G.O. conceived and designed research; M.B., Z.K., and G.O. analyzed data; M.B. and G.O. interpreted results of experiments; M.B., Z.K., and G.O. prepared figures; M.B. drafted manuscript; M.B. and G.O. edited and revised manuscript; G.O. approved final version of manuscript.

REFERENCES

Adrian ED. The impulses produced by sensory nerve endings: part I. *J Physiol* 61: 49–72, 1926. doi:10.1113/jphysiol.1926.sp002273.

Ahmed B, Anderson JC, Douglas RJ, Martin KA, Whitteridge D. Estimates of the net excitatory currents evoked by visual stimulation of identified neurons in cat visual cortex. *Cereb Cortex* 8: 462–476, 1998. doi:10.1093/cercor/8.5.462.

Amarasingham A, Geman S, Harrison MT. Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proc Natl Acad Sci USA* 112: 6455–6460, 2015. doi:10.1073/pnas.1506400112.

Anderson JS, Lampl I, Gillespie DC, Ferster D. The contribution of noise to contrast invariance of orientation tuning in cat visual cortex. *Science* 290: 1968–1972, 2000. doi:10.1126/science.290.5498.1968.

Archer EW, Koster U, Pillow JW, Macke JH. Low-dimensional models of neural population activity in sensory cortical circuits. *Adv Neural Inf Process Syst* 27: 343–351, 2014.

Azouz R, Gray CM. Cellular mechanisms contributing to response variability of cortical neurons in vivo. *J Neurosci* 19: 2209–2223, 1999.

Berkes P, Orbán G, Lengyel M, Fiser J. Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science* 331: 83–87, 2011. doi:10.1126/science.1195870.

Brette R, Gerstner W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J Neurophysiol* 94: 3637–3642, 2005. doi:10.1152/jn.00686.2005.

Britten KH, Shadlen MN, Newsome WT, Movshon JA. Responses of neurons in macaque MT to stochastic motion signals. *Vis Neurosci* 10: 1157–1169, 1993. doi:10.1017/S0952523800010269.

Carandini M. Amplification of trial-to-trial response variability by neurons in visual cortex. *PLoS Biol* 2: e264, 2004. doi:10.1371/journal.pbio.0020264.

Carandini M, Ferster D. Membrane potential and firing rate in cat primary visual cortex. *J Neurosci* 20: 470–484, 2000.

Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen MN. Variance as a signature of neural computations during decision making. *Neuron* 69: 818–831, 2011. doi:10.1016/j.neuron.2010.12.037.

Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T, Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani M, Shenoy KV. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13: 369–378, 2010. doi:10.1038/nn.2501.

Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nat Neurosci* 14: 811–819, 2011. doi:10.1038/nn.2842.

Cunningham JP, Yu BM. Dimensionality reduction for large-scale neural recordings. *Nat Neurosci* 17: 1500–1509, 2014. doi:10.1038/nn.3776.

de la Rocha J, Doiron B, Shea-Brown E, Josić K, Reyes A. Correlation between neural spike trains increases with firing rate. *Nature* 448: 802–806, 2007. doi:10.1038/nature06028.

Dorn JD, Ringach DL. Estimating membrane voltage correlations from extracellular spike trains. *J Neurophysiol* 89: 2271–2278, 2003. doi:10.1152/jn.000889.2002.

Ecker AS, Berens P, Cotton RJ, Subramaniyan M, Denfield GH, Cadwell CR, Smirnakis SM, Bethge M, Tolias AS. State dependence of noise correlations in macaque primary visual cortex. *Neuron* 82: 235–248, 2014. doi:10.1016/j.neuron.2014.02.006.

Ecker AS, Berens P, Keliris GA, Bethge M, Logothetis NK, Tolias AS. Decorrelated neuronal firing in cortical microcircuits. *Science* 327: 584–587, 2010. doi:10.1126/science.1179867.

Ecker AS, Denfield GH, Bethge M, Tolias AS. On the structure of neuronal population activity under fluctuations in attentional state. *J Neurosci* 36: 1775–1789, 2016. doi:10.1523/JNEUROSCI.2044-15.2016.

Faisal AA, Selen LPJ, Wolpert DM. Noise in the nervous system. *Nat Rev Neurosci* 9: 292–303, 2008. doi:10.1038/nrn2258.

Finn IM, Priebe NJ, Ferster D. The emergence of contrast-invariant orientation tuning in simple cells of cat visual cortex. *Neuron* 54: 137–152, 2007. doi:10.1016/j.neuron.2007.02.029.

Fiser J, Chiu C, Weliky M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* 431: 573–578, 2004. doi:10.1038/nature02907.

Froudarakis E, Berens P, Ecker AS, Cotton RJ, Sinz FH, Yatsenko D, Saggau P, Bethge M, Tolias AS. Population code in mouse V1 facilitates readout of natural scenes through increased sparseness. *Nat Neurosci* 17: 851–857, 2014. doi:10.1038/nn.3707.

Goris RLT, Simoncelli EP, Movshon JA. Origin and function of tuning diversity in macaque visual cortex. *Neuron* 88: 819–831, 2015. doi:10.1016/j.neuron.2015.10.009.

Gur M, Beylin A, Snodderly DM. Response variability of neurons in primary visual cortex (V1) of alert monkeys. *J Neurosci* 17: 2914–2920, 1997.

Gur M, Snodderly DM. High response reliability of neurons in primary visual cortex (V1) of alert, trained monkeys. *Cereb Cortex* 16: 888–895, 2006. doi:10.1093/cercor/bhj032.

Haefner RM, Berkes P, Fiser J. Perceptual decision-making as probabilistic inference by neural sampling. *Neuron* 90: 649–660, 2016. doi:10.1016/j.neuron.2016.03.020.

Haefner RM, Gerwinn S, Macke JH, Bethge M. Inferring decoding strategies from choice probabilities in the presence of correlated variability. *Nat Neurosci* 16: 235–242, 2013. doi:10.1038/nn.3309.

Haider B, Häusser M, Carandini M. Inhibition dominates sensory responses in the awake cortex. *Nature* 493: 97–100, 2013. doi:10.1038/nature11665.

Hennequin, G. and Lengyel, M. Characterizing variability in nonlinear recurrent neuronal networks (Preprint). arXiv:1610.03110 [q-bio.NC], 2016.

Hodgkin AL, Huxley AF. The components of membrane conductance in the giant axon of Loligo. *J Physiol* 116: 473–496, 1952. doi:10.1113/jphysiol.1952.sp004718.

Jazayeri M, Movshon JA. Optimal representation of sensory information by neural populations. *Nat Neurosci* 9: 690–696, 2006. doi:10.1038/nn1691.

Kenney JF, Keeping ES. Mathematics of Statistics: Part Two (2nd ed.). Princeton, NJ: Van Nostrand, 1951.

Kobak D, Brendel W, Constantinidis C, Feierstein CE, Kepecs A, Mainen ZF, Qi XL, Romo R, Uchida N, Machens CK. Demixed principal component analysis of neural population data. *eLife* 5: e10989, 2016. doi:10.7554/eLife.10989.

Lewandowski D, Kurowicka D, Joe H. Generating random correlation matrices based on vines and extended onion method. *J Multivariate Anal* 100: 1989–2001, 2009. doi:10.1016/j.jmva.2009.04.008.

Lin IC, Okun M, Carandini M, Harris KD. The nature of shared cortical variability. *Neuron* 87: 644–656, 2015. doi:10.1016/j.neuron.2015.06.035.

Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nat Neurosci* 9: 1432–1438, 2006. doi:10.1038/nn1790.

Ma WJ, Jazayeri M. Neural coding of uncertainty and probability. *Annu Rev Neurosci* 37: 205–220, 2014. doi:10.1146/annurev-neuro-071013-014017.

Machens CK, Romo R, Brody CD. Functional, but not anatomical, separation of "what" and "when" in prefrontal cortex. *J Neurosci* 30: 350–360, 2010. doi:10.1523/JNEUROSCI.3276-09.2010.

Macke JH, Buesing L, Cunningham JP, Yu BM, Shenoy KV, Sahani M. Empirical models of spiking in neural populations. *Adv Neural Inf Process Syst* 24: 1350–1358, 2011.

Mainen ZF, Sejnowski TJ. Reliability of spike timing in neocortical neurons. *Science* 268: 1503–1506, 1995. doi:10.1126/science.7770778.

McFarland JM, Cumming BG, Butts DA. Variability and correlations in primary visual cortical neurons driven by fixational eye movements. *J Neurosci* 36: 6225–6241, 2016. doi:10.1523/JNEUROSCI.4660-15.2016.

Moreno-Bote R, Renart A, Parga N. Theory of input spike auto- and cross-correlations and their effect on the response of spiking neurons. *Neural Comput* 20: 1651–1705, 2008. doi:10.1162/neco.2008.03-07-497.

Orbán G, Berkes P, Fiser J, Lengyel M. Neural variability and sampling-based probabilistic representations in the visual cortex. *Neuron* 92: 530–543, 2016. doi:10.1016/j.neuron.2016.09.038.

Paninski L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network* 15: 243–262, 2004. doi:10.1088/0954-898X_15_4_002.

Pillow JW. Likelihood-based approaches to modeling the neural code. In: *Bayesian Brain: Probabilistic Approaches to Neural Coding*, edited by Doya K, Ishii S, Pouget A, and Rao RPN. Cambridge, MA: MIT Press, 2007, p. 53–70.

Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP. Attention stabilizes the shared gain of V4 populations. *eLife* 4: e08998, 2015. doi:10.7554/eLife.08998.

Renart A, Machens CK. Variability in neural activity and behavior. *Curr Opin Neurobiol* 25: 211–220, 2014. doi:10.1016/j.conb.2014.02.013.

Sadtler PT, Quick KM, Golub MD, Chase SM, Ryu SI, Tyler-Kabara EC, Yu BM, Batista AP. Neural constraints on learning. *Nature* 512: 423–426, 2014. doi:10.1038/nature13665.

Schneidman E, Berry MJ II, Segev R, Bialek W. Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440: 1007–1012, 2006. doi:10.1038/nature04701.

Simoncelli EP, Paninski L, Pillow J, Schwartz O. Characterization of neural responses with stochastic stimuli. In: *The Cognitive Neurosciences*, edited by Gazzaniga M. Cambridge, MA: MIT Press, 2004, p. 327–338.

Singh A, Peyrache A, Humphries M. Task learning reveals signatures of sample-based internal models in rodent prefrontal cortex (Preprint). *bioRxiv*: 027102, 2016.

Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *J Neurosci* 28: 12591–12603, 2008. doi:10.1523/JNEUROSCI.2929-08.2008.

Softky WR, Koch C. The highly irregular firing of cortical cells is inconsistent with temporal integration of random EPSPs. *J Neurosci* 13: 334–350, 1993.

Tan AY, Chen Y, Scholl B, Seidemann E, Priebe NJ. Sensory stimulation shifts visual cortex from synchronous to asynchronous states. *Nature* 509: 226–229, 2014. doi:10.1038/nature13159.

Tkačik G, Prentice JS, Balasubramanian V, Schneidman E. Optimal population coding by noisy spiking neurons. *Proc Natl Acad Sci USA* 107: 14419–14424, 2010. doi:10.1073/pnas.1004906107.

Tolhurst DJ, Movshon JA, Dean AF. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision Res* 23: 775–785, 1983. doi:10.1016/0042-6989(83)90200-6.

Tolhurst DJ, Movshon JA, Thompson ID. The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Exp Brain Res* 41: 414–419, 1981.

Tomko GJ, Crapper DR. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain Res* 79: 405–418, 1974. doi:10.1016/0006-8993(74)90438-7.

Yu BM, Cunningham JP, Santhanam G, Ryu SI, Shenoy KV, Sahani M. Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *J Neurophysiol* 102: 614–635, 2009. doi:10.1152/jn.90941.2008.

Yu J, Ferster D. Membrane potential synchrony in primary visual cortex during sensory stimulation. *Neuron* 68: 1187–1201, 2010. doi:10.1016/j.neuron.2010.11.027.

Yu S, Yang H, Nakahara H, Santos GS, Nikolić D, Plenz D. Higher-order interactions characterized in cortical activity. *J Neurosci* 31: 17514–17526, 2011. doi:10.1523/JNEUROSCI.3127-11.2011.