Published in final edited form as:

Neural Comput. 2016 May ; 28(5): 849–881. doi:10.1162/NECO_a_00831.

Separating Spike Count Correlation from Firing Rate Correlation

Giuseppe Vinci¹, Valérie Ventura^{1,4}, Matthew A. Smith^{3,4}, and Robert E. Kass^{1,2,4}

¹Department of Statistics, Carnegie Mellon University, Pittsburgh, PA, USA

²Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA, USA

³Department of Ophthalmology, University of Pittsburgh, Pittsburgh, PA, USA

⁴Center for the Neural Basis of Cognition, Pittsburgh, PA, USA

Abstract

Populations of cortical neurons exhibit shared fluctuations in spiking activity over time. When measured for a pair of neurons over multiple repetitions of an identical stimulus, this phenomenon emerges as correlated trial-to-trial response variability via spike count correlation (SCC). However, spike counts can be viewed as noisy versions of firing rates, which can vary from trial to trial. From this perspective, the SCC for a pair of neurons becomes a noisy version of the corresponding firing-rate correlation (FRC). Furthermore, the magnitude of the SCC is generally smaller than that of the FRC, and is likely to be less sensitive to experimental manipulation. We provide statistical methods for disambiguating time-averaged drive from within-trial noise, thereby separating FRC from SCC. We study these methods to document their reliability, and we apply them to neurons recorded in vivo from area V4, in an alert animal. We show how the various effects we describe are reflected in the data: within-trial effects are largely negligible, while attenuation due to trial-to-trial variation dominates, and frequently produces comparisons in SCC that, because of noise, do not accurately reflect those based on the underlying FRC.

Keywords

attenuation effect; overdispersed spike counts; generalized tuning curve correlation; synchrony; hierarchical model; Poisson-lognormal distribution; macaque V4

1 Introduction

Analysis of spike count correlation has provided important neurophysiological insights, as summarized in Cohen and Kohn (2011), Gu et al. (2011), Harris et al. (2011), Jeanne et al. (2013), Mitchell et al. (2009), Smith and Kohn (2008), Smith and Sommer (2013), and references therein. However, as documented in these references, a striking feature of the interaction among neurons is that it can occur at multiple time scales. A neuron's spike count collected across a time interval of length T , such as $T = 1000$ ms, reflects both the time-averaged drive to that neuron and additional stochastic fluctuations, which produce irregular spiking at time scales smaller than T , usually described as point process variation (Shadlen and Newsome, 1998; Churchland et al., 2011). That is, the spike count involves both (a) a

relatively slowly varying input drive, at the time scale T , and (b) faster point-process fluctuations that determine when the neuron fires; across an interval of length T this point process variation determines the number of times each neuron fires for a given level of time-averaged input drive. Correspondingly, for a given pair of neurons, the *spike count correlation* (SCC) may reflect (A) the *firing-rate correlation* (FRC), i.e. the correlation in the time-averaged drive to these neurons, and (B) point process noise (Kass and Ventura, 2006; Staude et al., 2008; Goris et al., 2014). It also could involve (C) more precisely-timed within-trial correlation effects that might arise, for instance, from (lagged) synchronous spiking (Kelly and Kass, 2012; Harrison et al., 2013), which have been observed in some cortical areas (macaque V4, Smith and Kohn (2008); Smith and Sommer (2013)), and may be modulated by state transitions that are more prominent in some anesthetized preparations (Kelly et al., 2010; Ecker et al., 2014). Longer time scale effects, but still shorter than trial length, such as correlated up and down states (Steriade and Buzsaki, 1990; Cowan and Wilson, 1994; Timofeev et al., 2001; Steriade et al., 2001) can also produce within-trial spike count covariability.

In this paper, we describe ways in which the SCC may not accurately measure the FRC in (A) due to distortion caused by (B), and we present statistical methods for disambiguating (A) from (B). Other studies have developed related techniques: Churchland et al. (2011); Goris et al. (2014); Staude et al. (2008). Here we go further, using both parametric and nonparametric approaches, in order to provide a comprehensive set of tools which, as we show, perform well under conditions similar to those found with neural data. When (C) is negligible, statistical inference and interpretation are simplified. Therefore, we focus especially on methods for evaluating the relevance of (C) to the analysis of spike counts, including both statistical tests and estimates of magnitude. We then use the tools we have developed to analyze data recorded from area V4 while a monkey performed a fixation task, and we demonstrate that, in this setting, FRC behaves differently than SCC.

We begin by assuming that the data are observed across multiple trials of total elapsed time T , and that, for trials $r = 1, \dots, n$, each pair of observed spike counts (Y_{1r}, Y_{2r}) comes from the doubly stochastic correlated point process depicted in Figure 1. This process consists of two levels. The first level consists of a pair of latent (unobserved) variables (X_{1r}, X_{2r}) generated with some *across-trial* correlation (trial-to-trial correlation), where, on trial r , X_{ir} is the average firing rate FR_{ir} for neuron i multiplied by T , $X_{ir} = FR_{ir}T$. The variable X_{ir} may be interpreted as an average input drive to neuron i . The second level consists of observed spike counts (Y_{1r}, Y_{2r}) , generated jointly with expectation (X_{1r}, X_{2r}) . In addition, there may be some *within-trial* point-process correlation. Thus the conditional expectation of Y_{ir} , conditionally on (X_{1r}, X_{2r}) , is

$$E(Y_{ir} | X_{1r}, X_{2r}) = X_{ir} \quad (1)$$

We further assume that the conditional variance is

$$Var(Y_{ir} | X_{1r}, X_{2r}) = \phi_i X_{ir} \quad (2)$$

where $\phi_i \in \mathbb{R}^+$ controls the dispersion of Y_{ir} conditionally on X_{ir} (Kass and Ventura, 2006; Churchland et al., 2011; Goris et al., 2014). The case $\phi_i = 1$ holds when Y_{ir} follows a Poisson distribution, conditionally on X_{ir} .

A fundamental decomposition of variability is given by

$$Cov(Y_{1r}, Y_{2r}) = E[Cov(Y_{1r}, Y_{2r} | X_{1r}, X_{2r})] + Cov(E[Y_{1r} | X_{1r}, X_{2r}], E[Y_{2r} | X_{1r}, X_{2r}]). \quad (3)$$

The observed trial-to-trial *spike count correlation* (SCC) is the Pearson correlation of the counts Y_{1r} and Y_{2r} across trials, that is the normalized version of the left-hand side of eq.(3), while the unobserved trial-to-trial *firing-rate correlation* (FRC) is the Pearson correlation of FR_{1r} and FR_{2r} , or equivalently of X_{1r} and X_{2r} , across trials, that is the normalized version of the second term on the right-hand side. In Section 2 we show that, on a sequence of independent trials, the SCC can be seen as a perturbed version of FRC. We characterize this phenomenon by the simple formula

$$SCC = FRC \times ATT + \Gamma, \quad (4)$$

where we identify an *attenuation* component (ATT), mainly due to trial-to-trial variability, and a *deviation* component (Γ), which summarizes the impact of within-trial point process correlation (component C in our first paragraph, above) on SCC.

It is not possible to estimate the components in (4) separately, based on spike counts alone, without additional modeling assumptions (Amarasingham et al., 2015). In conjunction with equation (3), Churchland et al. (2011) imposed constraints on the variance of the spike counts, constraints estimated from the data, to extract method of moments estimates of correlations between single-neuron firing rates at different epochs of a trial. By assuming univariate Poisson-gamma models for single neuron spike counts, and then also assuming a sampling distribution for the values of SCC across stimuli, Goris et al. (2014) obtained estimates for FRC and for a quantity related to Γ . Because they used univariate models for the spike counts, they were only able to obtain a single estimate of FRC for all stimuli. However, as illustrated in Figure 3A, FRC of a neuron pair can (and often does) vary across stimuli. The approach we take allows for this. We build on previous work of Ventura et al. (2005) and Kass and Ventura (2006), who proved that spike count correlation necessarily increases with T in the presence of multiplicative trial-to-trial variation in time-averaged input drive, and Behseta et al. (2009), who used a bivariate hierarchical model as in Figure 1 to provide a correction for attenuation of correlation in the presence of point process noise. In Section 3 we derive nonparametric estimators of the components of (4) from spike trains (as opposed to the counts alone), by showing that certain estimators of intermediate quantities are unbiased; we also derive a nonparametric estimator of ϕ_i from spike trains, in contrast to Goris et al. (2014), who assume $\phi_i = 1$, and Churchland et al. (2011), who use ϕ_i like a tuning parameter and choose it such that variances and correlations are well defined. We further provide details on bivariate hierarchical models for spike counts that take the trial-to-trial distribution of (X_{1r}, X_{2r}) to be lognormal and the distribution of Y_{ir} to be Poisson with mean X_{ir} , and we suggest bootstrap methods to assess the fit of these models. Because analysis and interpretation are more straightforward when $\Gamma = 0$, we propose a jitter

test of this hypothesis. We show that the test is sufficiently powerful to find substantial departures from $\Gamma = 0$ under realistic scenarios and, because one typically must examine large numbers of pairs of neurons, we use methods that control the false discovery rate (FDR) and also the false non-discovery rate (FNR). The latter assesses the likelihood of error when a decision is made to assume $\Gamma = 0$ based on the jitter test. We also introduce generalized forms of tuning curve and tuning curve correlation (GTC and GTCC, Appendix H) that account for the trial-to-trial variability of the firing rates. In Section 4 we analyze the V4 data and find that the models fit well, and that they support the effects discussed in Section 2. In Section 5 we discuss implications of our findings.

Figure 2 summarizes the main features of our approach, in which we begin with nonparametric estimators both for exploratory purposes (e.g., making informative plots) and to formulate the jitter test. We then move on to model-based inference, and finally assess goodness-of-fit and sensitivity.

2 Spike count correlation as an attenuated firing-rate correlation

From equations (1) and (2) we have (see Appendix A)

$$E[Y_{ir}] = E[X_{ir}] \quad \text{and} \quad \text{Var}(Y_{ir}) = \phi_i E[X_{ir}] + \text{Var}(X_{ir}), \quad (5)$$

which implies that the spike counts are overdispersed ($\text{Var}(Y_{ir}) > E[Y_{ir}]$) if $\phi_i > 1 - F_i$, where $F_i = \text{Var}(X_{ir})/E[X_{ir}]$ is the Fano factor of X_{ir} .

Combining equations (1), (3), and (5), $\text{SCC} = \text{Cor}(Y_{1r}, Y_{2r})$ can be decomposed as

$$\text{SCC} = \text{FRC} \times \text{ATT} + \Gamma, \quad (6)$$

(see Appendix A) where

- $\text{FRC} = \text{Cor}(E(Y_{1r} | X_{1r}), E(Y_{2r} | X_{2r})) = \text{Cor}(X_{1r}, X_{2r})$ is the *trial-to-trial firing rate correlation*,

- $\text{ATT} = \prod_{i=1,2} (1 + \phi_i / F_i)^{-1/2}$ is the *attenuation coefficient*, $\text{ATT} \in (0, 1]$, and

- $\Gamma = \gamma / \prod_{i=1,2} (\text{Var}(Y_{ir}))^{1/2}$, where $\gamma = E[\text{Cov}(Y_{1r}, Y_{2r} | X_{1r}, X_{2r})]$ is the *trial-averaged within-trial covariance* of the spike counts.

The *attenuation effect* is defined as the case when SCC and FRC have the same sign and SCC has smaller magnitude ($|\text{SCC}| < |\text{FRC}|$). Figure 3A shows an example neuron pair (macaque V4 area, Section 4) with different attenuation strengths across stimuli, and Figure 3B shows that SCC is an attenuated version of FRC for the majority of neuron pairs in our dataset. Because $\text{ATT} \in (0, 1]$, attenuation occurs whenever Γ is small, specifically whenever Γ is between $\text{FRC} \times (1 - \text{ATT})$ and $-\text{FRC} \times \text{ATT}$. The special case $\Gamma = 0$ is of particular interest because eq.(6) simplifies to $\text{SCC} = \text{FRC} \times \text{ATT}$, so that the attenuation effect is completely specified by ATT (and the results in Kass and Ventura (2006), which assume $\Gamma = 0$, hold). In Section 3.3 we develop a *jitter* method to test the hypothesis $\Gamma = 0$

and in Section 4 we show that this hypothesis is reasonable for the majority of neuron pairs in the in vivo macaque V4 dataset we analyze.

3 Estimating sources of correlations

Our aim here is to estimate the population quantities in eqs. (5) and (6) that characterize the hierarchical structure defined in Section 2 and Figure 1. In Section 3.1 we propose bivariate parametric models for the spike counts of neuron pairs, which imply parametric estimates for the components of eqs. (5) and (6); this models assume $\varphi_i = 1$, allowing for the overdispersed spike counts that are typical of the visual cortex V4 data we analyze in Section 4 (Tomko and Crapper, 1974; Tolhurst et al., 1981; Cohen and Maunsell, 1999; Mitchell et al., 2009; Goris et al., 2014). We also provide methods to evaluate the fit of the models to data and assess their sensitivity to the assumptions. In Section 3.2, we propose nonparametric estimators of the population quantities in eqs. (5) and (6) based on binned spike counts, which only make assumptions about the time scale of within-trial dependencies between neurons. Because the interpretation of eq.(6) is more straightforward when $\Gamma = 0$, in Section 3.3 we develop a model-free jitter test of this hypothesis.

3.1 Parametric estimation

Goris et al. (2014) used a univariate Poisson-Gamma model – $Y_{ir} | X_{ir} \sim \text{Poisson}(X_{ir})$ where $X_{ir} \sim \Gamma(\alpha_i, \beta_i)$ – to model single neuron spike counts, but because it has no natural bivariate extension, we extend instead the univariate Poisson-lognormal model of Kass and Ventura (2006):

$$\begin{aligned} Y_{ir} | X_{ir} &\sim \text{Poisson}(X_{ir}), \quad i=1, 2 \\ \begin{pmatrix} \log X_{1r} \\ \log X_{2r} \end{pmatrix} &\sim N\left(\mu = \begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \Sigma = \begin{pmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{pmatrix}\right). \end{aligned} \quad (7)$$

This bivariate Poisson-lognormal model (PLN) allows for overdispersed spike counts since $\varphi_i = 1$ (Section 2 and Appendix C) but assumes that $Y_{1r} | X_{1r}$ and $Y_{2r} | X_{2r}$ are independent, which implies $\gamma = \Gamma = 0$. We therefore generalize it to allow $\gamma \neq 0$. We assume $\gamma \neq 0$ and let

$$\begin{aligned} Y_{ir} &= R_{0r} + R_{ir} \\ R_{0r} &\sim \text{Poisson}(\gamma) \\ R_{ir} | W_{ir} &\sim \text{Poisson}(W_{ir}), \quad i=1, 2 \\ \begin{pmatrix} \log W_{1r} \\ \log W_{2r} \end{pmatrix} &\sim N(\mu, \Sigma), \end{aligned} \quad (8)$$

where R_{0r} and $R_{ir} | W_{ir}$, $i = 1, 2$, are mutually independent, and the firing rates are $X_{ir} = \gamma + W_{ir}$. We call this model a bivariate **Doubly Correlated Poisson-lognormal** model (DC-PLN) because the spike-counts (Y_{1r}, Y_{2r}) given the rates have a correlated Poisson distribution, and the rates (W_{1r}, W_{2r}) are also correlated. Figure 8 shows tolerance regions of the joint distribution of (X_{1r}, X_{2r}) fitted by maximum likelihood (ML; Appendix C) to a neuron pair in macaque area V4 (Section 4), which we also call a *generalized tuning curve* (GTC, Appendix H). Model (8) is appropriate when the spike trains of two neurons share common spikes R_{0r} in addition to, and independently from, (R_{1r}, R_{2r}) , for example if they

receive common input from an adjacent third neuron that causes positive conditional within-trial covariance γ . This model does not accommodate the case $\gamma < 0$, which is acceptable for the dataset analyzed in Section 4, because over 95% of neurons-conditions pairs have $\gamma \geq 0$. When $\gamma = 0$ (equivalently $\Gamma = 0$), model (8) reduces to the PLN model in eq.(7). Section 3.3 details a test of this hypothesis.

Under this model, the spike count correlation formula (6) reduces to:

$$SCC = \underbrace{(e^{\sigma_{12}^2} - 1) \prod_{i=1,2} (e^{\sigma_i^2} - 1)^{-\frac{1}{2}}}_{FRC} \times \underbrace{\prod_{i=1,2} \left(1 + a_i + \left(e^{\mu_i + \sigma_i^2/2} (e^{\sigma_i^2} - 1) \right)^{-1} \right)^{-\frac{1}{2}}}_{ATT} + \Gamma \quad (9)$$

where $a_i = \gamma \left(e^{\mu_i + \sigma_i^2/2} \right)^2 (e^{\sigma_i^2} - 1)^{-1}$ and

$\Gamma = \gamma \prod_{i=1,2} \left[\gamma + e^{\mu_i + \sigma_i^2/2} + \left(e^{\mu_i + \sigma_i^2/2} \right)^2 (e^{\sigma_i^2} - 1) \right]^{-\frac{1}{2}}$. Replacing the parameters by their ML estimates (see Appendix C) yields ML estimates of the quantities of interest: $\hat{\Gamma}_{mle}$, $\widehat{FRC}_{mle}$, $\widehat{ATT}_{mle}$ and $\widehat{SCC}_{mle}$.

Alternatively, the parameters of model (8) could be estimated via Bayesian methods (Bernardo and Smith (2009); Kass et al. (2014), Ch.16), which require choices of priors for μ , Σ , γ ; assuming squared error loss, FRC would be estimated as the posterior mean from the marginal posterior of Σ .

3.1.1 Assessing the fit of the parametric model—We assess the global goodness-of-fit (GOF) of a parametric model to data using a discrepancy measure between a nonparametric quantity and its parametric analog under the fitted model. A large discrepancy suggests that the model is not a good representation of the true spike count generation process. We consider two discrepancies: the Kullback-Leibler (KL) divergence of the fitted model (univariate or bivariate) from the empirical distribution function (EDF) of the data, and in the bivariate case, the divergence between the sample SCC and the SCC predicted by the fitted model.

- KULLBACK-LEIBLER DIVERGENCE. Let $Z_r = Y_{ir}$ be the spike count of a neuron in trial r in the univariate case, and $Z_r = (Y_{1r}, Y_{2r})$ in the bivariate case. Let $P_n(z) = n^{-1} \sum_{r=1}^n I(\|Z_r - z\|_1 = 0)$ be their EDF, where $\|x\|_1 = \sum_j |x_j|$ is the L_1 -norm, and $Q_{\hat{\theta}}$ be a parametric model fitted to $(Z_1, \dots, Z_n)$. The Kullback-Leibler (KL) divergence between P_n and $Q_{\hat{\theta}}$ is:

$$\begin{aligned} KL(P_n \parallel Q) &= \sum_z \log(P_n(z)/Q(z)) P_n(z) \\ &= n^{-1} \sum_{r=1}^n \log(P_n(Z_r)/Q(Z_r)) \\ &= n^{-1} \sum_{r=1}^n \log P_n(Z_r) - n^{-1} \sum_{r=1}^n \log Q(Z_r) \\ &\geq 0, \end{aligned}$$

where $\sum_{r=1}^n \log Q_{\hat{\theta}}(Z_r) = \max_{\theta} \log L_Q(\theta; Z_1, \dots, Z_n)$ is the maximum of the log-likelihood of Q_{θ} . Hence the KL divergence can be written as:

$$KL(P_n \parallel Q_{\hat{\theta}}) = n^{-1} (EML - PML)$$

where *EML* stands for “empirical maximum log-likelihood” and *PML* for “parametric maximum log-likelihood”. We use $KL(P_n \parallel Q_{\hat{\theta}})$ to test the hypothesis $H_0: Z_r \sim Q_{\theta}$ and approximate its null distribution, F_{KL} , via parametric bootstrap (Efron and Tibshirani, 1994). We reject H_0 if $KL(P_n \parallel Q_{\hat{\theta}}) > F_{KL}^{-1}(1 - \alpha)$ for some significance level α .

- **SCC DIVERGENCE.** Let Q_{θ} be a bivariate parametric model for (Y_{1r}, Y_{2r}) . Define

$$D = f_Z(\widehat{SCC}_{samp}) - f_Z(\widehat{SCC}_{mle}) \quad (10)$$

where $\widehat{SCC}_{samp}$ is the sample SCC of (Y_{1r}, Y_{2r}) (eq.(18)), $\widehat{SCC}_{mle}$ is the estimate derived from model Q_{θ} fitted to the data (eq.(9)), and $f_Z(r) = 1/2 \cdot \log((1+r)/(1-r))$ is the Fisher transformation. We reject the hypothesis $H_0: Z_r \sim Q_{\theta}$ if $D < F_D^{-1}(\alpha/2)$ or $D > F_D^{-1}(1 - \alpha/2)$, where F_D is the null distribution of D that we approximate using a parametric bootstrap.

3.1.2 Sensitivity of FRC estimates to model assumptions—Goodness of fit tests do not always have high power to detect deviations between the data and the assumed parametric model, so it is useful to assess how sensitive the parametric estimate of FRC is to mild deviations from the DC-PLN model. We do that by simulating spike counts from the more general model specified in eq.(11), estimate FRC assuming that the data are DC-PLN distributed, and quantify the deviations of the estimate from its true value.

The **generalized bivariate spike count model** is defined as:

$$\begin{aligned} Y_{ir} &= R_{0r} + R_{ir}, \quad i=1, 2 \\ R_{0r} &\sim P_{\gamma, \phi} \\ R_{ir} | W_{ir} &\sim P_{W_{ir}, \phi}, \quad i=1, 2 \\ (W_{1r}, W_{2r})^T &\sim f_{b, \mu, \Sigma}, \end{aligned} \quad (11)$$

which has the same structure as model (8) but with the Poisson and lognormal distributions replaced by $P_{\lambda, \phi}$ and $f_{b, \mu, \Sigma}$. It includes as special case the DC-PLN model (8) when $\phi = 1$, $b = 0^+$, $\gamma = 0$, and the PLN model (7) when $\phi = 1$, $b = 0^+$, and $\gamma = 0$. The distributions $P_{\lambda, \phi}$ and $f_{b, \mu, \Sigma}$ are generalizations of the Poisson and lognormal distributions that allow different spread and skewness. They are defined as follows:

1. Under model (8), spike counts Y conditional on their rates X are Poisson, which implies $\phi = \text{Var}(Y|X)/E(Y|X) = 1$ (Section 2). Here we consider a generative model permitting $\phi \neq 1$:

$$P_{\lambda,\phi} = \begin{cases} \text{Binomial} \left(N = \left\lfloor \frac{\lambda}{1-\phi} \right\rfloor, p = 1 - \phi \right), & \phi \in (0, 1) \\ \text{Poisson}(\lambda), & \phi = 1 \\ \text{Negative Binomial} \left(r = \frac{\lambda}{1-\phi}, p = \frac{\phi-1}{\phi} \right), & \phi > 1, \end{cases} \quad (12)$$

where $\lfloor x \rfloor$ is the closest integer to x . This model has $E(Y|\lambda) \approx \lambda$ and $\text{Var}(Y|\lambda) \approx \lambda\phi$ (see Appendix D). Figure 4A illustrates the deviations of $P_{\lambda,\phi}$ with $\phi = 0.8$ and 1.2 from the Poisson distribution, i.e. $P_{\lambda,\phi}$ with $\phi = 1$.

2. Under model (8), (W_{1r}, W_{2r}) are bivariate lognormal. Here we generate them as follows: set $b > 0$ and (i) generate $(Z_{1r}, Z_{2r}) \sim \mathcal{N}(\mu, \Sigma)$, (ii) evaluate $W_{ir} = (1 + Z_{ir}b)^{1/b}$, $i = 1, 2$, and (iii) repeat i-ii until $W_{1r}, W_{2r} > 0$. When $b \rightarrow 0^+$, it is easy to show that (W_{1r}, W_{2r}) are bivariate lognormal. Figure 4B shows $f_{b,\mu,\Sigma}$ for $b = 0$ and $b = 1$, with μ, Σ chosen to match typical values in the V4 data analyzed in Section 4.

Simulations details and sensitivity results for the V4 data are in Section 4.2.

3.2 Nonparametric estimation

Let $S_i = \{S_{ir}\}_{r=1}^n$ be n independent spike trains of duration $[0, T]$ for neuron i . We divide $[0, T]$ into m bins of equal length and we let Y_{irj} denote the spike count in trial r and bin j , so that $\sum_{j=1}^m Y_{irj} = Y_{ir}$ is the total spike count for neuron i on trial r . We assume $E[Y_{irj}|X_{1r}, X_{2r}, Y_{1r}, Y_{2r}] = p_{ij}Y_{ir}$, where $\sum_j p_{ij} = 1$ and p_{ij} is the probability that, given that neuron i spikes, the spike occurs in bin j , independent of trial r ; in other words, the expected spike count in bin j in trial r is a portion p_{ij} of the total spike count Y_{ir} in trial r , with p_{ij} constant across trials. We let $0 < K < m - 1$ and define:¹

$$G_{m,K}(S_1, S_2) = \frac{\sum_{r=1}^n \sum_{|j-h| \leq K} (Y_{1rj} - p_{1j}Y_{1r})(Y_{2rh} - p_{2h}Y_{2r})}{n \left(1 - \sum_{|j-h| \leq K} p_{1j}p_{2h} \right)}. \quad (13)$$

Proposition 1—If the spikes of neurons 1 and 2 given their rates are uncorrelated whenever they are separated by more than K bins, that is they satisfy $\text{Cov}(Y_{irj}, Y_{krh}|X_{ir}, X_{kr}) = 0$ for $1 \leq i \leq k \leq 2$ whenever bins j and h are such that $K < |j - h| \leq m - 1$, then $G_{m,K}(S_1, S_2)$ and $G_{m,K}(S_i, S_i)$ are unbiased estimators of γ and $\phi E[X_{ir}]$, respectively.

The proof is in Appendix E. The probabilities p_{ij} are typically unknown but can be estimated by the corresponding observed proportions $\hat{p}_{ij} = \sum_{r=1}^n Y_{irj} / \sum_{r=1}^n Y_{ir}$, or smooth versions; in Section 4 we estimated p_{ij} based 50ms binwidths peri-stimulus time histograms. The

¹A sufficient condition for the denominator of eq.(13) to be strictly positive is $p_{ij} > 0$ for all $i = 1, 2$, since $K < m - 1$ implies $\sum_{|j-h| \leq K} p_{1j}p_{2h} < \sum_{j,h} p_{1j}p_{2h}$, where $\sum_{j,h} p_{1j}p_{2h} = 1$. The p_{ij} could be smoothed to achieve $p_{ij} > 0$ for low firing rate neurons.

unbiased property depends on K : given the trial length T and number of bins m , we need K such that $K/m \gg \delta/T$, where δ is such that the spikes of the two neurons are correlated only if they are within δ . If the spikes are still correlated when they are further than T apart, we cannot obtain unbiased estimates of γ and $\phi_i E[X_{ii}]$. Thus, Proposition 1 yields unbiased estimates only if within-trial covariabilities occur on time scales smaller than the trial duration T , so that $\delta < T$, and m is sufficiently large, specifically $m > T/(T - \delta)$ (see Appendix E). Within these constraints, we use the smallest value of K , that is $K/m \approx \delta/T$, because the variance of $G_{m,K}$ increases with K . For the data in Section 4, Smith and Sommer (2013) found that, based on jitter corrected cross-correlograms, short time scale correlation happened on a 5 – 10 ms half-width, so we use a half-width of $\delta = 20$ ms to be conservative. The variance of the estimators decreases with the number of bins m because: (i) the trial length T is finite so the cross terms of the extreme indices in the summand of eq.(13) are not available, so that the ends of the spike trains do not contribute to estimating γ ; using small bins (m large) therefore means using more data, which reduces the variance of the estimators; (ii) $(\delta/T) < 1$ and m large makes it possible for $K = \delta(T/m)$ to be an integer; otherwise, we must round K up to the next integer so we avoid biasing the estimators, with consequence that the cross terms for $|j - k|$ large in eq.(13) may contain less signal and more noise, and thereby increase the variance of the estimators; and (iii) small bins better localize the beginnings and ends of within-trial effects, which reduces the variance of eq.(13) for the same reason as in (ii). In Section 4 we used $m = 100$ bins, and thus $K = 2$; it would be statistically more efficient to use the largest m possible up to recording accuracy (in our case this would be $m = 1000$, since $T = 1000$ ms and spike time accuracy is in ms), but the computational burden to perform the within-trial covariability jitter tests (Section 3.3) would be prohibitive (Appendix G).

Plug-in estimators—We use Proposition 1 to define plugin estimators for the components of eqs. (5) and (6):

$$\widehat{ATT}_{plug} = \prod_{i=1,2} \left(1 + \frac{G_{m,K}(S_i, S_i)}{\widehat{Var}(Y_{ir}) - G_{m,K}(S_i, S_i)} \right)^{-1/2} \quad (14)$$

$$\hat{\Gamma}_{plug} = \frac{G_{m,K}(S_1, S_2)}{\sqrt{\widehat{Var}(Y_{1r}) \widehat{Var}(Y_{2r})}} \quad (15)$$

$$\widehat{FRC}_{plug} = \frac{\widehat{SCC}_{smp} - \hat{\Gamma}_{plug}}{\widehat{ATT}_{plug}} \quad (16)$$

$$\hat{\phi}_{iplug} = \frac{G_{m,K}(S_i, S_i)}{\bar{Y}_i} \quad (17)$$

where

$$\widehat{Cov}(Y_{1r}, Y_{2r}) = (n-1)^{-1} \sum_{r=1}^n (Y_{1r} - \bar{Y}_{1\cdot}) (Y_{2r} - \bar{Y}_{2\cdot}),$$

$$\bar{Y}_{i\cdot} = n^{-1} \sum_{r=1}^n Y_{ir}, \quad \widehat{Var}(Y_{ir}) = \widehat{Cov}(Y_{ir}, Y_{ir}),$$

and

$$\widehat{SCC}_{samp} = \frac{\widehat{Cov}(Y_{1r}, Y_{2r})}{\sqrt{\widehat{Var}(Y_{1r}) \widehat{Var}(Y_{2r})}} \quad (18)$$

are the sample estimates of the corresponding population quantities. When $\gamma = \Gamma = 0$, we set $\hat{\Gamma}_{plug} = 0$ (see Section 3.3). Note that the plug-in estimators do not always lie in the required domains (Churchland et al. (2011) encountered similar problems); for example, $\widehat{FRC}_{plug}$ could lie outside $[-1, 1]$. They are nevertheless useful estimates for exploratory data analysis.

We are able to estimate the components of eqs. (5) and (6) without specific assumptions on ϕ_i in contrast to the assumption that $\phi_i = 1$ in Section 3.1 and in Goris et al. (2014), and to the approach of Churchland et al. (2011), who use ϕ_i like a tuning parameter and choose it such that variances and correlations are well defined. Here we derive a nonparametric estimate of ϕ_i (eq.(17)) by using the information carried by spike counts at time scales smaller (i.e. spike counts in smaller bins) than the time scale of SCC and FRC.

3.3 Jitter test for within-trial correlation

In Section 2 we presented a hierarchical model where spike counts (Y_{1r}, Y_{2r}) are trial-to-trial correlated through the rates (X_{1r}, X_{2r}) , but also, conditionally on (X_{1r}, X_{2r}) , they might share *within-trial covariability*, measured by γ or Γ (Section 2). The case $\gamma = 0$ (equivalently $\Gamma = 0$) is of particular interest because then the attenuation of FRC is fully explained by the attenuation factor ATT (eq.(6)).

A parametric likelihood ratio test (LRT) of $H_0: \gamma = 0$ based on model (8) is problematic because it is limited to the one-sided alternative $H_a: \gamma > 0$, whereas γ could be negative, and the asymptotic chi-square null distribution of the LRT does not hold because $\gamma = 0$ is at the boundary of the parameter space; a bootstrap would be needed to obtain the null distribution. The parametric estimate of γ also turned out to be less efficient than the nonparametric estimate (eq.(13)) in a simulation study, where we repeatedly generated synthetic datasets similar to ours and estimated γ parametrically and nonparametrically: both estimators were unbiased, but the nonparametric estimator had smaller variance, which translates into more power in a testing context. The difference in efficiency happens because the two estimators are obtained from different reductions of the observed spike trains: the parametric estimate of γ was based on spike counts in bins of $T = 1$ s and the nonparametric estimate on spike counts in 10 ms bins. The spike trains contain more information about γ than spike counts (data processing inequality, Cover and Thomas (2012)), and presumably spike counts in

smaller bins contain more information than spike counts in larger bins, since they are closer to the original data.

We therefore prefer a model-free *jitter test* that accommodates the two-sided alternative $\gamma \neq 0$, does not depend on parametric assumptions, and is likely to be more powerful. We use the nonparametric estimator of γ in Section 3.2, $\tilde{\gamma} = G_{m,K}$ ($m = 100$ and $K = 2$ are appropriate for the data in Section 4), and we reject $H_0 : \gamma = 0$ if $\tilde{\gamma} < F^{-1}(\alpha/2)$ or $\tilde{\gamma} > F^{-1}(1 - \alpha/2)$, where F is the null distribution of $\tilde{\gamma}$ obtained using a jitter method (Harrison et al., 2013), as follows. Let $(Y_{ir1}, \dots, Y_{irm})$ denote the spike counts of neuron i in trial r and time-bin $j = 1, \dots, m$, such that $\sum_{j=1}^m Y_{irj} = Y_{ir}$ is the total spike count of neuron i in trial r . Let

$$(Y_{ir1}^*, \dots, Y_{irm}^*) | Y_{ir} \sim \text{Multinomial}(Y_{ir}; p_{i1}, \dots, p_{im}), \quad r=1, \dots, n$$

be *jittered spike counts* for neuron i for each of the n trials, where p_{ij} is the probability that, given that neuron i spikes, the spike occurs in bin j (the same p_{ij} 's appear in section 3.1). Let $\tilde{\gamma}^*$ be the nonparametric estimator of γ calculated from the jittered spike counts of the neuron pair. We repeat this B times to obtain B values of $\tilde{\gamma}^*$: their EDF approximates the null distribution F of $\tilde{\gamma}$. To obtain accurate p-values from an EDF requires a simulation size B in the thousands. However, a normal probability plot diagnostic revealed that the $\tilde{\gamma}^*$ were approximately normally distributed (this is not surprising given that eq.(13) is an average), so a faster approximation of F is the normal distribution with mean 0 and variance equal to the sample variance of B values of $\tilde{\gamma}^*$, where B can be much smaller, say $B = 100$ (see Canty et al. (2006)). The jitter test is a conditional parametric bootstrap test that is exact when the p_{ij} 's are known. In our application we must estimate the p_{ij} 's (Section 3.2), so the test is no longer exact. However, Figure 5 shows that its spurious detection rate is nevertheless close to the significance level α , so the test is reliable.

Statistical power—The power of a test is the probability to detect deviations from the null hypothesis when they exist. It is difficult to calculate analytically the power of the jitter test, so we estimated it by simulation. For values of $\gamma \in (0, 1.25)$, we simulated 1000 datasets of 60 trials of pairs of independent spike trains satisfying model (7) and induced within-trial covariability by adding an independent homogeneous Poisson spike train with rate γ to each spike train pairs, lagged by either 0, 10, or 20 ms. The resulting pairs of spike trains satisfy model (8) with parameters $\mu_i = 1.9$, $\sigma_i = 0.31$, $\rho = 0.51$; these parameter values match the mean estimated values in the data in Section 4. Then for each γ , we estimated the power as the proportion of times the jitter test rejected the hypothesis $\gamma = 0$ out of the 1000 datasets ($\delta = 20$ ms, $m = 100$, $K = 2$), and plotted this estimate in Figure 5A against γ and a standardized version of γ that takes value in $[-1, 1]$, the *expected conditional spike count correlation*:

$$ECC = E \left[\frac{\text{Cov}(Y_{1r}, Y_{2r} | X_{1r}, X_{2r})}{\prod_{i=1,2} \sqrt{\text{Var}(Y_{ir} | X_{ir})}} \right]. \quad (19)$$

Appendix F provides nonparametric and parametric estimates of ECC. The power of the jitter test increases with the amount of within-trial covariability γ and has very high power when $\gamma > 0.75$, $\text{ECC} > 0.1$.

False Discovery Rate control—When many tests are performed, it is useful to cap the expected proportion of false rejections among all rejections – the False Discovery Rate (FDR) – below some desired β . Benjamini and Hochberg (1995) proceed as follows: given M independent tests with null hypotheses $H_{0(1)}, \dots, H_{0(M)}$ and corresponding p -values $P_1, \dots, P_M$, they reject $H_{0(1)}, H_{0(2)}, \dots, H_{0(K)}$ where $K = \max\{i \text{ such that } P_{(i)} \leq i\beta/M\}$ and $P_{(j)}$ is the j -th largest p -value; then, in expectation, fewer than a proportion β of the K rejected tests correspond to true null hypotheses. Benjamini and Yekutieli (2001) show that this procedure also controls the FDR in the case of positive regression dependent tests, but advocate a more conservative FDR control for tests with general dependencies, which consists of replacing β with $\beta / \sum_{i=1}^M i^{-1}$.

The jitter tests are dependent, since they are applied to all neuron pairs, but they are not necessarily positive regression dependent, so we investigated the performances of both procedures for data like ours. We simulated $n = 60$ independent trials of spike trains from an N -variate Poisson-lognormal model with $N = 30$ (we extended model (7) from $N = 2$ to 30), where μ and Σ were assigned values that match the V4 data in Section 4, such that half of the neuron pairs had $\text{FRC} < 0$ and half $\text{FRC} > 0$. Neurons were randomly assigned to three groups of size $30/3 = 10$. For each trial r , the spike trains of all neurons in a group were contaminated by the same noise spike train with rate γ , and the noise spike trains were independent between groups; we used $\gamma = 0.5, 1$, and 2 for the three groups respectively, to match approximately the values calculated in the data. The proportion of tests with H_0 false

was therefore $1/3$. We performed the $\binom{N}{2} = 435$ jitter tests for all neuron pairs and recorded the observed FDR and False Non-discovery Rate (FNR), that is the expected proportion of false nulls among the non-rejections, for a range of nominal FDR β . We repeated the simulation 1000 times to obtain accurate mean estimates of the observed FDR and FNR. In Figure 5B we show that the Benjamini and Hochberg (1995) procedure provides reliable FDR control for data like ours (observed FDR is below β), perhaps because the tests were only weakly dependent, as measured by the Brownian covariance test of Székely and Rizzo (2009). We also observe $\text{FNR} < 10\%$ whenever $\beta \geq 5\%$, which means that most cases with $\Gamma = 0$ are detected by the jitter test with FDR control. The Benjamini and Yekutieli (2001) procedure for tests with general dependencies provided too conservative a control (smaller FDR and larger FNR; results not shown), so we did not use it.

4 Analysis of V4 data

We analyze the rhesus macaque monkey (*Macaca mulatta*) visual V4 data collected by Smith and Sommer (2013) during the following task (Figure 6): the animal fixated a small blue dot in the center of a computer screen while a sinusoidal grating covered the aggregate receptive field area of the recorded neurons; the grating's size and spatial and temporal frequencies were fixed, and the direction of drift was one of 12 equally spaced directions for each trial; n

= 61 independent trials of duration $T = 1$ sec were recorded for each of the 12 stimuli (directions). The neural data were acquired from a Utah array (100 electrodes, Blackrock Microsystems) implanted in V4. The electrode voltages were amplified and band-pass filtered (250-7.5 kHz) prior to spike sorting with custom software (Kelly et al., 2007). All procedures were approved by the Institutional Animal Care and Use Committee of the University of Pittsburgh and were in compliance with the guidelines set forth in the National Institutes of Health's Guide for the Care and Use of Laboratory Animals.

For our statistical analysis we retained only the 53 well-isolated single units with signal-to-noise ratio $\text{SNR} > 2.75$ (Kelly et al., 2007) and 1 Hz minimum trial-averaged firing rate, and

studied the relationships between all $\binom{53}{2} \times 2 = 16536$ neuron pairs-stimuli combinations.

We proceeded as outlined in Figure 2. We first applied the jitter test (Section 3.3) with a false discovery rate (FDR) of 10% to all neuron pairs-stimuli combinations to determine which pairs have $\Gamma = 0$: only a minority of the V4 neuron pairs exhibited significant within-trial covariance (two-sided test, 6.04% of cases reject the null hypotheses; one-sided test to find $\gamma > 0$, 10.07%), which suggests that within-trial correlations have a weak impact on SCC, and therefore that the attenuation effect is mainly characterized by ATT in eq.(6).

Figure 7B shows the distribution of the nonparametric estimates $\widehat{ATT}_{plug}$ (eq.(14)); it is concentrated around 0.35, which means that the magnitude of observed SCCs is on average only 35% of the magnitude of FRC.

Next we fitted the parametric models in Section 3.1 consistently with the jitter test results (the DC-PLN model (8) if $\Gamma > 0$, and the PLN model (7) otherwise) using the maximum likelihood Algorithm 2 in Appendix C, and obtained estimates of the components of eq.(6). A component of the model fitted to a neuron pair is shown in Figure 8, while goodness of fit diagnostics and sensitivity to parametric assumptions are summarized in Sections 4.1 and 4.2. Figure 3B shows the sample estimates of SCC in eq.(18) versus the model-based estimates of FRC in eq.(9) for all the 16536 neuron pairs-stimuli combinations: the SCC is substantially attenuated compared to the FRC, which is consistent with the nonparametric estimates of ATT being well below one in Figure 7B. Figures 7C and 7D further show that, on average, SCC and FRC decrease with pairwise *interneural distance* (distance between electrodes) and increase with *tuning curve correlation* (TCC, eq.(25)), which is a measure of similarity between the neural activity of two neurons. Smith and Sommer (2013) had noted these relationships for SCC, but not FRC. The sample estimate of ECC (eq.24) also decreases with interneural distance and increases with TCC, although it is hard to see by the naked eye because ECC is close to zero.

In Appendix H, we also display SCC, FRC, and ECC against the *generalized tuning curve correlation* (GTCC), which is a combined measure of tuning curve similarity and firing rate covariability between a neuron pair, which is motivated in Figure 8. This figure shows the estimated joint *generalized tuning curve* (GTC) of a neuron pair, that is the estimated joint distribution of their firing rates for all stimuli, and shows not only how the neurons' average firing rates vary jointly with the stimulus – the TCC provides a one number summary of that covariability across stimuli –, but also how differently these rates are correlated across trials

for different stimuli; for the neuron pair shown in Figure 8, these rates are virtually uncorrelated when stimulus 7 is presented, and very strongly positively correlated for stimuli 2, 9, and 12. The GTCC is a one number summary of both tuning curve similarity and firing rate covariability between a neuron pair.

4.1 Parametric model goodness of fit

Figures 3B, 7C and 7D display the parametric estimates of FRC. We justify here that the parametric models fit the data.

Figure 9A shows the sample variance versus the sample mean of the spike counts of the 53 neurons for all 12 stimuli: the spike counts are overdispersed, which is a feature of the Poisson-lognormal model. The KL divergence test (Section 3.1) confirms that the univariate Poisson-lognormal model (i.e. the univariate version of eq.(7)) fits single neuron spike counts well: only 5% of the tests were rejected, which matches the spurious detection rate α . To provide a contrasting result, we also performed KL divergence tests assuming Poisson spike counts. Their p-values are plotted in Figure 9B against the spike counts' Fano factors: 52% of the Poisson tests were rejected, including all the tests applied to data with $\log_e$ Fano factors above 0.6. The p-values of the univariate Poisson-lognormal KL divergence tests are also plotted and exhibit no relationship with the Fano factor.

The univariate Poisson-lognormal model fits single neurons' spike counts. Now we assess the fits of the PLN and DC-PLN models in eqs.(7) and (8) to pairs of spike counts: we randomly selected 2500 cases ($\approx 15\%$ of the 16536 neuron pairs-stimuli combinations) to which we applied the KL and SCC divergence tests (Section 3.1). Only about 4% of the neuron pairs did not pass the SCC divergence test and less than 1.5% did not pass the KL test, at the 5% level, which suggests that the bivariate parametric models fit pairs of spike counts well.

4.2 Sensitivity of FRC estimates to model assumptions

Figures 3B, 7C and 7D display the parametric estimates of FRC where in about 90% of cases, we imposed $\gamma = \Gamma = 0$ according to the results of the jitter test. This test has low false non-discovery rate (FNR, Figure 5B) and high power to detect deviations from $\gamma = 0$ when $\gamma > 0.75$ or $ECC > 0.1$ (Figure 5A) so we are confident that we fitted appropriate model, that is the PLN model, in such cases. However, we would like to know how the parametric estimate of FRC is affected by assuming $\gamma = 0$ when the jitter test fails to detect that γ is small and positive, and how sensitive it is to mild deviations from the Poisson and log-normal assumptions.

We applied the methods in Section 3.1.2. We generated $n = 60$ trials of duration $T = 1$ sec of spike count pairs (Y_{1T}, Y_{2T}) from the generalized bivariate spike count model (11) for (φ, b, γ) in $(0.8, 1.2) \times (0, 1) \times (0, 1)$, where $\mu(\varphi, b, \gamma)$ and $\Sigma(\varphi, b, \gamma)$ were calibrated for each (φ, b, γ) so that spike counts had means 7 and variances 12, and $FRC = 0.5$, to match typical values in the V4 data. We used these data to estimate FRC under model (7), and repeated the simulation 10,000 times. Figure 10A shows the mean of the repeat values of $\widehat{FRC}_{mle}$ and Figure 10B the standard deviation of $\widehat{FRC}_{mle}$ relative to what it should be if the spike

counts were generated from the estimation model (7), as functions of the generating model parameters, ϕ , b , and γ (the estimating model has $\phi = 1$, $b = 0^+$, and $\gamma = 0$). The deviations from lognormality ($b > 0$) have practically no effect on $\widehat{FRC}_{mle}$, and the deviations from $\phi = 1$ moderate effects: $\widehat{FRC}_{mle}$ is fairly robust to mild deviations from the Poisson-lognormal model. However, the deviations from $\gamma = 0$ increase the bias of $\widehat{FRC}_{mle}$ substantially and reduce its variance², so inferences are not reliable (for example, a confidence interval for the true FRC would be centered on the wrong value and have lower coverage than expected). The possibility of this happening in our data is small because the FNR is less than 10% (Figure 5B).

5 Discussion

We have provided a collection of methods for eliminating Poisson-like noise from the SCC to produce the FRC. As shown in 3B, FRC tends to be much larger in magnitude than SCC, and should be much more sensitive to experimental manipulations. A relatively simple bivariate hierarchical model may be introduced when the within-trial correlation is sufficiently small that its impact on FRC is negligible, and we proposed a jitter test of this assumption. We showed that the majority of pairs of V4 neurons failed to reject the null hypothesis of zero within-trial correlation. Careful analysis requires an evaluation of the power of the test and the sensitivity of the proposed parametric model to mild departures from assumptions, which we supplied, along with an evaluation of the expected false non-discovery rate (FNR) of the test, which we determined to be low. Together, these analyses strongly suggest that, in our data, within-trial correlation rarely induce distortions in the estimate of FRC when the jitter test does not reject the null hypothesis. For those pairs of neurons where the null hypothesis of zero within-trial correlation was rejected, we used a nonparametric estimate of within-trial correlation based on spike trains, and then applied a correlated Poisson hierarchical model. This model produces positive within-trial correlation. A more complete accounting of within-trial variation, which would include negative within-trial correlation, could be obtained from point process models (Snyder and Miller, 1991; Kass et al., 2014) and is the subject of future research.

Our finding that only a few V4 neuron pairs exhibited within-trial correlations differs from the result of Goris et al. (2014) in the macaque visual area V1. First, their SCCs are larger than in our V4 data; a similar discrepancy is also documented by the V1 analysis of Smith and Kohn (2008) and the V4 studies of Smith and Sommer (2013), Cohen and Maunsell (1999), and Mitchell et al. (2009). Second, the attenuation effect in Goris et al. (2014) is much weaker than ours, so that SCC is only weakly attenuated compared to FRC (there called “gain correlation”), while their point process covariability is much higher (ours is measured by Γ or ECC). These differences may be due to the different area under study. However, the different anesthetics used in the two experiments (our animal was alert) might also have induced different levels of up and down states (Steriade and Buzsaki, 1990; Cowan

²The FRC is overestimated likely because the positive within-trial covariance ($\gamma > 0$) causes an increase in SCC (see eq.(6)), which is erroneously attributed to the firing-rates covariance by the estimation model (7), which assumes $\text{Cov}(Y_{1r}, Y_{2r}) \equiv \text{Cov}(X_{1r}, X_{2r})$.

and Wilson, 1994; Timofeev et al., 2001; Steriade et al., 2001), which in turn could have affected within-trial covariabilities.

Our overall goal was to provide a useful suite of methods for examining the trial-to-trial correlated behavior of neurons at trial-length time scales. To the extent that neural implementations of maximum likelihood and Bayesian inference are possible (e.g., Ma et al. (2006)), downstream neural networks should be able to strip away noise from spike counts and use FRC to advantage, just as we have been able to do statistically. Whether they do so remains an open question.

Acknowledgments

This work was supported by grants from the NIH: RO1 MH064537. We used the R-PACKAGES `poilog`, by Vidar Grøtan and Steinar Engen; `energy`, by Maria L. Rizzo and Gabor J. Székely; and `locfit`, by Catherine Loader.

Appendix

A: Spike Count Correlation formula

Using the laws of total expectation and total variance, equations (1) and (2) imply

$$E[Y_{ir}] = E\{E[Y_{ir}|X_{ir}]\} = E[X_{ir}]$$

and

$$\begin{aligned} Var(Y_{ir}) &= E\{Var(Y_{ir}|X_{ir})\} + Var(E[Y_{ir}|X_{ir}]) \\ &= \phi_i E[X_{ir}] + Var(X_{ir}) \end{aligned}$$

Thus

$$\begin{aligned} Cor(Y_{1r}, Y_{2r}) &= \frac{\gamma + Cov(X_{1r}, X_{2r})}{\prod_{i=1,2} \sqrt{Var(Y_{ir})}} \quad \text{with} \quad \gamma = E[Cov(Y_{1r}, Y_{2r}|X_{1r}, X_{2r})] \\ &= \frac{Cov(X_{1r}, X_{2r})}{\prod_{i=1,2} \sqrt{(\phi_i E[X_{ir}] + Var(X_{ir}))}} + \frac{\gamma}{\prod_{i=1,2} \sqrt{Var(Y_{ir})}} \\ &= \frac{Cov(X_{1r}, X_{2r})}{\prod_{i=1,2} \sqrt{Var(X_{ir})(\phi_i/F_i + 1)}} + \Gamma \\ &= Cor(X_{1r}, X_{2r}) \prod_{i=1,2} (1 + \phi_i/F_i)^{-1/2} + \Gamma \end{aligned}$$

B: Lognormal distribution

Let $(W_{1r}, W_{2r}) = (\log X_{1r}, \log X_{2r}) \sim \mathcal{N}(\mu, \Sigma)$. Then:

$$E[X_{ir}] = E[e^{W_{ir}}] = e^{\mu_i + \sigma_i^2/2}$$

$$\begin{aligned}
Var(X_{ir}) &= E[e^{2W_{ir}}] - (E[e^{W_{ir}}])^2 \\
&= e^{2\mu_i + 2\sigma_i^2} - e^{2\mu_i + \sigma_i^2} \\
&= (e^{\mu_i + \sigma_i^2/2})^2 (e^{\sigma_i^2} - 1)
\end{aligned}$$

$$\begin{aligned}
Cov(X_{1r}, X_{2r}) &= E[e^{W_{1r} + W_{2r}}] - E[e^{W_{1r}}] E[e^{W_{2r}}] \\
&= e^{\mu_1 + \mu_2 + (\sigma_1^2 + \sigma_2^2)/2 + \sigma_{12}} - e^{\mu_1 + \mu_2 + (\sigma_1^2 + \sigma_2^2)/2} \\
&= e^{\mu_1 + \sigma_1^2/2} e^{\mu_2 + \sigma_2^2/2} (e^{\sigma_{12}} - 1)
\end{aligned}$$

$$\text{such that } F_i = e^{\mu_i + \sigma_i^2/2} (e^{\sigma_i^2} - 1) \text{ and } Cor(X_{1r}, X_{2r}) = (e^{\rho\sigma_1\sigma_2} - 1) \prod_{j=1,2} (e^{\sigma_j^2} - 1)^{-1/2}.$$

C: Doubly Correlated Poisson-lognormal distribution

Let $p_{\lambda}(y)$ be the p.f. of $\text{Poisson}(\lambda_i)$. Let $R_{ir} \sim \text{Poisson}(\lambda_i)$, $i = 0, 1, 2$, be independent, and define $Y_{ir} = R_{0r} + R_{ir}$, $i = 1, 2$. The joint distribution of (Y_{1r}, Y_{2r}) is:

$$\begin{aligned}
P(Y_{1r}=y_1, Y_{2r}=y_2) &= P(R_{0r} + R_{1r}=y_1, R_{0r} + R_{2r}=y_2) \\
&= P(R_{1r}=y_1 - R_{0r}, R_{2r}=y_2 - R_{0r}) \\
&= \sum_{z=0}^{\min\{y_1, y_2\}} P(R_{1r}=y_1 - R_{0r}, R_{2r}=y_2 - R_{0r}, R_{0r}=z) \\
&= \sum_{z=0}^{\min\{y_1, y_2\}} p_{\lambda_1}(y_1 - z) p_{\lambda_2}(y_2 - z) p_{\lambda_0}(z).
\end{aligned}$$

By integrating with respect to $(\log \lambda_1, \log \lambda_2)^T \sim \mathcal{N}(\mu, \Sigma)$ we obtain the p.f. of the DC-PLN model (8):

$$p_{\mu, \Sigma, \gamma}(y_1, y_2) = \sum_{z=0}^{\min\{y_1, y_2\}} q_{\mu, \Sigma}(y_1 - z, y_2 - z) p_{\lambda_0}(z) \quad (20)$$

where

$$q_{\mu, \Sigma}(y_1, y_2) = \int_{\mathbb{R}^2} p_{e^{x_1}}(y_1) p_{e^{x_2}}(y_2) g_{\mu, \Sigma}(x_1, x_2) dx_1 dx_2 \quad (21)$$

is the p.f. of the PLN model (7), and $g_{\mu, \Sigma}$ is the Gaussian p.d.f..

We can compute the ML estimators of μ , Σ , γ numerically using Algorithm 1, when only the spike counts are available, or Algorithm 2, when the spike trains are available. In data simulated as in Section 4.2, we found that Algorithm 2 typically yielded lower variance estimates of FRC. We therefore recommend Algorithm 2, keeping in mind that Algorithm 1

could provide better estimates of other quantities. Note that if we assume $\gamma = 0$, that is the PLN model, then we just use step 2 of Algorithm 2 with $\gamma = 0$.

D: Generalized count model

Let $Y|\lambda$ have the distribution in (12). Thus

- if $\phi \in (0, 1)$, then $E[Y|\lambda] = \left\lfloor \frac{\lambda}{1-\phi} \right\rfloor (1-\phi) = \left(\frac{\lambda}{1-\phi} + \epsilon \right) (1-\phi) = \lambda + \epsilon(1-\phi) \approx \lambda$, and $\text{Var}(Y|\lambda) = E[Y|\lambda]\phi = \lambda\phi + \epsilon(1-\phi)\phi \approx \lambda\phi$, where $\epsilon \in (-0.5, 0.5)$, such that $\text{Var}(Y|\lambda)/E[Y|\lambda] \approx \phi$,
- if $\phi = 1$, then $E[Y|\lambda] = \lambda$ and $\text{Var}(Y|\lambda) = \lambda$,
- and if $\phi > 1$, then $E[Y|\lambda] = \lambda\phi(1-\phi)^{-1} = \lambda$ and $\text{Var}(Y|\lambda) = \lambda\phi(1-\phi)^{-2} = \lambda\phi$.

E: Non-parametric estimators of γ and ϕ_i

Proof of Proposition 1. Let $X_r = (X_{1r}, X_{2r})$ and $Y_r = (Y_{1r}, Y_{2r})$. Exploiting the covariance decomposition as in eq. (3) and the fact that $E[Y_{ir}|X_r, Y_r] = p_{ij}Y_{ir}$ we get

$$\begin{aligned} \text{Cov}(Y_{1rj}, Y_{2rh}|X_r) &= E[\text{Cov}(Y_{1rj}, Y_{2rh}|X_r, Y_r)|X_r] + \text{Cov}(E[Y_{1rj}|X_r, Y_r], E[Y_{2rh}|X_r, Y_r]|X_r) \\ &= E\{E[(Y_{1rj} - p_{1j}Y_{1r})(Y_{2rh} - p_{2h}Y_{2r})|X_r, Y_r]|X_r\} + \text{Cov}(p_{1j}Y_{1r}, p_{2h}Y_{2r}|X_r) \\ &= E[(Y_{1rj} - p_{1j}Y_{1r})(Y_{2rh} - p_{2h}Y_{2r})|X_r] + p_{1j}p_{2h}\text{Cov}(Y_{1r}, Y_{2r}|X_r) \end{aligned}$$

which implies

$$E[(Y_{1rj} - p_{1j}Y_{1r})(Y_{2rh} - p_{2h}Y_{2r})|X_r] = \text{Cov}(Y_{1rj}, Y_{2rh}|X_r) - p_{1j}p_{2h}\text{Cov}(Y_{1r}, Y_{2r}|X_r) \quad (22)$$

Moreover, zero conditional covariance across time bins j, h s.t. $|j-h| > K$ implies

$$\text{Cov}(Y_{1r}, Y_{2r}|X_r) = \sum_{|j-h| \leq K} \text{Cov}(Y_{1rj}, Y_{2rh}|X_r) \quad (23)$$

Therefore, combining eqs. (22) and (23) we get

$$\sum_{|j-h| \leq K} E[(Y_{1rj} - p_{1j}Y_{1r})(Y_{2rh} - p_{2h}Y_{2r})|X_r] = \text{Cov}(Y_{1r}, Y_{2r}|X_r) \left(1 - \sum_{|j-h| \leq K} p_{1j}p_{2h}\right)$$

implying that

$$E[G_{m,K}(S_1, S_2)] = E[\text{Cov}(Y_{1r}, Y_{2r}|X_{1r}, X_{2r})].$$

The proof that $G_{m,K}(S_i, S_j)$ is an unbiased estimator of $\phi_i E[X_{ir}]$ is similar; all we need is to replace (Y_{1r}, Y_{2r}) and (X_{1r}, X_{2r}) by (Y_{ir}, Y_{jr}) and (X_{ir}, X_{jr}) .

Note that combining the conditions $K < m - 1$ and $K/m \ll T$ of Section 3.2 yields

$\frac{\delta}{T} \leq \frac{K}{m} < \frac{m-1}{m} < 1$, which implies $\frac{\delta}{T} < 1$ and $\frac{m-1}{m} > \frac{\delta}{T} \Rightarrow m > T / (T - \delta)$; these conditions are necessary to obtain unbiased estimators.

F: Expected Conditional spike-count Correlation

The *Expected Conditional spike-count Correlation* (ECC) in eq.(19) can be estimated non-parametrically analogously to (13):

$$\widehat{ECC}_{samp} = \frac{1}{|\mathcal{S}|} \sum_{r \in \mathcal{S}} \frac{\sum_{|j-h| \leq K} (Y_{1rj} - \hat{p}_{1j} Y_{1r}) (Y_{2rh} - \hat{p}_{2h} Y_{2r}) / \hat{Q}_{12}(K)}{\prod_{i=1,2} \sqrt{\sum_{|j-h| \leq K} (Y_{irj} - \hat{p}_{ij} Y_{ir}) (Y_{irh} - \hat{p}_{ih} Y_{ir}) / \hat{Q}_{ii}(K)}} \quad (24)$$

where $Q_{il}(K) = 1 - \sum_{|j-h| \leq K} p_{ij} p_{lh}$, and $\hat{p}_{ij}$'s are as discussed in Section 3.2, and $\mathcal{S} = \{r \in \{1, \dots, n\} : Y_{1r}, Y_{2r} > 0\}$, i.e. $\mathcal{S}$ is the set of trials where both neurons spiked at least once. Notice that if $\text{Cov}(Y_{1r}, Y_{2r} | X_{1r}, X_{2r})$ is constant across trials, then

$ECC = \gamma \times E \left[1 / \prod_{i=1,2} \sqrt{\text{Var}(Y_{ir} | X_{ir})} \right]$, so that $ECC = 0 \Leftrightarrow \gamma = 0$. Otherwise the equivalence is not necessarily true. For instance, under model (8) we have

$$ECC = \gamma \times E \left[1 / \sqrt{(\gamma + W_{1r})(\gamma + W_{2r})} \right].$$

G: Computational time

We implemented our estimation and testing procedures in R, by using personal codes and modifications of the package `poilog` to estimate the DC-PLN model (8). We used a laptop with an Intel® Core™ 2 Duo CPU T6600 ($2 \times 2.20\text{GHz}$), RAM 4GB (DDR 800MHz); and OS Ubuntu 15.04. Table 1 displays the time to compute $G_{m,K}$ (eq.(13)) and perform the jitter test on a neuron pair, and Table 2 to estimate the parametric model. The times required to process a dataset with N neurons and S stimuli can be obtained by multiplying the values in

the tables by $\binom{N}{2} \times S$. For example, in Section 4, the multiple jitter tests required about $\binom{53}{2} \times 12 \times 354 \text{ ms} \approx 97$ minutes.

Choosing m of $G_{m,K}$

In Section 3.1 we argued that the number of bins m should be large. However, the computational burden may be too large when we perform the jitter test (see Table 1): the number of elements in the summand in eq.(13) is $n \times M$, where $M = \text{card}(\{j, h : 1 \leq j, h \leq m, |j-h| \leq K\}) = m(2K+1) - K(K+1)$, so that for a fixed ratio $\pi = K/m = \delta/T$, the computational time of eq.(13) is proportional to $n \times M = O(n \times m^2)$. To set an upper bound

D on M given π , we use $m^*(\pi, D) = \frac{1-\pi + \sqrt{(1-\pi)^2 + 4(2\pi-\pi^2)D}}{2(2\pi-\pi^2)}$ and $K = \lceil \pi m^* \rceil$.

H: Generalized Tuning Curve Correlation

The *tuning curve* of a neuron i is generally defined as the spike count averaged across trials, $f_i(s) = E[Y_{it}(s)]$, which varies with the stimulus s when s modulates the neuron (Butts and Goldman, 2006). Under the assumptions of Section 2, the tuning curve reduces to $f_i(s) = E[X_{it}(s)]$, which contains information about the firing rate averaged across trials but none about the trial-to-trial variability of $X_{it}(s)$. Instead of reducing the tuning properties to the mean firing rate, we may retain in addition the full distribution of $X_{it}(s)$, say F_{is} , for each stimulus s : this ensemble is what we call the *generalized tuning curve* (GTC). Figure 11A displays the GTC for two example V4 neurons, assuming the Poisson-lognormal model (7) for their spike counts; Figure 8 shows the joint GTC for a neuron pair, which further displays their shared trial-to-trial covariability.

The *tuning curve correlation* (TCC) is widely considered an important measure of similarity of the neural activity between two neurons across stimuli. It can be defined as:

$$TCC = \text{Cor}(f_1(S), f_2(S)), \quad (25)$$

where the covariance is with respect to the random variable S which codes for the stimulus; for example, when all stimuli are presented with equal frequency, S is uniform over the stimuli. Consider the toy example where, for a stimulus s , $(X_{1r}(s), X_{2r}(s))$ are (trial-to-trial) random variables with generative model:

$$\begin{cases} X_{1r}(s) \sim F_{1s}, \\ X_{2r}(s) = (1+\beta) f_1(s) - \beta X_{1r}(s), \end{cases}$$

where we choose $\beta > 0$ to achieve $P(X_{2r}(s) \geq 0) = 1$. This model implies $E[X_{1r}(s)] = E[X_{2r}(s)]$ for all s , and in turn $TCC = 1$, from which one might conclude that the two neurons have proportional tuning curves. But the model also implies $\text{FRC}_s = \text{Cor}(X_{1r}(s), X_{2r}(s)) = -1$ for all s , which reveals that TCC only summarizes the similarity of the expected firing rates of the two neurons across stimuli, but does not convey the sign and magnitude of their shared trial-to-trial variability. This motivates us to define the *generalized tuning curve correlation*:

$$GTCC = \text{Cor}(X_{1r}(S), X_{2r}(S)), \quad (26)$$

where the correlation is with respect to the stimulus random variables S , as in eq.(25), but also to the trial-to-trial random firing rates $X_{it}(s)$. Hence, the TCC only captures the covariability of the expected firing rates across stimuli, and the GTCC also accounts for changes in their joint distribution, such as trial-to-trial firing rate correlation (FRC). That is, the GTCC summarizes the strength of the relationship between two neurons across both trials and stimuli.

Note that when the experiment involves only one stimulus, so that $S = s$ is not a random variable since it always takes the same value, then $GTCC = \text{FRC}$; and when the firing rates have no trial-to-trial variability, that is when $X_{it}(s) = E(X_{it}(s))$ for all r , then $GTCC = TCC$.

since $f_i(s) = X_{ir}(s)$ for all r . When both the stimulus and the firing rates are random, the GTCC can therefore be viewed as an approximate average of FRC across stimuli or as an approximate average of TCC across trials. More precisely, we can rewrite:

$$\begin{aligned} GTCC &= \frac{E[Cov(X_{1r}(S), X_{2r}(S)|S)] + Cov(f_1(S), f_2(S))}{\prod_{i=1,2} \sqrt{E[Var(X_{ir}(S)|S)] + Var(f_i(S))}} \\ &= TCC \times \frac{E[Cov(X_{1r}(S), X_{2r}(S)|S)] / Cov(f_1(S), f_2(S)) + 1}{\prod_{i=1,2} \sqrt{E[Var(X_{ir}(S)|S)] / Var(f_i(S)) + 1}} \quad (27) \\ &= TCC \times \Omega, \end{aligned}$$

where Ω depends on the joint distributions of the firing rates at different stimuli (Figure 8 shows such a distribution), so that GTCC can be viewed as a modified version of TCC. Figure 11B shows GTCC versus TCC for all the neuron pairs in the V4 data: GTCC has smaller magnitude and GTCC and TCC sometimes have different signs. Figure 11C shows SCC, FRC, and ECC as functions of GTCC, which resembles the plot of the same quantities versus TCC in Figure 7B, although the relationships appear stronger. This happens because GTCC is approximately an average of FRC across stimuli, as justified above, whereas SCC is an attenuated version of FRC.

References

- Amarasingham A, Geman S, Harrison MT. Ambiguity and nonidentifiability in the statistical analysis of neural codes. *Proceedings of the National Academy of Sciences*. 2015; 112(20):6455–6460.
- Behseta S, Berdyeva T, Olson CR, Kass RE. Bayesian correction for attenuation of correlation in multi-trial spike count data. *Journal of neurophysiology*. 2009; 101(4):2186–2193. [PubMed: 19129297]
- Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society. Series B (Methodological)*. 1995:289–300.
- Benjamini Y, Yekutieli D. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*. 2001:1165–1188.
- Bernardo, JM.; Smith, AF. Bayesian theory. Vol. 405. John Wiley & Sons.; 2009.
- Butts DA, Goldman MS. Tuning curves, neuronal variability, and sensory coding. *PLoS biology*. 2006; 4(4):e92. [PubMed: 16529529]
- Canty AJ, Davison AC, Hinkley DV, Ventura V. Bootstrap diagnostics and remedies. *Canadian Journal of Statistics*. 2006; 34(1):5–27.
- Churchland AK, Kiani R, Chaudhuri R, Wang XJ, Pouget A, Shadlen. Variance as a signature of neural computations during decision making. *Neuron*. 2011; 69(4):818–831. [PubMed: 21338889]
- Cohen MR, Maunsell JHR. Attention improves performance primarily by reducing interneuronal correlations. *Nature neuroscience*. 2009; 12(12)(2009):1594–1600. [PubMed: 19915566]
- Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nature neuroscience*. 2011; 14(7):811–819. [PubMed: 21709677]
- Cover, TM.; Thomas, JA. Elements of information theory. John Wiley & Sons; 2012.
- Cowan RL, Wilson CJ. Spontaneous firing patterns and axonal projections of single corticostriatal neurons in the rat medial agranular cortex. *Journal of neurophysiology*. 1994; 71(1):17–32. [PubMed: 8158226]
- Efron, B.; Tibshirani, RJ. An introduction to the bootstrap. Vol. 57. CRC press; 1994.
- Ecker AS, Berens P, Cotton RJ, Subramaniam M, Denfield GH, Cadwell CR, Tolias AS. State dependence of noise correlations in macaque primary visual cortex. *Neuron*. 2014; 82(1):235–248. [PubMed: 24698278]
- Goris RLT, Movshon JA, Simoncelli EP. Partitioning neuronal variability. *Nature neuroscience*. 2014

- Gu Y, Liu S, Fetsch CR, Yang Y, Fok S, Sunkara A, DeAngelis GC, Angelaki DE. Perceptual learning reduces interneuronal correlations in macaque visual cortex. *Neuron*. 2011; 71(4):750–761. [PubMed: 21867889]
- Harris KD, Thiele A. Cortical state and attention. *Nature reviews neuroscience*. 2011; 12(9):509–523. [PubMed: 21829219]
- Harrison, MT.; Amarasingham, A.; Kass, RE. Statistical identification of synchronous spiking.. In: Lorenzo, PD.; Victor, J., editors. *Spike timing: Mechanisms and function*. Taylor and Francis; London: 2013.
- Jeanne JM, Sharpee TO, Gentner TQ. Associative learning enhances population coding by inverting interneuronal correlation patterns. *Neuron*. 2013; 78(2):352–363. [PubMed: 23622067]
- Kass, RE.; Eden, UT.; Brown, EN. *Analysis of Neural Data*. Springer; 2014.
- Kass RE, Ventura V. Spike count correlation increases with length of time interval in the presence of trial-to-trial variation. *Neural computation*. 2006; 18(11):2583–2591. [PubMed: 16999571]
- Kelly RC, Smith MA, Samonds JM, Kohn A, Bonds AB, Movshon JA, Lee TS. Comparison of recordings from microelectrode arrays and single electrodes in the visual cortex. *The Journal of neuroscience*. 2007; 27(2):261–264. [PubMed: 17215384]
- Kelly RC, Smith MA, Kass RE, Lee TS. Local field potentials indicate network state and account for neuronal response variability. *Journal of computational neuroscience*. 2010; 29(3):567–579. [PubMed: 20094906]
- Kelly RC, Kass RE. A framework for evaluating pairwise and multiway synchrony among stimulus-driven neurons. *Neural computation*. 2012; 24(8):2007–2032. [PubMed: 22509967]
- Ma WJ, Beck JM, Latham PE, Pouget A. Bayesian inference with probabilistic population codes. *Nature neuroscience*. 2006; 9(11):1432–1438. [PubMed: 17057707]
- Mitchell JF, Sundberg KA, Reynolds JH. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron*. 2009; 63(6):879–888. [PubMed: 19778515]
- Shadlen MN, Newsome WT. The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *The Journal of neuroscience*. 1998; 18(10):3870–3896. [PubMed: 9570816]
- Smith MA, Kohn A. Spatial and temporal scales of neuronal correlation in primary visual cortex. *The Journal of Neuroscience*. 2008; 28(48):12591–12603. [PubMed: 19036953]
- Smith MA, Sommer MA. Spatial and temporal scales of neuronal correlation in visual area v4. *The Journal of Neuroscience*. 2013; 33(12):5422–5432. [PubMed: 23516307]
- Snyder, DL.; Miller, MI. *Random Point Processes in Time and Space*. Springer Science & Business Media; 1991.
- Staude B, Rotter S, Grün S. Can spike coordination be differentiated from rate covariation?. *Neural Computation*. 2008; 20(8):1973–1999. [PubMed: 18336077]
- Steriade M, Buzsáki G. Parallel activation of thalamic and cortical neurons by brainstem and basal forebrain cholinergic systems. *Brain cholinergic systems*. 1990:3–62.
- Steriade M, Timofeev I, Grenier F. Natural waking and sleep states: a view from inside neocortical neurons. *Journal of neurophysiology*. 2001; 85(5):1969–1985. [PubMed: 11353014]
- Székely GJ, Rizzo ML. Brownian distance covariance. *The annals of applied statistics*. 2009; 3(4):1236–1265.
- Timofeev I, Grenier F, Steriade M. Disfacilitation and active inhibition in the neocortex during the natural sleep-wake cycle: an intracellular study. *Proceedings of the National Academy of Sciences*. 2001; 98(4):1924–1929.
- Tolhurst DJ, Movshon JA, Thompson ID. The dependence of response amplitude and variance of cat visual cortical neurones on stimulus contrast. *Experimental brain research*. 1981; 41(3-4):414–419. [PubMed: 7215502]
- Tomko GJ, Crapper DR. Neuronal variability: non-stationary responses to identical visual stimuli. *Brain research*. 1974; 79(3):405–418. [PubMed: 4422918]
- Ventura V, Cai C, Kass RE. Trial-to-trial variability and its effect on time-varying dependency between two neurons. *Journal of neurophysiology*. 2005; 94(4):2928–2939. [PubMed: 16160096]

Hierarchical model

Figure 1. Hierarchical model for correlation sources

First level: the firing rates (FR_{1r}, FR_{2r}) have correlation FRC. Second level: the spike counts (Y_{1r}, Y_{2r}) are jointly generated with rates (FR_{1r}, FR_{2r}) and have correlation SCC. The first level correlation (which is unobserved) describes the relationship between the two neurons, but it is contaminated in the second level with Poisson-like noise.

Figure 2. Analysis roadmap
We first use nonparametric estimators (Section 3.2) to explore the data and implement hypothesis tests to determine which pairs of neurons have $\Gamma = 0$ (jitter test, Section 3.3). We then fit appropriate bivariate parametric models (we impose $\Gamma = 0$ if $\Gamma = 0$ was obtained from the jitter test; Section 3.1) and check their goodness of fit and sensitivity (Sections 3.1.1 and 3.1.2).

Figure 3. **SCC and FRC** at the 12 experimental stimuli (drift direction degrees) for a neuron pair (macaque V4 area; Section 4). **(B) SCC versus FRC** for the macaque V4 data analyzed in Section 4. The SCC is an attenuated version of the FRC; both generally have the same sign.

Figure 4.

(A) Deviations of the spike counts from Poisson. Top: *p.m.f* of model (12) for $\phi = 0.8, 1, 1.2$; $\phi = 1$ corresponds to the Poisson *p.m.f*. Bottom: relative % difference between model (12) for $\phi = 0.8$ and 1.2 and the Poisson *p.m.f*. (B) Deviations of the firing rates from Lognormality.

Figure 5.

(A) Power of the jitter test to detect lagged synchrony within-trial effects. The test has full power when $\gamma > 1$, *approximately*. **(B) FDR-FNR control:** the actual FDR is below the nominal FDR β so the FDR procedure of Benjamini and Hochberg (1995) is reliable. The FNR is below 10% for all $\beta > 5\%$: we expect that the FDR testing procedure will fail to detect less than 10% of existing within-trial effects in data like the V4 data of Section 4.

Figure 6.

(A) The experimental task. The dashed circle indicates the aggregate receptive field (RF) of all the V4 neurons recorded from the 100-electrode array. After the animal fixated on the central blue dot for 100 ms, a drifting grating stimulus (one of 12 directions, block randomized) was shown for 1 second. The central dot and grating were then extinguished, and a blue target dot appeared at a random location positioned along a circle 8 degrees from fixation. The animal was rewarded for making a successful saccade to this target dot within 300 ms. **(B) The experimental stimulus** consisted of a full-contrast drifting sinusoidal grating with 12 equally spaced directions of drift (30 degrees). The spatial frequency (1.3 cpd) and temporal frequency (6.25 Hz) were fixed, and the grating was positioned and sized to cover the aggregate RF of the neurons.

Figure 7.

(A) The estimates of $\hat{\Gamma}_{plug}$ (eq.(15)) concentrate around zero: within-trial correlations have a weak effect on SCC. (B) The estimates of $\widehat{ATT}_{plug}$ (eq.(14)) concentrate around 0.35: the magnitude of SCC is about 35% that of FRC on average. (CD) SCC, FRC, ECC vs interneural distance and TCC obtained by local linear regressions of the sample estimates of SCC and ECC (eqs. (18), (19), and (24)), and parametric and nonparametric estimates of FRC (eqs.(9) and (16); solid and dashed curves respectively), along with 95% simultaneous confidence bands: all correlations decrease with interneural distance and increase with TCC. The SCC and FRC curves are similar in shape, but SCC is attenuated compared to FRC; ECC is close to zero, which suggests that within-trial correlations have a small impact on spike-counts in 1 s bins.

Figure 8. Joint generalized tuning curve (GTC; Appendix H) for a macaque V4 neuron pair
 Tolerance regions (10% to 90% levels) of the estimated joint bivariate lognormal distributions for the rates (X_{1r}, X_{2r}) in model (8) for an example pair of neurons at 12 stimuli. The stimulus number is written in the top right corner of each panel. This plot shows not only how the neurons' average firing rates vary with the stimulus, but also how differently the rates are correlated for different stimuli.

Figure 9.

(A) Spike count log variance vs log mean: *the spike counts are overdispersed.* **(B)**

Kullback-Leibler divergence test p-values. *Using $\alpha = 5\%$, the Poisson model is rejected 52% of the times, and the Poisson-lognormal model 5% of the times, which matches the chosen spurious detection rate: the second model provides a better fit to the spike counts.*

Figure 10.

Sensitivity of $\widehat{FRC}_{mle}$ to deviations of the spike counts from the PLN model (7), i.e. to deviations from $\phi = 1$ (*horizontal axis*), $\gamma = 0$ (*different colors*), and $b = 0$ (*thickness of the curves*) in model (11), in terms of mean and relative standard deviation (SD). The true value of FRC is 0.5. The mean and SD of $\widehat{FRC}_{mle}$ are not sensitive to changes in $b \in (0, 1)$, somewhat sensitive to changes in dispersion of the spike counts $\phi \in (0.8, 1.2)$, and very sensitive to changes in within-trial covariability.

Figure 11.

(A) Generalized tuning curves of two V4 neurons with 0.05, 0.15, . . . , 0.95 quantiles, estimated assuming the Poisson-lognormal model (7); the red curves are the standard tuning curves. **(B) GTCC vs TCC:** they can have different signs, and GTCC appears more attenuated than TCC. **(C) SCC, FRC, ECC vs GTCC** (local linear regression with 95% simultaneous confidence bands): all correlations increase with GTCC; the relationships are more pronounced than when plotted versus TCC in Figure 7.

Algorithm 1

ML estimator of DC-PLN (spike counts)

1. Obtain the profile MOM estimates of μ, Σ , given γ :

$$\hat{\sigma}_i^2(\gamma) = \log \left(1 + \frac{s_{Y_i}^2 - \bar{Y}_i}{(\bar{Y}_i - \gamma)^2} \right); \hat{\mu}_i(\gamma) = \log(\bar{Y}_i - \gamma) - \frac{\hat{\sigma}_i^2}{2}; \hat{\sigma}_{12}(\gamma) = \log \left(1 + \frac{s_{Y_1, Y_2} - \gamma}{(\bar{Y}_1 - \gamma)(\bar{Y}_2 - \gamma)} \right)$$

2. Find $\tilde{\gamma} = \arg \max_{\gamma} L(\hat{\mu}(\gamma), \hat{\Sigma}(\gamma), \gamma; y_{11}, \dots, y_{2n})$

3. Find $(\mu^*, \Sigma^*) \arg \max_{\mu, \Sigma} L(\mu, \Sigma, \tilde{\gamma}; y_{11}, \dots, y_{2n})$

4. Update $\tilde{\gamma} = \arg \max_{\gamma} L(\mu^*, \Sigma^*, \gamma; y_{11}, \dots, y_{2n})$

5. Iterate steps 3-4.

Algorithm 2

ML estimator of DC-PLN (spike trains)

-
1. Estimate $\tilde{\gamma}$ nonparametrically using eq.(13)
 2. Find $(\mu^*, \Sigma^*) = \arg \max_{\mu, \Sigma} L(\mu, \Sigma, \tilde{\gamma}; y_{11}, \dots, y_{2n})$
-

Table 1

Time (ms) to compute $G_{m,K}$ (eq.(13)) for a neuron pair for different values of n and m ($K = \lceil m\delta/T \rceil$). The implementation time of the jitter test based on $G_{m,K}$ (Section 3.3) with 100 bootstraps and 20-bin psth is indicated in parentheses.

			Number of trials n		
			δT	m	
$G_{m,K}$	0.01	50	0.3 (207)	0.3 (275)	0.4 (346)
		100	0.8 (322)	0.9 (426)	1.1 (526)
		200	2.1 (603)	2.5 (867)	3.4 (976)
	0.02	50	0.6 (235)	0.6 (289)	0.8 (357)
		100	1.2 (354)	1.3 (492)	1.8 (581)
		200	3.3 (744)	4.3 (914)	5.5 (1202)

Table 2

Computational time (ms) to estimate the parameters of the DC-PLN parametric model (8) for a neuron pair. Time for step 1 of Algorithm 2 can be found in Table 1.

	Number of trials n		
	60	80	100
Estimate μ, Σ assuming $\gamma = 0$	240	293	383
Estimate μ, Σ, γ (Algorithm 1)	4516	6364	7494
Estimate μ, Σ, γ (Algorithm 2, step 2)	2257	2769	3721