

Robust Inference of Neuronal Correlations from Blurred and Noisy Spiking Observations

Anuththara Rupasinghe

*Department of Electrical & Computer Engineering
The Institute for Systems Research
University of Maryland
College Park, MD, USA
raar@umd.edu*

Behtash Babadi

*Department of Electrical & Computer Engineering
The Institute for Systems Research
University of Maryland
College Park, MD, USA
behtash@umd.edu*

Abstract—Emerging large-scale neuronal recording technologies, such as two-photon calcium imaging, typically provide blurred and noisy surrogates of spiking activity. Extracting the underlying neuronal correlations, which are key to understanding neural function and circuitry, from such data is thus a challenging task. Though deconvolution techniques are often applied to such data to recover spiking activity, they require high temporal resolution and signal-to-noise ratio conditions to be effective. In addition, their solutions are biased towards obtaining accurate first-order statistics (i.e., spike detection) via spatiotemporal priors, which may be detrimental to recovering second-order statistics (i.e., correlations). Existing methods for inferring neuronal correlations from two-photon data thus suffer from significant bias and variability. In this work, we propose an algorithm to directly estimate neuronal correlations from ensemble two-photon imaging data, by integrating techniques from point process modeling and variational Bayesian inference, with no recourse to intermediate spike deconvolution. We demonstrate through simulation studies that the proposed method outperforms existing approaches in accurately capturing the underlying neuronal correlations.

Index Terms—Neuronal correlations, Variational inference, Point process models, Two-photon calcium imaging

I. INTRODUCTION

Recent advances in optical imaging have significantly enhanced neural data acquisition throughout by allowing to simultaneously record the activity of hundreds of neurons *in vivo* [1], [2]. These breakthroughs are hallmarked by two-photon calcium imaging, in which calcium ion concentrations resulting from spiking activity are recorded as fluorescence traces. In particular, these data have facilitated the investigation of functional network properties of neuronal ensembles. A popular characterization of these functional networks is given by neuronal correlations, which are crucial to understanding how populations of neurons encode information and how they interact to coordinate networked activity [3]–[5]. Extracting neuronal correlations is thus key to deciphering the computations carried out by neuronal populations.

Two-photon calcium imaging data, however, provide blurred and noisy measurements of neuronal spiking activity, due

to the slow time constant and variability of the underlying biochemical processes. In order to recover spiking activity from these data, spike deconvolution has become a popular and well-established area of research [6]–[12]. While existing results utilize spike deconvolution to further infer connectivity [13], [14] or Granger causal dynamics [15], [16] of the underlying networks, they require high temporal resolution and signal-to-noise ratio conditions to be effective in identifying the spike trains. In addition, their solutions are biased towards obtaining accurate first-order statistics (i.e., spike detection) via spatiotemporal priors, which may be detrimental to recovering second-order statistics (i.e., correlations).

Extracting the underlying correlations of spiking neurons directly from two-photon imaging data is thus a challenging task. In existing work [17]–[19], neuronal correlations are typically computed by directly averaging the two-photon imaging data across trials and computing Pearson correlations [5], with the purpose of identifying the key latent covariates that govern ensemble spiking activity. In this approach, the two-photon imaging data are assumed to be the direct measurements of spiking activity, which may result in highly biased and variable estimates of the neuronal correlations. Even if the ground truth spiking data were available, evaluating the correlations by direct averaging undermines the well-known nonlinearities that relate neural covariates to spike trains. Point processes and generalized linear models have been utilized to address this issue by accounting for the nonlinear effect of the latent processes that govern spiking activity [20], [21]. A unified framework for inferring the underlying correlations of these latent processes directly from two-photon imaging data is thus lacking.

In this paper, we close this gap by proposing a method to directly estimate neuronal correlations from high-dimensional calcium imaging data, by integrating techniques from point process modeling and Bayesian inference. We consider the intracellular calcium concentrations to be linearly related to the fluorescence observations, and characterize the exponential decay of calcium traces and their relationship to spiking activity by a first-order autoregressive model. Next, we model the spiking activity as a Bernoulli process that is related to a latent process through a logistic link. The correlations between

the different components of the latent process can then be recovered through parameter estimation. Considering the hierarchy of latent processes in the model, we develop an iterative parameter estimation method by integrating elements from variational inference [22], [23], Pólya-Gamma augmentation [24], iteratively re-weighted least square estimation [25], and fixed interval smoothing [26]. We demonstrate the utility of the proposed method using simulated calcium imaging data, which reveals significant gains over existing methodologies in terms of robustness.

II. PRELIMINARIES AND PROBLEM FORMULATION

Throughout the paper, we use upper-case bold-face, lower-case bold-face, lower-case letters to denote matrices, vectors and scalars, respectively. Furthermore, $(\mathbf{v})_i$ denotes the i^{th} element of a vector $\mathbf{v}$ and $(\mathbf{M})_{i,j}$ represents the $(i,j)^{\text{th}}$ element of a matrix $\mathbf{M}$.

We consider an observation duration of T frames of duration Δ each, corresponding to L independent trials, from J neurons. Thus, the fluorescence traces of the l^{th} trial, at the t^{th} time frame, $\mathbf{y}_t^{(l)}$ is a J -variate vector. We model $\mathbf{y}_t^{(l)}$ as a noisy linear function of the corresponding intracellular calcium concentration $\mathbf{z}_t^{(l)} \in \mathbb{R}^J$, and relate $\mathbf{z}_t^{(l)}$ to the spiking activity $\mathbf{n}_t^{(l)} \in \mathbb{R}^J$ by the first-order autoregressive model:

$$\begin{aligned} \mathbf{y}_t^{(l)} &= \mathbf{A} \mathbf{z}_t^{(l)} + \mathbf{w}_t^{(l)}, \quad \mathbf{w}_t^{(l)} \sim \mathcal{N}(\mathbf{0}, \Sigma_w), \\ \mathbf{z}_t^{(l)} &= \alpha \mathbf{z}_{t-1}^{(l)} + \mathbf{n}_t^{(l)}. \end{aligned} \quad (1)$$

We consider the time bin Δ to be small enough so that the probability of having two or more spikes within an interval of Δ is negligible [20]. The binary random variable $(\mathbf{n}_t^{(l)})_j$ in Eq. (1) indicates the spiking activity of the j^{th} neuron at time bin t , during the l^{th} trial. The Conditional Intensity Function [20] of a point process $N(t\Delta)$, with a spiking history H_t is defined as:

$$\lambda(t\Delta|H_t) := \lim_{\Delta \rightarrow 0} \frac{P[N(t\Delta + \Delta) - N(t\Delta) = 1|H_t]}{\Delta}.$$

Accordingly, the discretized point process can be modeled by a Bernoulli process with success probability $\lambda_t := \lambda(t\Delta|H_t)\Delta$. Following this point process framework, we model the binary process $(\mathbf{n}_t^{(l)})_j$, by a Bernoulli process with success probability $(\lambda_t)_j$. Further, we assume a logistic link between the success probability $\lambda_t \in \mathbb{R}^J$ and a latent process $\mathbf{x}_t \in \mathbb{R}^J$ governing spiking activity:

$$\begin{aligned} (\mathbf{n}_t^{(l)})_j &\sim \text{Bernoulli}((\lambda_t)_j), \\ (\lambda_t)_j &= \text{logistic}((\mathbf{x}_t)_j) \\ &= 1/(1 + \exp(-(\mathbf{x}_t)_j)), \\ \mathbf{x}_t &\sim \mathcal{N}(\boldsymbol{\mu}_x, \Sigma_x), \end{aligned}$$

where we assume $\mathbf{x}_t$ to be a Gaussian random vector with mean $\boldsymbol{\mu}_x$ and covariance Σ_x .

Based on this model, the covariance of the latent process Σ_x can be identified as the underlying neuronal covariance, which in turn can be used to compute the neuronal correlation matrix $\mathbf{N} \in \mathbb{R}^{J \times J}$ as pairwise Pearson correlations:

$$(\mathbf{N})_{i,j} = \frac{(\Sigma_x)_{i,j}}{\sqrt{(\Sigma_x)_{i,i}(\Sigma_x)_{j,j}}}. \quad (2)$$

Our goal is therefore to estimate Σ_x directly from the observed two-photon data $\mathbf{y} := \{\mathbf{y}_t^{(l)}\}_{t,l=1}^{T,L}$. To this end, we need to make additional assumptions to alleviate the ill-posed nature of this problem, which we will discuss in the forthcoming section.

III. PROPOSED ESTIMATION PROCEDURE

Given the temporal sparsity of spiking activity, suitable prior assumptions are necessary for robust parameter estimation. First, we assume an Inverse Wishart prior over Σ_x :

$$\Sigma_x \sim \text{InvWish}_J(\boldsymbol{\psi}_x, \rho_x),$$

which is indeed the conjugate prior in our model. Second, in order to simplify the exposition, we assume that the constants α , $\mathbf{A}$, Σ_w and $\boldsymbol{\mu}_x$ are either known or can be consistently estimated from training data.

Considering the complexity of the model and the hierarchy of latent variables, we propose a method based on Variational Inference [22], [23], for parameter estimation. Variational inference is widely used in Bayesian estimation in order to approximate complicated posterior densities via optimization, and can be thought of as an alternative strategy to Markov Chain Monte Carlo sampling [27].

Furthermore, we employ Pólya-Gamma latent variables [24] to decouple the logistic function, transforming the likelihood into an analytically convenient form. Following the Pólya-Gamma augmentation scheme [24], the complete data log-likelihood takes the form:

$$\begin{aligned} \log p(\mathbf{y}, \mathbf{z}, \mathbf{x}, \boldsymbol{\omega}, \Sigma_x) &= -\frac{1}{2} \left\{ (T + \rho_x + J + 1) \log(|\Sigma_x|) + \text{trace}(\boldsymbol{\psi}_x \Sigma_x^{-1}) \right\} \\ &+ \sum_{t=1}^T \left\{ \sum_{l=1}^L \left\{ -\frac{1}{2} (\mathbf{y}_t^{(l)} - \mathbf{A} \mathbf{z}_t^{(l)})^\top \Sigma_w^{-1} (\mathbf{y}_t^{(l)} - \mathbf{A} \mathbf{z}_t^{(l)}) \right. \right. \\ &+ \sum_{j=1}^J \left\{ \left((\mathbf{z}_t^{(l)})_j - \alpha (\mathbf{z}_{t-1}^{(l)})_j - 1/2 \right) (\mathbf{x}_t)_j - (\boldsymbol{\omega}_t)_j ((\mathbf{x}_t)_j)^2 / 2 \right. \\ &\left. \left. + \log p((\boldsymbol{\omega}_t)_j) \right\} \right\} - \frac{1}{2} (\mathbf{x}_t - \boldsymbol{\mu}_x)^\top \Sigma_x^{-1} (\mathbf{x}_t - \boldsymbol{\mu}_x) \left. \right\} + C. \end{aligned} \quad (3)$$

where $(\boldsymbol{\omega}_t)_j \sim \text{PG}(1, 0)$ for $j = 1, \dots, J$ and $t = 1, \dots, T$ and C account for terms not depending on $\mathbf{y}, \mathbf{z}, \mathbf{x}, \boldsymbol{\omega}$, and Σ_x . Next, we apply variational inference for inferring the random variables $\mathbf{x} = \{\mathbf{x}_t\}_{t=1}^T$, $\boldsymbol{\omega} = \{\boldsymbol{\omega}_t\}_{t=1}^T$ and Σ_x , under the mean field assumptions [23], resulting in the overall variational distribution:

$$q(\mathbf{x}, \boldsymbol{\omega}, \Sigma_x) = q(\Sigma_x) \prod_{t=1}^T \left(q(\mathbf{x}_t) \prod_{j=1}^J q((\boldsymbol{\omega}_t)_j) \right). \quad (4)$$

We employ the Coordinate Ascent Variational Inference algorithm [23], [28] to derive the optimal variational densities. Accordingly, we see that the optimal variational densities in Eq. (4) that maximize the log-likelihood in Eq. (3) take the forms:

$$\begin{aligned}
q^*(\mathbf{x}_t) &\sim \mathcal{N}(\mathbf{m}_{\mathbf{x}_t}, \mathbf{Q}_{\mathbf{x}_t}), \\
q^*((\omega_t)_j) &\sim \text{PG}(1, (\mathbf{c}_t)_j), \quad j = 1, 2, \dots, J, \\
q^*(\Sigma_x) &\sim \text{InvWish}_J(\mathbf{P}_x, \gamma_x).
\end{aligned}$$

The explicit expressions of the optimal variational parameters $\mathbf{m}_{\mathbf{x}_t}$, $\mathbf{Q}_{\mathbf{x}_t}$, $\mathbf{c}_t$, $\mathbf{P}_x$ and γ_x are outlined in Algorithm 1.

Note that even though $\mathbf{z} = \{\mathbf{z}_t^{(l)}\}_{t,l=1}^{T,L}$ is also an unknown variable in our model, we have not applied variational inference to $\mathbf{z}$. Given that $\mathbf{z}$ includes variables with temporal dependencies due to the underlying state-space model, imposing variational distributions under the mean field assumption is not straightforward. Thus, we propose an alternative strategy to estimate $\mathbf{z}$ using the derived variational distribution $q^*(\mathbf{x}, \omega, \Sigma_x)$, while treating $\mathbf{z}$ as an unknown parameter.

Note that the likelihood in Eq. (3) is decoupled in l , as a result of the independence of the realizations, for $l = 1, \dots, L$. Hence, we can derive independent updates for $\mathbf{z}^{(l)} = \{\mathbf{z}_t^{(l)}\}_{t=1}^T$ for $l = 1, \dots, L$. We propose to estimate $\mathbf{z}^{(l)}$ by,

$$\hat{\mathbf{z}}_t^{(l)} = \underset{\mathbf{z}_t^{(l)}}{\text{argmax}} \mathbb{E}_{q^*(\mathbf{x}, \omega, \Sigma_x)} [\log p(\mathbf{y}, \mathbf{z}, \mathbf{x}, \omega, \Sigma_x)], \quad (5)$$

under the constraints $0 \leq (\mathbf{z}_t^{(l)})_j - \alpha(\mathbf{z}_{t-1}^{(l)})_j \leq 1$, for $t = 1, \dots, T$ and $j = 1, \dots, J$. These constraints are a direct consequence of $(\mathbf{n}_t^{(l)})_j = (\mathbf{z}_t^{(l)})_j - \alpha(\mathbf{z}_{t-1}^{(l)})_j$ being a Bernoulli random variable with $\mathbb{E}[(\mathbf{n}_t^{(l)})_j] \in [0, 1]$.

However, this constrained optimization problem is intractable and solving for $\mathbf{z}_t^{(l)}$ directly from Eq. (5) is not straightforward. We thus consider an alternative unconstrained optimization problem by relaxing these constraints. We relax the constraint $\mathbf{z}_t^{(l)} - \alpha\mathbf{z}_{t-1}^{(l)} \leq \mathbf{1}$ and capture the effect of the constraint $\mathbf{z}_t^{(l)} - \alpha\mathbf{z}_{t-1}^{(l)} \geq \mathbf{0}$ by adding penalty terms proportional to $|(\mathbf{z}_t^{(l)})_j - \alpha(\mathbf{z}_{t-1}^{(l)})_j|$ to the cost function. Thus, the alternative problem can be formulated as:

$$\begin{aligned}
\underset{\{\mathbf{z}_t^{(l)}\}_{t=1}^T}{\text{minimize}} \quad & \sum_{t=1}^T \left\{ \frac{1}{2} (\mathbf{y}_t^{(l)} - \mathbf{A}\mathbf{z}_t^{(l)})^\top \Sigma_w^{-1} (\mathbf{y}_t^{(l)} - \mathbf{A}\mathbf{z}_t^{(l)}) \right. \\
& \left. + \sum_{j=1}^J (\nu_t)_j \left| (\mathbf{z}_t^{(l)})_j - \alpha(\mathbf{z}_{t-1}^{(l)})_j \right| \right\}, \quad (6)
\end{aligned}$$

where $\nu_t = \beta |\mathbf{m}_{\mathbf{x}_t}|$, with $\beta \geq 1$ being a hyper-parameter. Note that due to the temporal sparsity of neuronal spiking, we may assume $|\mathbf{m}_{\mathbf{x}_t}| = -(\mathbf{m}_{\mathbf{x}_t})$.

We employ a procedure similar to that in [10] to solve for the optimal $\mathbf{z}_t^{(l)}$ in Eq. (6), based on Iteratively Re-weighted Least Squares (IRLS) [25] and Fixed Interval Smoothing (FIS) [26] algorithms. Incorporating IRLS with ϵ -perturbation of the absolute value function as in [10], we see that the solution to Eq. (6) coincides with the FIS solution for the Gaussian state-space model:

$$\begin{aligned}
\mathbf{y}_t^{(l)} &= \mathbf{A}\mathbf{z}_t^{(l)} + \mathbf{w}_t^{(l)}, \quad \mathbf{w}_t^{(l)} \sim \mathcal{N}(0, \Sigma_w) \\
\mathbf{z}_t^{(l)} &= \alpha\mathbf{z}_{t-1}^{(l)} + \mathbf{v}_t^{(l)}, \quad \mathbf{v}_t^{(l)} \sim \mathcal{N}(0, \Sigma_{\mathbf{v}_t^{(l)}}) \quad (7)
\end{aligned}$$

where $\Sigma_{\mathbf{v}_t^{(l)}} \in \mathbb{R}^{J \times J}$ is a diagonal matrix with $(\Sigma_{\mathbf{v}_t^{(l)}})_{j,j} := \sqrt{((\hat{\mathbf{z}}_t^{(l)})_j - \alpha(\hat{\mathbf{z}}_{t-1}^{(l)})_j)^2 + \epsilon^2} / (\nu_t)_j$. Thus, we derive the current

updates for $\{\mathbf{z}_t^{(l)}\}_{t=1}^T$ by applying the FIS algorithm [26] to the model in Eq. (7). The overall iterative procedure of alternatively estimating the variational parameters and calcium concentrations is outlined in Algorithm 1. Following convergence, we derive the final estimate of the covariance Σ_x by the mean of the corresponding variational density:

$$\hat{\Sigma}_x = \mathbb{E}_{q^*(\Sigma_x)}[\Sigma_x] = \frac{\mathbf{P}_x}{\gamma_x - J - 1}.$$

Finally, the neuronal correlation matrix $\hat{\mathbf{N}}$ can be computed using Pearson correlations as in Eq. (2).

Algorithm 1 Proposed Iterative Procedure for Estimating Σ_x

Inputs: Ensemble of fluorescence measurements $\{\mathbf{y}_t^{(l)}\}_{t,l=1}^{T,L}$, Known variables $\alpha, \mathbf{A}, \Sigma_w$ and μ_x , hyper-parameters ψ_x, ρ_x, β and ϵ , tolerance at convergence δ

Output: Estimate of the covariance matrix, $\hat{\Sigma}_x$
Initialization: Initial choice of $\Sigma_{\mathbf{v}_t^{(l)}}$, $\hat{\Omega}_t$, $\hat{\Sigma}_x$ and $\hat{\Sigma}_x^{-1}$, $\text{res} = 10\delta$, $\gamma_x = \rho_x + T$

```

1: while res  $\geq \delta$  do
    Estimate calcium concentrations using FIS
2:   for  $l = 1, \dots, L$  do
    Forward filter:
3:     for  $t = 1, \dots, T$  do
4:        $\mathbf{z}_{t|t-1}^{(l)} = \alpha\mathbf{z}_{t-1|t-1}^{(l)}$ 
5:        $\mathbf{P}_{t|t-1}^{(l)} = \alpha^2\mathbf{P}_{t-1|t-1}^{(l)} + \Sigma_{\mathbf{v}_t^{(l)}}$ 
6:        $\mathbf{B}_t^{(l)} = \mathbf{P}_{t|t-1}^{(l)}\mathbf{A}^\top(\mathbf{A}\mathbf{P}_{t|t-1}^{(l)}\mathbf{A}^\top + \Sigma_w)^{-1}$ 
7:        $\mathbf{z}_{t|t}^{(l)} = \mathbf{z}_{t|t-1}^{(l)} + \mathbf{B}_t^{(l)}(\mathbf{y}_t^{(l)} - \mathbf{A}\mathbf{z}_{t|t-1}^{(l)})$ 
8:        $\mathbf{P}_{t|t}^{(l)} = (\mathbf{I} - \mathbf{B}_t^{(l)}\mathbf{A})\mathbf{P}_{t|t-1}^{(l)}$ 
9:     end for
    Backward smoother:
10:    for  $t = T-1, \dots, 1$  do
11:       $\hat{\mathbf{z}}_t^{(l)} = \mathbf{z}_{t|t}^{(l)} + \alpha\mathbf{P}_{t|t}^{(l)}\mathbf{P}_{t+1|t}^{(l)-1}(\hat{\mathbf{z}}_{t+1}^{(l)} - \mathbf{z}_{t+1|t}^{(l)})$ 
12:    end for
13:  end for

  Update variational parameters
14:  for  $t = 1, \dots, T$  do
15:     $\mathbf{Q}_{\mathbf{x}_t} = (L\hat{\Omega}_t + \hat{\Sigma}_x^{-1})^{-1}$ 
16:     $\mathbf{m}_{\mathbf{x}_t} = \mathbf{Q}_{\mathbf{x}_t} \left( \hat{\Sigma}_x^{-1}\mu_x + \sum_{l=1}^L \left\{ \hat{\mathbf{z}}_t^{(l)} - \alpha\hat{\mathbf{z}}_{t-1}^{(l)} - \frac{1}{2}\mathbf{1} \right\} \right)$ 
17:     $\nu_t = \beta |\mathbf{m}_{\mathbf{x}_t}|$ 
18:    for  $j = 1, \dots, J$  do
19:       $(\mathbf{c}_t)_j = \sqrt{(\mathbf{Q}_{\mathbf{x}_t})_{j,j} + ((\mathbf{m}_{\mathbf{x}_t})_j)^2}$ 
20:       $(\hat{\Omega}_t)_{j,j} = \frac{1}{2(\mathbf{c}_t)_j} \tanh\left(\frac{(\mathbf{c}_t)_j}{2}\right)$ 
21:    end for
22:  end for
23:   $\mathbf{P}_x = \psi_x + \sum_{t=1}^T \{ \mathbf{Q}_{\mathbf{x}_t} + \mathbf{m}_{\mathbf{x}_t}\mathbf{m}_{\mathbf{x}_t}^\top - \mu_x\mathbf{m}_{\mathbf{x}_t}^\top - \mathbf{m}_{\mathbf{x}_t}\mu_x^\top + \mu_x\mu_x^\top \}$ 
24:   $\hat{\Sigma}_x^{-1} = \gamma_x\mathbf{P}_x^{-1}$ 

  Update IRLS covariance approximation
25:  for  $l = 1, \dots, L, t = 1, \dots, T$ , and  $j = 1, \dots, J$  do
26:     $(\Sigma_{\mathbf{v}_t^{(l)}})_{j,j} = \frac{\sqrt{((\hat{\mathbf{z}}_t^{(l)})_j - \alpha(\hat{\mathbf{z}}_{t-1}^{(l)})_j)^2 + \epsilon^2}}{(\nu_t)_j}$ 
27:  end for

  Update the convergence criterion and output
28:   $(\hat{\Sigma}_x)_{\text{prev}} = \hat{\Sigma}_x, \hat{\Sigma}_x = \frac{\mathbf{P}_x}{\gamma_x - J - 1}$ 
29:   $\text{res} = \|(\hat{\Sigma}_x)_{\text{prev}} - \hat{\Sigma}_x\|_2 / \|(\hat{\Sigma}_x)_{\text{prev}}\|_2$ 
30: end while
31: Return  $\hat{\Sigma}_x$ 

```

IV. SIMULATION RESULTS

A. Simulated Two-Photon Imaging Data

We consider a time duration of $T = 10^5$ frames, $J = 10$ neurons and $L = 10$ independent trials per neuron for simulation purposes. We set $\alpha = 0.98$, $\mathbf{A} = 0.1 \times \mathbf{I}$, $\boldsymbol{\mu}_x = -5.6 \times \mathbf{1}$ and $\boldsymbol{\Sigma}_w = 10^{-4} \times \mathbf{I}$ ($\mathbf{I} \in \mathbb{R}^{10 \times 10}$ is the identity matrix and $\mathbf{1} \in \mathbb{R}^{10}$ is the vector of all ones) when generating the fluorescence traces $\{\mathbf{y}_t^{(l)}\}_{t,l=1}^{T,L}$, so that the SNR of simulated data is in the same range as that of experimentally-recorded data. We simulate the spike trains based on a Poisson process [29] using the discrete time re-scaling procedure [29], [30]. Following the assumptions of [30], we use an exponential link to get:

$$(\boldsymbol{\lambda}_t)_j = \exp((\mathbf{x}_t)_j), \quad (\mathbf{n}_t^{(l)})_j \sim \text{Poisson}((\boldsymbol{\lambda}_t)_j).$$

Note that the data are simulated using a different model than that used in our inference framework (i.e., Bernoulli process with a logistic link), in order not to bias the performance in favor of our proposed method. As an example, Fig. 1 shows the simulated fluorescence trace $((\mathbf{y}_t^{(1)})_1)$, estimated calcium concentrations $((\hat{\mathbf{z}}_t^{(1)})_1)$ and the estimated spike train $((\hat{\mathbf{n}}_t^{(1)})_1 = (\hat{\mathbf{z}}_t^{(1)})_1 - \alpha(\hat{\mathbf{z}}_{t-1}^{(1)})_1)$ corresponding to the first trial of the first neuron.




Fig. 1. Top: simulated fluorescence trace from the first neuron during the first trial, $(\mathbf{y}_t^{(1)})_1$. Bottom: estimates of the calcium signal $(\hat{\mathbf{z}}_t^{(1)})_1$ (blue) and spike train $(\hat{\mathbf{n}}_t^{(1)})_1$ (red) using the proposed method.

B. Performance Evaluation and Comparison

We compare the performance of our proposed estimation framework with three other techniques, benchmarked by the true correlation matrix used in simulating the data:

1) *Oracle Estimate*: Given the actual random process $\mathbf{x}_t$ that was used to simulate the data, we can directly compute the Pearson correlations from its sample covariance, which we refer to as the Oracle Estimate, as if an oracle would provide the true latent process that underlies neuronal spiking activity.

2) *FCSS Estimate*: In this method, we first obtain the Fast Compressible State-Space (FCSS) estimates of the spike trains from the simulated fluorescence traces using the deconvolution developed in [10]. We then smooth the estimated spike trains with a Gaussian kernel to obtain a continuous process, followed by computing the empirical covariances of the smoothed FCSS estimates for $l = 1, \dots, L$. Then, we

compute the average covariance across trials, from which the Pearson correlations are obtained.

3) *Direct Estimate*: The direct estimate is based on the conventional methods for estimating neuronal correlations [4], [5]. As in [17]–[19], we compute the neuronal correlations by first evaluating the empirical covariances of the two-photon observations $\{\mathbf{y}_t^{(l)}\}_t^T$ for $l = 1, \dots, L$, and then computing the Pearson correlations of the average covariance across trials.

Fig. 2 shows the ground truth correlation matrix along with the Oracle, proposed, FCSS, and direct estimates. For the sake of comparison, we have normalized the correlation matrices and have set the diagonal elements to zero in Fig. 2. It can be observed from Fig. 2 that our proposed estimate (Fig. 2C) closely resembles the Oracle estimate (Fig. 2B) and the ground truth (Fig. 2A), while the direct (Fig. 2E) and FCSS (Fig. 2D) estimates exhibit multiple spurious correlations and thus provide a highly biased characterization of the underlying correlation structure of the latent process.

We further quantify these observations by comparing each estimate to the ground truth using the disparity metric:

$$D_{\text{frob}}(\mathbf{X}, \mathbf{Y}) = \|\mathbf{X} - \mathbf{Y}\|_F,$$

where $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{J \times J}$ and $\|\cdot\|_F$ is the Frobenius matrix norm. To this end, for each estimate $\hat{\mathbf{N}}$, we evaluate $D_{\text{frob}}(\hat{\mathbf{N}}, \mathbf{N})$, to quantify its similarity to the ground truth $\mathbf{N}$. Table I summarizes these comparisons. As expected, the Oracle estimate is the closest to the ground truth. Among the estimates that use the fluorescence traces, our proposed estimate exhibits the closest performance to the Oracle estimate. In accordance with the foregoing visual comparisons from Fig. 1, the performance of the FCSS and direct estimates is far from the Oracle estimate, due to their significant biases.

TABLE I
PERFORMANCE COMPARISON

Estimation method	$D_{\text{frob}}(\hat{\mathbf{N}}, \mathbf{N})$
Oracle Estimate ($\hat{\mathbf{N}}_{\text{Oracle}}$)	0.0666
Proposed Estimate ($\hat{\mathbf{N}}$)	0.7535
FCSS Estimate ($\hat{\mathbf{N}}_{\text{FCSS}}$)	1.9887
Direct Estimate ($\hat{\mathbf{N}}_{\text{direct}}$)	2.3182

V. CONCLUSIONS

The advent of two-photon calcium imaging has paved the way to simultaneous data acquisition from large-scale neuronal ensembles. Extracting neuronal correlations from these data is key to understanding the functional properties of neuronal populations. This problem, however, is challenging due to the blurred and noisy nature of two-photon measurements in comparison to spike recordings. Existing methods either treat two-photon recordings as the underlying latent processes that govern spiking activity, or are based on spike deconvolution techniques, both of which result in a biased characterization of the underlying neuronal correlations. In this work, we addressed this challenge by developing a variational inference




Fig. 2. Ground truth and estimated neuronal correlations: (A) Ground truth correlations ($\mathbf{N}$), (B) Oracle estimate ($\hat{\mathbf{N}}_{\text{Oracle}}$), (C) Proposed estimate ($\hat{\mathbf{N}}$), (D) FCSS estimate ($\hat{\mathbf{N}}_{\text{FCSS}}$) and (E) Direct estimate ($\hat{\mathbf{N}}_{\text{direct}}$).

framework to extract neuronal correlations directly from two-photon fluorescence observations. Through simulation studies, we demonstrated that the proposed method accurately characterizes the neuronal correlations that govern spiking activity, while significantly outperforming several existing methods. Future work includes application of this methodology to experimentally-recorded two-photon data, as well as accounting for the effect of external stimuli.

REFERENCES

- [1] C. Stosiek, O. Garaschuk, K. Holthoff, and A. Konnerth, “In vivo two-photon calcium imaging of neuronal networks,” *Proceedings of the National Academy of Sciences*, vol. 100, no. 12, pp. 7319–7324, 2003.
- [2] K. Svoboda and R. Yasuda, “Principles of two-photon excitation microscopy and its applications to neuroscience,” *Neuron*, vol. 50, no. 6, pp. 823–839, 2006.
- [3] B. B. Averbeck, P. E. Latham, and A. Pouget, “Neural correlations, population coding and computation,” *Nature Reviews Neuroscience*, vol. 7, no. 5, pp. 358–366, 2006.
- [4] M. Cohen and A. Kohn, “Measuring and interpreting neuronal correlations,” *Nature neuroscience*, vol. 14, pp. 811–9, 06 2011.
- [5] D. R. Lyamzin, S. J. Barnes, R. Donato, J. A. Garcia-Lazaro, T. Keck, and N. A. Lesica, “Nonlinear transfer of signal and noise correlations in cortical networks,” *Journal of Neuroscience*, vol. 35, no. 21, pp. 8065–8080, 2015.
- [6] J. T. Vogelstein, A. M. Packer, T. A. Machado, T. Sippy, B. Babadi, R. Yuste, and L. Paninski, “Fast nonnegative deconvolution for spike train inference from population calcium imaging,” *Journal of Neurophysiology*, vol. 104, no. 6, pp. 3691–3704, 2010.
- [7] T. Deneux, A. Kaszas, G. Szalay, G. Katona, T. Lakner, A. Grinvald, B. Rózsa, and I. Vanzetta, “Accurate spike estimation from noisy calcium signals for ultrafast three-dimensional imaging of large neuronal populations in vivo,” *Nature Communications*, vol. 7, pp. 12 190 EP –, Jul 2016, article.
- [8] E. Pnevmatikakis, D. Soudry, Y. Gao, T. A. Machado, J. Merel, D. Pfau, T. Reardon, Y. Mu, C. Lacefield, W. Yang, M. Ahrens, R. Bruno, T. M. Jessell, D. Peterka, R. Yuste, and L. Paninski, “Simultaneous denoising, deconvolution, and demixing of calcium imaging data,” *Neuron*, vol. 89, no. 2, pp. 285–299, 2016.
- [9] J. Friedrich, P. Zhou, and L. Paninski, “Fast online deconvolution of calcium imaging data,” *PLOS Computational Biology*, vol. 13, no. 3, pp. 1–26, 03 2017.
- [10] A. Kazemipour, J. Liu, K. Solarana, D. A. Nagode, P. O. Kanold, M. Wu, and B. Babadi, “Fast and stable signal deconvolution via compressible state-space models,” *IEEE Transactions on Biomedical Engineering*, vol. 65, no. 1, pp. 74–86, 2018.
- [11] M. Pachitariu, C. Stringer, and K. D. Harris, “Robustness of spike deconvolution for neuronal calcium imaging,” *Journal of Neuroscience*, vol. 38, no. 37, pp. 7976–7985, 2018.
- [12] S. W. Jewell, T. D. Hocking, P. Fearnhead, and D. M. Witten, “Fast nonconvex deconvolution of calcium imaging data,” *Biostatistics*, 02 2019.
- [13] Y. Mishchenko, J. Vogelstein, and L. Paninski, “A Bayesian approach for inferring neuronal connectivity from calcium fluorescent imaging data,” *Annals of Applied Statistics*, vol. 5, 07 2011.
- [14] D. Soudry, S. Keshri, P. Stinson, M.-h. Oh, G. Iyengar, and L. Paninski, “Efficient ‘shotgun’ inference of neural connectivity from highly subsampled activity data,” *PLOS Computational Biology*, vol. 11, no. 10, pp. 1–30, 10 2015.
- [15] F. De Vico Fallani, M. Corazzoli, J. R. Sternberg, C. Wyart, and M. Chavez, “Hierarchy of neural organization in the embryonic spinal cord: Granger-causality graph analysis of in vivo calcium imaging data,” *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, vol. 23, no. 3, pp. 333–341, 2015.
- [16] A. Sheikhattar, S. Miran, J. Liu, J. B. Fritz, S. A. Shamma, P. O. Kanold, and B. Babadi, “Extracting neuronal functional network dynamics via adaptive Granger causality analysis,” *Proceedings of the National Academy of Sciences*, vol. 115, no. 17, pp. E3869–E3878, 2018.
- [17] G. Rothschild, I. Nelken, and A. Mizrahi, “Functional organization and population dynamics in the mouse primary auditory cortex,” *Nature Neuroscience*, vol. 13, pp. 353 EP –, Jan 2010.
- [18] D. E. Winkowski, S. Bandyopadhyay, S. A. Shamma, and P. O. Kanold, “Frontal cortex activation causes rapid plasticity of auditory cortical processing,” *Journal of Neuroscience*, vol. 33, no. 46, pp. 18 134–18 148, 2013.
- [19] N. Francis, D. Winkowski, A. Sheikhattar, K. Armengol, B. Babadi, and P. Kanold, “Small networks encode decision-making in primary auditory cortex,” *Neuron*, vol. 97, 02 2018.
- [20] W. Truccolo, U. T. Eden, M. R. Fellows, J. P. Donoghue, and E. N. Brown, “A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects,” *Journal of Neurophysiology*, vol. 93, no. 2, pp. 1074–1089, 2005.
- [21] L. Paninski, “Maximum likelihood estimation of cascade point-process neural encoding models,” *Network: Computation in Neural Systems*, vol. 15, no. 4, pp. 243–262, 2004.
- [22] M. I. Jordan, Z. Ghahramani, T. S. Jaakkola, and L. K. Saul, “An introduction to variational methods for graphical models,” *Machine Learning*, vol. 37, no. 2, pp. 183–233, Nov 1999.
- [23] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *Journal of the American Statistical Association*, vol. 112, no. 518, pp. 859–877, 2017.
- [24] N. G. Polson, J. G. Scott, and J. Windle, “Bayesian inference for logistic models using Pólya-Gamma latent variables,” *Journal of the American Statistical Association*, vol. 108, no. 504, pp. 1339–1349, 2013.
- [25] D. Ba, B. Babadi, P. L. Purdon, and E. N. Brown, “Convergence and stability of iteratively re-weighted least squares algorithms,” *IEEE Transactions on Signal Processing*, vol. 62, no. 1, pp. 183–195, 2014.
- [26] H. E. Rauch, C. T. Striebel, and F. Tung, “Maximum likelihood estimates of linear dynamic systems,” *AIAA Journal*, vol. 3, no. 8, pp. 1445–1450, aug 1965.
- [27] W. K. Hastings, “Monte Carlo sampling methods using Markov chains and their applications,” *Biometrika*, vol. 57, no. 1, pp. 97–109, 04 1970.
- [28] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Berlin, Heidelberg: Springer-Verlag, 2006.
- [29] A. C. Smith and E. N. Brown, “Estimating a state-space model from point process observations,” *Neural Comput.*, vol. 15, no. 5, pp. 965–991, May 2003.
- [30] E. N. Brown, R. Barbieri, V. Ventura, R. E. Kass, and L. M. Frank, “The time-rescaling theorem and its application to neural spike train data analysis,” *Neural Comput.*, vol. 14, no. 2, pp. 325–346, Feb. 2002.