

Granger Causal Inference from Spiking Observations via Latent Variable Modeling

Sahar Khosravi, Anuththara Rupasinghe, and Behtash Babadi

Department of Electrical & Computer Engineering

Institute for Systems Research

University of Maryland, College Park, MD

E-mails: saharkh@umd.edu, raar@umd.edu, behtash@umd.edu

Abstract—Extracting directional connectivity in a neuronal ensemble from spiking observations is a key challenge in understanding the circuit mechanisms of brain function. Existing methods proceed in two stages, by first estimating the latent processes that govern spiking, followed by characterizing connectivity using said estimates. As such, the extracted networks in the second stage are highly sensitive to the accuracy of the estimates in the first stage. In this work, we introduce a framework to directly extract Granger causal links from spiking observations, without requiring intermediate time-domain estimation, by explicitly modeling the endogenous and exogenous latent processes that underlie spiking activity. Our proposed method integrates several techniques such as point processes, state-space modeling and Pólya-Gamma augmentation. We demonstrate the utility of our proposed approach using simulated data and application to real data from the rat brain during sleep.

I. INTRODUCTION

Recent advances in neural data acquisition techniques such as high-density arrays and two-photon calcium imaging have paved the way to recording the simultaneous activity of large neuronal populations [1], [2]. Extracting the networks underlying such population activity is key to studying the mechanisms of cognitive and sensory processing [3]–[8]. Granger causality is a well-established method for characterizing such directional connectivity: considering two processes X_t and Y_t , if the knowledge of process X_t , significantly improves the prediction of Y_{t+1} we say there is a Granger causal (GC) link from X to Y [9]. This notion of directed connectivity has been widely used in the analysis of neuroimaging data as well as spiking neuronal ensembles [8], [10], [11].

While point process modeling [11], [12] provides a principled framework to analyze spiking data, existing methods for inferring GC networks from spiking observations have several shortcomings. First, commonly used approaches proceed in two stages, where the underlying latent processes driving spiking activity are first estimated from the spiking data, followed by GC inference from said estimates [13]–[15]. As a result, biases incurred in the estimation of the latent processes propagate to the subsequent GC inference stage. While recent results use history-dependent point process models to directly extract GC links from spiking observations

[8], [16], the neuronal dynamics in these models are limited to the binary-valued history of the ensemble and may not fully account for the rich range of endogenous (i.e., spontaneous) and exogenous (i.e., stimulus-driven) processes that drive spiking activity. On the other hand, latent variable models that are used to capture neural dynamics from spiking activity, consider instantaneous correlations and do not account for delayed interactions between the neurons [17], [18]. Thus, a unified modeling and estimation framework to extract GC links directly from spiking data by explicitly modeling the effect of exogenous and endogenous processes is lacking.

We here address these issues and propose a methodology to directly extract GC links from spiking data using an inverse problem shown in Fig. 1. We integrate point processes and multivariate autoregressive modeling to relate the spiking observations to the underlying network, while explicitly taking into account the spontaneous and stimulus-driven activity. We combine several methods such as Pólya-Gamma augmentation [19] and Expectation-Maximization [20] to directly solve the inverse problem and recover the underlying GC connectivity. Furthermore, we assess the significance of the detected GC links using Youden’s J-statistics [8], [21]. We compare the performance of our method to existing approaches using a simulation study, which reveals significant improvements in hit and false alarm rates. Finally, we apply our method to experimentally-recorded data from the rat cortical neurons during sleep to identify network-level changes across different sleep stages.

II. METHODS

A. Problem Formulation

Suppose that we observe the spiking activity of a population of neurons across multiple repeated trials. Let $n_{t,l}^{(j)}$ denote the spiking of the j^{th} neuron at the t^{th} time bin and l^{th} trial. We assume that spiking observations are Bernoulli realizations primarily driven by two independent covariates, an endogenous process $x_{t,l}^{(j)}$ accounting for spontaneous activity and an exogenous component capturing the activity driven by a stimulus $s_{t,l}$. We assume that these covariates change over a slower time scale than the observed spiking, i.e., spiking observations within a small window of length W are driven by the same processes. Accordingly, we have:

$$n_{t,l}^{(j)} \sim \text{Bernoulli} \left(\text{logistic} \left(\mu_l^{(j)} + x_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l} \right) \right) \quad (1)$$

This work has been supported in part by the National Science Foundation awards no. 1552946, 1807216, and 2032649, and the National Institutes of Health award no. 1U19NS107464-01.




Fig. 1: A schematic depiction of the generative model and the inverse problem. The spiking activity of networked neurons in an ensemble is driven by latent endogenous processes (red traces) as well as an exogenous stimulus (green traces). Black arrows show the GC connectivity. Our contribution is to solve the inverse problem by estimating the underlying GC connectivity directly from the spiking observations.

where $k = \text{ceil}(\frac{t}{W})$, $\text{logistic}(\cdot) = \frac{\exp(\cdot)}{1 + \exp(\cdot)}$, $\mathbf{d}_j$ is the receptive field of neuron j encoding the effect of stimulus $\mathbf{s}_{k,l}$, and $\mu_l^{(j)}$ is the baseline activity.

To investigate the underlying GC connectivity, we model the latent variables $x_{k,l}^{(1:J)}$ as a multivariate autoregressive (AR) process. Defining $\mathbf{x}_{k,l} := [x_{k,l}^{(1)}, \dots, x_{k,l}^{(J)}]^\top \in \mathbb{R}^J$, we have:

$$\mathbf{x}_{k,l} = \sum_{p=1}^P \mathbf{A}_p \mathbf{x}_{k-p,l} + \mathbf{w}_{k,l}, \quad \mathbf{w}_{k,l} \sim \mathcal{N}(\mathbf{0}, \Sigma) \quad (2)$$

where $\mathbf{A}_p \in \mathbb{R}^{J \times J}$ represents the AR coefficients of the p^{th} time lag and the noise covariance $\Sigma \in \mathbb{R}^{J \times J}$ is a diagonal matrix with the j^{th} diagonal entry given by σ_j^2 . Further, we regularize the AR coefficients using the ℓ_2 norm, to mitigate over-fitting in the regime of large P .

Given spiking observations of J neurons over a time period T (i.e., $K = \text{ceil}(\frac{T}{W})$ windows) and L trials, $n_{1:T,1:L}^{(1:J)}$, we aim at extracting the GC links among the neurons in the population. To this end, we first estimate the AR coefficients and the state covariance matrix without any restrictions, namely the *full* model. Then, to assess the GC link from neuron i to j , we remove the contribution of neuron i to neuron j in the AR model and re-estimate the AR coefficients and state covariance matrix to get the *reduced* model parameters. If including $x_{k,l}^{(i)}$ improves the prediction of $x_{k,l}^{(j)}$, the *full* model is preferred over the *reduced* model, and we say that there exists a GC link from neuron i to j . Thus, by assessing the significance of the difference in the log-likelihoods of these two models, we can test for the existence of a GC link.

B. Parameter Estimation via Expectation-Maximization

To estimate the unknown parameters $\theta := \{\mathbf{A}_p, \Sigma, \mathbf{d}_j, |p \in \{1, 2, \dots, P\}, j \in \{1, 2, \dots, J\}\}$ we use Expectation-Maximization (EM) [20]. To illustrate this procedure, first consider the full model. The joint log-likelihood of the observations, latent processes and parameters can be written as:

$$\begin{aligned} \log p(n_{1:T,1:L}^{(1:J)}, x_{1:K,1:L}^{(1:J)}, \theta) \\ = \sum_{j,k,l,w=1}^{J,K,L,W} \left(n_{(k-1)W+w,l}^{(j)} (\mu_l^{(j)} + x_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l}) \right. \\ \left. - \log(1 + \exp(\mu_l^{(j)} + x_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l})) \right) \\ - \sum_{k,l=1}^{K,L} \frac{1}{2} \left(\mathbf{x}_{k,l} - \sum_{p=1}^P \mathbf{A}_p \mathbf{x}_{k-p,l} \right)^\top \Sigma^{-1} \left(\mathbf{x}_{k,l} - \sum_{p=1}^P \mathbf{A}_p \mathbf{x}_{k-p,l} \right) \\ - \frac{KL}{2} \log |\Sigma| - \lambda \sum_{p=1}^P \|\mathbf{A}_p\|_2^2 + C, \end{aligned} \quad (3)$$

where C represents terms that are not functions of $n_{1:T,1:L}^{(1:J)}$, $x_{1:K,1:L}^{(1:J)}$ or θ , and λ represents the ℓ_2 -regularization parameter.

1) *E-Step*: We need to compute the so-called Q-function given by:

$$Q(\theta | \hat{\theta}^{(r)}) := \mathbb{E} \left[\log(p(n_{1:T,1:L}^{(1:J)}, x_{1:K,1:L}^{(1:J)})) | n_{1:T,1:L}^{(1:J)}, \hat{\theta}^{(r)} \right],$$

where $\hat{\theta}^{(r)}$ is the parameter estimates at iteration r . This requires evaluation of $\mathbb{E}[\mathbf{x}_{k,l} | n_{1:T,1:L}^{(1:J)}, \hat{\theta}^{(r)}]$ and $\mathbb{E}[\mathbf{x}_{k,l} \mathbf{x}_{k-p,l}^\top | n_{1:T,1:L}^{(1:J)}, \hat{\theta}^{(r)}]$, for $p = 1, \dots, P$. Using a standard state augmentation, we define:

$$\begin{aligned} \begin{bmatrix} \mathbf{x}_{k,l} \\ \mathbf{x}_{k-1,l} \\ \vdots \\ \mathbf{x}_{k-(P-1),l} \end{bmatrix} &= \mathbf{A} \begin{bmatrix} \mathbf{x}_{k-1,l} \\ \mathbf{x}_{k-2,l} \\ \vdots \\ \mathbf{x}_{k-P,l} \end{bmatrix} + \begin{bmatrix} \mathbf{w}_{k,l} \\ \mathbf{0}_{(P-1)J \times 1} \end{bmatrix}, \\ \underbrace{\begin{bmatrix} \mathbf{x}_{k,l} \\ \mathbf{x}_{k-1,l} \\ \vdots \\ \mathbf{x}_{k-(P-1),l} \end{bmatrix}}_{\tilde{\mathbf{x}}_{k,l}} &= \underbrace{\mathbf{A} \begin{bmatrix} \mathbf{x}_{k-1,l} \\ \mathbf{x}_{k-2,l} \\ \vdots \\ \mathbf{x}_{k-P,l} \end{bmatrix}}_{\tilde{\mathbf{x}}_{k-1,l}} + \underbrace{\begin{bmatrix} \mathbf{w}_{k,l} \\ \mathbf{0}_{(P-1)J \times 1} \end{bmatrix}}_{\tilde{\mathbf{w}}_{k,l}}, \\ \mathbf{A} &= \begin{bmatrix} \mathbf{A}_1 & \mathbf{A}_2 & \dots & \mathbf{A}_P \\ & \mathbf{I}_{(P-1)J \times (P-1)J} & & \mathbf{0}_{(P-1)J \times J} \end{bmatrix}, \end{aligned}$$

with $\tilde{\mathbf{w}}_{k,l} \sim \mathcal{N}(\mathbf{0}, \tilde{\Sigma})$ where

$$\tilde{\Sigma} = \begin{bmatrix} \Sigma & \mathbf{0}_{J \times (P-1)J} \\ \mathbf{0}_{(P-1)J \times J} & \mathbf{0}_{(P-1)J \times (P-1)J} \end{bmatrix}.$$

The complete log-likelihood of Eq. (3) can be expressed as:

$$\begin{aligned} \sum_{j,k,l,w=1}^{J,K,L,W} \left(n_{(k-1)W+w,l}^{(j)} (\mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l}) \right. \\ \left. - \log(1 + \exp(\mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l})) \right) \\ - \sum_{k,l=1}^{K,L} \frac{1}{2} (\tilde{\mathbf{x}}_{k,l} - \mathbf{A} \tilde{\mathbf{x}}_{k-1,l})^\top \tilde{\Sigma}^{-1} (\tilde{\mathbf{x}}_{k,l} - \mathbf{A} \tilde{\mathbf{x}}_{k-1,l}) \\ - \frac{KL}{2} \log(|\tilde{\Sigma}|) - \lambda \sum_{p=1}^P \|\mathbf{A}_p\|_2^2 + C. \end{aligned} \quad (4)$$

Given that the Bernoulli log-likelihood is not quadratic in $\mathbf{x}_{k,l}$, computing the expectation of Eq. (4) is intractable. To address this issue, we use Pólya-Gamma (PG) augmentation, which leverages the following identity [17], [19]:

$$a\psi - \log(1 + \exp(\psi)) = \left(a - \frac{1}{2}\right)\psi - \log 2 + \int_0^\infty \left(-\frac{\omega}{2}\psi^2 + \log p_{\text{PG}(1,0)}(\omega)\right) d\omega, \quad (5)$$

where $p_{\text{PG}(1,0)}(\omega)$ is the PG density. Comparing the first term in Eq. (4) with the left hand side of Eq. (5), we see that by setting $a := n_{(k-1)W+w,l}^{(j)}$, $\psi := \mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l}$, the PG augmented log-likelihood can be expressed as:

$$\begin{aligned} & \log \left(p \left(n_{1:T,1:L}^{(1:J)}, x_{1:K,1:L}^{(1:J)}, \omega_{1:K,1:L}^{(1:J)}, \boldsymbol{\theta} \right) \right) \\ &= \sum_{j,k,l,w=1}^{J,K,L,W} \left(\left(n_{(k-1)W+w,l}^{(j)} - \frac{1}{2} \right) \left(\mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l} \right) \right. \\ & \quad \left. - \frac{\omega_{k,l}^{(j)}}{2} \left(\mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l} \right)^2 + \log P_{\text{PG}(1,0)} \left(\omega_{k,l}^{(j)} \right) \right) \\ & \quad - \sum_{k,l=1}^{K,L} \frac{1}{2} (\tilde{\mathbf{x}}_{k,l} - \mathbf{A} \tilde{\mathbf{x}}_{k-1,l})^\top \tilde{\boldsymbol{\Sigma}}^{-1} (\tilde{\mathbf{x}}_{k,l} - \mathbf{A} \tilde{\mathbf{x}}_{k-1,l}) \\ & \quad - \frac{KL}{2} \log(|\boldsymbol{\Sigma}|) - \lambda \sum_{p=1}^P \|\mathbf{A}_p\|_2^2 + C \end{aligned} \quad (6)$$

introducing PG distributed i.i.d. latent random variable $\omega_{k,l}^{(j)} \sim \text{PG}(1,0)$, for $1 \leq k \leq K$, $1 \leq l \leq L$ and $1 \leq j \leq J$. Note that the augmented log-likelihood in Eq. (6) is quadratic in $\tilde{\mathbf{x}}_{k,l}$, hence the conditional distribution $p(\tilde{\mathbf{x}}_{k,l} | n_{1:T,1:L}^{(1:J)}, \hat{\omega}_{1:K,1:L}^{(1:J)}, \hat{\boldsymbol{\theta}}^{(r)})$ is indeed Gaussian. Thus, we can use fixed interval smoothing [22] and covariance smoothing [23] to efficiently compute the conditional expectations of $\tilde{\mathbf{x}}_{k,l}$, $\tilde{\mathbf{x}}_{k,l} \tilde{\mathbf{x}}_{k,l}^\top$, and $\tilde{\mathbf{x}}_{k-1,l} \tilde{\mathbf{x}}_{k,l}^\top$ given $n_{1:T,1:L}^{(1:J)}, \hat{\omega}_{1:K,1:L}^{(1:J)}, \hat{\boldsymbol{\theta}}^{(r)}$.

From the augmented log-likelihood in Eq. (6), we note that the conditional distribution $p(\omega_{k,l}^{(j)} | x_{1:K,1:L}^{(1:J)}, \boldsymbol{\theta}) \sim \text{PG}(1, c_{k,l}^{(j)})$, where $c_{k,l}^{(j)} = |\mu_l^{(j)} + \tilde{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l}|$. Thus, we estimate $\hat{\omega}_{k,l}^{(j)} = \mathbb{E}[\omega_{k,l}^{(j)} | x_{1:K,1:L}^{(1:J)}, \hat{\boldsymbol{\theta}}^{(r)}] = \frac{1}{2c_{k,l}^{(j)}} \tanh\left(\frac{c_{k,l}^{(j)}}{2}\right)$. To increase the accuracy of latent variable estimation, we iterate between estimating the moments of $\tilde{\mathbf{x}}_{k,l}$ and $\omega_{k,l}^{(j)}$, before advancing to the M-step.

2) *M-Step*: In the M-step we update the values of $\mathbf{A}$, $\boldsymbol{\Sigma}$ and $\mathbf{D}$ maximizing the Q-function. The update rules for the j th row of $\mathbf{A}$, $\boldsymbol{\Sigma}$ and $\mathbf{d}_j$ are given by:

$$\begin{aligned} \hat{\mathbf{a}}_j^{(r+1)} &= \left(\sum_{k,l}^{K,L} \mathbb{E} \left[\mathbf{x}_{k,l} \tilde{\mathbf{x}}_{k-1,l}^\top \right] \right) \left(\sum_{k,l}^{K,L} \mathbb{E} \left[\tilde{\mathbf{x}}_{k-1,l} \tilde{\mathbf{x}}_{k-1,l}^\top \right] + \lambda \sigma_j^2 \mathbf{I} \right)^{-1}, \\ \hat{\boldsymbol{\Sigma}}^{(r+1)} &= \frac{1}{KL} \left(\sum_{k,l}^{K,L} \mathbb{E} \left[\tilde{\mathbf{x}}_{k,l} \tilde{\mathbf{x}}_{k,l}^\top \right] - \mathbb{E} \left[\tilde{\mathbf{x}}_{k,l} \tilde{\mathbf{x}}_{k-1,l}^\top \right] \mathbf{A}^\top \right. \\ & \quad \left. - \mathbf{A} \mathbb{E} \left[\tilde{\mathbf{x}}_{k-1,l} \tilde{\mathbf{x}}_{k,l}^\top \right] + \mathbf{A} \mathbb{E} \left[\tilde{\mathbf{x}}_{k-1,l} \tilde{\mathbf{x}}_{k-1,l}^\top \right] \mathbf{A}^\top \right), \end{aligned}$$

$$\begin{aligned} \hat{\mathbf{d}}_j^{(r+1)} &= \left[W \sum_{k,l}^{K,L} \hat{\omega}_{k,l}^{(j)} \mathbf{s}_{k,l} \mathbf{s}_{k,l}^\top \right]^{-1} \left[\sum_{k,l,w}^{K,L,W} \left(n_{(k-1)W+w,l}^{(j)} - \frac{1}{2} \right) \mathbf{s}_{k,l} \right. \\ & \quad \left. - \mu_l^{(j)} \hat{\omega}_{k,l}^{(j)} \mathbf{s}_{k,l} - \hat{\omega}_{k,l}^{(j)} \mathbf{s}_{k,l} \mathbb{E} \left[x_{k,l}^{(j)} \right] \right], \end{aligned}$$

where we have dropped the conditioning of the expectations by $n_{1:T,1:L}^{(1:J)}, \hat{\omega}_{1:K,1:L}^{(1:J)}, \hat{\boldsymbol{\theta}}^{(r)}$ for notational convenience. To check for the convergence of the EM algorithm, we compare the normalized relative error of the estimated AR coefficients and the state covariance matrix with a pre-specified threshold (here, we used a threshold of 10^{-2}).

We follow the same EM procedure when estimating the parameters of the reduced model, with slight modifications to the update rules. Note that in the reduced model, we remove the contributions of the source neuron to the activity of the target neuron in Eq. (2), by setting the corresponding AR coefficients to zero. In the reduced model for testing the GC link from neuron i to neuron m , the update rules for $\hat{\mathbf{a}}_j$ are the same as before, for $j \neq m$. However, for $j = m$, we use the reduced vector $\tilde{\mathbf{x}}_{k-1,l} \in \mathbb{R}^{(J-1)P}$ by removing the elements of $\tilde{\mathbf{x}}_{k-1,l}$ corresponding to neuron i , and then update $\hat{\mathbf{a}}_m^{(r+1)} \in \mathbb{R}^{(J-1)P}$ as

$$\hat{\mathbf{a}}_m^{(r+1)} = \left(\sum_{k,l}^{K,L} \mathbb{E} \left[\mathbf{x}_{k,l}^{(m)} (\tilde{\mathbf{x}}_{k-1,l}^{\text{red}})^\top \right] \right) \left(\sum_{k,l}^{K,L} \mathbb{E} \left[\tilde{\mathbf{x}}_{k-1,l}^{\text{red}} (\tilde{\mathbf{x}}_{k-1,l}^{\text{red}})^\top \right] + \lambda \sigma_m^2 \mathbf{I} \right)^{-1},$$

and set the coefficients corresponding to neuron i to zero. Similarly, when deriving the reduced model to test for the GC effect of stimulus to neuron m , we set $\mathbf{d}_m = \mathbf{0}$.

C. Statistical Testing of Granger Causal Links

To assess the significance of the GC links, we use the deviance difference between the reduced and full models [8], [24] defined as:

$$D^{(i \mapsto j)} = 2(\ell^f - \ell^r), \quad (7)$$

where ℓ^f and ℓ^r denote the log-likelihood of the observations in the full and reduced models, respectively. Since a direct evaluation of the observation log-likelihood is intractable, we simplify it using Bayes' theorem and the sampling scheme of [25] as follows:

$$\begin{aligned} P \left(n_{1:T,1:L}^{(1:J)} \right) &= \frac{p(n_{1:T,1:L}^{(1:J)}, \tilde{\mathbf{x}}_{1:K,1:L})}{p(\tilde{\mathbf{x}}_{1:K,1:L} | n_{1:T,1:L}^{(1:J)})} \\ &= \frac{\prod_{j,t,l}^{J,T,L} P(n_{t,l}^{(j)} | \tilde{\mathbf{x}}_{1:K,1:L}) \prod_{k=2,L}^{K,L} p(\tilde{\mathbf{x}}_{k,l} | \tilde{\mathbf{x}}_{k-1,l}) p(\tilde{\mathbf{x}}_{1,l})}{\prod_{l=1}^L p(\tilde{\mathbf{x}}_{K,l} | n_{1:T,1:L}^{(1:J)}) \prod_{k,l=1}^{K-1,L} p(\tilde{\mathbf{x}}_{k,l} | \tilde{\mathbf{x}}_{k+1:K,l}, n_{1:T,1:L}^{(1:J)})}. \end{aligned} \quad (8)$$

While the numerator in Eq. (8) has a closed-form expression, the denominator does not, due to the coupling of $\mathbf{x}_{k,l}$ and the Bernoulli observation model. Thus, we use the Laplace approximation to simplify the denominator in Eq. (8). Specifically, for $k = 1, \dots, K$ and $l = 1, \dots, L$ we assume:

$$p(\tilde{\mathbf{x}}_{k,l} | \tilde{\mathbf{x}}_{k+1:K,l}, n_{1:T,1:L}^{(1:J)}) \sim \mathcal{N}(\hat{\mathbf{x}}_{k|K,l}, \mathbf{P}_{k|K,l}),$$

where $\hat{\mathbf{x}}_{k|K,l}$ and $\mathbf{P}_{k|K,l}$ satisfy:

$$\hat{\mathbf{x}}_{k|K,l} = (\mathbf{I} - \mathbf{C}_{k+1,l} \hat{\mathbf{A}}) \hat{\mathbf{x}}_{k|k,l} + \mathbf{C}_{k+1,l} \tilde{\mathbf{x}}_{k+1,l},$$

$$\mathbf{P}_{k|K,l} = (\mathbf{I} - \mathbf{C}_{k+1,l} \hat{\mathbf{A}}) \mathbf{P}_{k|k,l},$$

$$\mathbf{C}_{k+1,l} = \mathbf{P}_{k|k,l} \hat{\mathbf{A}}^\top (\hat{\mathbf{A}} \mathbf{P}_{k|k,l} \hat{\mathbf{A}}^\top + \hat{\Sigma})^{-1},$$

with $\tilde{\mathbf{x}}_{k|k,l}$ and $\mathbf{P}_{k|k,l}$ being the outputs of the Kalman filter in the E-step, respectively given by $\mathbb{E}[\tilde{\mathbf{x}}_{k,l} | n_{1:kW,l}^{(1:J)}, \hat{\boldsymbol{\theta}}]$ and $\mathbb{E}(\tilde{\mathbf{x}}_{k,l} - \tilde{\mathbf{x}}_{k|k,l})(\tilde{\mathbf{x}}_{k,l} - \tilde{\mathbf{x}}_{k|k,l})^\top | n_{1:kW,l}^{(1:J)}, \hat{\boldsymbol{\theta}}]$.

We use the fact that the deviance difference $D^{(i \mapsto j)} \sim \chi^2(M)$ in the absence of a GC link and $D^{(i \mapsto j)} \sim \chi^2(M, \nu^{(i \mapsto j)})$ in the presence of a GC link, where $\chi^2(M)$ and $\chi^2(M, \nu^{(i \mapsto j)})$ are respectively central and non-central chi-square distributions with M degrees of freedom and non-centrality parameter $\nu^{(i \mapsto j)}$ [8]. The degrees of freedom M is equal to the number of parameters removed in the reduced model. In the analysis of GC links between the neurons $M = P$, and assessing the GC effect of the stimulus on a neuron, M is equal to the length of the receptive field vector $\mathbf{d}_j$. The non-centrality parameter can be estimated as $\hat{\nu}^{(i \mapsto j)} = \max\{D^{(i \mapsto j)} - M, 0\}$ [8].

To summarize the strength of a GC link at a significance level α , we use the J-statistic for the link $i \mapsto j$ defined as:

$$J^{(i \mapsto j)} = 1 - \alpha - \eta^{(i \mapsto j)}(\alpha), \quad (9)$$

where $\eta^{(i \mapsto j)}(\alpha) := \mathcal{F}_{\chi^2(M, \hat{\nu}^{(i \mapsto j)})}(\mathcal{F}_{\chi^2(M)}^{-1}(1 - \alpha))$ with $\mathcal{F}_{\chi^2(M, \hat{\nu}^{(i \mapsto j)})}(\cdot)$ and $\mathcal{F}_{\chi^2(M)}(\cdot)$ denoting the cumulative distribution functions of the non-central and central chi-square distributions with M degrees of freedom and non-centrality $\hat{\nu}^{(i \mapsto j)}$. The J-statistic summarizes the overall performance of a diagnostic test by accounting for both type-I and type-II errors. When $J^{(i \mapsto j)} \approx 0$ there is no strong evidence to reject the null hypothesis, whereas when $J^{(i \mapsto j)} \approx 1$, there is strong evidence that the alternative hypothesis is true. Thus $J^{(i \mapsto j)} \in [0, 1]$ is an indicator of the strength of the detected GC link. Finally, to account for the multiple comparisons performed to detect the GC links in a network, we use the Benjamini-Hochberg [26] (BH) procedure to correct the false discovery rate (FDR).

III. RESULTS

We demonstrate the utility of our proposed method using a simulation study and application to real data. In the simulations, we compare our method with two existing approaches:

1) *Two-Stage AR Estimation Method*: This method first estimates the conditional intensity of the spiking activity through a logistic link as follows:

$$n_{t,l}^{(j)} \sim \text{Bernoulli} \left(\text{logistic} \left(y_{k,l}^{(j)} \right) \right),$$

where $y_{k,l}^{(j)}$ is the logit of the conditional intensity and is modeled as an AR(1) process to enforce temporal smoothness [27]. Then, the GC links are assessed by fitting a multivariate AR model using Eq. (2) and assuming a linear model $y_{k,l}^{(j)} = \mathbf{x}_{k,l}^{(j)} + \mathbf{d}_j^\top \mathbf{s}_{k,l} + v_{k,l}^{(j)}$, where $v_{k,l}^{(j)} \sim \mathcal{N}(0, \sigma_v^2)$.

2) *Two-stage PSTH Method*: This method is a simplification of the Two-Stage AR Method, where the latent process $y_{k,l}^{(j)}$ is estimated by the peri-stimulus time histogram (PSTH) of the spiking activity $\hat{y}_{k,l}^{(j)} := \frac{1}{W} \sum_{w=1}^W n_{(k-1)W+w,l}^{(j)}$

A. Simulation Study

We consider $J = 5$ neurons and a stimulus given by an AR(2) process. There are 4 GC links as shown in Fig. 2-A, where red (resp. blue) arrows show facilitative (resp. suppressive) GC links, reflecting the positive (resp. negative) sign of the aggregate AR coefficients or receptive fields. The spiking activity was simulated over bin sizes of 10 ms over $L = 3$ trials for total duration of $T = 30$ s. The stimulus and the endogenous AR process (with $P = 2$ lags) were generated over 100 ms sampling intervals, corresponding to $W = 10$. Fig. 2-B shows the ground truth GC links, as well as those estimated by our proposed method and the two-stage approaches. While our proposed method accurately detects all the true GC links, with only one false detection, the two-stage methods significantly deviate from the ground truth. Also, our proposed method is the only one that correctly identifies the GC link from the stimulus neuron 4. To further quantify this observation, we repeated this experiment 10 times and computed the average hit and false alarm rates as shown in Fig. 2-C (FDR controlled at 10^{-4}). Our method exhibits a hit rate close to 1, while also resulting in a lower false alarm rate as compared to the other methods.




Fig. 2: Result of the simulation study. (A) Schematic depiction of the simulated network with 5 neurons and an external stimulus, (B) GC links, from left to right: ground truth, proposed method, Two-stage AR method, and Two-stage PSTH method. (C) Detection performance of the three methods, with FDR controlled at 10^{-4} .

B. Application to Experimentally-Recorded Data

Finally, we applied our proposed method to experimentally recorded data from the rat's brain during the nonREM and REM phases of sleep (dataset from [28], publicly available in the Collaborative Research in Computational Neuroscience data sharing website [29]). This dataset includes the spiking activity of $J = 2$ neural populations: putative pyramidal cells (pE) and putative interneurons (pI). Since there is no external stimulus, we set $\mathbf{s}_{k,l} = \mathbf{0}$. We considered two neurons from each population and assumed their spike trains to be independent trials $L = 2$, segmented the data into 30 s episodes sampled at $f_s = 50$ Hz and considered a stationary window of length $W = 10$ samples. We assessed the GC links between the pE and pI populations over 47 nonREM and 29 REM episodes, and compared the changes in connectivity




Fig. 3: Application to experimentally-recorded data from the rat brain during sleep. The GC links from pE to pI populations significantly decrease from nonREM to REM stage (t-test, $*p < 0.05$).

across the two sleep stages as shown in Fig. 3. We found that the average number of GC links from pE to pI neurons significantly decreased during REM as compared to nonREM sleep, while there was no significant change in the GC links from pI to pE neurons.

IV. CONCLUSION

In this work, we developed a framework to infer Granger causal links from spiking activity of a neuronal population using a latent variable model accounting for the endogenous and exogenous effects. By explicitly modeling the latent endogenous processes that govern spontaneous activity and the exogenous effect of external stimuli on the observed spiking data, we developed a direct inference framework to identify the GC connectivity in a neuronal ensemble. We used multivariate AR modeling, Pólya-Gamma augmentation, and point process modeling to address the shortcomings of existing methods. Using simulated data, we demonstrated the superior performance of our method compared to existing approaches. Application of our proposed method to real data from the rat brain during sleep revealed changes in the connectivity of neuronal populations across different sleep stages.

REFERENCES

- [1] F. L. e. a. Kashkoush A.I., Gaunt R.A., "Recording single- and multi-unit neuronal action potentials from the surface of the dorsal root ganglion," *Sci Rep*, vol. 9, 2019.
- [2] V. V. e. a. Akemann W., Wolf S., "Fast optical recording of neuronal activity by three-dimensional custom-access serial holography," *Nature Meth*, vol. 19, pp. 100–120, 2022.
- [3] F. Najafi, G. F. Elsayed, R. Cao, E. Pnevmatikakis, P. E. Latham, J. P. Cunningham, and A. K. Churchland, "Excitatory and inhibitory subnetworks are equally selective during decision-making and emerge simultaneously during learning," *Neuron*, vol. 105, no. 1, pp. 165–179.e8, 2020.
- [4] "Functional clustering of dendritic activity during decision-making," *Elife*, vol. 8, no. e46966, 2019 Oct 30.
- [5] R. N. Ramesh, C. R. Burgess, A. U. Sugden, M. Gyetvan, and M. L. Andermann, "Intermingled ensembles in visual association cortex encode stimulus identity or predicted outcome," *Neuron*, vol. 100, no. 4, pp. 900–915.e9, 2018.
- [6] 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.
- [7] D. Yatsenko, K. Josi, A. S. Ecker, E. Froudarakis, R. J. Cotton, and A. S. Tolias, "Improved estimation and interpretation of correlations in neural circuits," *PLOS Computational Biology*, vol. 11, no. 3, pp. 1–28, 03 2015.

- [8] 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.
- [9] S. L. Bressler and A. K. Seth, "Wienergranger causality: A well established methodology," *NeuroImage*, vol. 58, no. 2, pp. 323–329, 2011.
- [10] A. K. Seth, A. B. Barrett, and L. Barnett, "Granger causality analysis in neuroscience and neuroimaging," *Journal of Neuroscience*, vol. 35, no. 8, pp. 3293–3297, 2015.
- [11] S. Kim, D. Putrino, S. Ghosh, and E. N. Brown, "A granger causality measure for point process models of ensemble neural spiking activity," *PLOS Computational Biology*, vol. 7, no. 3, pp. 1–13, 03 2011.
- [12] B. E. Smith AC, "Estimating a state-space model from point process observations," *Neural Comput*, vol. (5), pp. 965–91, May 15 2003.
- [13] K. Sameshima and L. A. Baccalá, "Using partial directed coherence to describe neuronal ensemble interactions," *Journal of neuroscience methods*, vol. 94, no. 1, pp. 93–103, 1999.
- [14] A. G. Nedungadi, G. Rangarajan, N. Jain, and M. Ding, "Analyzing multiple spike trains with nonparametric granger causality," *Journal of computational neuroscience*, vol. 27, no. 1, pp. 55–64, 2009.
- [15] M. Krumin and S. Shoham, "Multivariate autoregressive modeling and granger causality analysis of multiple spike trains," *Computational intelligence and neuroscience*, vol. 2010, 2010.
- [16] M. Eichler, R. Dahlhaus, and J. Dueck, "Graphical modeling for multivariate hawkes processes with nonparametric link functions," *Journal of Time Series Analysis*, vol. 38, no. 2, pp. 225–242, 2017.
- [17] A. Rupasinghe, N. Francis, J. Liu, Z. Bowen, P. O. Kanold, and B. Babadi, "Direct extraction of signal and noise correlations from two-photon calcium imaging of ensemble neuronal activity," *eLife*, vol. 10, p. e68046, jun 2021.
- [18] L. Aitchison, L. Russell, A. M. Packer, J. Yan, P. Castonguay, M. Hausser, and S. C. Turaga, "Model-based bayesian inference of neural activity and connectivity from all-optical interrogation of a neural circuit," *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [19] N. G. Polson, J. G. Scott, and J. Windle, "Bayesian inference for logistic models using polya-gamma latent variables," 2013.
- [20] R. H. Shumway and D. S. Stoffer, "An approach to time series smoothing and forecasting using the EM algorithm," *Journal of time series analysis*, vol. 3, no. 4, pp. 253–264, 1982.
- [21] W. Youden, "Index for rating diagnostic tests," in *Cancer*, vol. 3, 1950, pp. 32–35.
- [22] 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.
- [23] P. D. Jong and M. J. Mackinnon, "Covariances for smoothed estimates in state space models," *Biometrika*, vol. 75, no. 3, pp. 601–602, 1988.
- [24] B. Soleimani, P. Das, J. Kulasingham, J. Z. Simon, and B. Babadi, "Granger causal inference from indirect low-dimensional measurements with application to meg functional connectivity analysis," in *2020 54th Annual Conference on Information Sciences and Systems (CISS)*, 2020, pp. 1–5.
- [25] S. Frühwirth-Schnatter, "Data augmentation and dynamic linear models," Department of Statistics and Mathematics, WU Vienna University of Economics and Business, Vienna, Forschungsberichte / Institut für Statistik 28, 1992. [Online]. Available: <https://epub.wu.ac.at/392/>
- [26] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: A practical and powerful approach to multiple testing," *Journal of the Royal Statistical Society. Series B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [27] A. C. Smith and E. N. Brown, "Estimating a state-space model from point process observations," *Neural Computation*, vol. 15, no. 5, pp. 965–991, 2003.
- [28] B. O. Watson, D. Levenstein, J. P. Greene, J. N. Gelinas, and G. Buzsáki, "Network homeostasis and state dynamics of neocortical sleep," *Neuron*, vol. 90, no. 4, pp. 839 – 852, 2016.
- [29] B. O. Watson, D. Levenstein, J. P. Greene, J. N. Gelinas, and G. Buzsáki, "Multi-unit spiking activity recorded from rat frontal cortex (brain regions mPFC, OFC, ACC, and M2) during wake-sleep episode wherein at least 7 minutes of wake are followed by 20 minutes of sleep," *CRCNS.org*, 2016. [Online]. Available: <http://dx.doi.org/10.6080/K02N506Q>