Functional Connectivity: Continuous-Time Latent Factor Models for Neural Spike Trains

Meixi Chen (meixi.chen@uwaterloo.ca)
Department of Statistics and Actuarial Science, University of Waterloo, 200 University Ave. W., Waterloo, ON N2L 3G1, Canada

Martin Lysy (mlysy@uwaterloo.ca)
Department of Statistics and Actuarial Science, University of Waterloo, 200 University Ave. W., Waterloo, ON N2L 3G1, Canada

David Moorman (moorman@umass.edu)
Department of Psychological and Brain Sciences, University of Massachusetts Amherst, 639 N. Pleasant St., Amherst, Massachusetts 01003-9298, USA

Reza Ramezan (rramezan@uwaterloo.ca) – corresponding author
Department of Statistics and Actuarial Science, University of Waterloo, 200 University Ave. W., Waterloo, ON N2L 3G1, Canada

Abstract

Modelling the dynamics of interactions in a neuronal ensemble is an important problem in functional connectivity research. One popular framework is latent factor models (LFMs), which have achieved notable success in decoding neuronal population dynamics. However, most LFMs are specified in discrete time, where the choice of bin size significantly impacts inference results. In this work, we present what is, to the best of our knowledge, the first continuous-time multivariate spike train LFM for studying neuronal interactions and functional connectivity. We present an efficient parameter inference algorithm for our biologically justifiable model which (1) scales linearly in the number of simultaneously recorded neurons and (2) bypasses time binning and related issues. Simulation studies show that parameter estimation using the proposed model is highly accurate. Applying our LFM to experimental data from a classical conditioning study on the prefrontal cortex in rats, we found that coordinated neuronal activities are affected by (1) the onset of the cue for reward delivery, and (2) the sub-region within the frontal cortex (OFC/mPFC). These findings shed new light on our understanding of cue and outcome value encoding.

Keywords: multivariate point processes; spike trains; latent factor models; functional connectivity; neural correlation, classical conditioning

Introduction

An important question in neuroscience research is understanding the functional connectivity between neurons in different parts of the brain. Spike trains based on simultaneously recorded neurons provide information about population coding and neuronal interaction. Both model-free and model-based spike-train analysis tools have been developed to answer this question. While model-free methods (Perkel et al., 1967; Ventura et al., 2005; Fujisawa et al., 2008; Humphries, 2011; Lopes-dos-Santos et al., 2013) are typically more efficient and convenient to implement, they often fail to uncover more complex underlying neuronal relationships beyond correlation at the level of observed data. In contrast, latent factor models (LFMs) are able to discover patterns which model-free algorithms cannot, thanks to their ability to specify different structures in the latent layers. The remarkable success of spike-train LFMs in predictive tasks (Yu et al., 2008; Gao et al., 2016; Wu et al., 2017; Pandarinath et al., 2018) and neuronal clustering (Buesing et al., 2014; Wei et al., 2022) motivates our work in this paper to study the functional connectivity between neurons. As all the aforementioned LFMs are specified in discrete time, they are commonly applied upon binning the experimental time to obtain spike counts, which are then modelled using, e.g., a Poisson likelihood. However, the bin size is often chosen arbitrarily– despite this having a significant impact on parameter estimation (Nelson, 2002; Kass & Ventura, 2006; Ramezan et al., 2014). Although there are Poisson process models set in continuous time without binning (e.g., Duncker & Sahani, 2018; Williams et al., 2020), these models cannot be readily applied to the functional connectivity analysis of multiple neurons.

In this work we present a continuous-time multivariate point process LFM to study neuronal interactions based on simultaneously recorded spike trains in a neural population. To the best of our knowledge, this is the first continuous-time LFM proposed for the analysis of neural spike trains. In our model, the activities of a neuronal population are described by correlated Wiener processes with resetting. Each of these processes is viewed as a proxy of the evolving membrane potential of a neuron which resets after reaching a threshold. Crucially, we assume that the high-dimensional multivariate latent process can be summarized by a small number of dynamic factors. Not only does this factor analysis framework provide an interpretable low-dimensional representation of neural activities, it also serves as a means of dimension reduction for studying large neuronal populations. Our model generalizes the limiting case of the multivariate Skellam point process with resetting of Ramezan et al. (2022). However, by passing to the Brownian limit, we are able to develop an efficient algorithm for parameter inference, which both circumvents the choice of bin size (it can be made so small as to approximate the continuous time process arbitrarily well) and scales linearly in the number of neurons in the analysis. The applicability of the proposed model and inference procedure is demonstrated in simulated and experimental data analyses.

Model

A graphical summary of the proposed model is presented in Figure 1. Let $\mathbf{Y}_t = (y_{1,t}, \ldots, y_{q,t}) \in \{0,1\}^q$ denote the observed binary spike trains of a population of q neurons at time t, and $\mathbf{X}_t = (x_{1,t}, \ldots, x_{q,t}) \in \mathbb{R}^q$ be the unobserved latent dynamic processes of the neuronal population. We assume that the neuronal activities are governed by $\mathbf{X}_t$, which are modelled as correlated Wiener processes with resetting. Each process i is reset to its initial value whenever it crosses some neuron-specific threshold, and then a spike occurs, i.e., $y_{i,t} = 1$ is observed. We let $\boldsymbol{\theta}$ denote the threshold parameters and the Wiener process drift vector. The threshold-crossing and resetting processes mimic the spike-generating mechanism and the refractoriness of the neuron, respectively. We further assume that the q dynamic processes can be represented as $\mathbf{X}_t = \boldsymbol{\Lambda}\mathbf{F}_t + \boldsymbol{\varepsilon}_t$, where $\mathbf{F}_t = (f_{1,t}, \ldots, f_{d,t})$ are d $(d \ll q)$ dynamic factors, and $\boldsymbol{\varepsilon}_t = (\varepsilon_{1,t}, \ldots, \varepsilon_{q,t})$ are neuron-specific idiosyncrasies. The loading matrix $\boldsymbol{\Lambda}$ can be interpreted within the traditional factor analysis framework– that is, it identifies a small number of factors driving the neuronal dynamic processes. It is also used to model the correlation matrix between latent neuron processes, i.e., $\mathrm{cor}(\mathbf{X}_t) = \boldsymbol{\Sigma} = \boldsymbol{\Lambda}\boldsymbol{\Lambda}' + \boldsymbol{\Psi}$, where $\boldsymbol{\Psi}$ is a diagonal matrix with elements determined by $\boldsymbol{\Lambda}$ and the fact that $\boldsymbol{\Sigma}$ has unit diagonal.

Inference

Model inference is carried out in two steps. First, the drift of the Wiener process and the threshold parameters, sum-

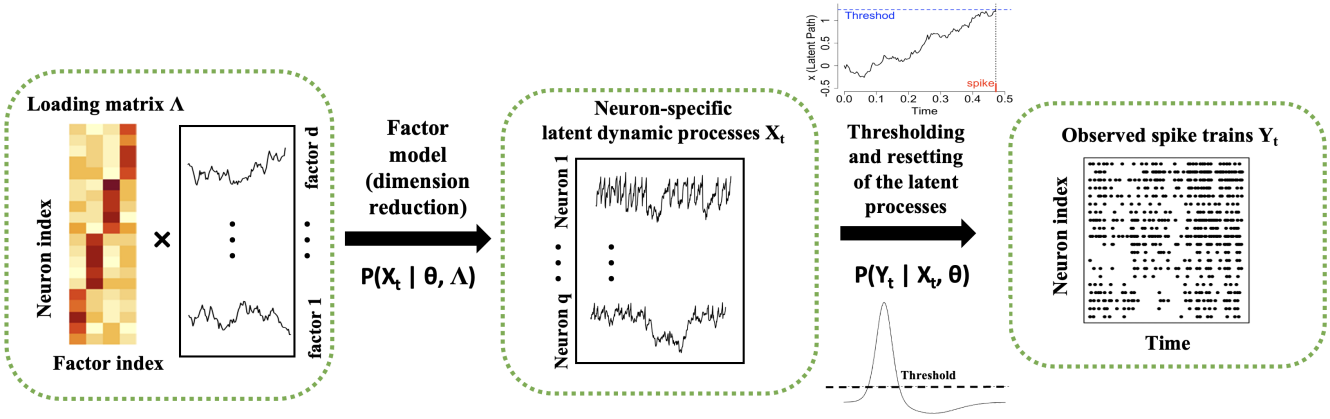

Figure 1: Illustration of the proposed model.

marized by $\boldsymbol{\theta}$, can be estimated analytically since the first passage time of a Wiener process follows an Inverse Gaussian distribution (Brown, 2005). However, closed-form estimation of $\boldsymbol{\Lambda}$ is not available. One common solution is to employ MCMC sampling on $p(\mathbf{X}_{1:T}, \boldsymbol{\Lambda} \mid \mathbf{Y}_{1:T})$, where time has been discretized to an arbitrarily fine grid. However, this can be prohibitively slow in high dimensions. Instead, the latent processes $\mathbf{X}_{1:T}$ are integrated out via the Laplace approximation (Tierney & Kadane, 1986; Skaug & Fournier, 2006; Koyama et al., 2010), i.e., $p_{\mathrm{LA}}(\mathbf{Y}_{1:T} \mid \boldsymbol{\Lambda}) \approx \int p(\mathbf{Y}_{1:T}, \mathbf{X}_{1:T} \mid \boldsymbol{\Lambda})\, d\mathbf{X}_{1:T}$, so that an estimate of $\boldsymbol{\Lambda}$ can be obtained by maximizing the Laplace-approximated marginal likelihood $p_{\mathrm{LA}}(\mathbf{Y}_{1:T} \mid \boldsymbol{\Lambda})$. An efficient gradient-based nested optimization algorithm for this is implemented in our R/C++ library `fastr` (Chen et al., 2023). The algorithm involves repeatedly solving linear systems of the form $\boldsymbol{\Sigma}\mathbf{x} = \mathbf{b}$, which due to the factor structure $\boldsymbol{\Sigma} = \boldsymbol{\Lambda}\boldsymbol{\Lambda}' + \boldsymbol{\Psi}$, scales linearly in the number of neurons q for fixed number of factors d.

Data Analysis

To assess the performance of the proposed model, we first applied it to a biologically plausible simulated dataset generated in NEURON (Hines & Carnevale, 1997). We found that the K-means clustering results based on $\hat{\boldsymbol{\Lambda}}$ are far more accurate than those based on (convolved) multiple neuron spike trains $\mathbf{Y}_{1:T}$ (see Figure 2(a)-(c)). Therefore, the estimate $\hat{\boldsymbol{\Lambda}}$ uncovers the underlying associations between neurons beyond what is observed at the data level. Next, we applied our model on neuronal ensembles recorded from a rat's medial prefrontal cortex (mPFC) and orbitofrontal cortex (OFC) during a classical conditioning experiment. The rat was trained to recognize a cue predictive of an appetitive outcome (sucrose). We modelled neuronal population dynamics in relation to encoding appetitive outcomes and functional connectivity, and found that coordinated activities are modulated by brain areas mPFC and OFC. The estimated correlation matrices $\hat{\boldsymbol{\Sigma}}$ in Figure 3 also show a decrease in overall neuronal interactions after the onset of the cue.

Discussion

We have proposed a novel continuous-time multivariate point process latent factor model for simultaneously recorded spike

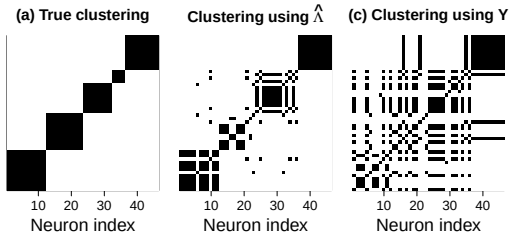

Figure 2: Clustering performance. The (i, j)-th entry of the matrix is black if neuron i and j are assigned to the same cluster, and white otherwise.

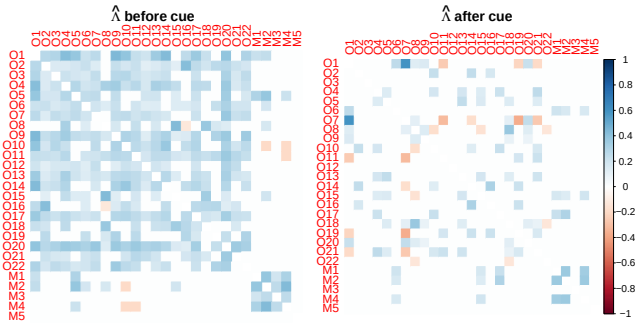

Figure 3: Estimated correlation between neurons. "O" stands for OFC neurons and "M" stands for mPFC neurons. The numbers are the neuron indices. For visual presentation, non-significant values and the diagonal elements are set to zero.

trains. Downstream analyses using the proposed model can reveal neuronal clustering and estimate correlations between neurons. Computational challenges in model inference are addressed by carefully designing and implementing an efficient model-fitting procedure based on the Laplace approximation. We have confirmed, via simulation studies, that our algorithm achieves satisfactory accuracy and precision in parameter estimation. Finally, compared to black-box deep learning models for neural spike trains, our proposed model is able to provide more reliable statistical inference results with uncertainty quantification, which is vital for making scientifically sound conclusions. One immediate future direction is to apply our model to appetitive and aversive classical conditioning outcomes (available within the same experimental data) to investigate reward-value coding, and to identify value signals for the most relevant contextual features.

Acknowledgments

This work was supported by the Natural Sciences and Engineering Research Council of Canada, grant numbers RGPIN-2018-04376 (Ramezan), DGECR-2018-00349 (Ramezan) and RGPIN-2020-04364 (Lysy).

References

Brown, E. N. (2005). Theory of point processes for neural systems. In *Methods and models in neurophysics* (p. 691-727). Elsevier.

Buesing, L., Machado, T. A., Cunningham, J. P., & Paninski, L. (2014). Clustered factor analysis of multineuronal spike data. In *Advances in neural information processing systems* (Vol. 27).

Chen, M., Ramezan, R., & Lysy, M. (2023). *fastr: **F**actor **A**nalysis of **S**pike **T**rains in **R**.* `https://github.com/meixichen/fastr`. GitHub.

Duncker, L., & Sahani, M. (2018). Temporal alignment and latent Gaussian process factor inference in population spike trains. In *Advances in neural information processing systems.*

Fujisawa, S., Amarasingham, A., Harrison, M. T., & Buzsáki, G. (2008). Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nature Neuroscience*, *11*, 823–833.

Gao, Y., Archer, E., Paninski, L., & Cunningham, J. P. (2016). Linear dynamical neural population models through nonlinear embeddings. In *30th conference on neural information processing systems.*

Hines, M. L., & Carnevale, N. T. (1997). The neuron simulation environment. *Neural Computation*, *9*, 1179-1209.

Humphries, M. D. (2011). Spike-train communities: Finding groups of similar spike trains. *Journal of Neuroscience*, *31*(6), 2321-2336.

Kass, R. E., & Ventura, V. (2006). Spike count correlation increases with length of time interval in the presence of trial-to-trial variation. *Neural Computation*, *18*(11), 2583-2591.

Koyama, S., Eden, U. T., Brown, E., & Kass, R. (2010). Bayesian decoding of neural spike trains. *Annals of the Institute of Statistical Mathematics*, *62*, 37–59.

Lopes-dos-Santos, V., Ribeiro, S., & Tort, A. B. L. (2013). Detecting cell assemblies in large neuronal populations. *Journal of Neuroscience Methods*, *220*, 149-166.

Nelson, M. E. (2002). Multiscale spike train variability in primary electrosensory afferents. *Journal of Physiology-Paris*, *96*(5), 507-516.

Pandarinath, C., O'Shea, D. J., Collins, J., Józefowicz, R., Stavisky, S. D., Kao, J. C., ... Sussillo, D. (2018). Inferring single-trial neural population dynamics using sequential auto-encoders. *Nature Methods*, *15*(10), 805-815.

Perkel, D. H., Gerstein, G. L., & Moore, G. P. (1967). Neuronal spike trains and stochastic point processes. ii. simultaneous spike trains. *Biophysical journal*, *7*(5), 419-40.

Ramezan, R., Chen, M., Lysy, M., & Marriott, P. (2022). A multivariate point process model for simultaneously recorded neural spike trains. In *Conference on cognitive computational neuroscience.*

Ramezan, R., Marriott, P., & Chenouri, S. (2014). Multiscale analysis of neural spike trains. *Statistics in Medicine*, *33*, 238–256.

Skaug, H. J., & Fournier, D. A. (2006). Automatic approximation of the marginal likelihood in non-gaussian hierarchical models. *Computational Statistics & Data Analysis*, *51*(2), 699–709.

Tierney, L., & Kadane, J. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, *81*, 82-86.

Ventura, V., Cai, C., & Kass, R. E. (2005). Trial-to-trial variability and its effect on time-varying dependency between two neurons. *Journal of neurophysiology*, *94*(4), 2928-2939.

Wei, G., Stevenson, I. H., & Wang, X. (2022). *Bayesian clustering of neural activity with a mixture of dynamic Poisson factor analyzers.* arXiv. Retrieved from `https://arxiv.org/abs/2205.10639`

Williams, A. H., Degleris, A., Wang, Y., & Linderman, S. W. (2020). Point process models for sequence detection in high-dimensional neural spike trains. In *Advances in neural information processing systems.*

Wu, A., Roy, N. A., Keeley, S., & Pillow, J. W. (2017). Gaussian process based nonlinear latent structure discovery in multivariate spike train data. In *Advances in neural information processing systems* (Vol. 30).

Yu, B. M., Cunningham, J. P., Santhanam, G., Ryu, S. I., Shenoy, K. V., & Sahani, M. (2008). Gaussian-process factor analysis for low-dimensional single-trial analysis of neural population activity. *Journal of Neurophysiology*, *102*(1), 614-635.