Inferring functional connectivity with priors on network topology

Scott W. Linderman and Ryan P. Adams
Harvard University

1 Summary

Neural computation is manifested in functional relationships between neurons. Reverse engineering this computation can be approached from the bottom up, looking for patterns in neural recordings, and from the top down, hypothesizing models of functional connectivity that could explain observed phenomena. We construct a Bayesian approach to modeling the latent functional network, using data from observed spike trains and representing top-down hypotheses via random graph models.

We build on the GLM framework [4], extending the model to explicitly incorporate prior distributions on the topology and weight of the functional connections, as well as a saturating nonlinear link function. Since many interesting priors lack convexity, we depart from the traditional MLE and MAP approaches, opting instead for a Gibbs sampling algorithm with an efficient GPU implementation.[1]

Our method is demonstrated on a synthetic spike train generated from an integrate-and-fire network with clusters of highly-connected excitatory neurons suppressing one another through an inhibitory pathway, thereby forming stable attractor states. A prior for a densely connected functional network, as implicitly assumed in the GLM, claims direct inhibition between the two clusters due to their negative correlation, whereas a prior biased toward blocks of functionally similar neurons recovers the inhibitory pathway.

This is but a proof of concept for a very general Bayesian framework. We can build priors for spatially proximal connectivity or time-varying weights, models for hierarchical or distributed representations of information among neurons, or models where neurons represent states in complex dynamical systems. More importantly, we have an algorithm for efficiently performing inference given spike train recordings, and for comparing models in a principled framework.

2 Model

We decompose the impulse response from neuron n onto neuron n' as $h(\Delta t; n, n') = A_{n,n'} W_{n,n'} g_\theta(\Delta t)$, where $A_{n,n'}$ is an entry in the directed adjacency matrix $\mathbf{A} \in \{0,1\}^{N \times N}$, $\mathbf{W} \in \mathbb{R}^{N \times N}$ is the corresponding weight matrix, and $g_\theta(\Delta t)$ is a parameterized, normalized function governing the temporal dependence of the filter. We use a logistic transfer function (i.e. logit link function) scaled by λ_n^{max}, reflecting the unique maximum conditional intensities of each neuron.[2] The stimulus-dependent component is taken to be constant in the experiments below, but could easily be modeled as, for example, a stimulus-dependent Gaussian process.

Network Priors. By explicitly separating $A_{n,n'}$ and $W_{n,n'}$ from the temporal component of the impulse response we can place prior distributions on the topology of the network. Similar work [3] has considered priors on $\mathbf{W}$, enforcing sparsity through a bias for low weights, but does not directly model the topology of $\mathbf{A}$. For example, we may hypothesize that neurons act in functionally similar blocks, and the relations between neurons are determined by their block affiliations. A stochastic block model (SBM) [2] prior on $\mathbf{A}$ is a natural way to represent this intuition. To capture the idea that neurons are either excitatory or inhibitory, we place a prior on $\mathbf{W}$ that requires all rows to have the same sign. These models fall under the Aldous-Hoover representation, which contains all random graph models invariant to permutations of the node labels [1]. Any network prior in this class can be directly incorporated into our inference algorithm.

Inference We use Gibbs sampling to obtain samples from the posterior distribution. To efficiently compute the log-likelihood, which contains an intractable integral and a sum over per-spike instantaneous rates, we make a piecewise linear approximation to the conditional intensity. We can then use trapezoidal quadrature

[1] This code has a simple interface and will be made publicly available.

[2] Clearly $\boldsymbol{\lambda}^{max}$ is not a hard limit on the minimum interspike interval, but it does allow the model to explain limited firing rates despite strong excitatory influence.

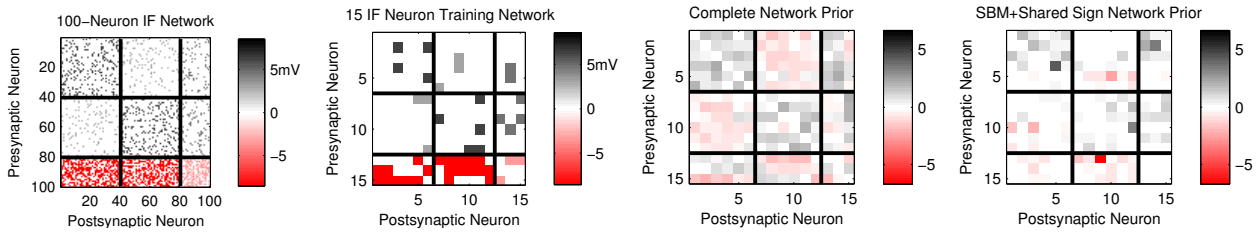

(a) (b) (c) (d)

Figure 1: (a) Weighted adjacency matrix of 100-node IF network. Black lines delineate clusters. Excitatory clusters preferentially connect to themselves and to inhibitory neurons; the inhibitory cluster provides global suppression. (b) True network of a subset of 15 IF neurons whose spike trains were used for training. (c) Average inferred weight matrix ($\mathbf{A} \circ \mathbf{W}$, 10K samples) with a complete network prior, the effective prior in the GLM. (d) Average inferred network with an SBM, shared sign prior ($\mathbf{A} \circ \mathbf{W}$, 10K samples) exhibiting the expected block structure. Note that neurons 5 and 10 were, on average, misclassified as members of an inhibitory block. [(c,d) units of change in activation].

and linear interpolation, leveraging GPUs to parallelize the computation. Gibbs updates are then relatively straightforward: the maximum firing rates λ_n^{max} are conditionally conjugate with a gamma prior; Gaussian variables such as the weights and baseline activation are suited to Elliptical Slice Sampling; and parameters of g_θ can be sampled with Hybrid Monte Carlo.

3 Experimental Results

We test our approach on synthetic data from a deterministic integrate-and-fire (IF) network. We create two excitatory clusters of 40 neurons each. The clusters compete, suppressing one another through a population of 20 inhibitory neurons, thereby creating two stable "activity states," c.f. Figure 1a. Spiking variability arises due to the chaotic regime induced by balanced excitation and inhibition. We train our model on the spike trains of a subset of 15 neurons, simulating a potential tetrode recording, c.f. Figure 1b. Despite limited structural connections, we expect functional relationships due to similarity with unobserved neurons.

First consider a completely connected network $\mathbf{A}$, as in an unregularized GLM. As shown in Figure 1c, the completely connected model identifies the negative correlation between the excitatory clusters as an inhibitory functional connection. This is a valid interpretation, but it uses many weak relationships when we know there exists a compact explanation in terms of an inhibitory pathway. If we instead use an SBM prior where each block is constrained to be either excitatory or inhibitory, we recover the distinct functional blocks we know to exist, c.f. Figure 1d.

4 Conclusion

Statistics and machine learning have provided us with a wealth of network models, but they rarely conform to the convexity constraints of the GLM framework. We provide a method for combining network and spiking models, enabling Bayesian model comparison for different network hypotheses, and laying a foundation for future research on connectivity models specifically tailored to populations of neurons.

References

[1] David J Aldous. Representations for partially exchangeable arrays of random variables. *Journal of Multivariate Analysis*, 11(4):581–598, 1981.

[2] Krzysztof Nowicki and Tom A B Snijders. Estimation and prediction for stochastic blockstructures. *Journal of the American Statistical Association*, 96(455):1077–1087, 2001.

[3] Ian H Stevenson, James M Rebesco, Nicholas G Hatsopoulos, Zach Haga, Lee E Miller, and Konrad P Kording. Bayesian inference of functional connectivity and network structure from spikes. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 17(3):203–213, 2009.

[4] Wilson Truccolo, Uri T. Eden, Matthew R. Fellows, John P. Donoghue, and Emery N. Brown. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects. *Journal of Neurophysiology*, 93(2):1074–1089, 2005.