

A Bayesian approach to inferring functional connectivity and structure from spikes

IH Stevenson (1), JM Rebesco (1), NG Hatsopoulos (2), LE Miller (1,3), KP Körding (1,4)

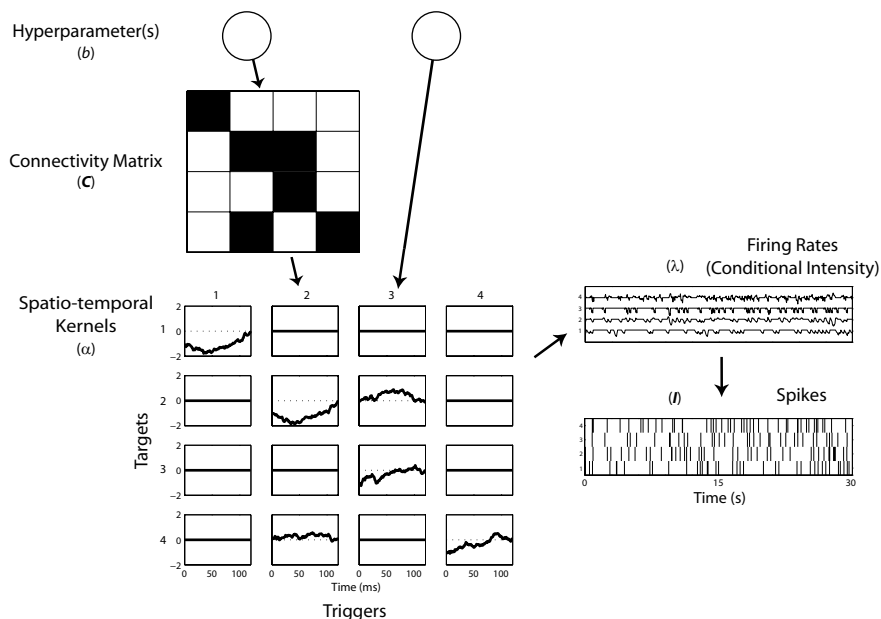
(1) Department of Physiology, Feinberg School of Medicine, Northwestern University

(2) Department of Organismal Biology and Anatomy, University of Chicago

(3) Department of Biomedical Engineering, Northwestern University

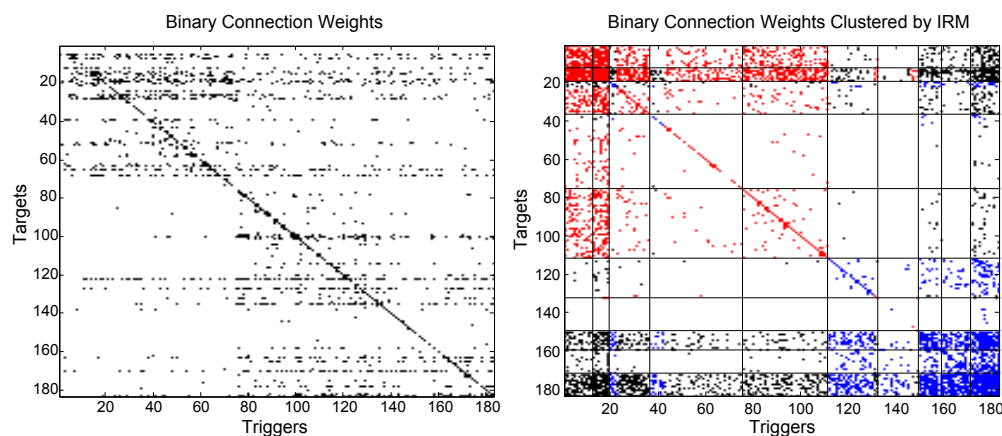
(4) Department of Physical Medicine and Rehabilitation, Rehabilitation Institute of Chicago

Current multi-electrode techniques enable the simultaneous recording of spikes from hundreds of neurons. To study neural processing, plasticity, and hierarchical structure, such as in the visuomotor system, it is desirable to infer the underlying functional connectivity between the recorded neurons. One challenge in inferring these connections is that a large number of parameters, which characterize how each neuron influences the other neurons, are often estimated from relatively little data. Such problems can be improved by using Bayesian methods that combine information from the recorded spikes (likelihood) with prior beliefs about functional connectivity (prior). Recent studies have used a likelihood function that captures the Poisson-like firing of neurons [2] and a prior that captures the sparseness of neural connections [3]. Here we extend these methods to include a prior that captures the empirical finding that interactions tend to vary smoothly in time. A generative model for this problem is shown in Figure 1. To describe the connections between neurons we follow a network likelihood model. Each neuron has a spontaneous firing rate, and connections between neurons are parameterized as a set of spatio-temporal kernels, α . These parameters generate a Poisson conditional intensity function for each neuron which, in turn, generates spikes. A likelihood function for this model has been previously derived using the point-process framework [1]. To implement a sparseness priors in this framework we introduce a set of latent variables, C , drawn i.i.d from an exponential distribution, to describe the strength of each connection. To implement a prior over the smoothness of connections we assume that the time course of each spatio-temporal kernel is a Gaussian random walk with variance determined by the latent variable, $C_{i,j}$. This incorporates our intuitive notion that most connections will be very weak and that fluctuations in spatio-temporal kernels will scale with connection strength.



To estimate the spatio-temporal kernels we find the parameters α and latent variables C that maximize the log-posterior probability of the model given the observed spike trains (MAP estimation). A single free parameter b allows us to vary the relative importance of smoothness and sparseness. Estimations are fairly robust to the choice of this parameter, and in practice, we can optimize it using cross-validation.

Using simulated data we show that this method, as well as a method using a linear approximation of the likelihood function, can successfully infer connectivity patterns from spikes. In datasets with many neurons (> 100) the use of priors over smoothness and sparseness can compensate, to some extent, for a less accurate likelihood function. Applying our method to real data, we infer functional connections from spike data recorded from primary motor (M1) and premotor (PMd) cortices of a macaque monkey. We use cross-validation to verify our estimates. Spatio-temporal kernels describing the interactions between neurons were similar to those found by previous studies. However, recording from two anatomically distinct areas (M1 and PMd) allows us to analyze the larger structure of these functional connections. We use a recently introduced Bayesian clustering algorithm, the infinite relational model (IRM), to make sense of the inferred functional networks and to study their structure. Unlike k-means or many other clustering algorithms, the complexity of the clustering (e.g., the number of clusters) is not predetermined with this algorithm. Using a prior over the number of clusters, IRM automatically determines the number of clusters and cluster assignments that best fit the data. Groups of neurons in M1 and PMd show distinct patterns of input and output that may correspond to functional assemblies (similar to cortical columns). A typical clustering result is shown in Figure 2. Clusters were composed, almost exclusively, of neurons from either area M1 (blue) or PMd (red). However, incorporating a model of external covariates, such as hand position [4] or gaze, is required to make finer functional distinctions. Although, these results are from motor cortex, this method has potential applications in analyzing functional connections anywhere in the visuomotor system and in the brain.



References

- [1] EN Brown, R. Barbieri, UT Eden, and LM Frank. *Likelihood methods for neural data analysis*, pages 253–286. London: Chapman and Hall, 2003.
- [2] Murat Okatan, Matthew A Wilson, and Emery N Brown. Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Comput*, 17(9):1927–1961, Sep 2005.
- [3] JW. Pillow, J. Shlens, L. Paninski, A. Sher, AM. Litke, E.J. Chichilnsiky, and E.P. Simoncelli. Deciphering correlations: Bayesian decoding of multi-neuronal spike trains in primate retina. *Computational and Systems Neuroscience (Cosyne)*, poster, II-77., 2007.
- [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. *J Neurophysiol*, 93(2):1074–1089, Feb 2005.