Neural network structure inference from its time series

Tony Hyun Kim[1, *]

[1] *Department of Electrical Engineering, Stanford University*
385 Serra Mall, Stanford, CA 94305

(Dated: November 13, 2013)

We explore the task of neural network connectivity inference from a set of temporal measurements at each node (neuron). Note: This document is a work in progress and much of the text is missing. On the other hand, I made a point of including key figures in the milestone, so that one can get a sense of the direction of the project. More information can be found in my project blog at cs224w.blogspot.com. I also realize that I may need to narrow the scope for the course project.

I. INTRODUCTION

A. (Brief) Motivation

I am interested in the task of network connectivity inference based on a measurement of the temporal dynamics of a set of nodes. Specifically, as part of my PhD research, I am involved in the design and implementation of optical neural recording devices, through which it is now possible to record the functional activity of hundreds (and soon, thousands) of neurons in a living animal with single-cell resolution [GBC+11].

The basic premise of the optical neural recording is as follows: fluorescent proteins whose fluorescence is modulated by neural activity [CWS+13] is introduced to the brain of a laboratory animal under study (mice, in my case). Portions of the skull (and possibly some brain tissue [JMA+04]) is surgically removed in order to provide optical access to the fluorescent neurons. This patch of neural tissue can subsequently be imaged over period of months using various imaging methodologies, that yield a movie of cellular populations whose brightness level encodes the neural activity. The video data can be segmented to yield optically-derived "activity traces" for each neuron [MNS09]. The activity dataset is studied in the context of the biological experiment. In the future, I intend on performing network analysis (e.g. recognition of network motifs, etc. [BS09]) on optical brain data at the cellular level.

For my CS224w term project, I am interested in developing a reference computational model for future analysis of experimental data. Generally, my intent is to perform "connectivity analysis" on a simulated dataset of neuron populations, to explore the use of various coupling metrics on time series data for network inference (*e.g.* linear cross-correlation, Granger causality, etc.) and to use the computational framework to explore several basic questions relevant to the experimental practice of optical neural recording.

* kimth@stanford.edu

B. Statement of problem

Consider Fig. 1, which shows a graph of n nodes (solid circles) that represent a single neuron in brain tissue, and directed edges (solid arrows) that represent synaptic connectivity between neurons. Each neuron can influence the spiking activity of other nodes through the synaptic contacts. In this work, only excitatory connections (those that tend the post-synaptic neuron to fire action potentials) are considered. Note that the n observed neurons may also be influenced by neurons that are not in the observed set (dashed circles and dashed edges).

In experimental neuroscience, there exist techniques to record the "functional" activity of a population of neu-

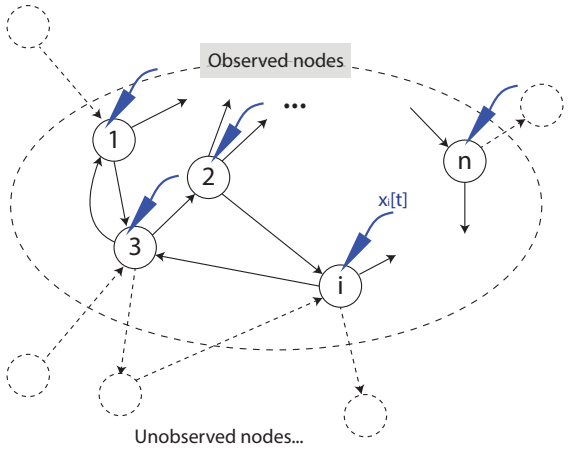

FIG. 1. The problem of neural network structure inference from its time series. We observe n neurons (solid circles) whose synaptic connections are represented by solid arrows. Note that the n observed neurons can be influenced by nodes that are not in the observed set (dashed circles). The blue arrows represent a single-neuron recording device (*e.g.* single-cell voltage-measuring electrodes) that yield time-series data $x_i[t]$ for each observed node. The basic task of "neural network inference" is to reconstruct the (observable) synaptic connections by computational analysis of the set of single-node measurements $\{x_i[t]\}$.

rons. For example, one can insert electrodes into a living cell, in order to measure the membrane voltage as a function of time, or use (as in my case) functional fluorescent proteins to "read out" the cell optically. Such single-neuron recording devices are represented by the blue arrows in Fig. 1. Regardless of the particular experimental method, the general feature is that one can obtain a set of time series data for each cell $\{x_i[t]\}$ that relates some information about the synaptic structure of the observed neurons. The deduction of the synaptic connectivity from a set of single-cell recordings is the task of "neural network inference from its time series."

In this work, I am interested in developing a statistically rigorous framework for neural network inference from its time series. The framework will have the ability to perform inference of both undirected and directed edges. Based on the network inference framework, I will then explore several questions relevant to the experimental practice of optical neural recording.

II. GENERAL QUESTIONS THAT MOTIVATE THE WORK

In this section, various questions relevant to the practice of optical neural recording are listed, that motivate the computational study. I do not expect that I address all questions in the CS224w project.

A. Comparison of various coupling metrics

The task of network inference from a set of time series recordings typically begins by computing a "coupling score" for all potential pairings in a network. For example, given two signals $x_i[t]$ and $x_j[t]$ of length n, a commonly used approach is to consider the cross-correlation $C_{ij}[\tau]$, defined by:

$$C_{ij}[l] = \frac{1}{\hat{\sigma}_i \hat{\sigma}_j (n - |l|)} \sum_{t=1}^{n-l} (x_i[t] - \bar{x}_i)(x_j[t+\tau] - \bar{x}_j) \quad (1)$$

where $\bar{x}$ and $\hat{\sigma}$ represent the mean and the standard deviation of the two signals, and the "lag" variable l accounts for possible temporal delays in the interaction between nodes i and j. Given the cross-correlation function, we may define the coupling score between $x_i[t]$ and $x_j[t]$ to be $\mathcal{C}_{ij} = \max_l |C_{ij}[l]|$ [KECK09]. Once the cross-correlation score has been computed between all potential pairs in the network, the inference algorithm then decides whether to assign an edge between the nodes.

It is clear that the inference result will be highly dependent on the coupling metric chosen to identify the interactions. For instance, because the cross-correlation is symmetric with respect to the node ordering ($\mathcal{C}_{ij} = \mathcal{C}_{ji}$), it can only infer undirected edges. Also, because the cross-correlation metric is purely bivariate (it does not consider the information provided by $x_k[t]$ where $k \neq i, j$), it is (as will be shown later) ill-suited to the task of inferring dense neural networks.

Therefore, in this work I am interested in comparing the performance and robustness of multiple coupling metrics for the inference of the underlying neural network from time series measurements.

B. Effect of temporal filtering on network inferrability

Currently, a major limitation of the optical recording method is that the fluorescent proteins involved in signaling functional activity have a slow temporal response – with time constants of 100s of ms [CWS$^+$13] – compared to the underlying voltage dynamics (the duration of an action potential is typically a few ms). In effect, the optical recording paradigm yields a low-pass filtered version of the underlying voltage dynamics.

How does the inferrability of the neural network vary as the temporal data is degraded? How sensitive is the inference performance to the time constant of the low-pass filter? How does the inference fail – with increases in false positives or increases in false negatives? Is the temporal susceptibility comparable between different coupling metrics or do certain measures perform better than others? Are there methods (*e.g.* Section II D) that may help recover the underlying network despite limitations in temporal resolution?

C. Inference of complex network topology

I am interested in simulating complex network topologies (rather than just random edges) such as $n_1 \to n_2 \to n_3$ and considering the performance of network inference on such structures (*e.g.* how often does the algorithm ascribe an erroneous edge $n_1 \to n_3$)? For instance, I suspect that the cross-correlation metric will be likely to falsely close the triad in the above case, whereas other multivariate measures (such as the Granger causality) would be less likely to make such a mistake.

More generally, for a given coupling metric, are there local graph structures that lead to errors (such as the $n_1 \to n_2 \to n_3$ example above)? By compiling a suite of such difficult-to-infer structures it may be possible to perform a more fine grained performance comparison of the different coupling metrics.

Also, I am curious whether different neural network generation processes (*e.g.* G_{nm} random graph, small world, preferential attachment) yield networks that are more difficult to infer.

D. Connection to cascades

[This may be beyond the scope of the course project.]

The concept of inferring network edges based on temporal data has conceptual similarities to information cascades in graphs, as discussed in lecture. On the other hand, to define a particular cascade instance, one has to "identify the contagion (*i.e.* the idea, information, virus, disease)" [GRLK10]. Unfortunately, when observing the spiking activity of an ensemble of neurons, it is difficult to decompose the overall activity into a superposition of independent spike cascades.

In the field of experimental neuroscience, a complementary technique to optical recording is optogenetics, where the electrical activity of neurons can be directly modulated by the experimenter via light illumination [PYG+12] (the effect can be both excitatory or inhibitory). What is the additional improvement in network inference if we are given the ability to drive the activity of a particular subset of neurons deterministically? Is there a method to characterize the resulting cascade even in the presence of endogeneous activity in the remainder of the population? Can we adapt methods of cascade-based network inference, in order to improve the performance of our neural network structure inference?

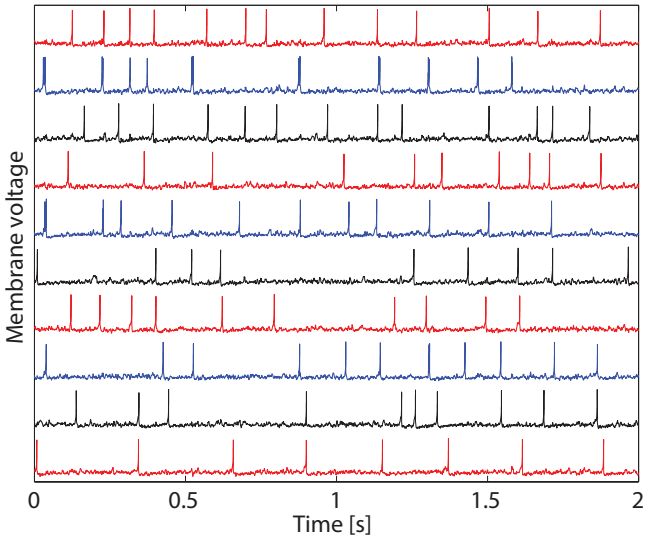

FIG. 2. Typical neural data as generated by the Izhikevich numerical model. For clarity, ten traces are shown. Most experiments in this study were based on a network consisting of $N = 100$ nodes.

III. ANALYSIS FRAMEWORK – PRIOR WORK

A. Numerical model of neuronal dynamics – model of Izhikevich [Izh03]

Fig. 2 shows a set of ten neural traces, as simulated by the Izhikevich numerical model. For later parts of this work, I typically used a network of $N = 100$ neurons, which is comparable to the number of cells captured by typical optical recordings. [More description later.]

B. Significance-based network inference from multivariate time series [KECK09]

Consider again the cross-correlation (Eq. 1), which is a typical metric for the coupling strength between two time series. With real and finite measurements, one can expect to find a distribution of nonzero correlations even in the case of truly independent processes. As a result, in experimental analysis, a common practice is to introduce a threshold and declare connectivity between nodes if the measured correlation exceeds that threshold. I find the thresholding approach to be unsatisfactory because: (1) the conclusions drawn from the data analysis is dependent on an arbitrary parameter, and (2) the approach is mathematically unprincipled.

I am thus drawn to the work by Kramer, *et. al.* [KECK09] describing a rigorous method of edge inference based on the concept of the "false discovery rate" (FDR) [BH95]. In essence, the FDR-based inference method makes conservative estimates regarding the presence of an edge, by bounding the expected fraction of false edges in the inferred set (such events are termed "false discovery"). Kramer's work applies FDR concepts to the problem of network inference.

I am conceptually drawn to the idea of using FDR control (rather than thresholding) for edge inference. However, Kramer's paper does not give a convincing summary of the expected performance of the method. (Two simulated examples, with known ground truths, are presented but only a single inference instance is presented.) In this work, I performed repeated FDR-based network inferences, to verify the performance of the method.

C. Coupling metrics for directed edge inference

1. *Granger causality*

IV. VALIDATION OF FDR-BASED EDGE INFERENCE

Fig. 3 illustrates the individual steps involved in FDR-based edge inference.

Fig. 4 shows an explicit validation of the FDR method.

V. EXPERIMENTS WITH UNDIRECTED NETWORK INFERENCE

A. Inference of dense networks

I observed that the cross-correlation metric for undirected edge inference fails when the edge density in the ground truth network is increased. I have a hypothesis for why this may be the case, which will be tested here.

B. Effect of temporal filtering

Fig. 5 shows the effect of the low-pass filter with time constant τ on the neural traces and their cross-correlation. (There is probably an easy closed-form expression for how the temporal linear filter affects the cross-correlation.)

Fig. 6 shows the degradation in performance when the neural traces are processed by a low-pass filter with a time constant of τ. In Fig. 6(a), it can be observed that the FDR framework is still valid (the actual FDR is approximately bounded by the FDR level q). However, Fig. 6(b) shows that the ROC curve has suffered due to low-pass filtering of the neural traces.

VI. EXPERIMENTS WITH DIRECTED NETWORK INFERENCE

A. Initial results

Fig. 7 shows preliminary results for directed edge inference in a network with $N = 30$ neurons based on Granger causality as the coupling metric.

[BH95] 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*, 57:289–300, 1995.

[BS09] E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neuro*, 10:186–198, 2009.

[CWS+13] T.-W. Chen, T. J. Wardill, Y. Sun, S. R. Pulver, S. L. Renninger, A. Baohan, E. R. Schreiter, R. A. Kerr, M. B. Orger, V. Jayaraman, L. L. Looger, K. Svoboda, and D. S. Kim. Ultrasensitive fluorescent proteins for imaging neuronal activity. *Nature*, 499:295–300, 2013.

[GBC+11] K. K. Ghosh, L. D. Burns, E. D. Cocker, A. Nimmerjahn, Y. Ziv, A. El Gamal, and M. J. Schnitzer. Miniaturized integration of a fluorescence microscope. *Nature Methods*, 8:871–878, 2011.

[GRLK10] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring network of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data*, 2010.

[JMA+04] J. C. Jung, A. D. Mehta, E. Aksay, R. Stepnoski, and M. J. Schnitzer. In vivo mammalian brain imaging using one- and two-photon fluorescence microscopy. *J Neurophysiol*, 92:3121–3133, 2004.

[KECK09] M.A. Kramer, U. T. Eden, S. S. Cash, and E. D. Kolaczyk. Network inference with confidence from multivariate time series. *Phys. Rev. E*, 79:061916, 2009.

[MNS09] E. A. Mukamel, A. Nimmerjahn, and M. J. Schnitzer. Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron*, 63:747–760, 2009.

[PYG+12] R. Prakash, O. Yizhar, B. Grewe, C. Ramakrishnan, N. Wang, I. Goshen, A. M. Packer, D. S. Peterka, R. Yuste, M. J. Schnitzer, and K. Deisseroth. Two-photon optogenetic toolbox for fast inhibition, excitation and bistable modulation. *Nature Methods*, 9:1171–1179, 2012.

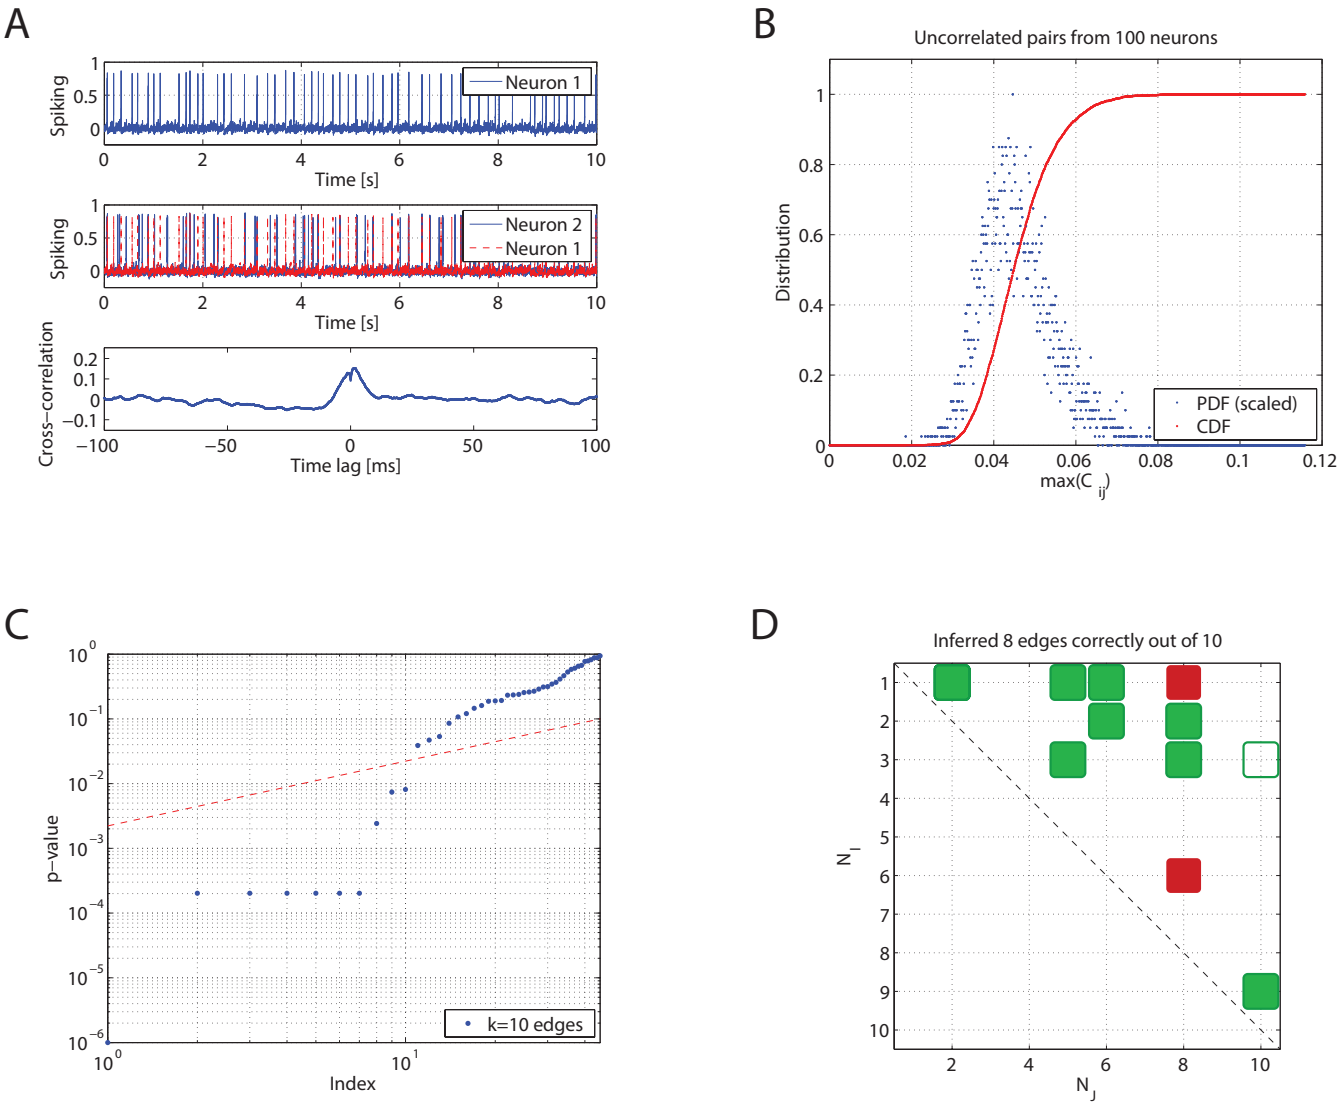

FIG. 3. Individual steps in the FDR-based edge inference mechanism, shown here for a ground truth neural network with $N = 10$ neurons and $M = 10$ directed edges. [A] For every pair of neurons in the network, compute the coupling score. In this case, there is a notable coupling between neurons 1 and 2, as measured by the cross-correlation $\mathcal{C}_{12} \approx 0.15$. [B] Compare the measured cross-correlation against the "null distribution," *i.e.* the distribution of $\mathcal{C}_{ij}$ for neurons i, j that are not connected. The comparison of $\mathcal{C}_{12}$ against the null distribution yields a significance value (a p-score) to be associated with the score. Here, the measurement of $\mathcal{C}_{12} \approx 0.15$ is highly significant, given that the null distribution PDF has most of its weight around $0.03 < \mathcal{C}_{ij} < 0.08$. [C] Identify the neuron pairs with sufficiently significant p-values according to the FDR-procedure of Benjamini and Hochberg (see text). [D] The results of the particular inference instance. Filled green squares represent true positives (the algorithm inferred a true edge), filled red squares represent false positives (the algorithm inferred a false edge), and the unfilled green square represent false negatives (the algorithm failed to infer a true edge). Note that we plot all results in the upper right triangle of the adjacency matrix, since the cross-correlation can only infer undirected edges.

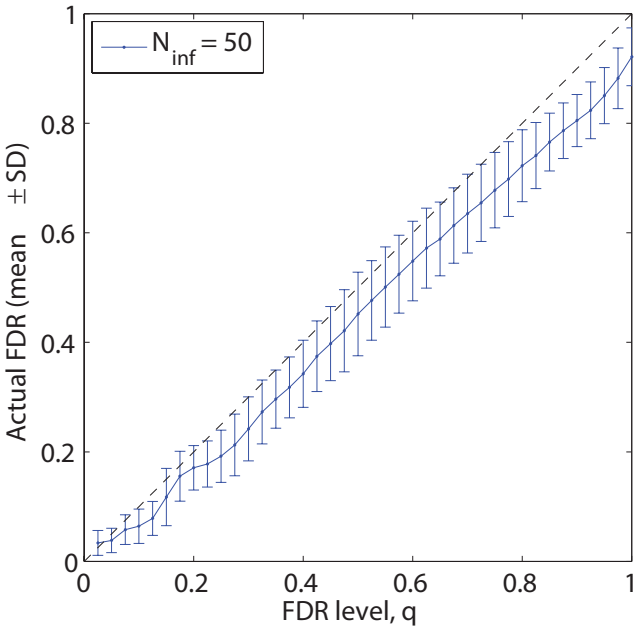

FIG. 4. Demonstration of the validity of the FDR method. For each FDR level q, we simulate 50 instances of the inference run on a graph with $N = 100$ and $M = 100$. We then measure the actual FDR (the fraction of false positives in the inferred set). The mean and standard deviation of the actual FDR is presented, which shows the mean to be well-bounded by the FDR level (dashed diagonal).

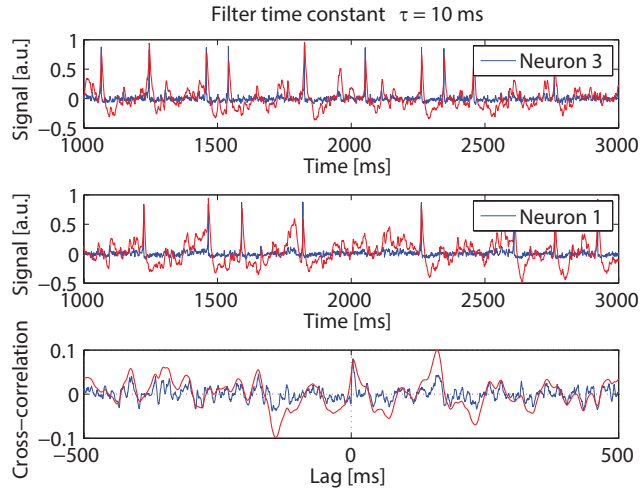

FIG. 5. The effect of low pass filtering (here, $\tau = 10$ ms) on the neural traces and the cross-correlation. The unfiltered quantities are plotted in blue, and the filtered results are shown in red.

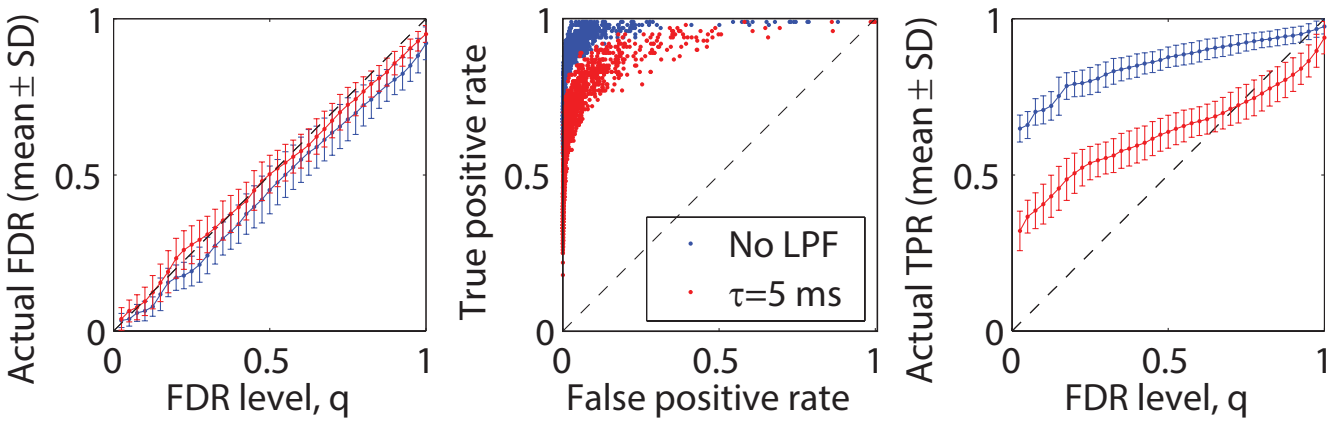

FIG. 6. Variation in the performance of undirected edge inference, when the neural signals are low-pass filtered prior to FDR-based inference. Unfiltered quantities are shown in blue, results with a $\tau = 5$ ms LPF are shown in red. (More will be added later.) [Left] The FDR-based inference method remains valid even after low-pass filtering the time series; the actual FDRs are bound by the FDR level q, independently of filtering. [Center] The degradation in inference performance as measured on a standard ROC plot. [Right] Significant loss in performance is shown due to filtering of the time traces. For a fixed FDR level q, the algorithm makes significantly fewer inferences (in order to bound the false discovery rate) which results in a lower true positive rate (TPR) compared to the unfiltered case.

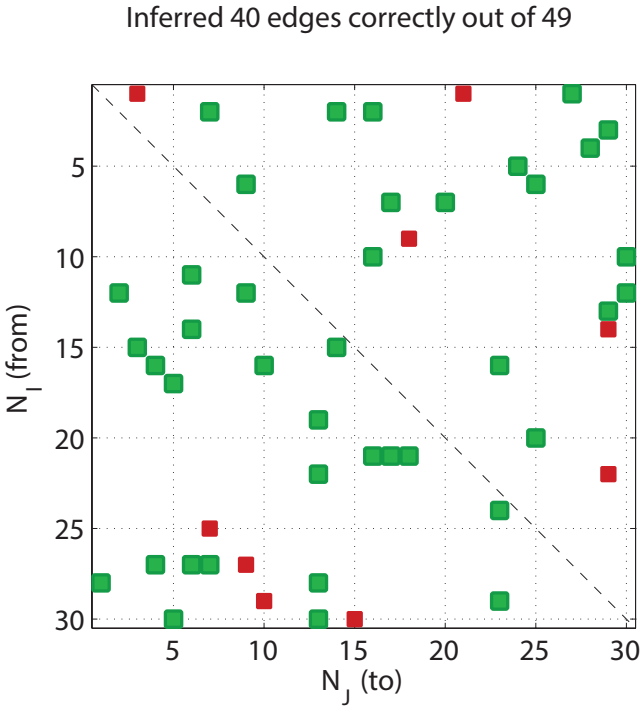

FIG. 7. Preliminary results in directed edge inference using the Granger causality as a coupling metric.