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: December 9, 2013)

I explore the task of neural network connectivity (*i.e.* graph) inference from a set of temporal measurements at each neuron (nodes). In this work, all experiments are based on numerically simulated neuronal populations, with known ground-truth graph structures. I develop a statistically principled framework for edge inference based on the concept of "false discovery rate" (FDR) as outlined in Ref. [1], and apply the FDR framework for undirected (based on max cross-correlation of temporal traces) and directed (pairwise Granger causality) edge inference. I then explore the variation in inference performance as the temporal measurements are degraded with low-pass filters. The filter models limitations in the currently available biological tools for neuronal readout.

I. INTRODUCTION

A. (Brief) Motivation

I am interested in inferring network connectivity from the temporal dynamics of its nodes. Such a task is relevant to my PhD research, in which I design and apply optical neural recording devices that observe the functional activity of hundreds (and soon, thousands) of neurons in a living animal with single-cell resolution [2].

The basic premise of the optical neural recording is as follows: fluorescent proteins whose fluorescence is modulated by neural activity [3] is introduced to the brain of a laboratory animal under study (mice, in my case). Portions of the skull (and possibly some brain tissue [4]) 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 movies of cellular populations whose brightness level as a function of time encodes the neural activity. The video data can be segmented to yield optically-derived "activity traces" for each neuron [5]. 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. [6]) on optical brain data at the single neuron level.

For my CS 224w term project, I am interested in developing the computational framework for future analysis of experimental data. In this work, my intent is to perform connectivity analysis (*i.e.* graph inference) 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 basic questions relevant to the experimental practice of optical neural recording.

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).
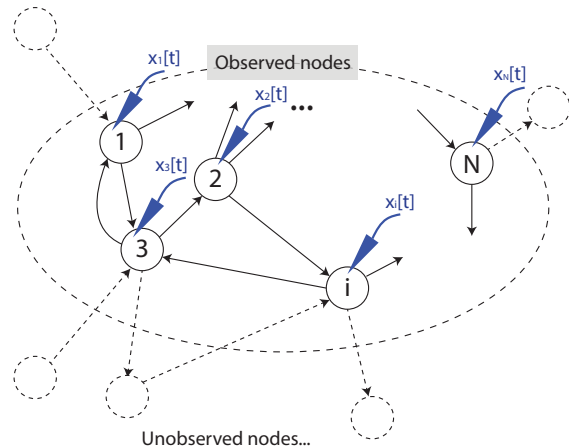

FIG. 1. The problem of neural network structure inference from its time series. We observe N neurons (solid circles) whose synaptic connections are represented by solid arrows. Note that the N observed neurons can be influenced by nodes that are not in the observed set (dashed circles). The blue arrows 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 synaptic connections by computational analysis of the set of single-node measurements $\{x_i[t]\}$.

* kimth@stanford.edu

In experimental neuroscience, there exist a variety of techniques to record the "functional" activity of a population of neurons. 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-level 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 may relay 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."

II. GENERAL QUESTIONS THAT MOTIVATE THE WORK

In this section, various questions relevant to the practice of optical neural recording are listed, that motivate my computational study. The current paper addresses Sections II A-II B, whereas the remaining sections outline future directions.

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 possible edges in a network. For example, given two signals $x_i[t]$ and $x_j[t]$ of length n, a commonly used coupling score is the cross-correlation $C_{ij}[\tau]$:

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

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 $\mathcal{C}_{ij} = \max_l |C_{ij}[l]|$. 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 based on $\{\mathcal{C}_{ij}\}$.

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 ill-suited to the task of inferring dense neural networks (more in Section II C).

In this work two different coupling metrics will be employed 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 [3] – compared to the underlying voltage dynamics. (The duration of an action potential is typically a few ms.) Today, the optical recording paradigm in effect 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? 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

Ultimately I am interested in simulating complex network topologies (rather than just random, isolated 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 infer a false edge $n_1 \to n_3$?). I suspect that bivariate measures such as cross-correlation and pairwise Granger causality (metrics that only consider pairs of nodes) will be likely to falsely close the triad in the above case, whereas more sophisticated multivariate measures (e.g. the conditional 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.

D. Connection to cascades

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)" [7]. Unfortunately, when observing the spiking activity of an ensemble of neurons, it is not obvious how 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 (the effect can be both excitatory or inhibitory) [8].
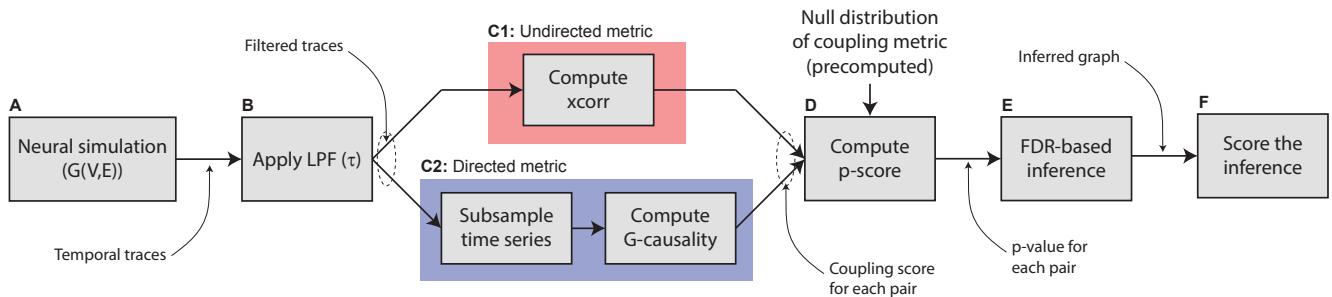
FIG. 2. Overall workflow for the numerical experiments performed in this paper. [A] A neural population of N cells with M excitatory connections is simulated. [B] Optional application of LPF on the neural time traces. [C1,C2] Computation of coupling metric for every node pair. [D] The coupling value for each pair is converted into a p-value. [E] The list of p-values is used by the FDR method to infer signifcant edges. [F] The inferred graph is scored against the ground truth.

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 identify "optogenetically-driven cascades" even in the presence of endogeneous activity (*e.g.* by patterning the drive input)? If so, can we adapt methods of cascade-based network inference for neural network determination? How does the performance of such a technique compare to typical "coupling-score"-based paradigms (*i.e.* the methods implemented in this paper)?

III. ANALYSIS FRAMEWORK

Fig. 2 outlines the overall workflow for the numerical experiments performed in this paper. First, we simulate a neural population of N cells with M excitatory connections among themselves (step A). This yields temporal traces (voltage dynamics) for each neuron. We then optionally apply a first-order LPF with time constant τ on the temporal traces (B). Next, we compute the coupling metric for all pairs. For undirected edge inference, we use the cross correlation metric (C1; Eq. 1); for directed edge inference, we use the pairwise Granger causality (C2; Section III B). Note that we subsample the time series for directed inference for computational reasons. The observed coupling values are converted into p-values by comparison against a precomputed null distribution (D). The list of p-values is then fed into the FDR-based method for edge inference (E). Finally, we score the inferred graph against the ground truth (F).

I implemented all steps of Fig. 2 in Matlab.

A. Numerical model of neuronal dynamics – Izhikevich neuron model

We use the neural simulation model of Izhikevich [9], which is a relatively popular model in computational neuroscience owing to its numerical simplicity [1]. In the Izhikevich model, we define the ground truth network

with N neurons and M directed edges, and simulate the voltage dynamics of the coupled neural network. The dynamical system is driven by the random inputs that represent the activity of unobserved neurons in the "background" (the dashed nodes in Fig. 1).

Fig. 3 shows a set of ten neural traces, as simulated by the Izhikevich numerical model. Each neural trace consists of spikes (action potentials) that indicate the "firing" of the neuron.

For undirected inferences, I used networks of $N = 100$ neurons, which is comparable to the number of cells captured by typical optical recordings. Traces were 10 s long. For directed inferences, I simulated much smaller networks of $N = 20$ neurons for 5 s due to the higher computational cost. In all cases, I selected $M = N$ directed edges at random.
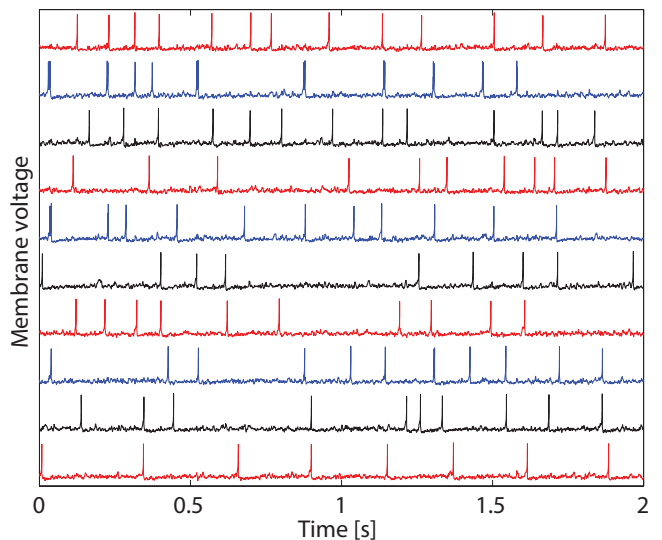

FIG. 3. Typical neural data as generated by the Izhikevich numerical model. For clarity, ten traces are shown from a simulation of $N = 100$ neurons.

B. Coupling metric for directed edge inference – pairwise Granger causality

Granger causality is a model-based implementation of the idea that a time series $x_i[t]$ can be said to "cause" another $x_j[t]$, if knowing the past values of x_i helps in predicting future values of x_j. In the past decade, Granger causality has emerged as a popular directed coupling metric for neuroscience applications [10].

The pairwise Granger causality ("G-causal") score from node i to node j (*i.e.* $i \rightarrow j$) is computed as follows. First, we fit an autoregressive model of order p (not to be confused with the p-value) to the temporal trace $x_j[t]$ (the target node). That is, we attempt to express x_j as:

$$x_j[t] = \sum_{k=1}^{p} a_k \cdot x_j[t-k] + \epsilon_j[t], \qquad (2)$$

where $\{a_k\}$ are fitted model parameters and ϵ_j is the residual. If x_j were a true autoregressive process, ϵ_j would be a white noise process characterized by some variance $\sigma_j = \mathrm{Var}(\epsilon_j)$.

Next, we perform a bivariate autoregressive fit for x_j using also the past values of x_i:

$$x_j[t] = \sum_{k=1}^{p} a_k' \cdot x_j[t-k] + \sum_{k=1}^{p} b_k' \cdot x_i[t-k] + \epsilon_{j|i}[t], \quad (3)$$

where $\{a_k', b_k'\}$ are model parameters and $\epsilon_{j|i}$ is the residual with variance $\sigma_{j|i} = \mathrm{Var}(\epsilon_{j|i})$.

Finally, the G-causal score $G_{i \rightarrow j}$ is defined by:

$$G_{i \rightarrow j} = \log\left(\frac{\sigma_j}{\sigma_{j|i}}\right). \qquad (4)$$

Note that the bivariate model (Eq. 4) is a strict superset of the univariate model (Eq. 3). Thus, we expect (in the absence of model fitting difficulties) $\sigma_{j|i} \leq \sigma_j$ and hence $G_{i \rightarrow j} \geq 0$. A large value for $G_{i \rightarrow j}$ indicates that the past information in x_i is useful for predicting the future values of x_j under the autoregressive assumption, and i can be said to "Granger cause" j.

Note: During my implementation of Granger causality analysis, I observed that the Izhikevich neural traces (Fig. 3) are not adequately fitted by an autoregressive model. In particular, I found that the residuals of the fit are definitely non-white and show "spikes" that coincide with the spiking in the neural trace. This implies to me that Granger causality, despite their prevalence in the literature, is not a particularly appropriate tool for inferring functional connectivity in neural traces.

C. FDR-based network inference from multivariate time series

Whatever the coupling metric, given real and finite measurements, one can expect to find a distribution of nonzero coupling scores even in the case of truly independent processes. As a result, a common practice in experimental analysis is to introduce an arbitrary threshold and declare connectivity between nodes if the measured coupling exceeds that threshold. However, 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.* [1] describing a rigorous method of edge inference based on the concept of the "false discovery rate" (FDR) [11]. 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 discoveries") to a specified "FDR level" q. Kramer's work applies FDR concepts to the problem of network inference.

In order to employ FDR-based edge inference, we begin by computing the coupling score for all pairs of neurons in the network (step C in Fig. 2). Next, we convert each coupling score into a p-value (step D) by comparing the observed value to the distribution of the coupling metric under the null hypothesis (*i.e.* the distribution of the coupling metric for independent neuron pairs). The p-value represents the probability of obtaining a coupling value at least as high as the one observed, assuming that the neurons are independent (*i.e.* the "null hypothesis"). Low p-values indicate significant coupling.

Then, FDR-based inference proceeds as follows (step E in Fig. 2; originally by Benjamini and Hochberg [11], adapted for network inference by Kramer, *et. al.* [1]). Sort the list of p-values in order of decreasing significance, *i.e.* $p_1 \leq p_2 \leq \cdots \leq p_m$ where m is the total number of potential edges. Choose a desired FDR level q, which represents an upper bound for the expected fraction of false positives in the set of inferred edges. Compare each p_i to the value $q \cdot i/m$ and find the maximum i (call it k) such that $p_k \leq q \cdot k/m$. Declare (infer) edges for the pairs corresponding to $p_1, p_2, \cdots, p_k$.

Fig. 4 illustrates a step-by-step walkthrough of the above process where the coupling metric is the cross-correlation (Eq. 1) and the ground truth network has $N = 10$ neurons and $M = 10$ edges. In panel A, we compute the cross-correlation between neurons 1 and 2, yielding a coupling score of $\mathcal{C}_{12} \approx 0.15$. In panel B, the observed $\mathcal{C}_{12}$ is compared against the null distribution, and the corresponding p-value is obtained. (The p-value is the CCDF of the null distribution evaluated at $\mathcal{C}_{12}$.) Panel C shows the sorted p-values for all $m = N \cdot (N-1)/2 = 45$ potential undirected edges in the network. The red dashed line indicates $q \cdot i/m$ for $q = 0.1$. A total of $k = 10$ p-values are deemed significant by the FDR method. Finally, panel D shows the comparison of the inference result against the ground truth network. See figure text for the detailed legend.

The FDR-based method can also make directed edge inferences, by substituting a directed coupling met-
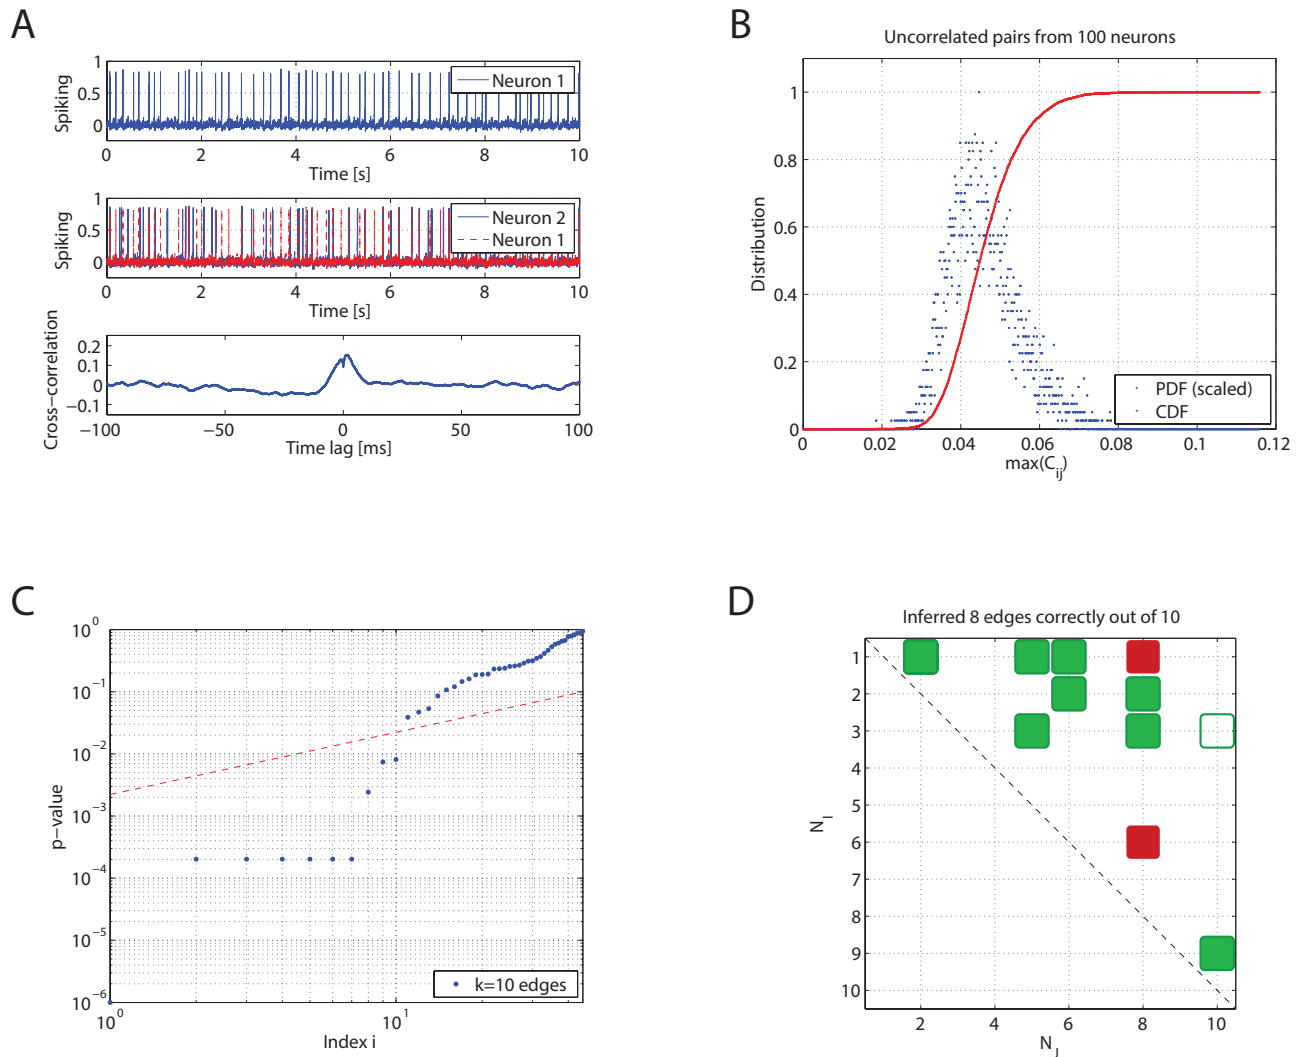
FIG. 4. Individual steps in the FDR-based mechanism for undirected edge inference, shown here for a ground truth neural network with $N = 10$ neurons and $M = 10$ directed edges. [A] For every pair of neurons in the network, compute the coupling score. In this case, there is a notable coupling between neurons 1 and 2, as measured by the max 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 observed 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). The red dashed line is $q \cdot i/m$ for FDR level $q = 0.1$. [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 an edge falsely), 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 (where multiple ground truth edges may overlap), since the cross-correlation can only infer undirected edges.

ric such as Granger causality in place of the cross-correlation.

D. Validation of FDR-based edge inference

I performed explicit validation of the FDR-based inference method, namely of the claim that the expected number of false positives in the inferred set is bounded by the FDR level q. In Fig. 5, I performed 50 inference instances on a ground truth network of $N = M = 100$, and measured the actual number of false positives in the inferred result. It can be observed that the actual FDR is indeed well-bounded by the specified FDR level. Overall, I find the FDR-based method to be a significant improvement over naive thresholding.
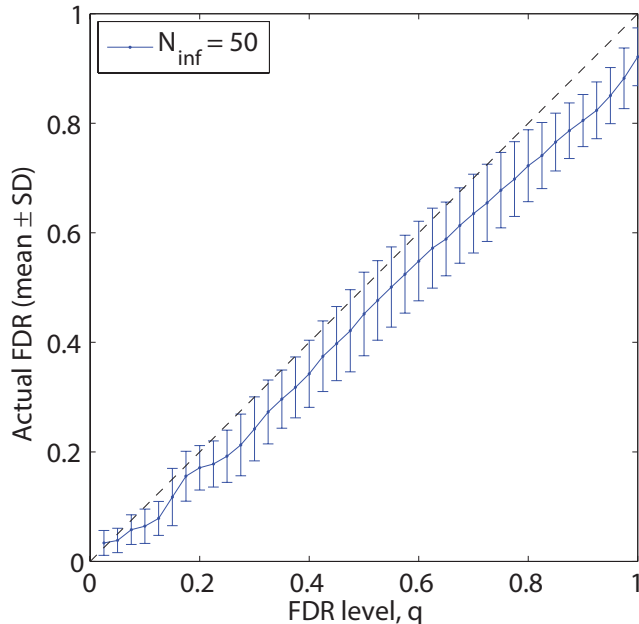
FIG. 5. Explicit validation of the FDR method. For each FDR level q, we simulate 50 instances of the inference run on a graph with $N = 100$ and $M = 100$. We then measure the actual FDR (the fraction of false positives in the inferred set). The mean and standard deviation of the actual FDR is presented, which shows the distribution of actual FDRs to be well-bounded by the specified FDR level (dashed diagonal).

IV. THE EFFECT OF TEMPORAL FILTERING ON EDGE INFERENCE

Having implemented the basic workflow for FDR-based edge inference, we now consider the effect of temporal lowpass filtering on overall inference performance. As shown in Fig. 2 (step B), a first-order lowpass filter with time constant τ is applied on the voltage traces from the Izhikevich model, which are then fed into the remaining inference pipeline. Note that the null distribution of the coupling metric (used in step D) must be recomputed for each value of τ.

A. Undirected inference

Fig. 6 shows the effect of applying a LPF to the neuronal traces. In the top two panels, the original temporal signal (blue) is compared to the filtered output (red; scaled to normalize for peak height). The LPF has two notable effects: (1) the width of each spike is increased from roughly 1 ms to the filter time constant τ; and (2) the signal-to-noise ratio is diminished in the filtered result, since a portion of the signal power is cut off by the LPF. (The spike waveform has high frequency components.) The bottom panel of Fig. 6 shows the effect of temporal filtering on the cross-correlation of the two
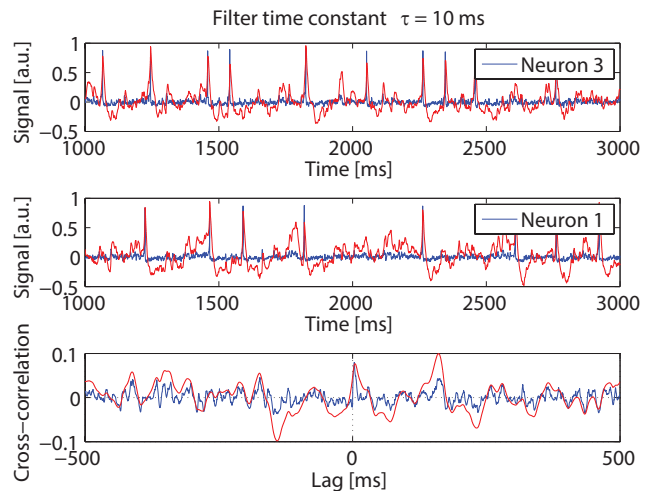

FIG. 6. The effect of low pass filtering (here, $\tau = 10$ ms) on the neural traces and the cross-correlation. The two neurons (1 and 3) share a synaptic connection in the ground truth network. The unfiltered quantities are plotted in blue, and the filtered results are shown in red.

signals. Compared to the original case, the filtered cross-correlation exhibits notable oscillations that makes the metric a less reliable indicator of the synaptic relationship between the neuron pair.

Fig. 7 summarizes the effect of temporal filtering on undirected edge inference. Firstly, the leftmost panel shows that the FDR-based mechanism is still valid after filtering the temporal traces; that is, the actual proportion of the false positives in the inferred set remains bounded by the FDR level q even in the presence of temporal filtering. Next, the center panel shows the effect of temporal filtering on inference performance in a standard ROC plot, where each data point represents an inference instance at a particular value of q. Even with relatively modest filtering ($\tau = 5, 10$ ms), the ROC curve is appreciably degraded from baseline performance.

The rightmost panel of Fig. 7 quantifies the drop in inference performance as a function of increasing τ in terms of the true positive rate (TPR). The TPR is a better indicator of the inference performance than the ROC plot, since the ground truth graphs have low edge densities at $N = M$. Here, we observe a dramatic degradation in performance. Consider a fixed FDR level $q = 0.1$ (dashed vertical line). When the inference is based on unfiltered traces, the FDR-based output correctly identifies 70% of the true synaptic connections on average. However, when the traces are filtered by a LPF of $\tau = 5$ ms, the TPR decreases to 45%. At $\tau = 10$ ms, the TPR is 20%. The interpretation is as follows: as the quality of the temporal data is degraded, the FDR-based mechanism becomes "more conservative," in order to bound the number of false positives in the inferred set. In the current experiment, a LPF with $\tau = 10$ ms is already able to render the inference nearly useless.

B.   Directed inference

Similarly, Fig. 8 shows the variation in the performance of directed edge inference as a function of LPF τ, using the pairwise Granger causality as the coupling metric. Note that the autoregressive model order p was set to $\max(1, n)$ where n is the number of signal samples corresponding to a temporal window of 5 ms. (This means that the autoregressive model considers past information of at least 5 ms prior to the current time.)

Interestingly, the inference performance based on Granger causality is robust up to $\tau = 10$ ms, as measured by the TPR as a function of q. In fact, the performance with $\tau = 5$ ms is even better than the unfiltered case, which could be due to the reduced model order used in the numerical analysis after the LPF has been applied.

On the other hand, there is a precipitous drop in performance at $\tau = 15$ ms. Admittedly, the $\tau = 15$ ms result is somewhat suspect, since the measured FDR greatly exceeds the specified FDR level. However, taken at face value, the experiment suggests that Granger causality-based inferences also fail with modest lowpass filtering of temporal data.

V.   CONCLUSIONS AND FUTURE WORK

In this work, I developed a workflow for network inference based on FDR, *i.e.* bounding the expected number of false positives in the inferred set. I then performed both directed (based on cross-correlation) and directed (pairwise Granger causality) edge inferences. Finally, I examined the degradation in inference performance as the voltage traces were preprocessed by first-order LPF of modest time constants ($\tau \leq 15$ ms; roughly $10\times$ the duration of a typical action potential spike).

Based on the LPF experiments, I found that even slight filtering of the temporal traces can have disastrous effects on inference performance. (I am confident about the numerical accuracy of the cross-correlation results; less so with Granger causality.) Given that current fluorescent optical probes have time constants of order 1 s, it is unlikely that one can perform functional network reconstruction by relying on optical readout alone. Thus, there is a strong need to consider alternative techniques – such as optogenetic perturbation of the neural circuit under study (Section II D) – in order to extract connectivity information in spite of poor temporal resolution.

The FDR-based inference method is compatible with arbitrary coupling metrics, and in this paper I demonstrated the use of cross-correlation and pairwise Granger causality. These metrics will not be sufficient to infer complex network topologies at higher edge densities than the ones presented here, as they are purely bivariate metrics (Section II C). Furthermore, I believe that the Granger causality is not an appropriate metric for the current application, because the underlying voltage dynamics is not adequately modeled by the vector autoregressive assumption underlying G-causality analysis. In the future, I am interested in applying Granger-causality models specifically adapted for spiking processes [12], and model-free information-theoretic directed measures such as the transfer entropy [13].

Finally, in order to apply more complex (*e.g.* multivariate) coupling metrics to longer simulated and experimental data, I will need to improve the computational performance of the inference mechanism. In particular, I am interested in developing an FPGA-based coprocessor for hardware-accelerated computation of coupling metrics from time series data.

[1] M. Kramer, U. T. Eden, S. S. Cash, and E. D. Kolaczyk, Phys. Rev. E **79**, 061916 (2009).

[2] K. K. Ghosh, L. D. Burns, E. D. Cocker, A. Nimmerjahn, Y. Ziv, A. El Gamal, and M. J. Schnitzer, Nature Methods **8**, 871 (2011).

[3] 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, Nature **499**, 295 (2013).

[4] J. C. Jung, A. D. Mehta, E. Aksay, R. Stepnoski, and M. J. Schnitzer, J Neurophysiol **92**, 3121 (2004).

[5] E. A. Mukamel, A. Nimmerjahn, and M. J. Schnitzer, Neuron **63**, 747 (2009).

[6] E. Bullmore and O. Sporns, Nat Rev Neuro **10**, 186 (2009).

[7] M. Gomez-Rodriguez, J. Leskovec, and A. Krause, ACM Transactions on Knowledge Discovery from Data (2010).

[8] 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, Nature Methods **9**, 1171 (2012).

[9] E. M. Izhikevich, IEEE Transactions on Neural Networks **14**, 1569 (2003).

[10] A. K. Seth, Journal of Neuroscience Methods **186**, 262 (2010).

[11] Y. Benjamini and Y. Hochberg, Journal of the Royal Statistical Society. Series B **57**, 289 (1995).

[12] S. Kim, D. Putrino, S. Ghosh, and E. N. Brown, PLoS Computational Biology **7** (2011).

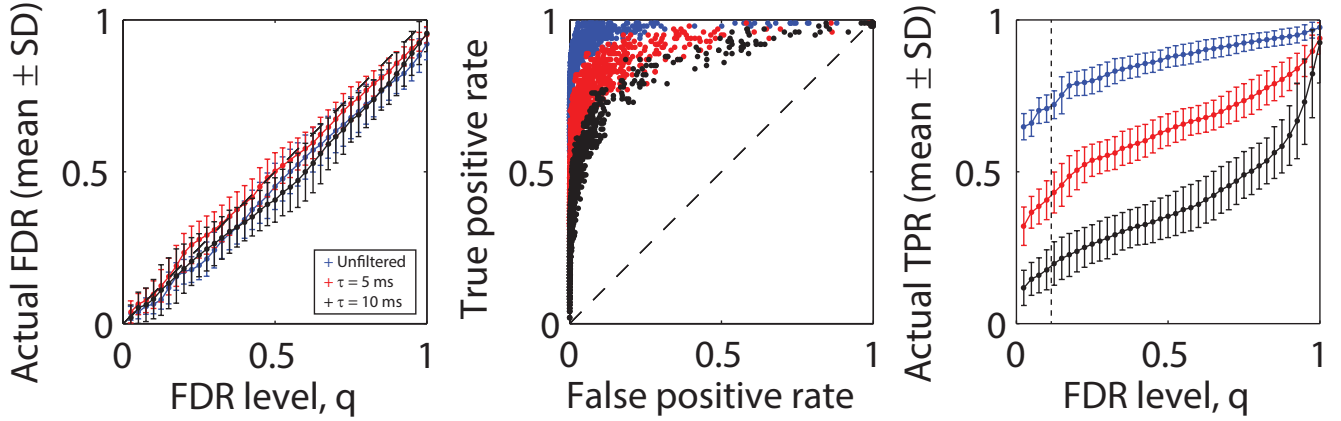[13] T. Schreiber, Phys. Rev. Lett. **85**, 461 (2000).

FIG. 7. Variation in the performance of undirected edge inference, when the neural signals are low-pass filtered prior to FDR-based inference. Unfiltered results are shown in blue, results with $\tau = 5$ ms LPF are shown in red, and results with $\tau = 10$ ms are shown in black. [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 overall true positive rate (TPR).
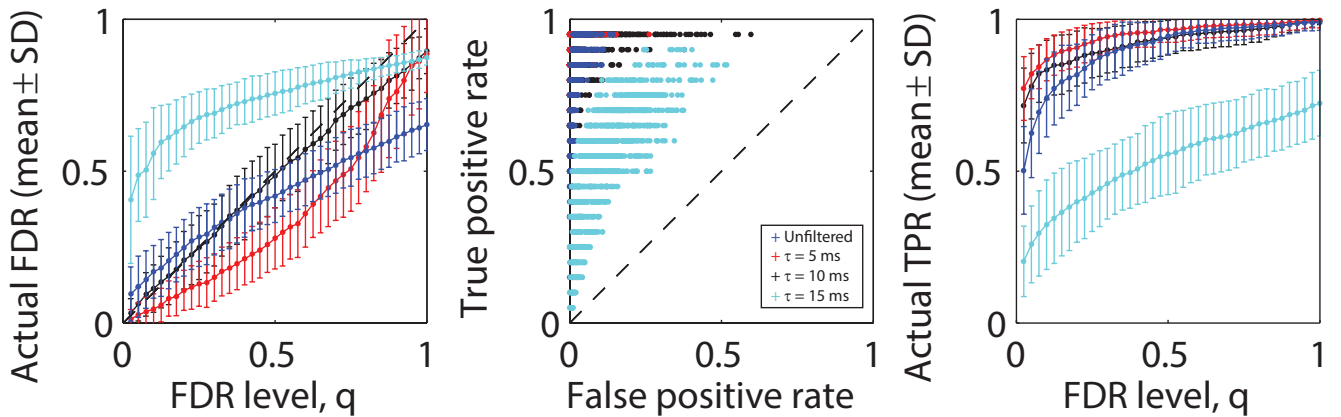

FIG. 8. Variation in the performance of directed edge inference, when the neural signals are low-pass filtered prior to FDR-based inference. (See the caption of Fig. 7 for details.)