

CS 224w Project Proposal: Neural network inference from its time series

Tony Hyun Kim

October 20, 2013

1 Reaction paper / Related work

1.1 Motivation

I am interested in the task of *network inference* based on 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 very soon thousands, given the pace of the field) 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 am interested in 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.¹ Generally, my intent for the project is to perform “connectivity analysis” on a simulated data of neuron populations, to explore the use of various coupling metrics on time series data (e.g. linear cross-correlation, Granger causality, etc.), and to use the computational framework to address several basic questions relevant to the experimental practice of optical neural recording.

Next, I describe specific papers relevant to my task.

1.2 Network inference with confidence from multivariate time series [KECK09]

Consider the cross-correlation, which is a standard basic metric for the “coupling strength” between two variables (e.g. two time-series $x_1[t]$ and $x_2[t]$). 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

¹Actually, I had intended to perform analysis on my original experimental data, but decided that it would be beyond the scope of the course project. Perhaps in CS341.

if the 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 unprincipled.

I am thus drawn to the work by Kramer, *et. al.* [KECK09] describing a rigorous foundation for connectivity inference based on the concept of “false discovery rate” (FDR) [BH95]. In essence, the FDR-based network inference makes conservative estimates regarding the existence 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, and provides “analytical shortcuts” when the choice of the coupling metric is linear cross-correlation.

I am conceptually drawn to the idea of using FDR control (over thresholding) for edge inference. However, Kramer’s paper does not give a convincing summary of the expected performance of the method. (Two simulated examples are provided, with known ground truths, but only a single instance of the inference is presented.) In Section 2.2.1, I performed repeated FDR-based inference on a toy network under various conditions, to verify the performance of the method.

In the current work, I also intend to generalize Kramer’s method by utilizing alternative coupling metrics, including ones that assign a directionality to the edge. Kramer alludes to several alternative metrics, *e.g.* the synchronization likelihood, wavelet coherence, Granger causality, and the directed transfer function.

1.3 Simple model of spiking neurons [Izh03]

My work will be based on simulations of neural dynamics. There are several choices for the model, including [Izh03, Ger95]. Kramer’s work on FDR-based edge inference (as well as other studies, *e.g.* [MPF05]) use the simulated neural model from Izhikevich [Izh03]. Based on this observation, I will be basing my simulations on the Izhikevich model. However, for generality, I will seek to demonstrate the conclusions of my work on multiple numerical models.

I will describe the Izhikevich model in more detail in Section 2.2.2.

1.4 Miscellaneous references

1.4.1 Relation to information cascades [GRLK10]

The concept of inferring network edges via a coupling measure inferred from temporal data has conceptual similarities to *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 (voltage) 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, through which the electrical activity of neurons can be directly modulated by the experimenter via light illumination [PYG⁺12] (the effect can be excitatory or inhibitory). In contrast to the endogenous neural activity, it may be possible to identify the result of an optogenically-driven activity as a cascade. I am interested in the possibility of capturing this idea (identification of unique cascades) in my simulations, and to quantify the additional information – if any – that the externally-driven cascade reveals about the network over network analysis of endogenous activity. More generally, I would like to evaluate whether cascade analysis algorithms can be integrated into the FDR-based network inference method.

1.4.2 Application of network theory to brain science [BS09]

The review paper by Bullmore and Sporns gives many references to the applications of graph theory to neuroscience, though much of the prior work is not at the individual cellular level.

2 Proposal

2.1 Project goals

2.1.1 Basic machinery

I will begin the work by building the machinery for FDR-based network inference, based on the work of [KECK09]. I will begin with the simplest coupling metric of linear cross-correlation. Using the cross-correlation, I will apply the inference method to simulated data of neural activity according to the Izhikevich model [Izh03]. The performance of the FDR-based inference method will be observed, by comparison of the predictions against the ground truth.

2.1.2 Robustness of network inference to degraded temporal signals

Currently, a major limitation of the optical recording method is that the fluorescent proteins involved in signaling functional activity have a slow temporal response – at roughly 10 Hz – than the voltage dynamics, which occurs at 1 kHz scale. In effect, the optical recording paradigm described in Section 1.1 yields a low-pass filtered version of the underlying voltage dynamics.

As the Izhikevich numerical model yields simulated traces that represent voltage dynamics with a physical scale (mV voltage; ms time scale), I am interested in the robustness (or lack thereof) of the FDR inference method to filtering of the temporal signals.

2.1.3 Complex network topology

The physical neuronal network (as well as its simulation model) implements directed edges, whereas the cross-correlation metric infers an undirected network.

I will simulate more complex network topologies (rather than just random edges) such as $n_1 \rightarrow n_2 \rightarrow n_3$, and consider the performance of the undirected network inference on such structures (*e.g.* how often does it ascribe an edge to $n_1 \rightarrow n_3$?). I would also consider the case in which portions of the underlying network are hidden from the inference analysis (*e.g.* hide n_2). This models the experimental fact that, even with the state-of-the-art recording methodologies, we have no *a priori* reason that our recording site reveals information over all neurons relevant to a particular biological function. It is thus interesting to examine how the performance of network inference varies when portions of the neural circuits are “hidden” from analysis.

2.1.4 Higher-order network analysis

On one hand, we do not expect that a set of neurons observable at a particular recording site captures all neurons relevant to a biological function. On the other hand, as we improve the capture efficiency of our

recording devices, we also expect that many neurons *unrelated* to a biological task at hand would also be present in the dataset.

Thus, I would like to evaluate the performance of the inference mechanism not only by the correctness of each inferred edge, but also by its ability to identify, for example, the connected components (so that we can disregard unrelated disconnected components), and other graphical structures.

2.1.5 Alternate coupling metrics

I am interested in exploring the use of other coupling metrics, especially ones that ascribe a directionality to an edge. Based on my early review of the literature, Granger causality appears to be a widely-accepted method for inferring directed edges. With alternate metrics, I am interested in re-examining the questions discussed above.

2.1.6 Connection to cascades

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? The resulting “cascade” can be quantified in the presence or the absence of endogenous activity in the remainder of the population. Can we adapt methods on cascade-based network inference to our FDR-based network inference, in order to improve the performance?

This question has implications for experimental neuroscience, where there are efforts underway to integrate recording and perturbation (optogenetic) methods [PYG⁺12].

2.2 Initial ground work

The following section describes some early work that I have performed. Note: I often find it useful to maintain a blog that describes progress on projects such as this one. I am making a record of my progress at: cs224w.blogspot.com.

2.2.1 FDR performance on a toy network

I recreated the FDR-based inference model of [KECK09]. As mentioned previously, the original publication did not give a complete sense of the performance of the method, as it gave the results of only a single inference instance. I am interested in the distribution of results over multiple inference instances.

As in [KECK09], I modeled a toy network of $N = 9$ nodes and $M = 9$ directed edges. The temporal signal $n_i[t]$ at each node is generated as follows:

$$n_i[t] = w_i[t] + \alpha \cdot \sum_{(j,i) \in E} w_j[t] \quad (1)$$

where $\{w_i\}$ are independent white noise signals for each node i , and $\alpha = 0.4$ represents the coupling constant.

I performed $N_{\text{inf}} = 10^5$ instances FDR-based edge inference with an FDR-level $q = 0.1$ (which means that, in expectation, 10% of our inferred edges will be incorrect). The results are shown in Fig. 1.

Some interesting features of Fig. 1:

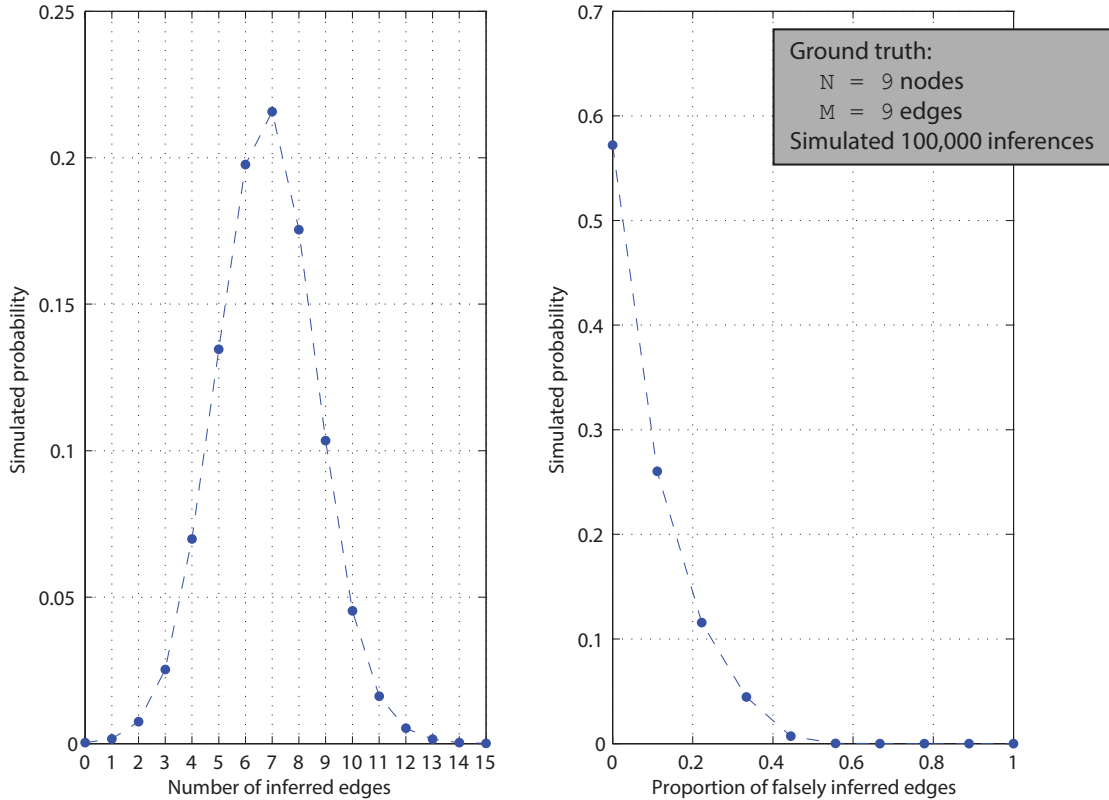


Figure 1: Results of applying FDR-based network inference method of [KECK09] on a toy network of $N = 9$ nodes and $M = 9$ randomly-chosen, directed edges, with FDR-level $q = 0.1$. The number of inferences is $N_{\text{inf}} = 10^6$. [Left] The distribution of the number of edges in the inferred network. [Right] The distribution of the false positives in the inferred network.

- On average, the inferred network has $M_{\text{inf}} \approx 7 < M = 9$ edges. In the toy network, the inference method tends to be conservative.
- The mean of the proportion of falsely inferred edges is $q_{\text{inf}} \approx 0.07 < q = 0.1$ as advertised, which means that roughly 7% of the inferred edges are false positives, in expectation.

2.2.2 Numerical simulation of neuron dynamics

Fig. 2 shows the basic implementation of the Izhikevich neuron model. Only a subset of the $N = 1000$ neurons are shown, whose connectivity structure was randomly generated.

Based on my initial work with the Izhikevich model, few issues were apparent:

1. Depending on choice of the numerical parameters, the dynamical model for a single neuron can show differing firing characteristics [Izh03]. I will have to make a simplifying choice for the type of neuron that I wish to model. (In the neural simulation example of [KECK09], the “simplest” type of neurons – known as “regularly spiking” (RS) neurons – are simulated.)

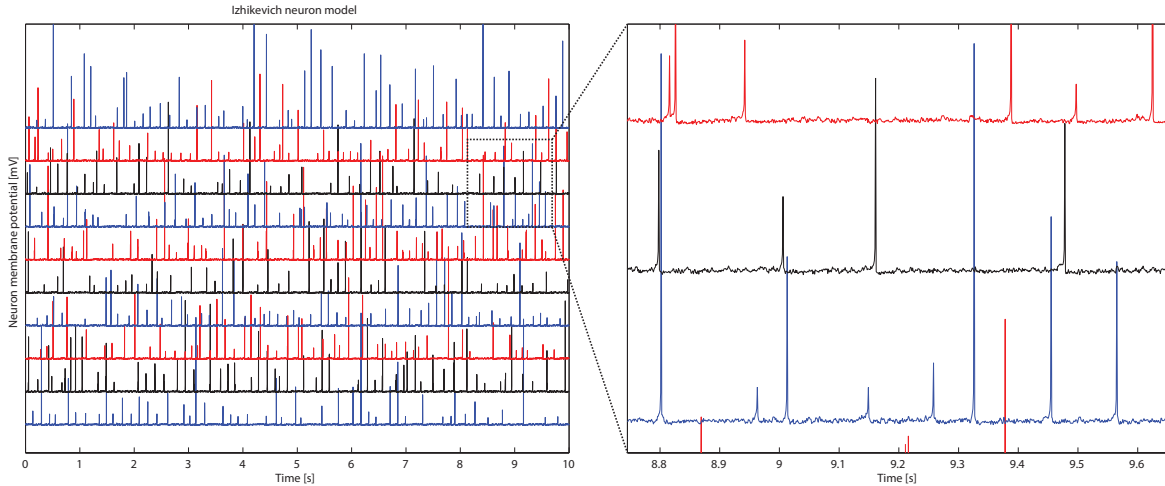


Figure 2: Example instance of the Izhikevich neuron model. Only a subset of the $N = 10^3$ simulated neurons are shown. The (directed) connection strengths between the neurons was generated randomly.

2. Izhikevich’s presentation of the model is very mathematical, and does not give an assessment of the reasonable range of parameters for the connection strengths. I have chosen to interpret the connection strengths in the Izhikevich model in terms of the amplitude of the excitatory postsynaptic potential that it generates – and determined a range of values that are biologically plausible.
3. I find that the simple numerical integration method proposed by Izhikevich is unstable over short time frames. In particular, as shown in Fig. 2, the height of the action potential is often too high (and is a result of numerical noise). This numerical instability perturbs the calculation of the analog cross-correlations. I will have to implement the integration at finer time steps.

Another possibility is to base the time-series analysis not on the voltage traces as shown in Fig. 2, but on digitized traces (*i.e.* $n_i[t] = 1$ if neuron i fired at time t , otherwise $n_i[t] = 0$). Performing analysis at the “digital trace” level would have natural correspondence to the conventional method of analyzing optical neural measurements, which digitizes the measurements.

References

- [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.

- [Ger95] W. Gerstner. Time structure of the activity in neural network models. *Phys. Rev. E*, 51:738–748, 1995.
- [GRLK10] M. Gomez-Rodriguez, J. Leskovec, and A. Krause. Inferring network of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data*, 2010.
- [Izh03] E. M. Izhikevich. Simple model of spiking neurons. *IEEE Transactions on Neural Networks*, 14:1569–1572, 2003.
- [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.
- [MPF05] V. A. Makarov, F. Panetsos, and O. Feo. A method for determining neural connectivity and inferring the underlying network dynamics using extracellular spike recordings. *Journal of Neuroscience Methods*, 144:265–279, 2005.
- [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.