A Granger Causality Measure for Point Process Models of Neural Spiking Activity

Diego Mesa
PhD Student - Bioengineering
University of California - San Diego
damesa@eng.ucsd.edu

Abstract

A network-centric view of brain function that is becoming more widely accepted would benefit from the directional interaction information that occurs between multiple neurons. Granger causality has been used previously to address this need, but previous methods can only operate on continuos-value data sets. This prevents it from being directly applied directly to neural spike train data so previous attempts have involved smoothing, binning, etc. to address this. Recently a point process framework that enables Granger causality to be applied to point process data was proposed. The newly proposed framework was used to investigate network interactions proposed in previous assignments to verify that the correct minimal generative model graph could be recovered from just the raw simulated neural spike train data. The interactions and non interactions present in the simulated networks were indeed recovered.

1 Introduction

Previously help beliefs about neural function maintained that different areas of the brain were task specific; that the structural connectivity local to a certain area somehow dictated the function of that piece. Collecting work that has been performed over many years, there has been a move to a different, network-centric approach to describing information flow in the brain [1]. Explanation of function is beginning to include the concept of networks existing at different levels and throughout different locations in the brain. The behavior of these networks are best described by non-deterministic processes that are occurring all the time. That is to say that given the same input stimulus, you will not get the same output from the network. The dynamics of these networks are governed by probabilities so it is more useful to start thinking in terms of stochastic (random) processes that capture these kinds of dynamics between different areas of the brain.

Different methods of obtaining some measure of information flow from the firing activities of a neuron and its surrounding ensemble has been explored in the past, but they are limited in the kinds of conclusions you can draw and tell you little about the directional flow of information, to what degree, and how it can change with time [2]. One method that addresses most of these issues is Granger causality. It involves whether or not your prediction was improved based on the inclusion of a set on information. In other words, if the past of X contains information that helps us better predict the future of Y, then we can say that X has a causal influence on Y. Previous granger-causality methods could only operate on continuous-valued data so the analysis of neural spike train recordings involved some kind of transformations that ultimately altered the stochastic properties of the data, indirectly altering the validity of the conclusions that could be drawn from it [2]. Recently however, a new general-purpose granger-causality framework was proposed that could directly operate on any modality, including neural-spike trains [4].

2 Methods

For this project, three simple networks previously proposed and implemented in class were re-implemented and simulated as point-processes for 100 seconds. They were then joined together to form one data set. This was done in an attempt to recover both connections within the simple networks themselves, as well as no connections between the networks.

2.1 Point Processes

The theory of point processes and their use in the proper representation of neural data has been previously explored in great detail [4]. Only a basic representation was used to model these networks, as explained below.

A temporal point process is a stochastic time-series of binary events that occurs in continuous time [5]. It can only take on two values at each point in time, indicating whether or not an event has actually occurred. This type of binary-valued representation of information suits the activity of neural populations because a single neurons action potential has a typical waveform. In this way, what carries the actual information being output from a neuron is the occurrence of a spike, as well as the time between successive spikes. Using this approach one could abstract the flow of information in a neural-network to be simply the spiking times for each neuron through an observation period. As can be seen in Figure 1, a point-process can be represented either by the timing of the spikes themselves, the waiting times between spikes, using a counting process, or, if time is discretized small enough to ensure that in each window only one event has the possibility of occurring, that is to say one time bin can only contain one event, as a set of 1s and 0s, very similar to binary.

One of the simplest types of neural-spiking models is the Poisson process [5]. This however, is limited in that it is memory-less. It does not account for any spiking history when calculating the current probability of firing. Neurons, however, exhibit a fundamental (biophysical) history dependence by way of its relative and absolute refractory periods. To address this we use a conditional intensity function to represent the probability of a neuron spiking, conditioned on its own history. The conditional intensity function expresses the instantaneous firing probability and implicitly defines a complete probability model for the point process [6]. It defines a probability per unit time. So if this unit time is taken small enough to ensure that only one spike could occur in that time window, then our conditional intensity function completely specifies the probability that a given neuron will fire in a certain time.

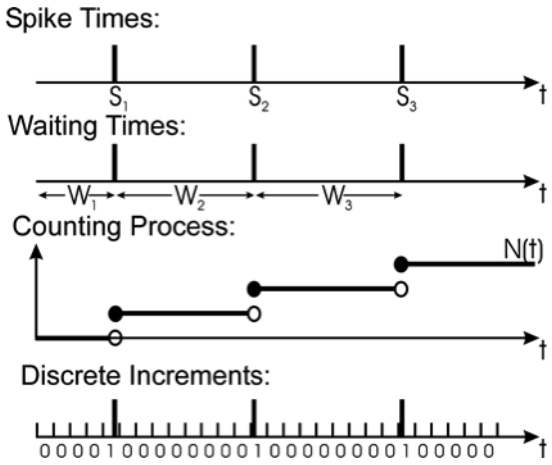

Figure 1: Point Processes

2.2 Generalized Linear Model

One may wish however to account for a number of covariates when analyzing neural data, not just a neurons own history. These covariates could include the spiking history of surrounding neurons or the onset of some external stimulation. Using a generalized linear model, one could construct a parametric model of the conditional intensity functions of these point processes. This allows us to intuitively model relationships and the biophysical properties we think are relevant in a certain situation. Under the GLM, the log of the conditional intensity function is modeled as a linear combination of functions of the postulated covariates. [5]

The GLM model is and has been used in countless numbers of papers and has been studied and optimized extensively, [2] so it also suitable for large network analysis. It should also be noted that the linearity of the model refers to the relationships between the covariates, not the functions of the covariates themselves. There are many different families of generalized linear models [7], and different techniques that have been proposed to optimize statistically and computationally. For the purposes of this paper, a very simple model was proposed that relates a neurons spiking probability to its own history and that of the ensemble. To determine the model order, or rather, how far back one should look in a neurons own history as well as those of its ensemble, several different orders were estimated and the Akaikes standard information criterion [8] was calculated to select the best model order.

2.3 Granger Causality

As mentioned previously, the granger causality framework was recently revisited [8] to generalize it across almost any modality. If we start from Grangers basic formulation, we have X causes Y if I can predict the future of Y given the past of Y and the past, present of X better than I can predict the future of Y given only the past of Y. From the methods proposed above, it is easy to see how we can predict the future of Y, given a set of information we think is relevant. The question now becomes, how exactly can we determine the effect of X on our prediction, and how can we say we can predict better with/without it.

To determine the effect of X on our prediction of Y, we use a point process likelihood function to fit the parametric conditional intensity function and analyze Granger causality between them. We determine if a causal relationship from X to Y exists by first calculating the likelihood of Y producing a particular set of spike trains, if the spiking history of neuron X is excluded. We then compare that with the calculated likelihood if all of the available covariates were used and calculate a reduction in the likelihood. The sign of the relationship can distinguish excitatory and inhibitory influences where a positive result is indicative of an excitatory effect and a negative result indicates an inhibitory effect. Zero indicates that no causal interactions were detected. [8]

2.4 Significance Testing

While a great many varieties of statistical significance tests exist, and anyone is free to choose whatever analysis they deem appropriate, for the purposes of this paper, a simple pair-wise hypothesis test was performed, controlling the False Discovery Rate at 0.01.

2.5 Network Simulation

The three networks in Figure 2 were described and implemented in project 4 earlier in the semester, so they were simulated using a simple GLM. These 9 neurons were simulated according to these simple graphs for 100 seconds, resulting in 100,000 samples per neuron. The raw data was then analyzed using the granger causality framework, without specifying absolutely anything about their connection. These three simple networks were analyzed as one large network in order to both verify that we could directly pick up directional information within the simple networks themselves, as well as the absence of any connection between them. In other words, as far as the granger causality analysis is concerned, the analysis was performed on recordings from 9 neurons for 100 seconds, without knowing anything about them. These neurons are then numbered from 1- 9, with network A being 1-3, network B being 4-6, and network C being 7-9.

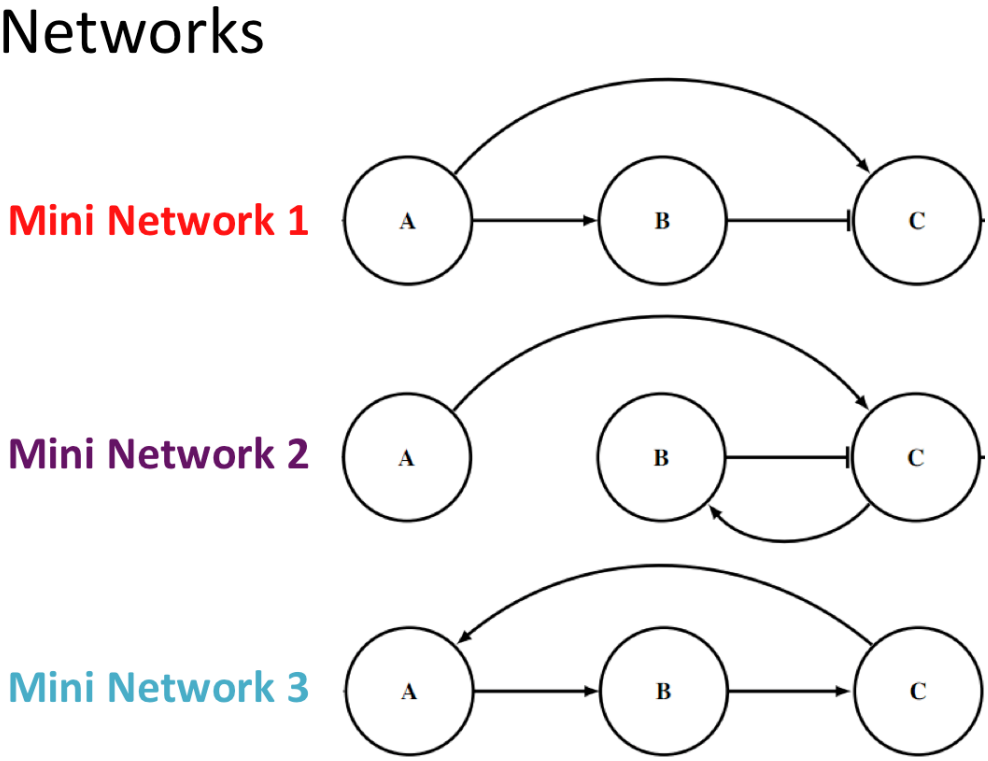

Figure 2: Network Simulations

3 Results

As an intermediary sanity check, any kind of GLM model we estimate to fit the data should correctly show how much each other neuron weighs in determining whether or not it would fire. As we can see in Figure 3, for neurons belonging to mini network 3, we do indeed correctly guess as to the weights each neuron has on these neurons firing probability. Using 2 parameters for each neuron (a model order determined by the AIC described above), we can see that each neuron shows a negative weight (represented by negative parameters) when it comes to itself. This is representative of the fact that each neuron was modeled with a self-inhibition representative of its absolute and relative refractory periods.

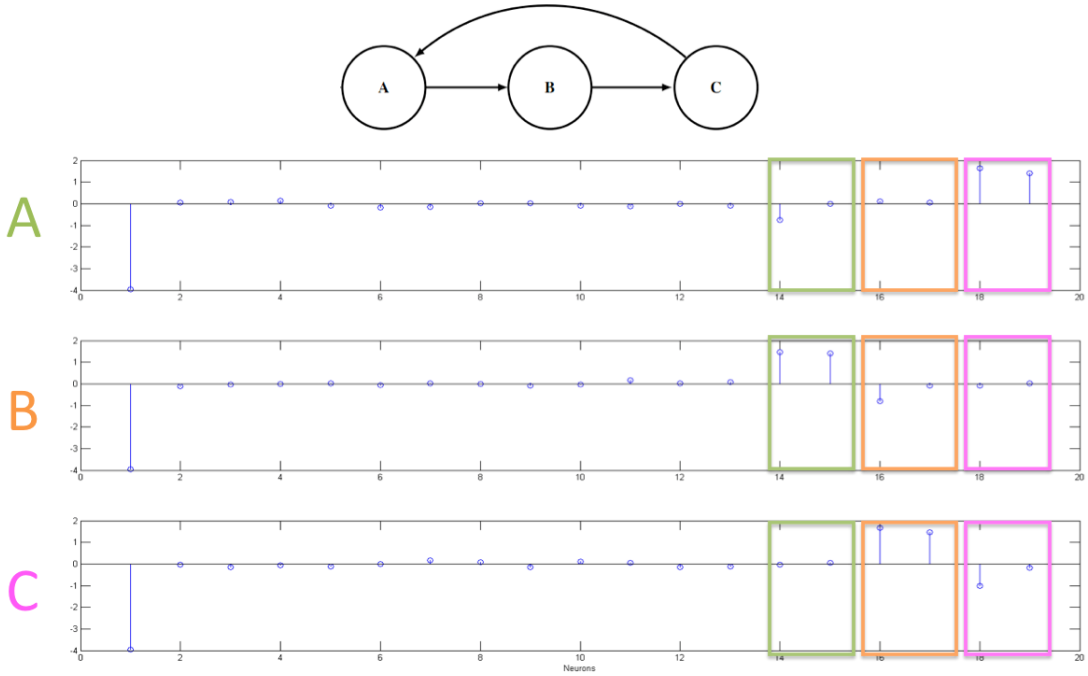

Figure 3: Parameter Estimations

Also, looking at Neuron A, the only neuron in the entire network of 9 that should have any direct effect on its firing probability is neuron C, and indeed, we see this by positive parameter values, representative of the excitatory connection between them. Neuron B also correctly shows that Neuron A is the only other neuron who affects its firing probability, and it does so positively, reflective of its excitatory connection. Neuron C also correctly shows that neuron B is the only neuron that affects it. Similar figures could be generated for the two other networks for verification.

After performing the aforementioned granger causality analysis, we are left with Figure 4, showing relationships from causal source neurons to causal sinks. That is to say, a causal source neuron acts as the start of a causal, directed arrow, where a causal source neuron acts as its end point. While this figure shows the strength of causal relationships, it has not yet had any significance testing performed on it.

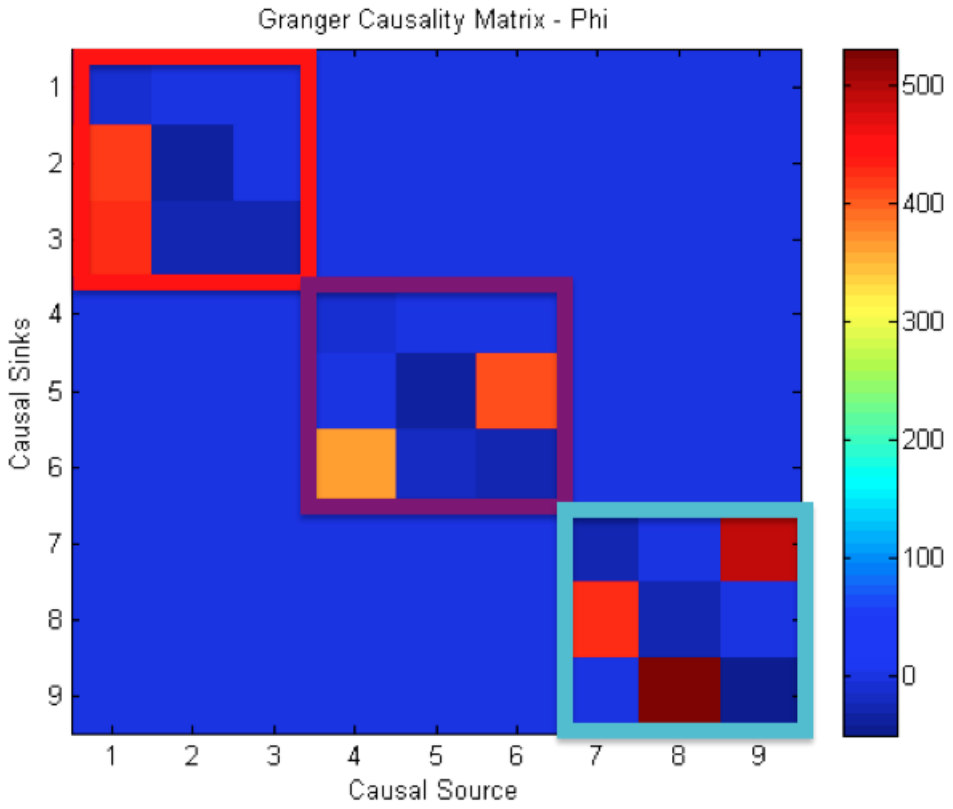

Figure 4: Granger Causality Matrix

After the hypothesis testing was performed, controlling the FDR at 0.01, we are left with Figure 5. As mentioned previously, if you start on a column, and draw an arrow from the column neuron to each row neuron where a red(excitatory) or blue(inhibitory) box resides, (green is no connection), you will generate the minimal generative model graph for this data set. As we can clearly see, we do indeed correctly direct connections within the networks, and no connections are detected between the networks.

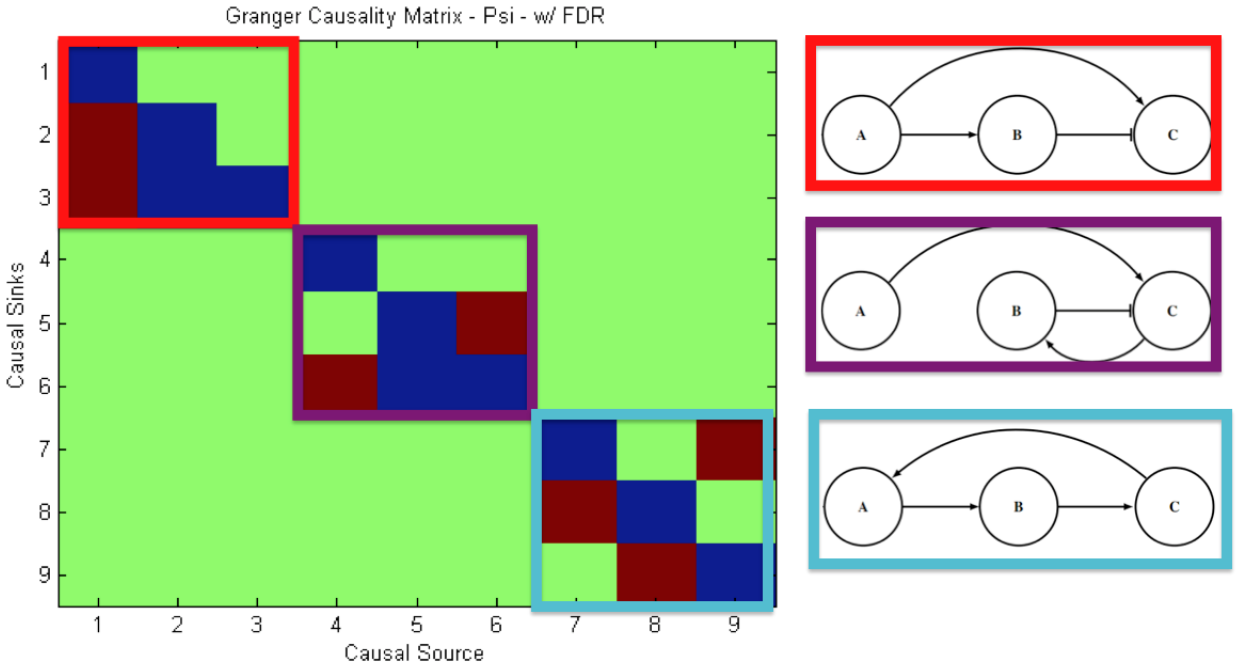

Figure 5: Granger Causality Matrix - With Statistical Significance Testing and FDR

4 Discussion

As we can see from the results mentioned above, the granger causality framework operating directly on the neural spike trains does indeed correctly reconstruct the networks used to simulate the data. This is a very satisfying result indeed, because from nothing more than the spike train recordings of each neuron, we were able to correctly reconstruct the rules used to generate them.

To continue developing a robust, concise description of statistical dynamics underlying neural activity, the development and testing of this framework needs to be continued and expanded to include not only other GLM families, but other point process models as well to test its robustness.

The analysis also needs to be continuously optimized so that it becomes more feasible to run this analysis across larger and larger data sets. Also, as it becomes more and more optimized, we can run the analysis on smaller time windows, giving us a picture of how causality can change through time. For example, in the analysis ran above, we generated what can be called the "average causality graph" for this network. In our case it worked out wonderfully because the data was generated using exactly those rules. That is to say, the rules used to generate the data did not change with time. In a real neural network however, that might not actually be the case. We therefore need to run the analysis with less samples, so that we could examine how the causality relationships evolve through time. Statistically however, less samples limits the strength of our assertions. To get around that we could average across a large number of trials of the same experiment. For example, suppose the same stimulate-measure experiment was conducted 10,000 times. We could then average all 20 second time intervals across all 10,000 trials, and run the analysis on each 20 second time window. We could then visualize this using some animation that shows how these directed arrows change through time, giving us a powerful visualization of the data, allowing us to make new and exciting statements about the nature of stochastic neural networks.

Acknowledgments

I would like to thank Sanggyun Kim, Todd Coleman, and Jeff Bush for their help with this project.

References

[1] R. T. Knight, ”Neural Networks Debunk Phrenology”, Science, 2010

[2] Brown, Kass, Mitra - ”Multiple Neural Spike Train Data Analysis: State of the Art and Future Challenges”, Nature Neuroscience, 2004

[3] Quinn, Coleman, Kiyavash, Hatsopoulos - ”Estimating the directed information to infer causal relationships in ensemble neural spike train recordings”, Journal of Computational Neuroscience, 2011

[4] Truccolo, Eden, Fellows, Donoghue, Brown - ”A Point Process Framework for Relating Neural Spiking Activity to Spiking History, Neural Ensemble, and Extrinsic Covariate Effects” - Journal of Neurophysiology, 2005

[5] Eden - ”Point Process Models for Neural Spike Trains” - Society of Neuroscience Short Course, 2008

[6] Daley, Vere-Jones - ”An Introduction to the Theory of Point Processes” Springer, 2003

[7] Zanos, Courellis, Berger - ”IEEE Transactions On Neural Systems and Rehabilitation Engineering”, 2008

[8] Kim, Putrino, Ghosh, Brown - A Granger Causality Measure for Point Process Models of Ensemble Neural Spiking Activity., PLoS Computational Biology, 2011