ESANN'2007 proceedings - European Symposium on Artificial Neural Networks
Bruges (Belgium), 25-27 April 2007, d-side publi., ISBN 2-930307-07-2.

Causality and communities in neural networks

Leonardo Angelini, Daniele Marinazzo, Mario Pellicoro, Sebastiano Stramaglia

TIRES-Center for Signal Detection and Processing - Università di Bari, Bari, Italy
Dipartimento Interateneo di Fisica, Bari, Italy
Istituto Nazionale di Fisica Nucleare, Sezione di Bari, Bari, Italy

Abstract. A recently proposed nonlinear extension of Granger causality is used to map the dynamics of a neural population onto a graph, whose community structure characterizes the collective behavior of the system. Both the number of communities and the modularity depend on transmission delays and on the learning capacity of the system.

1 Introduction

Identifying causal relations among simultaneously acquired signals is an important problem in computational time series analysis. One major approach to analyze causality between two time series is to examine if the prediction of one series could be improved by incorporating information of the other, as proposed by Granger [1] in the context of linear regression models of stochastic processes. In particular, if the prediction error of the first time series is reduced by including measurements from the second time series in the linear regression model, then the second time series is said to have a causal influence on the first time series.

In this paper we propose the use of causal interactions, detected as described above, to map a network of spiking neurons with synaptic plasticity, which mimics a sensorial cortex, onto a graph, in order to characterize collective behavior, and to mine patches of interdependent neurons by means of the search of community structures in the graph [2].

2 Discovering Causal Communities

It is well known how synaptic plasticity models the connections between the neurons on the basis of the activity [3]: that is the dynamics and the functional properties of the networks are tightly connected with the architecture. In our view causality relations are the key of collective behavior, and constitute a suitable way to detect functional communities [4]. We thus first look for causal interactions, then we use this information to discover community structures in the neural network.

2.1 Granger Causality

We recall here the key concepts, while referring the reader to [5, 6] for all the details. As stated in the Introduction, the notion of Granger causality between two time series examines if the prediction of one series could be improved by

incorporating information of the other. The interactions between individual neurons in a network are strongly nonlinear: we thus address the problem of nonlinear Granger causality by means of the generalization of a radial basis function approach. It is important to point out that not all nonlinear prediction schemes are suitable to evaluate causality between two time series, since they should be invariant if statistically independent variables are added to the set of input variables.

2.2 Community Discovery Algorithm

A hierarchical structure of modules characterizes the topology of most of real-world networked systems [7, 8]. In social networks, for instance, these modules are densely connected groups of individuals belonging to social communities. *Modules* (called also *community structures*) are defined as tightly connected subgraphs of a network, i.e. subsets of nodes *within* which the density of links is very high, while *between* which connections are much sparser. These tight-knit modules constitute units that separately (and in parallel) contribute to the collective functioning of the network. For instance, the presence of subgroups in *biological* and *technological* networks is at the basis of their functioning. Hence the issue of detecting and characterizing module structures in networks received considerable amount of attention .

Many methods have been proposed to partition a group into communities. A central quantity used by hierarchical clustering techniques is the modularity, which was introduced by Girvan and Newman [2, 9, 10] in connection to an algorithm able to produce a hierarchy of subdivisions in a network, from a single agglomerate to isolated nodes. The modularity Q (see [2] for the mathematical definition) is a measure of the correlation between the probability of having an edge joining two sites and the fact that the sites belong to the same modules. It is clear that, if a partition in a fixed number of subgroups of the networks is requested, the best solution is obtained minimizing the number of edges connecting vertices belonging to different subgroups (or minimizing the number of vertices belonging to the same subgroup). A good division into communities is the one in which the number of edges between vertices belonging to the same group is significantly greater than the number expected by a random distribution of edges. As an example of a network, the three-community structure resulting from the analysis of a reduced version of the population of excitatory and inhibitory neurons described in the next sessions is reported in figure 1.

We then address the problem of detecting communities in a network in the following way: to find the optimal number of clusters the modularity Q is calculated at different values of n_c, the proposed number of modules. Once n_c is determined, the optimal clusterization is found maximizing an objective function called *ratio association*, as described in [11].

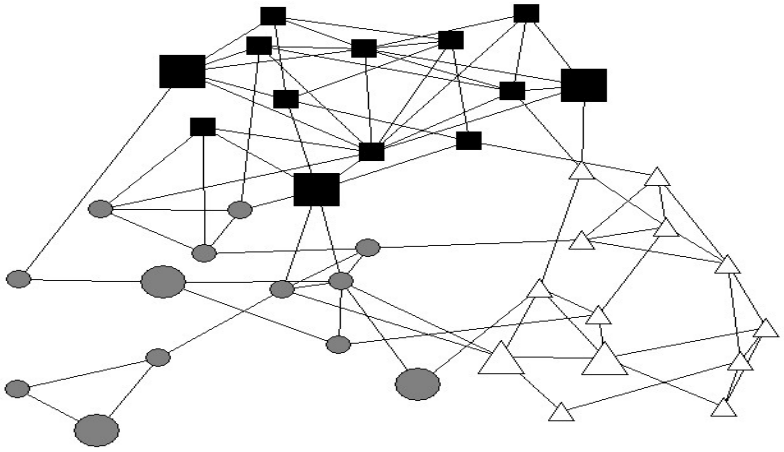

Fig. 1: Visualization of a reduced network of neurons, divided into three communities. Larger nodes indicate inhibitory neurons.

3 Application: Network with delays and spike-time dependent plasticity

We test our procedure on data obtained with the neural network model proposed by Izhikevich [12]. It consists of 1000 randomly connected neurons with spike-timing-dependent plasticity (STDP) and conduction delays. We introduced in the model the possibility of fixing the maximum delay, varying it from 0 to 30 ms. As shown by Izhikevich, the rhythmic activity of the spiking model evolves with time, presenting delta frequency oscillations at the very beginning. Then, as synaptic weights evolve according to STDP, uncorrelated Poissonian activity, and then gamma frequency oscillations, appear. The mechanism generating these rhythms is called PING (pyramidal-interneuron network gamma) [13]. This is explained by individual excitatory cells which fire at or near gamma frequency, and their active participation is crucial: the excitatory cells drive and synchronize the inhibitory ones, which in turn gate and synchronize the excitatory ones. It is clear how in this case there is a whole activity of functional communities and causal interactions going on.

For each of the three cases listed above, we extracted time series of length 1000 points, from the PSP of any presynaptic neuron, and from the membrane potential of the corresponding postsynaptic neurons. Doing this we take also into account the effects of the external input on the subthreshold behavior of the neurons [14]. The series were checked for covariance stationarity with a Dickey-Fuller test ($p < 0.01$). The Granger causality algorithm was then applied. We identified statistically relevant interactions performing an F-test (Levene test) of the null hypothesis that the error on the prediction of one series is not decreased when information on the other series is added to the model. The

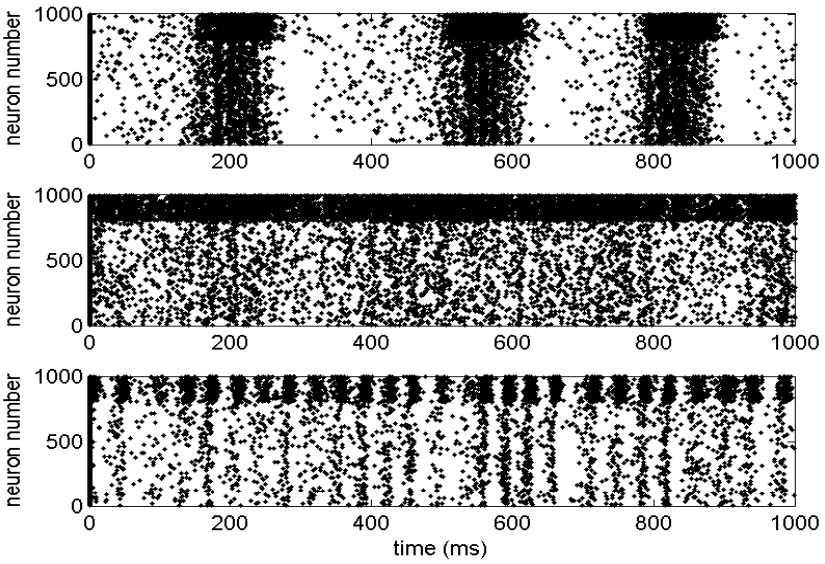

Fig. 2: The three phases displayed by the model as time, and thus STDP, evolves. Top: delta rhythm. Middle: uncorrelated Poissonian activity. Bottom: emergency of gamma rhythm.

neurons connected by significant causal interactions were marked as connected in our causal adjacency matrix, which was then analyzed by the community discover algorithm described in the previous section.

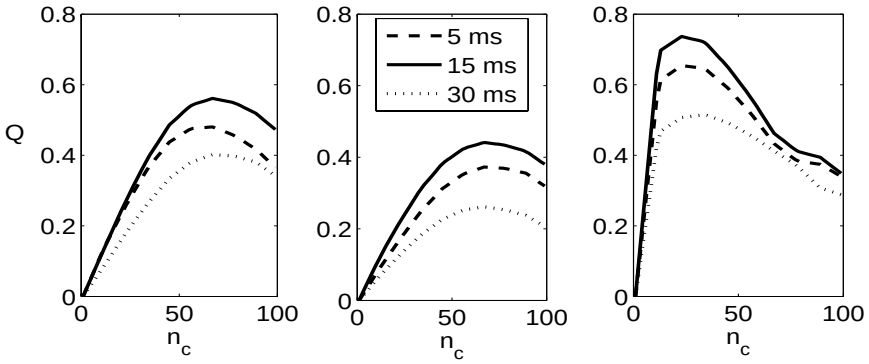

Fig. 3: The averaged modularity as a function of the proposed number of communities, for different values of the transmission delays, for the three phases. Left: delta rhythm. Center: uncorrelated Poissonian activity. Right: gamma rhythm.

Figure 3 displays the average modularity as a function of the number of communities. The number of communities which maximizes the modularity varies with time. At the beginning, when there is a dominant delta rhythm, the optimal number of causal communities is around 60. Then, as the synaptic weights evolve according to STDP, this number reaches 70. But, when gamma rhythms begin to emerge, it goes back to around 30. This is because the increased synchronization of inhibitory and excitatory neurons leads to the formation of a lower number of clusters with common activity. Furthermore, we can see that the number of communities is not affected by the maximum delay in the interactions. This is explained by the fact that our model for detecting causality spans many past values of the time series, thus detecting causal interactions also within polychronous groups, that is groups of neurons which maintain a fixed, though not synchronous, pattern of firing times. On the other hand, the modularity is affected by the internal delays, and it displays a maximum around $D = 15ms$. This behavior can be explained if one considers that, for very small delays, we have no polychronous groups, and not all the possibly relevant interactions are taken into account. On the other hand, for large delays, we have many interactions at different times, thus adding interactions which don't contribute to form new communities but add some "noise" to the adjacency matrix. This preferred value of internal delays has also been observed when causal interactions were investigated when external inputs with different spatiotemporal characteristics were presented to the neural network.

Linear Granger causality has been successfully applied to networks of neurons in a "brain-based device" [15]. But when, as in the present case, the presynaptic and postsynaptic activity of the neurons is simulated with more biologically plausible equations (leaky integrate and fire, or conductance-based), we believe that the strong nonlinear character of the neural behavior should not be neglected. Indeed the application of linear Granger causality to the model described above still gave rise to a plausible network structure, but with significantly lower values of modularity and a less defined profile for different values of delays and community number.

4 Conclusions

We have shown that nonlinear Granger causality measures may be successfully applied to map a complicated dynamical system, such as a population of inhibitory and excitatory neurons, onto a graph whose edges connect pairs of neurons with significant causal relations. By subsequently analyzing the community structures in the graph obtained, we found interesting insights on the collective behavior of the system. In particular, we found that the delays don't change the number of communities, but influence the modularity, that is to what extent these communities are interconnected. On the other hand, the number of communities evolves with time as the process of learning models the synaptic weights.

Our findings confirm the potentiality of the notion of causality in the analysis

of complex systems and the importance of the nonlinear extension of Granger causality to model them.

References

[1] C W J Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–38, July 1969.

[2] M Girvan and M E J Newman. Community structure in social and biological networks. *Proc. Natl. Acad. Sci. USA*, 99:7821, 2002.

[3] L F Abbott and S B Nelson. Synaptic plasticity: taming the beast. *Nat Neurosci*, 3 Suppl:1178–1183, Nov 2000.

[4] C R Shalizi, M Camperi, and K L Klinkner. Discovering functional communities in neural networks. In *proceedings of the ICML 2006 workshop on Statistical Network Analysis: Models, Issues and New Directions, to appear*. Springer, http://arxiv.org/abs/q-bio.NC/0609008.

[5] N Ancona, D Marinazzo, and S Stramaglia. Radial basis function approach to nonlinear Granger causality of time series. *Phys Rev E*, 70:056221, Nov 2004. Evaluation Studies.

[6] D Marinazzo, M Pellicoro, and Stramaglia S. Nonlinear parametric model for granger causality of time series. *Phys. Rev. E*, 73(6):066216, 2006.

[7] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393(6684):440–442, June 1998.

[8] S Boccaletti, V Latora, Y Moreno, M Chavez, and D U Hwang. Complex networks: Structure and dynamics. *Phys. Rep.*, 424:175, 2006.

[9] M E J Newman and M Girvan. Finding and evaluating community structure in networks. *Phys Rev E*, 69:026113, 2004.

[10] M E J Newman. Modularity and community structure in networks. *Proc. Natl. Acad. Sci. USA*, 103:8577, 2006.

[11] L Angelini, S Boccaletti, D Marinazzo, M Pellicoro, and S Stramaglia. Fast identification of network modules by optimization of ratio association. *Submitted*.

[12] E M Izhikevich. Polychronization: Computation with spikes. *Neural Computation*, 18(2):245–282, 2006.

[13] M A Whittington, R D Traub, N Kopell, B Ermentrout, and E H Buhl. Inhibition-based rhythms: experimental and mathematical observations on network dynamics. *Int J Psychophysiol*, 38(3):315–336, Dec 2000.

[14] W Gerstner, R Kempter, J L van Hemmen, and H Wagner. A neuronal learning rule for sub-millisecond temporal coding. *Nature*, 383(6595):76–81, Sep 1996.

[15] A K Seth and G M Edelman. Distinguishing causal interactions in neural populations. *Neural Computation*, in press.