Journal of Neuroscience Methods 236 (2014) 92–106

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

journal homepage: www.elsevier.com/locate/jneumeth

Computational Neuroscience

Revealing cell assemblies at multiple levels of granularity

Yazan N. Billeh[a,*,1], Michael T. Schaub[b,**,1], Costas A. Anastassiou[c], Mauricio Barahona[b], Christof Koch[c]

[a] Computation and Neural Systems Program, California Institute of Technology, Pasadena, CA 91125, USA
[b] Department of Mathematics, Imperial College London, London SW7 2AZ, UK
[c] Allen Institute for Brain Science, Seattle, WA 98103, USA

HIGHLIGHTS

- Modern experiments enable the simultaneous monitoring of the activity of hundreds of neurons.
- We present a dynamics-driven methodology to detect neural assemblies from spike-train data.
- Our biophysically inspired measure extracts directed functional relations between both excitatory and inhibitory neurons.
- Community detection is used to reveal groups at all levels of granularity, without prior knowledge of their expected size.
- We extensively assess our method with synthetic, simulated, and experimental spike-train data.

GRAPHICAL ABSTRACT

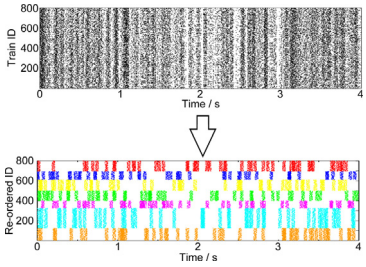

ARTICLE INFO

Article history:
Received 29 April 2014
Received in revised form 6 August 2014
Accepted 8 August 2014
Available online 26 August 2014

Keywords:
Biophysically inspired and directed spike-train association metric
Spike-train communities
Community detection
Clustering

ABSTRACT

Background: Current neuronal monitoring techniques, such as calcium imaging and multi-electrode arrays, enable recordings of spiking activity from hundreds of neurons simultaneously. Of primary importance in systems neuroscience is the identification of cell assemblies: groups of neurons that cooperate in some form within the recorded population.
New method: We introduce a simple, integrated framework for the detection of cell-assemblies from spiking data without *a priori* assumptions about the size or number of groups present. We define a biophysically-inspired measure to extract a directed functional connectivity matrix between both excitatory and inhibitory neurons based on their spiking history. The resulting network representation is analyzed using the Markov Stability framework, a graph theoretical method for community detection across scales, to reveal groups of neurons that are significantly related in the recorded time-series at different levels of granularity.
Results and comparison with existing methods: Using synthetic spike-trains, including simulated data from leaky-integrate-and-fire networks, our method is able to identify important patterns in the data such as hierarchical structure that are missed by other standard methods. We further apply the method to experimental data from retinal ganglion cells of mouse and salamander, in which we identify cell-groups that correspond to known functional types, and to hippocampal recordings from rats exploring a linear track, where we detect place cells with high fidelity.

* Corresponding author at: Caltech, 1200 East California Boulevard, MC 136-93, Pasadena, CA 91125, USA.
** Corresponding author at: Imperial College London, South Kensington Campus, Huxley Building 6M34, London SW7-2AZ, UK.
E-mail addresses: ybilleh@caltech.edu (Y.N. Billeh), michael.schaub09@imperial.ac.uk (M.T. Schaub).
[1] These authors contributed equally to this work.

http://dx.doi.org/10.1016/j.jneumeth.2014.08.011

Conclusions: We present a versatile method to detect neural assemblies in spiking data applicable across a spectrum of relevant scales that contributes to understanding spatio-temporal information gathered from systems neuroscience experiments.

1. Introduction

As capabilities for parallel recordings from large neuronal populations continue to improve (Ahrens et al., 2013; Buzsaki, 2004) experimentalists are now able to probe neural population encoding in ever more detail. These experimental advances allow the study of the intricate links between topology and dynamics of neural interactions, which underpin the functional relationships within neural populations. One such example is the activity of cell assemblies. The problem is to identify groups of neurons (termed cell assemblies) within a large number of simultaneously recorded neurons where, due to functional cooperativity, each cell in an assembly is more similar in its temporal firing behavior to members of its own group than to members of other groups. Such strongly intertwined activity patterns are believed to underpin a wide range of cognitive functions (Hebb, 1949; Harris, 2005; Buzsaki, 2010). However, the reliable identification of cell assemblies remains challenging.

Here we introduce a technique to identify such neuron assemblies directly from multivariate spiking data, based on two steps: the definition of a simple biophysically-inspired similarity measure obtained from the observed spiking dynamics, followed by its analysis using a recent framework for multiscale community detection in weighted, directed graphs. A variety of techniques have been proposed to cluster spike-train groups to date, and have shown promising results in particular settings (Fellous et al., 2004; Feldt et al., 2009; Humphries, 2011; Lopes-Dos-Santos et al., 2011, 2013; Quiroga and Panzeri, 2009; Abeles and Gat, 2001; Laubach et al., 1999; Peyrache et al., 2010; Gansel and Singer, 2012). In contrast to these techniques, our methodology provides a dynamics-based framework, in which both the similarity measure and the community detection method are geared toward incorporating key features of neural network dynamics. The framework is purposely designed to be simple, yet capturing a breadth of features not present concurrently in other methods.

Our similarity measure evaluates the association between neuron pairs based on their spiking history and integrates three features that are key for a network-based analysis of neurophysiological data: (i) an intuitive biophysical picture, allowing a simple interpretation of the computed associations; (ii) a measure that is directed in time, hence asymmetric in the sense that spike-time dependent information is retained (*e.g.*, spiking of neuron A precedes that of neuron B); (iii) excitatory and inhibitory interactions are both included yet treated differently, inspired by their distinct effects on post-synaptic cells.

The detected dynamic associations are interpreted as an induced functional network, which is used to identify neuronal assemblies using a directed version of the recently introduced Markov Stability framework for community detection in graphs (Delvenne et al., 2010). Unlike other approaches, this framework allows us to analyze directed networks and search for cell assemblies at all levels of granularity, from fine to coarse levels of resolution, extracting relevant, possibly hierarchical groupings in spike-trains without *a priori* assumptions about the groups present. In the following, we present our framework and evaluate it on a series of examples, including synthetic spike-trains and leaky-integrate-and-fire network models. We also apply it to experimental datasets from retinal ganglion cells and hippocampal pyramidal neurons.

2. Materials and methods

Most existing methods to detect groups in spike-train neuronal population data are based on the following generic paradigm (Fellous et al., 2004; Feldt et al., 2009; Humphries, 2011; Lopes-Dos-Santos et al., 2011). First, a metric is defined to quantify the relationship between all neuron pairs leading to a $N \times N$ association matrix, where N is the number of observed neurons. We call this the *functional connectivity matrix* (FCM) hereafter. Every (i, j) entry in this matrix is a non-negative number that indicates how similar the spike trains of neurons i and j are over the observed time. Second, the FCM is clustered, *i.e.*, partitioned into different groups (Newman, 2004; Fortunato, 2010; Aggarwal and Reddy, 2014).

Here we introduce a simple framework that addresses both of these steps in a consistent and integrated manner, focusing on the dynamical relations between neurons: a new directed ('causal') biophysically-inspired measure is introduced to calculate the FCM, which is then analyzed using the recently introduced dynamics-based technique of Markov Stability for community detection (Delvenne et al., 2010, 2013; Lambiotte et al., 2009; Schaub et al., 2012) to identify cell assemblies at multiple scales in the neuronal population.

The numerics are performed in MATLAB (2011b or later versions). Code implementing the algorithm for spike-train analysis is available upon request and available at github.com/CellAssembly/Detection.

2.1. Biophysically-inspired causal measure of spike-train similarity

A plethora of metrics exists to describe the relationship between two signals, ranging from generic measures, such as cosine similarity and Pearson or Spearman correlation coefficients, to specialized measures designed for spike-train analysis (Kreuz et al., 2013; Lyttle and Fellous, 2011; Victor and Purpura, 1996; Fellous et al., 2004; Schreiber et al., 2003; van Rossum, 2001; Okatan et al., 2005; Vincent et al., 2012). Although these methods can be well suited in particular contexts, they only partially account for three important features for network-driven analyses of neural recordings. First, most current metrics are based on statistical arguments lacking a simple biophysical interpretation that would allow the use of relevant biophysical characteristics of neuronal dynamics. Second, most commonly used measures are distance metrics, *i.e.*, symmetric by construction, and thus neglect spike-timing information contained in the ordering of events. Finally, to the best of our knowledge, all measures ignore whether the neurons under consideration are excitatory or inhibitory. While an even finer characterization of neuronal subtypes could be of further interest, the distinction between excitatory and inhibitory neurons underpins fundamental balances in neuronal network dynamics and should be reflected in the analysis of data. Here, we propose a similarity measure that incorporates these three ingredients in a simple, intuitive form (see Fig. 1).

Consider first an excitatory neuron A connected to neuron B. The action potentials of A induce excitatory postsynaptic potentials (EPSPs) in neuron B, increasing the likelihood of neuron B