ELSEVIER

Journal of Neuroscience Methods 157 (2006) 303–316

JOURNAL OF NEUROSCIENCE METHODS

www.elsevier.com/locate/jneumeth

Spike sorting: Bayesian clustering of non-stationary data

Aharon Bar-Hillel[a,*], Adam Spiro[b,1], Eran Stark[c,2]

[a] The Interdisciplinary Center for Neural Computation, The Hebrew University of Jerusalem, Jerusalem 91904, Israel
[b] School of Computer Science and Engineering, The Hebrew University of Jerusalem, Jerusalem 91904, Israel
[c] Department of Physiology, Hadassah Medical School, The Hebrew University of Jerusalem, Jerusalem 91120, Israel

Received 5 September 2005; received in revised form 19 March 2006; accepted 23 April 2006

Abstract

Spike sorting involves clustering spikes recorded by a micro-electrode according to the source neurons. It is a complicated task, which requires much human labor, in part due to the non-stationary nature of the data. We propose to automate the clustering process in a Bayesian framework, with the source neurons modeled as a non-stationary mixture-of-Gaussians. At a first search stage, the data are divided into short time frames, and candidate descriptions of the data as mixtures-of-Gaussians are computed for each frame separately. At a second stage, transition probabilities between candidate mixtures are computed, and a globally optimal clustering solution is found as the maximum-a-posteriori solution of the resulting probabilistic model. The transition probabilities are computed using local stationarity assumptions, and are based on a Gaussian version of the Jensen–Shannon divergence. We employ synthetically generated spike data to illustrate the method and show that it outperforms other spike sorting methods in a non-stationary scenario. We then use real spike data and find high agreement of the method with expert human sorters in two modes of operation: a fully unsupervised and a semi-supervised mode. Thus, this method differs from other methods in two aspects: its ability to account for non-stationary data, and its close to human performance.

Keywords: Clustering; Jensen–Shannon divergence; Mixture of Gaussians; Monkey recordings; Non-stationary data; Semi-supervised learning; Spike sorting

1. Introduction

In extra-cellular recording of brain activity, a micro-electrode usually picks up the activity of multiple neurons. Spike sorting is the task of finding a clustering of the spike data such that each cluster is comprised of the spikes generated by a different neuron. Currently, this task is mostly done manually. It is a tedious mission, requiring hours of human work for a single recording session. Many algorithms were proposed in order to automate this process (see Lewicki, 1998 for a review) and some were implemented to provide a helping tool for manual sorting (Harris et al., 2000; Lewicki, 1994; Yu and Kreiman, 2000). However, the ability of suggested algorithms to replace the human worker has been quite limited. One of the main obstacles for a successful spike sorting application is the non-stationary nature of the data (Lewicki, 1998). The primary source of this non-stationarity is slight movements of the recording electrode. Small drifts of the electrode's location, which are almost inevitable, result in changes of the shapes of recorded spikes over time. Other sources of non-stationarity are variable background noise, and, perhaps, changes of the characteristic spike shape generated by a neuron. The increasing usage of multiple electrode recordings turns non-stationarity into an important issue, since electrodes are being placed in a single location for long durations.

Manual sorting often uses the first two principal component analysis (PCA) coefficients (PCs) to represent spike data (see Hulata et al., 2002; Quiroga et al., 2004 for other, wavelet-based representations). Such a representation preserves up to 93% of the variance of the recorded spikes (Abeles and Goldstein, 1977). A human often clusters spikes by visual inspection of the projected spikes in small time frames, in which non-stationary effects are insignificant. Problematic scenarios which can appear due to non-stationarity are demonstrated in Fig. 1 and include: (1) Movements and considerable shape changes of clusters over time. Such movement causes a cluster to be smeared when observed at a large time window. (2) Two clusters which are initially well-separated may move until they merge into one. A split of a single cluster is possible in the same manner. We would

* Corresponding author. Tel.: +972 2 658 4079; fax: +972 3 535 7155.
E-mail addresses: aharonbh@cs.huji.ac.il (A. Bar-Hillel), adams@cs.huji.ac.il (A. Spiro), eranstark@md.huji.ac.il (E. Stark).
[1] Tel.: +972 3 5238584.
[2] Tel.: +972 2 6758381.

0165-0270/$ – see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2006.04.023

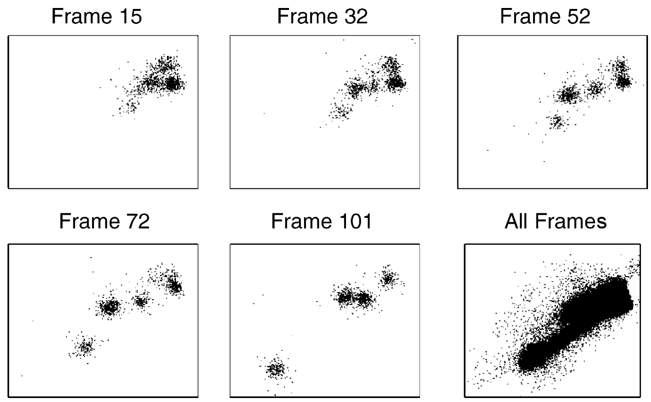

Fig. 1. Example of several time frames from a single electrode recording. Five time frames from a sequence of 116 are presented with their chronological frame numbers (1000 spikes per frame). The bottom right panel shows all the data points in a single frame (115,782 spikes). Spikes were projected on two principal components. In this recording, several clusters move constantly. In the process, a big fuzzy cluster splits into distinguishable clusters and later, two of those become indistinguishable again. While single time frames may have several plausible clustering interpretations, the sequence as a whole is usually less ambiguous. Clearly, the "right" clustering cannot be seen when all data points are observed in a single frame.

like to distinguish between different clusters whenever possible, and to state that a cluster is a multi-unit cluster when separation is not possible anymore.

Most spike sorting algorithms do not address the presented difficulties at all, as they assume full stationarity of the data. Some methods do try to cope with the lack of stationarity by grouping data points into many small clusters and then identifying clusters that can be combined to represent the activity of a single unit. For instance, Fee et al. (1996) use inter-spike interval (ISI) information and density of points on the borders between clusters to decide which clusters should be combined. Snider and Bonds (1998) point out that an ISI criterion is not valid for correlated neurons and base the cluster unification decision on the density of points between clusters. Shoham et al. (2003) suggest handling non-stationary clusters by modeling clusters using a *t*-distribution. This distribution can partially describe asymmetric smeared clusters. None of these methods partition the data into time frames and hence none can cope with complicated cases requiring this information. For example, clusters of similar spike shapes existing at distant time frames are always unified by such methods, while tracking the cluster through time may reveal that it is not the same cluster.

Emondi et al. (2004) suggested a semi-automated method in which the data are divided into small time frames. Each time frame is clustered manually, and then the correspondence between clusters in consecutive time frames is established automatically. In order to establish the correspondence, match scores between clusters are computed using a heuristic metric, and a greedy procedure is then used to choose matching pairs. No splits or merges of clusters are considered in that method.

In this paper we suggest a fully automated method for solving the clustering problem offline in a Bayesian framework. A preliminary version of the method was presented in a conference paper (Bar-Hillel et al., 2005). Trying to mimic human experts, we focus here on two-dimensional spike representations, although our algorithm can be applied to higher dimensions as well. In brief, the spike data are divided into short time frames in which stationarity is a reasonable assumption. At a first stage, good descriptions of the data as mixtures-of-Gaussians are computed for each time frame independently. Many possible mixture descriptions are considered for each time frame. In a second stage, transition probabilities between consecutive mixtures are computed, and the complete dataset is modeled as a chain of local mixture models. The final clustering, as well as the correspondence between clusters in consecutive time frames, is found as the maximum-a-posteriori (MAP) solution of the chain model. This global optimization allows the algorithm to successfully disambiguate problematic time frames in a way similar to the behavior of a human expert. The model and its optimization are presented in Section 2.

An important contribution to the success of the method is made by the derivation and computation of the transition probabilities. The assumptions and algorithms used to derive these probabilities are described in Section 3. In essence, we assume that the data are approximately stationary in two consecutive time frames, and we look for a hidden mixture model that explains well the Gaussian components in both time frames. We frame the search for the "right" hidden mixture model as a constrained optimization problem with a Jensen–Shannon (JS) based score. The problem can be optimally solved if the number of clusters is assumed to be constant, and we suggest an efficient suboptimal strategy in a general case in which splits and merges of clusters are allowed.

Empirical validation of the proposed method is presented in Section 4. We first test the performance of the method using synthetically generated spike data. In this case we know the correct source neuron for each spike, and hence we can verify the correctness of sorting in various scenarios and compare our method with other recently proposed spike sorting methods. The agreement between the clustering of the proposed method and the correct labels is found to be higher than that of other methods in most scenarios. Second, we test the performance of the method operating in an unsupervised mode using electrode recordings from premotor cortices of Macaque monkeys (Stark et al., 2006). The clustering is evaluated by computing its agreement with manual clustering done by a human expert, and high agreement rates are found in most cases. In another experiment, we compare the agreement between the algorithm's and human clustering solutions to the agreement among several expert human sorters. The algorithm is comparable to the human experts in this test. Finally, we test the algorithm in a semi-supervised mode, in which a human user clusters part of the time frames, and the algorithm has to fill in the rest. In this mode, the clustering results highly agree with the intentions of the human expert, while human labor is substantially reduced.

2. A chain of Gaussian mixtures

A single recording session may contain tens or hundreds of thousands of spikes recorded over several hours (Section 4). While spike clusters are not necessarily stationary over such