Journal of Neuroscience Methods 222 (2014) 47–55

Contents lists available at ScienceDirect

Journal of Neuroscience Methods

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

Computational neuroscience

A unified framework and method for automatic neural spike identification

Chaitanya Ekanadham[a,*], Daniel Tranchina[a,b], Eero P. Simoncelli[a,b,c]

[a] Courant Institute of Mathematical Sciences, New York University, United States
[b] Center for Neural Science, New York University, United States
[c] Howard Hughes Medical Institute, United States

ARTICLE INFO

Article history:
Received 19 December 2012
Received in revised form
30 September 2013
Accepted 2 October 2013

Keywords:
Spike sorting
Action potential
Neural spike identification
Clustering
Spike detection
Multi-electrode

ABSTRACT

Automatic identification of action potentials from one or more extracellular electrode recordings is generally achieved by clustering similar segments of the measured voltage trace, a method that fails (or requires substantial human intervention) for spikes whose waveforms overlap. We formulate the problem in terms of a simple probabilistic model, and develop a unified method to identify spike waveforms along with continuous-valued estimates of their arrival times, even in the presence of overlap. Specifically, we make use of a recent algorithm known as Continuous Basis Pursuit for solving linear inverse problems in which the component occurrences are sparse and are at arbitrary continuous-valued times. We demonstrate significant performance improvements over current state-of-the-art clustering methods for four simulated and two real data sets with ground truth, each of which has previously been used as a benchmark for spike sorting. In addition, performance of our method on each of these data sets surpasses that of the best possible clustering method (i.e., one that is specifically optimized to minimize errors on each data set). Finally, the algorithm is almost completely automated, with a computational cost that scales well for multi-electrode arrays.

The problem of detection, time-estimation, and cell classification of neural action potentials from extracellular electrode measurements is fundamental to experimental neuroscience. Electrode(s) are embedded in neural tissue, and a voltage trace is recorded as a function of time. When a neuron in the vicinity of the electrode fires an action potential, a stereotypical waveform is superimposed onto the recorded voltage (Lewicki, 1998; Sahani et al., 1997; Wehr et al., 1999). The shape of this waveform depends on the cell's morphology and position, as well as the filtering properties of the medium and the electrode(s). The "spike sorting" problem consists of detecting the occurrence of these individual waveforms and estimating their corresponding times of occurrence.

Despite the ubiquity and succinct formulation of the problem, there is no *de facto* standard for spike sorting. Traditionally, experimentalists manually position a single electrode and define threshold triggers to identify the spikes of individual cells (Rodieck, 1967). However, this becomes substantially more difficult when recording from several cells simultaneously, and is infeasible for multi-electrode arrays. Computer-assisted solutions have converged on a general methodology that we will refer to as "clustering," consisting of three steps (Lewicki, 1998), illustrated in Fig. 1: (1) detection of temporal segments of the voltage trace that are likely to contain spikes, (2) estimation of a set of features for each segment, and (3) classification of the segments according to these features. A variety of methods exist for solving each step (e.g., (1) thresholding based on absolute value (Obeid and Wolf, 2004), squared values (Rutishauser et al., 2006), Teager energy (Choi et al., 2006), or other nonlinear operators (Rebrik et al., 1999), (2) features such as peak-to-peak width/amplitude, projections onto principal components (Lewicki, 1998), or wavelet coefficients (Quiroga et al., 2004; Kwon and Oweiss, 2011), and (3) classification methods such as *K*-means (Lewicki, 1998), mixture models (Sahani, 1999; Shoham et al., 2003), or superparamagnetic methods (Quiroga et al., 2004)).

Although methods exist for solving each of the three steps in isolation, it is unclear how to relate the sequential application of these steps to the optimization of a single objective function, making it difficult to state the assumptions and operating conditions needed for success. Since each step does not take into account errors introduced in previous steps, errors tend to accumulate. In addition, many of these methods require human supervision (especially for the classification step), which is not only costly, but generally inaccurate (Harris et al., 2000) and highly variable (Wood et al., 2004). The lack of a standard automated methodology makes it difficult to compare results of scientific studies.

* Corresponding author. Tel.: +1 9146216070.
E-mail address: chaitue@gmail.com (C. Ekanadham).

0165-0270/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jneumeth.2013.10.001

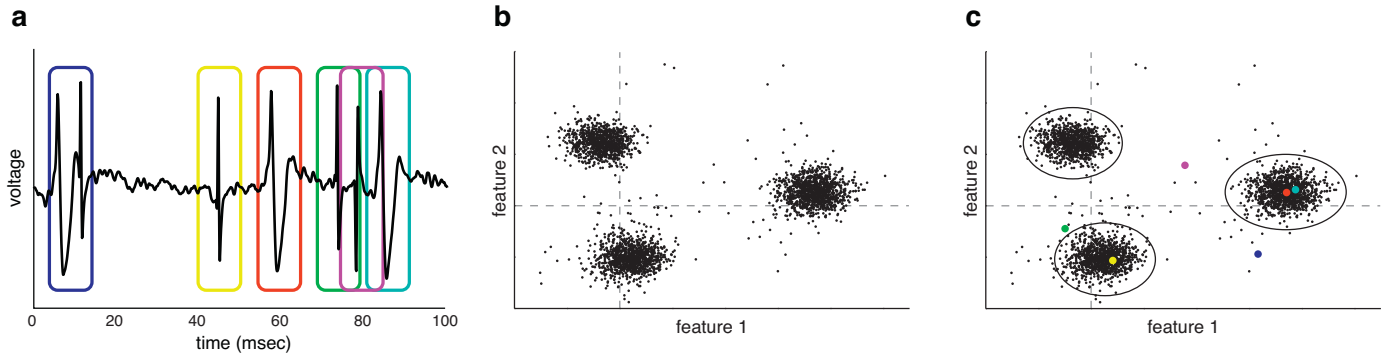

Fig. 1. Schematic of 3-step procedure common to most current spike sorting methods. (a) Thresholding/windowing. Voltage peaks are detected (by comparison to a threshold), and their occurrence times are estimated. Temporal segments of the voltage trace that lie within a fixed-duration window around each peak (colored rectangles) are gathered. (b) Feature estimation. Segments are projected into a low-dimensional feature space. Here, we plot the projection of each segment onto the first two principal components of the full set of segments. (c) Classification. Segments are grouped within the feature space, typically using an automatic clustering method such as *K*-means or estimation of a Gaussian mixture model. Colored points correspond to the windowed segments in (a). Note that several of these are mis-classified because they contain a superposition of more than one spike. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of the article.)

Most importantly, the conventional three-step procedure mishandles overlapping spikes. If two or more cells fire near-synchronously, their respective waveforms will overlap in the voltage trace, creating a shape that differs from either spike in isolation (Lewicki, 1994; Sahani et al., 1997; Wehr et al., 1999; Pillow et al., 2013). If the waveform shapes partially cancel, the initial detection stage may miss the spikes altogether. Even if the segment is detected, its appearance will depend on the time delay between the two spikes (Pillow et al., 2013). If this is significantly different from that of either spike in isolation, it will be misidentified as a fictitious third cell or discarded as an outlier, as illustrated in Fig. 1. Even in the best case scenario, only one of the two spikes can be correctly identified by a clustering method, and the other discarded.

Failure to resolve overlapping spikes can have serious consequences: Basic measurements, such as mean firing rates and cross-correlations, can be heavily biased due to spike sorting artifacts (Bar-Gad et al., 2001; Pazienti and Grn, 2006; Pillow et al., 2013). Properly handling this bias is crucial when studying a neural population where there is a high level of synchronous activity or when the study itself focuses on the correlation of firing patterns (Mastronarde, 1989; Devries, 1999; Schnitzer and Meister, 2003; Shlens et al., 2008; Pillow et al., 2008). Such studies are more frequent with the advent of multi-electrode array recordings, which allow the simultaneous recording of large populations of neurons (Meister et al., 1994; Gerstein and Clark, 1964; Brown et al., 2004; Pillow et al., 2008; Shlens et al., 2009).

There have been several proposed methods to augment the clustering approach to account for overlapping spikes (Atiya, 1992; Lewicki, 1994; Segev et al., 2004; Zhang et al., 2004; Vargas-Irwin and Donoghue, 2007; Pillow et al., 2008; Chen et al., 2011; Prentice et al., 2011; Pillow et al., 2013). However, these methods generally rely on brute-force examination of all combinations of spike waveforms at all time separations (impractical for simultaneous recordings of many cells), or "greedy" algorithms that iteratively subtract the waveform of the best-fitting cell until the residual amplitude is within the range expected for noise. A notable exception is the family of ICA-based spike sorting methods (Takahashi et al., 2003; Takahashi and Sakurai, 2005; Franke et al., 2009), which bear some resemblance to our approach, but have not been developed or implemented in the context of a unified probabilistic model for the voltage measurements, and have not been extensively tested and compared to traditional clustering methods.

In this paper, we present a method for estimation of the most probable spike patterns given the observed voltage trace, which is assumed to be a noisy linear superposition of spike waveforms shifted to their respective spike times, corrupted by additive noise. We use a recently developed method in sparse signal decomposition, known as *Continuous Basis Pursuit* (Ekanadham et al., 2011a), as the basis for an accurate and efficient approximation of the solution. The resulting method provides a unified procedure for the estimation of continuous-valued spike times that operates correctly in the presence of overlapping spikes, and does not rely on any auxiliary heuristic pre-processing or post-processing such as alignment of spike segments or searching for spike combinations. An initial version of this work was presented in (Ekanadham et al., 2011b). A software implementation is available at http://www.cns.nyu.edu/~lcv/spikeSorting.html

1. Methods

1.1. Spike sorting using Continuous Basis Pursuit (CBP)

Our method is derived from a simple generative model for the observed voltage trace (Sahani, 1999; Pillow et al., 2008, 2013), as illustrated in Fig. 2. A spike from the nth neuron, occurring at time $\{\tau_{ni}\}$, is assumed to produce a temporally localized waveform $a_{ni}W_n(t-\tau_{ni})$, where $W_n(t)$ has unit norm, and a_{ni} represents the (root-mean-squared) spike amplitude. These time-shifted and scaled waveforms are then added together with noise to form the electrode voltage trace:

$$V(t) = \sum_{n=1}^{N}\sum_{i=1}^{C_n} a_{ni}W_n(t-\tau_{ni}) + \epsilon(t) \quad (1)$$

In the case of multi-electrode recordings, $V(t)$ and $W_n(t)$ are vector-valued with as many dimensions as electrodes, but for notational convenience, the derivation below is written for the scalar case.

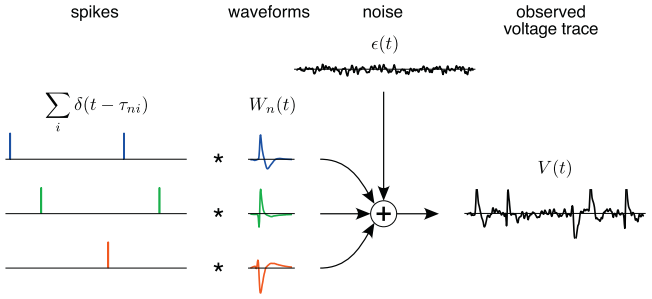

Fig. 2. Measurement model, illustrated for three neurons and a single electrode. Each cell generates a voltage trace containing time-shifted copies of its spike waveform, and the observed voltage trace is assumed to be a sum of these and noise.