Large-scale Multi Electrode Array Spike Sorting Algorithm Introducing Concurrent Recording and Stimulation

Gonzalo Mena
Statistics Department
Columbia University

Lauren Grosberg
Department of Neurosurgery
Hansen Experimental Physics Laboratory
Stanford University

Frederick Kellison-Linn
Columbia University

E.J. Chichilnisky
Department of Neurosurgery
Hansen Experimental Physics Laboratory
Stanford University

Liam Paninski
Statistics Department
Center for Theoretical Neuroscience
Grossman Center for the Statistics of Mind
Columbia University

1 Introduction

In recent years, the use of multi-electrode arrays (MEAs) for simultaneous electrical stimulation and recording has shown to be a useful technique in systems neuroscience and neural engineering applications: for example, for the assessment of functional connectivity in neural circuits by targeted activation of specific cells while recording from others [1], to give better experimental access to axonal function [2], and for the design of retinal prosthetic devices [3]. At the same time, the big data era in neuroscience has brought about new scientific questions in large-scale settings, opening the doors for new uses of this technique. Specifically, simultaneous stimulation and recording in large-scale MEAs could be a powerful tool for online control and perturbation of neural networks and lead to much higher retinal prosthesis resolution through closed-loop feedback control.

However, for this to be possible, an important technical hurdle have to be overcome: most of the aforementioned applications rely on an intermediate spike sorting stage from which voltage activations are matched to their underlying neuronal identities. Although in the past decades many different spike sorting methodologies have been developed [4, 5, 6, 7, 8], all of them fail to assign neuronal identity when electrical stimulation is introduced experimentally, primarily because recordings with stimulation are corrupted with stimulus-induced distortions or artifacts. Indeed, in the context of retinal stimulation the only method for spike identification that can be used heavily relies on human intervention [3], making it unsuitable for the study of even modestly sized neural populations.

In this article, a method for automated spike sorting for MEA recordings with electrical stimulation is presented, and its effectiveness is demonstrated with experimental data from the primate retina. The algorithm is based on a probabilistic generative model of the data: it is assumed that recorded voltage traces following electrical stimulation are made up by the sum of electrical signals from

cells with known spike waveforms, stimulus artifact with certain regularities, and gaussian noise. Priors on the shapes of both the response probability of neurons and the artifact as functions of stimulus current amplitude are used to constrain the model and improve algorithm performance. Moreover, these priors are used for post-processing purposes, to come up with diagnostic measures that assess the plausibility of solutions, and to construct an artifact resampling device to correct erroneous solutions. We begin with a description of the method, and then proceed to demonstrate its effectiveness in analyzing large-scale multi-electrode data.

2 Methods

Electrophysiology was recorded on a 512-electrode MEA capable of electrical stimulation [9] from primate retinas as described previously [3]. Data is comprised of voltage recordings following electrical stimulation arranged in amplitude series. We refer to an *amplitude series (AS)* as a collection of responses to J increasing current amplitudes repeatedly applied on one or more electrodes. At the j-th amplitude (also called *condition* j) a number of I_j voltage traces, or trials, are recorded over a time window $t = 1, \ldots, T$ and on a collection of $e = 1, \ldots, E$ recording electrodes. An AS is the minimal data unit for which the algorithm can be applied, and define the underlying experimental design for data collection. Given an AS, the goal is to identify spikes of N neurons provided their action potential waveforms (templates) on all electrodes. In the present case, templates are obtained separately, from data collected in the presence of visual stimuli rather than electrical stimuli.

2.1 Generative Model

Let $Y_{t,e}^{i,j}$ denote the observed voltage for trial i of condition j at time t and electrode e. Then the model is:

$$Y_{t,e}^{i,j} = A_{t,e}^{j} + \sum_{n=1}^{N} (K_n s_n^{i,j})_{t,e} + \epsilon_{t,e}^{i,j} \ , \quad \frac{\epsilon_{t,e}^{i,j}}{\sigma_{e,j}} \sim \mathcal{N}(0,1) \quad \textit{i.i.d.} \tag{1}$$

Here $A_{t,e}^{j}$ is the artifact at time t, electrode e and condition j (fixed over trials i), and $s_n^{i,j}$ is a binary spiking vector: $s_n^{i,j}(l) = 1$ if a spike of neuron n occurs at time t_l on trial i of condition j. We impose that the sum of any of these vectors is at most one; that is, at most one spike occurs per trial. K_n is a convolution matrix whose rows contain copies of template of neuron n as recorded in all electrodes but with spike onset aligned at all different possible spike times.

We also consider a vectorized version of equation 1 to state the generative model as a gaussian likelihood for the data vector Y in terms of artifact A, suitable artifact covariates X, template matrix K, spiking vector s and covariance matrix Σ:

$$p(Y|A, s, \Sigma) \propto \exp\left(-\frac{1}{2}(Y - XA - Ks)^t \Sigma^{-1} (Y - XA - Ks)\right). \tag{2}$$

We impose structure on artifact by penalizing squared time and condition differences, to represent a continuity requirement in these two dimensions. This structure can be expressed as a gaussian prior: to see this, notice that penalized differences can be written as quadratic forms$A_e^t D_{e,k} A_e$, and the amount of penalization can be set by some parameters $\lambda_{e,k}$. The prior is, then:

$$p(A|\lambda) \propto \exp\left(-\frac{1}{2}\sum_{e,k} \lambda_{e,k} A_e^t D_{e,k} A_e\right) \quad \text{or} \quad A_e|\lambda_e \sim \mathcal{N}\left(0, \left(\sum_k \lambda_{e,k} D_{e,k}\right)^{-1}\right). \tag{3}$$

Spike probabilities increase smoothly as a function of condition [9], and these physiologically characteristic sigmoidal activation curves are used to further constrain the model. To this end we use a logistic regression prior for the activation curve of each neuron (variables $r_n^{i,j}$ indicate presence or absence of spikes):

$$p(r_n^{i,j} = 1|\alpha_n) = \frac{1}{1 + \exp\left(-\alpha_n^0 - j\alpha_n^1\right)} \ . \tag{4}$$

Finally, for the entries of the diagonal of the matrix (i.e. the variances $\sigma^2_{e,j}$) we consider a non-informative prior [10], only to exploit conjugacies:

$$p(\Sigma) = p(\sigma^2_{e,j}, e = 1 \dots E, j = 1 \dots J) \propto \prod_{e,j} \frac{1}{\sigma^2_{e,j}} \ . \tag{5}$$

2.2 Algorithm

Our algorithm is based on a thorough exploration of the parameter space based on the posterior distribution formed from the product of the data likelihood (equation 2) and the parameter priors (equations 3-5):

$$p(s, A, \Sigma, \alpha | Y, \lambda) \propto p(Y|s, A, \Sigma) p(\Sigma) p(r|\alpha) p(A|\lambda). \tag{6}$$

Computing the MAP solution here can be intractable: since the variables sni,j are integer valued, maximization of 6 is a non-concave and multimodal problem. Also, using maximization of the posterior as the unique suitable device for doing spike sorting can be unsatisfying because of model misspecification: in practice, correct spike sorting solutions correspond to one high probability posterior mode, but there can be nonsensical solutions associated with modes of even higher probability.

This motivates the pursuit of an algorithm possessing the following features: i) it explores the parameter space looking for regions of high posterior probability ii) it is computationally tractable, iii) it can assess the plausibility of solutions and in the negative case propose an improved one. To do so, our algorithm is made up of an initialization step to obtain a reasonable first guess, a Gibbs sampling stage that iterates until convergence to a spike sorting solution, and a post-processing stage where changes to the current possibly erroneous solutions are proposed.

The initialization has two steps. In the first, a surrogate quadratic program (QP) is solved. Essentially, it corresponds to a convex relaxation of a much less complex problem that nevertheless captures the structure of the original one. In the QP the objective function is the RSS, spiking vectors are allowed to belong to the probability simplex, the artifact prior is replaced by a polynomial structure and logistic regression priors for activation curves are replaced by a requirement of increasingness in spike probabilities as a function of condition, which is a simple linear constraint. The solution of the QP leads to initial estimates of all the values, except the hyperparameters λ. These hyperparameters are found in the second step, by maximizing the likelihood of having obtained the artifact initial solution $A_{0,e}$, that is:

$$\begin{aligned} \min_{\lambda_e} \ & \frac{1}{2} A^t_{0,e} \left(\sum_k \lambda_{e,k} D_{e,k} \right) A_{0,e} - \log \left| \sum_k \lambda_{e,k} D_{e,k} \right| \\ \text{s.t.} \ & \lambda_e \geq 0, \quad \sum_k \lambda_{e,k} D_{e,k} \succ 0 \ . \end{aligned} \tag{7}$$

The Gibbs sampler stage is straightforward and alternates between sampling conditional spikes (multinomial), artifact (gaussian), variances (inverse gamma) and logistic regression parameters given the rest of the variables and data. Sampling is performed until iterations no longer lead to changes in s(convergence to local optimum).

Finally, in the post-processing stage the agreement between the empirical activation curves and their logistic regression counterparts is assessed. If lack of fit is diagnosed (based on residual statistics, which follow a χ^2 distribution) the artifact is interpolated at the set of conditions that contribute the most to the mismatch, based on the rest of the conditions. This interpolation is implemented via resampling from the gaussian conditional distributions that arise from equation 6. After resampling, the Gibbs stage is executed again to converge to a new solution, and the procedure is repeated until no lack of fit is detected.

3 Results

Analysis was based on 710 AS, coming from eight retinal preparations, making up a total of 924,118 trials. For each AS, an individual neuron was targeted for spike sorting based on the voltage recordings of its closest electrode. To assess performance, results of the algorithm were compared to human sorted data, which serves as ground "truth", applying the usual accuracy, sensitivity and sensitivity measures. Comparisons were based on two events: spikes (from individual trials) and activation (whether or not the activation curve surpassed a threshold of 50%, obtained from individual AS). Results are summarized in table 1, and overall are satisfactory.

	Occurrence	Performance		
Type	(%)	Accuracy (%)	Sensitivity (%)	Specificity (%)
Spikes	9.66	$98.16 \pm .02$	$90.03 \pm .19$	$99.03 \pm .02$
Activation	44.1	93.4 ± 1.8	93.3 ± 2.8	93.5 ± 2.4

Table 1: Spike-by-spike and AS-by-AS results.

We also assessed performance in a more practical context: for the assessment of selective activation strategies in the development of retinal prosthesis, it is of particular interest to create spatial sensitivity maps: images that depict to which extent electrodes in the MEA can elicit spiking on a set of targeted neurons. These maps are made by doing spike sorting on a large set of AS, one for each electrode, and subsequently finding the corresponding activation thresholds. Information contained in the maps is useful for the determination of stimulation patterns across electrodes that will elicit activity in some neurons but not in others. Figure 1 shows that human and algorithm sorting are in close agreement.

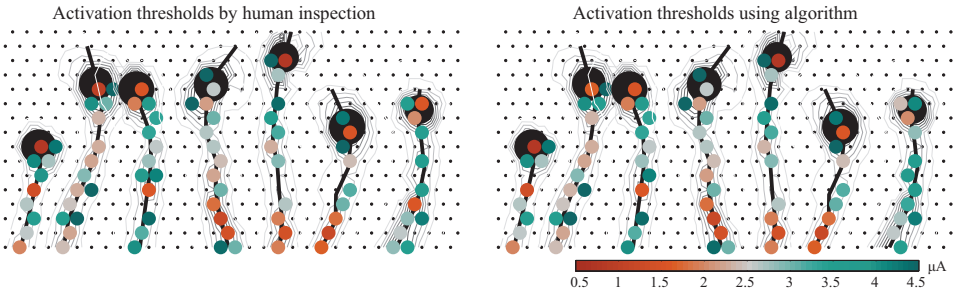

Figure 1: Comparison of spatial sensitivity maps, human v.s. algorithm. Individual neurons are shown in black overlaid with the MEA. Colored dots represent electrodes that activated the nearest neuron, and the color scale indicates the current amplitude required for activation.

4 Discussion and further work

We have demonstrated the plausibility of automating spike sorting in electrical stimulation paradigms, opening a possibility for the development of closed-loop, online data analysis during experiments. However, further improvements are required to permit the simultaneous sorting of many nearby cells, and to reduce the computational costs. Future implementations will heavily benefit from state-of-the-art machine learning methods.

References

[1] Lauren H. Jepson, Pawel Hottowy, Keith Mathieson, Deborah E. Gunning, Wladyslaw Dabrowski, Alan M. Litke, and E. J. Chichilnisky. Spatially patterned electrical stimulation to enhance resolution of retinal prostheses. *J Neurosci.*, 34(14):487–4881, 2014.

[2] Douglas J. Bakkum, Urs Frey, Milos Radivojevic, Thomas L. Russell, Jan Muller, Michele Fiscella, Hirokazu Takahashi, and Andreas Hierlemann. Tracking axonal action potential propagation on a high-density microelectrode array across hundreds of sites. *Nature Communications*, 4(2181), 2013.

[3] Lauren H Jepson, Pawel Hottowy, Keith Mathieson, Deborah E Gunning, Wladyslaw Dabrowski, Alan M Litke, and EJ Chichilnisky. Focal electrical stimulation of major ganglion cell types in the primate retina for the design of visual prostheses. *The Journal of Neuroscience*, 33(17):7194–7205, 2013.

[4] M. S. Lewicki. A review of methods for spike sorting: the detection and classification of neural action potentials. *Network: Computation in Neural Systems*, 9(4):53–78, 1998.

[5] Emery N Brown, Robert E Kass, and Partha P Mitra. Multiple neural spike train data analysis: state-of-the-art and future challenges. *Nature neuroscience*, 7(5):456–461, 2004.

[6] Jonathan W. Pillow, Jonathon Shlens, E. J. Chichilnisky, and Eero P. Simoncelli. A model-based spike sorting algorithm for removing correlation artifacts in multi-neuron recordings. *PLoS ONE*, 8(5):e62123, 05 2013.

[7] Chaitanya Ekanadham, Daniel Tranchina, and Eero P. Simoncelli. A unified framework and method for automatic neural spike identification. *Journal of Neuroscience Methods*, 222(0):47 – 55, 2014.

[8] Hernan Gonzalo Rey, Carlos Pedreira, and Rodrigo Quian Quiroga. Past, present and future of spike sorting techniques. *Brain Research Bulletin*, 2015. In press.

[9] Pawel Hottowy, Andrzej Skoczen, Deborah E Gunning, Sergei Kachiguine, Keith Mathieson, Alexander Sher, Piotr Wikacek, Alan M Litke, and Wladyslaw Dabrowski. Properties and application of a multichannel integrated circuit for low-artifact, patterned electrical stimulation of neural tissue. *Journal of neural engineering*, 9(6):066005, 2012.

[10] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian data analysis*, volume 2. Taylor & Francis, 2014.