Large-scale spike sorting for the analysis of electrical stimulation and a first application

Gonzalo E. Mena, Lauren E. Grosberg, Sasidhar Madugula
Paweł Hottowy, Alan Litke, John Cunningham,E.J. Chichilnisky & Liam Paninski

1 Abstract

Simultaneous electrical stimulation and recording using multi-electrode arrays (MEAs) can provide a valuable technique for studying circuit connectivity and engineering neural interfaces. However, interpreting these recordings is challenging because the spike sorting process (identifying and segregating action potentials arising from different neurons) for use with large scale MEAs is greatly complicated by electrical stimulation artifacts, which can exhibit complex and nonlinear waveforms. To our knowledge, no current data-analysis methods are available for large scale applications. Also, current approaches are often based on highly restrictive assumptions.

We provide a principled solution to this problem by explicitly stating a generative model to which we perform inferences. This model linearly decomposes neural activity and artifact contributions to data, and impose a structured Gaussian Process (GP) prior on the artifact for a parsimonious representation of its properties. We develop a scalable algorithm based on this structured GP to jointly infer the neural activity and the artifact, a nuisance parameter.

The effectiveness of our method is demonstrated in both real and simulated 512-electrode recordings with electrical stimulation in the peripheral primate retina. Also, we include the first application of the algorithm to show — based on a comprehensive large-scale analysis — that differential electrophysiological excitability characteristics of RGC are seen depending on the locus of stimulation. Our technology may be helpful in the design of future high-resolution retinal prostheses, and for closed-loop neural stimulation at a much larger scale than currently possible.

2 Additional Information

Our model is based on the decomposition $Y = A + s + \epsilon$, where Y accounts for the observed traces over the array (indexed by time, space, stimulus strength and repetition), s is the neural activity that we want to infer, A is the stimulation artifact and ϵ is a gaussian noise term. Our goal is to infer s, which is done jointly with A and s. We consider a structured GP prior for $A \sim GP(0, K)$ with

$$K = \rho K_t \otimes K_e \otimes K_s + \phi^2 I_d.$$

Here, K_t, K_e and K_s are kernels that account for the observed temporal, spatial and stimulus-wise non-stationarities, and are constructed as product of Matérn kernels with gamma-like envelope functions that achieve their maximum values in the regions of higher variability. Underlying the use of the Kronecker product $\otimes$ it is the assumption that variability in each of the dimensions is independent of the other. Although arguably, this is a common practice in the Machine Learning community, as it enables tractable large-scale inferences [1]. Also, ρ and ϕ^2 are scale and noise parameters, respectively.

Out algorithm is based on coordinate ascent: it iterates between estimating the artifact given an estimate of neural activity (based on standard GP inference, and estimating the neural activity given the artifact, which is done by trying to match typical neural templates (assumed known beforehand) at the times where they would explain the residuals the best.

Figure 1 contains sample traces for which spike sorting is required, along with estimated artifact and neural activity estimates. Figure 2 shows an example of the kind of large-scale analysis that can now routinely be done with our technology, perhaps in closed-loop paradigms. Figure 3 shows how the algorithm can be used to investigate relevant a neurophysiological question by aiding and extending the role of the experimenter. Indeed, for that later figure, available human-curated data was scarce, and results provided with the algorithm were based on a many times larger volume of data.

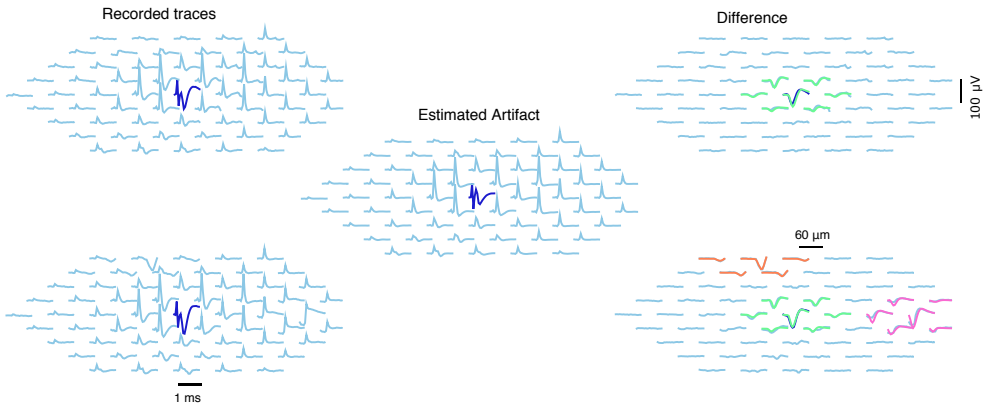

Figure 1: Example of neural activity and artifact inference in a neighborhood of the stimulating electrode. *Left:* Two recordings in response to a 2.01 μA stimulus. *Center:* estimated artifact (as the stimulus doesn't change, it is the same for both trials). *Right:* Difference between raw traces and estimated artifact, with inferred spikes in color. In one case (above) three spiking neurons were detected, while in the other (below) there was only one. The algorithm can correctly identify the neural activity in over 99% of cases (704,984 available trials)

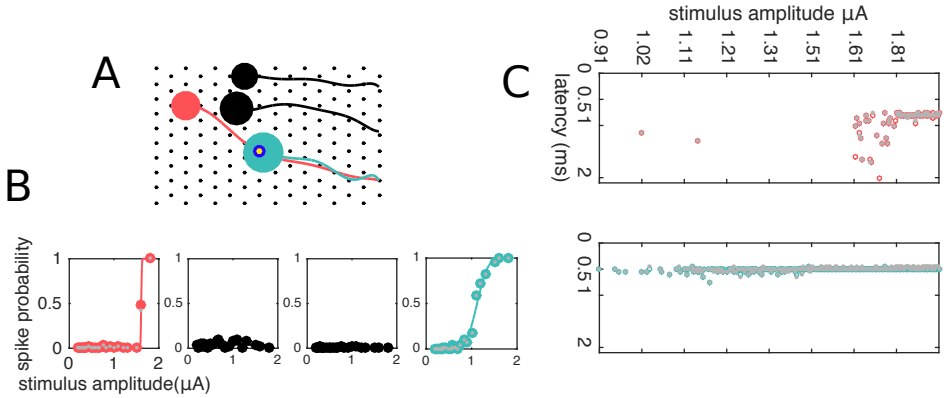

Figure 2: Analysis of responses of neurons in a neighborhood of the stimulating electrode. A) Spatial configuration: stimulating electrode (blue/yellow) and four neurons on its vicinity. Soma of green neuron and axon of pink neuron overlap with stimulating electrode. **B)** Activation curves (solid lines) along with with human and algorithm inferred spike probabilities (gray and colored circles, respectively) of all the four cells. Stimulation elicited activation of green and pink neurons; however, the two other neurons remained inactive. **C)** Rasters for the activated cells, with responses sorted by stimulation strength in the y axis. Ground truth and algorithm inferred latencies are indicated by gray and colored circles, respectively. Notice that direct somatic activation of the green neuron leads to a different pattern of activation that of the pink neuron, that is activated through its axon.

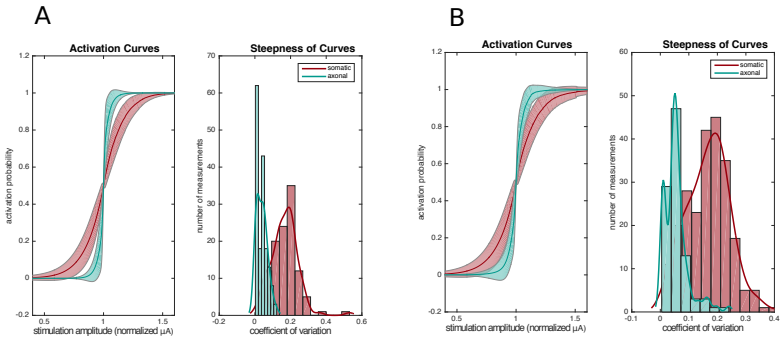

Figure 3: Different characteristics of activation depending on the type of stimulation (axonal or somatic) A) human, **B**). *Left*: Comparison of normalized activation curves (all are re-scaled to occur in the same stimulation range). For axonal stimulation activation curves tend to be steeper, whereas for somatic stimulation this curve is shallower *Right*: The coefficient of variation (σ/μ of the normal CDF fit to the activation curves, a sensible parameter to quantify the steepness of curves) exhibits a clearly stimulation-type dependent distribution.

References

[1] Elad Gilboa, Yunus Saatçi, and John P Cunningham. Scaling multidimensional inference for structured gaussian processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 37(2):424–436, 2015.