COLUMBIA UNIVERSITY
IN THE CITY OF NEW YORK

Introducing large-scale spike sorting for simultaneous electrical stimulation and recording

Gonzalo E. Mena[1*], Lauren E. Grosberg [2*], Frederick Kellison-Linn [2], E.J. Chichilnisky[2] and Liam Paninski[1,3], 1. Department of Statistics, Columbia University; 2. Department of Neurosurgery and Hansen Experimental Physics Laboratory, Stanford University d; 3. Center for Theoretical Neuroscience and Grossman Center for the Statistics of Mind, Columbia University. (*these authors contributed equally to this work) ,

Summary

Overarching goal:

- To develop methods for the analysis of large-scale electrophisiological recordings (MEAs) with electrical stimulation.
- Potentially a powerful tool for closed-loop feedback control and perturbation of neural networks.

Context: retina

- Interest in understanding how retinal ganglion cells (RGC) respond to electricity.
- Useful for the development of high-resolution prosthetic devices [2].
- Here we aim to extend the model-based spike sorting [1] framework to the context of electrical stimulation.
- Despite the context specificity, developed methods should extend to other settings.

Challenge

- Electrical stimulation induces transient distortions, or artifacts, on the recordings.
- With the presence of artifacts, it is difficult to tell which are the spike and artifact contributions to data.
- Existing artifact subtraction methods cannot handle with short latency or low time variability in spikes, the norm in this context.
- Current method: hybrid human-automatic, does not scale.

Strategy

- Algorithm is based on a generative model of data. Artifact structure and response properties of RGC are encoded in suitable priors.
- We leverage an experimental design and from the availability of action potentials shapes obtained from spike sorting in the absence of electrical stimuli.

Results

- We demonstrated the plausibility of automating spike sorting in electrical stimulation paradigms.
- Next step: extend this framework to account for spatial structure.

Experimental Setup

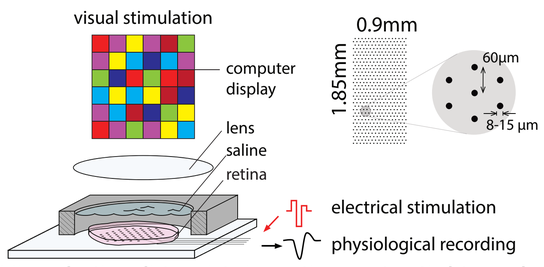

- Visual stimuli to obtain action potential templates.
- Triphasic current pulses with relative amplitudes of 2:-3:1 and phase widths of 50s were applied to electrodes. Currents ranged between $0.20\mu A$ and $4.79\mu A$
- In some experiments tetrodotoxin (TTX) was perfused into the retina to inhibit all action potentials in order to directly measure the stimulus artifact in a retinal preparation.
- Experimental design: given an stimulating electrode, for each current amplitude $j = 1, \dots J$ (also called conditions) there are I_j recordings of responses, or trials. Consecutive amplitudes increase: $a_{j+1} = fa_j$.
- We refer to this collection of trials as an amplitude series (AS), the minimal data unit for which the algorithm is applied.

Templates and electrical image from visual stimuli

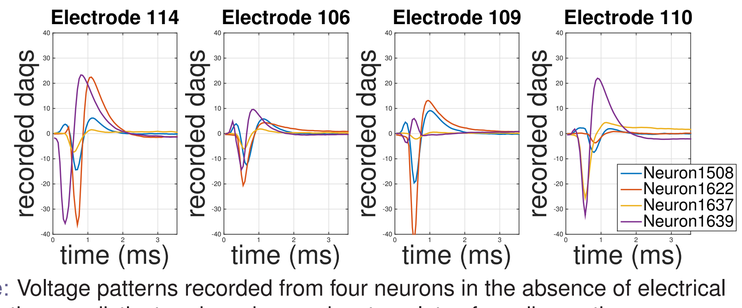

Figure: Voltage patterns recorded from four neurons in the absence of electrical stimulation are distinct and can be used as templates for spike sorting

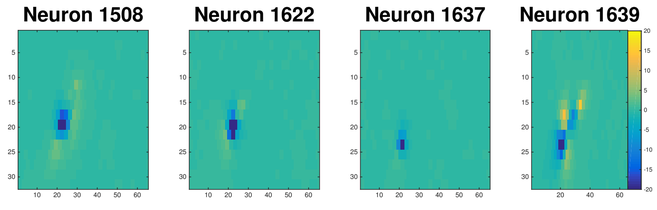

Figure: Electrical image of action potential at a fixed time

Stimulus Artifact

- Any contribution to data beyond noise and spikes of targeted RGCs.
- Properties: continuous in time, magnitude increases with amplitude of stimulus, stabilizes eventually. Also, it affects primarily the stimulating electrode and the effect decreases with distance to the stimulating electrode.
- Made up by hardware [3] and axonal contributions.

Figure: Artifact estimates (mean voltage across trials). Each colored trace corresponds to a different amplitude of stimulation (equivalently, condition)

Activation curves

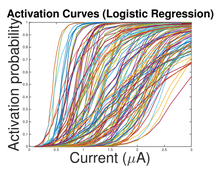

Figure: Logistic Regression fit of activation curves for many AS. Spike probabilities increase smoothly as a function of stimulus amplitude.

Why spike sorting is hard: an example

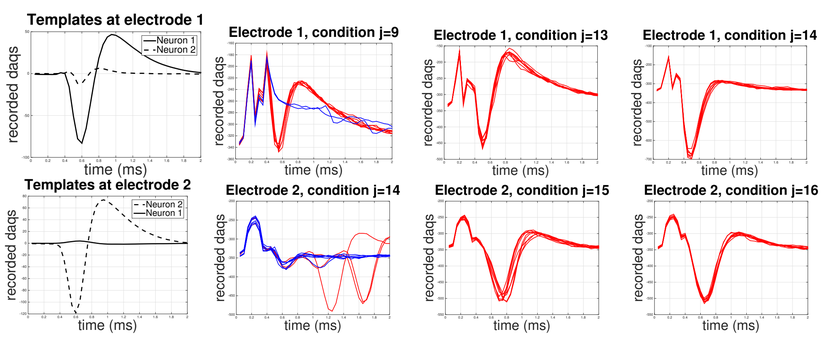

Figure: Examples of recordings for an AS with two electrodes (one stimulating) and two neurons. First and second rows show recordings in the stimulating and non-stimulating electrodes, respectively. Only red traces contain spikes

Generative Model

- Observed voltage $Y^{i,j}_{t,e}$

$$Y^{i,j}_{t,e} = A^j_{t,e} + \sum_{n=1}^{N} (K_n s^{i,j}_n)_{t,e} + \epsilon^{i,j}_{t,e}, \quad \frac{\epsilon^{i,j}_{t,e}}{\sigma_{e,j}} \sim \mathcal{N}(0,1) \quad i.i.d.$$

- t =time, e =electrode, j =condition(stimulus), i =trial, A =Artifact. s_n = are binary vectors indicating spike times for neuron n. K_n encodes action potentials at all possible times

$$s = (0,0,0,0) \xRightarrow{Ks}$$
$$s = (0,1,0,0) \xRightarrow{Ks}$$
$$s = (0,0,1,0) \xRightarrow{Ks}$$
$$s = (0,0,0,1) \xRightarrow{Ks}$$

- Equivalently,

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

- Gaussian prior for the artifact to account for its structure (D is a Brownian-motion-like inverse-covariance matrix) and to exploit conjugacies:

$$p(A|\lambda) \propto \exp\left(-\frac{1}{2}\sum_e \lambda_e A_e^t D_e A_e\right)$$

- Spike probability increases smoothly with stimulus amplitude ($r^{i,j}_n \equiv \sum_l s^{i,j}_n(l)$ indicate presence or absence of spikes):

$$p(r^{i,j}_n = 1|\alpha_n) = \frac{1}{1 + \exp\left(-\alpha^0_n - j\alpha^1_n\right)}$$

- Non-informative prior for the variances:

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

Algorithm Description

Main idea

- Inferences based on the (multimodal) posterior:

$$p(s, A, \Sigma, \alpha|Y, \lambda) \propto p(Y|s, A, \Sigma)p(\Sigma)p(r|\alpha)p(A|\lambda)$$

Step One: Initialization

- Solve a similar quadratic problem, a 'relaxation' where spikes can be probability vectors and Artifact has a polynomial structure.
- Compute λ maximizing the likelihood of obtaining the above 'initial' artifact.

Step Two: Gibbs sampler

- Sample from the conditionals of s, A, Σ and α given the rest and data.
- Repeat until s does not change (local maximum of the posterior)

Step Three: Post-processing

- Evaluate plausibility of the current solution based on residuals (e.g. compare 'empirical' activation curves with logistic regression.
- If lack of fit is detected, artifact is interpolated at the conditions where fit is the worst. Then, Gibbs sample again.

Algorithm Example

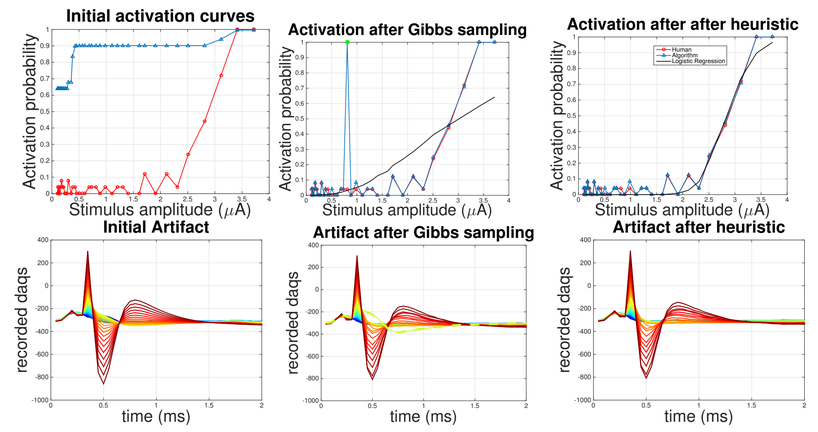

Figure: Top: activation curves after different stages. Bottom: Artifact estimates at different stages of algorithm. Green circle denotes detected false positive, corrected via artifact interpolation (yellow-ish trace)

Results

	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: Spike-by-spike and AS-by-AS results. Dataset: 710 AS, 924,118 trials, eight retinal preparations. One neuron and electrode per AS

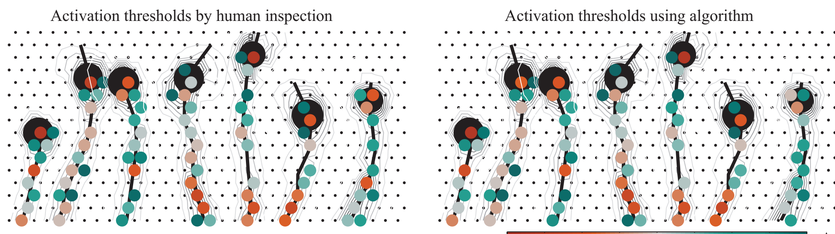

Figure: Large-scale application: RGC spatial sensitivity to electrical stimuli

Extension: Gaussian process artifact model

- Artifact does possess a clearly defined spatial structure. We can use that shared structure to estimate artifact.
- Instead of the 'naïve' brownian motion prior consider suitable Kronecker products (to mantain tractability [4]) of non-stationary Kernels.
- For example, $K = K_t \otimes K_r \otimes K_j + \sigma^2 I_d$. K_t, K_r can be non-stationary extensions of the usual Matérn Kernels.

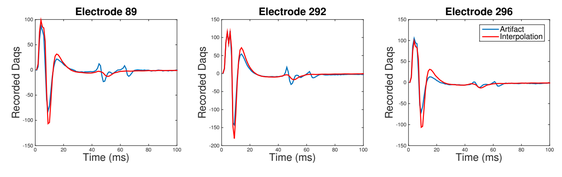

References

- Pillow J. et al., **A Model-Based Spike Sorting Algorithm for Removing Correlation Artifacts in Multi-Neuron Recordings**, *PloS ONE*, 2013.
- Jepson, L. et al. **Focal electrical stimulation of major ganglion cell types in the primate retina for the design of visual prostheses**, *The Journal of Neuroscience*, 2013.
- Hottowy, P. et al., **Properties and application of a multichannel integrated circuit for low-artifact, patterned electrical stimulation of neural tissue**, *J. neural. eng.* 2012.
- Gilboa, E. et al., **Scaling multidimensional inference for structured Gaussian processes**, *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 2015.

Acknowledgements

L.E.G. was funded by NIH Grant 1F32EY025120 and E.J.C. by NEI grant EY021271. Thanks to John Cunningham for useful discussions.

Email: gem2131@columbia.edu