Optimization of fully differentiable ODE neurons using gradient descent

Ilenna S. Jones (ijones@fas.harvard.edu)
Harvard Kempner Institute, 150 Western Ave
Allston, Boston, Massachusetts 02139 USA

Konrad P. Kording (kording@upenn.edu)
University of Pennsylvania Neuroscience and Bioengineering, 3400 Hamilton Ave
Philadelphia, PA 19104 USA

Abstract

Neuroscientists fit simulations of single neurons to data. Fitting morphologically and biophysically detailed neuron models is computationally expensive as typical gradient-free approaches, such as evolutionary algorithms, converge slowly for neurons with many parameters. Here we introduce a gradient-based algorithm using differentiable ODE solvers, a class of models that scales well to high-dimensional problems. We employ GPUs to efficiently run many morphologically detailed neuron simulations in parallel and thus fit heterogeneously distributed ion channel densities. We use this efficient optimization algorithm to provide a proof of concept by fitting models analogous to specific experimental conditions in less than 4 hours on 1 GPU. We find that individually stimulating all dendritic compartments of the model produces outputs that lead to identifiable models. It reliably converges, even when limited numbers of recording sites. However, limiting stimulation sites reduces the reliability of this optimization method. Our approach makes model fitting efficient with the potential to allow models to have many parameters. Differentiable neuron models promise a new era of optimizable neuron models with many free parameters, a key feature of real neurons.

Keywords: optimization; neuron modelling; conductance model; gradient descent

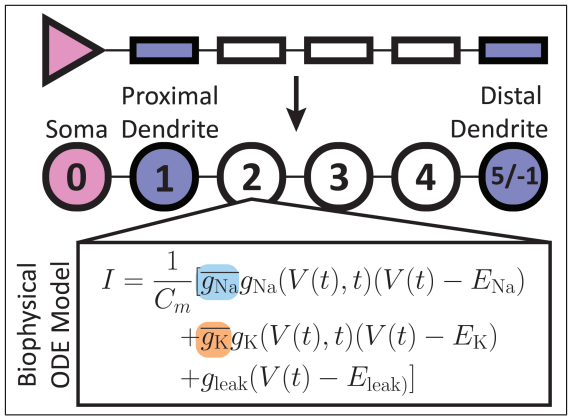

Figure 1: 6-compartment conductance-based neuron model with 2 maximal conductance parameters each.

In systems neuroscience, explanation of multi-neuronal systems often relies on single-neuron models. Depending on the purpose or scientific question, the form of a neuron model is determined by the judgment of the modeler. Point neurons are often the neuron model of choice in multineuronal systems, which abstracts away dendritic detail and separates these models from biological reality (Schutter, 2008). This and other dimensionality-reducing assumptions are made because fitting morphologically and biophysically detailed versions of these neuron models is a computationally expensive high-dimensional optimization problem (Amsalem et al., 2020).

Method

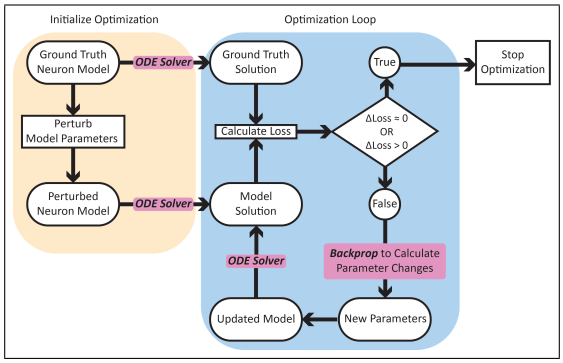

Figure 2: Optimization protocol with differentiable ODE solver and backprop

In order to more efficiently optimize neuron models, we turned to the field of machine learning and sought to use the gradient-based backpropagation of error (backprop) algorithm to optimize simple multicompartment neuron models in a supervised learning optimization protocol. The backprop algorithm is used to optimize artificial neural networks in linear time with respect to the parameters, as opposed to quadratic time in other gradient-based methods (Gegenfurtner, 1992; Carnevale & Hines, 2006), allowing the optimization of artificial neural networks with several magnitudes more parameters than typical neuron models (Bianco, Cadene, Celona, & Napoletano, 2018; Gidon et al., 2020). Backprop requires that every operation in the "forward" calculation of the model algorithm is differentiable. We implemented a system of ordinary differential equations (ODEs) describing a conductance-based neuron model (Gidon et al., 2020) using the python

library PyTorch (Paszke et al., 2019) and used a fully differentiable ODE solver (Chen, Rubanova, Bettencourt, & Duvenaud, 2019) to produce a voltage trace solution. This differentiable "forward" calculation allows use of the backprop algorithm to optimize a biophysically detailed ODE neuron model with dendrites.

The ODE neuron model has 2 learnable maximal conductance parameters per compartment, corresponding to sodium ($\overline{g_{Na}}$) and potassium ($\overline{g_K}$) voltage gated ion channels (Fig. 1). These parameters correspond to the density distribution of ion channels across the dendritic tree of the neuron model, which determines the intrinsic excitability of the neuron. We initialized the optimization protocol by producing a "ground truth" (GT) model with randomized values for $\overline{g_{Na}}$ and $\overline{g_K}$ parameters and generated a set of 100 voltage trace outputs for 100 different inputs. Using a single GPU, we simulated 100 identical optimizable models receiving these inputs in parallel for each optimization loop epoch, and followed the optimization protocol detailed in Figure 2. After optimization, we then calculated the difference between the GT parameter values and the trajectories of the inferred parameters during optimization to observe if the optimization protocol helped the model approach ground truth. Measuring the approximation to the ground truth parameter values allows us to validate the effectiveness of this optimization method in multiple simulation contexts.

Results

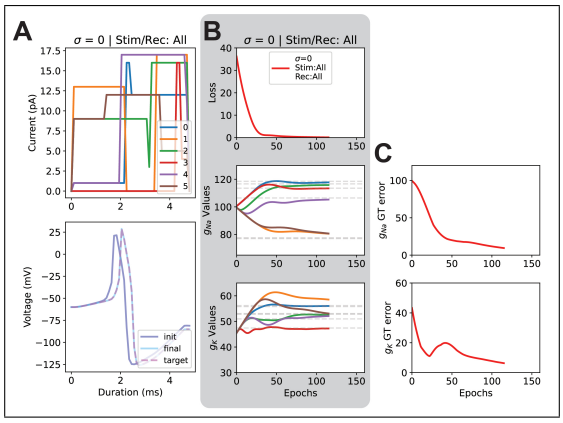

Figure 3: Optimization using input/output of all compartments.

To test if the optimization method is effective under different experimental design decisions, we tested the impact of limits to the number of patched compartments. We first tested if the optimization method could work in the best conditions: stimulating and recording from all 6 compartments of the neuron model. Figure 3AB shows that the optimization protocol successfully optimizes the neuron model voltage trace output with a loss that monotonically converges close to zero. Figure 3B shows the trajectories of each parameter over optimization and Figure 3C shows that the distance from GT approaches zero before early stopping. The model was effectively optimized in less than 4 hours on a single GPU. This demonstrates that optimization of a neuron model is effective at fitting the model output to GT output and inferring the original GT parameters.

We then stress-tested the optimization method by limiting patched compartment stimulation or recording. This is a loose analogy to glutamate imaging, glutamate uncaging, and voltage dye imaging (Ellis-Davies, 2011; Fino et al., 2009; Aggarwal et al., 2023; Aseyev, Ivanova, Balaban, & Nikitin, 2023). The loss curves and ground truth error in Figure 4F shows that limiting patched compartment stimulation to combinations of soma, proximal dendrite, and distal dendrite locations reduces the effectiveness of the method to fit GT output and infer GT parameters. However, Figure 4L shows that stimulating all compartments and limiting compartment recordings leads to successful model fitting. Notably, soma-only stimulation, (Fig. 4E), a common experimental design condition, with full morphology recordings completely fails to fit the GT voltage trace output or begin to fit the GT parameter values (Fig. 4F). However, we found using the combination of soma, proximal dendrite, and distal dendrite patch stimulation had the best limited stimulation condition performance (Fig. 4B), which shows that increased dendritic recordings produces helpful data for optimization using the gradient descent. We confirm that stimulating dendrites makes identification of distributed parameters of dendrites far easier than when using somatic current inputs only.

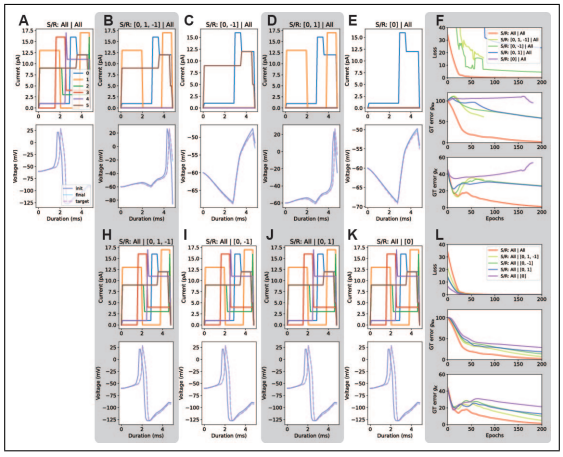

Figure 4: Optimization using input/output of limited compartments

Conclusion

This efficient method was implemented using a single GPU, which could improve accessibility to model tuning. This optimization method has the potential to make tuning high-parameter biologically realistic neuron models feasible, which could allow exploration of the heterogeneity, complexity, and mechanistic detail of neurons in multineuronal systems in both theoretical and systems neuroscience.

Acknowledgments

We would like to acknowledge the members of the Kording Lab, specifically Jordan Matelsky and Richard Lange, for help in the development of this project. This work was funded by grants from the National Institute of Health, National Science Foundation, and Howard Hughes Medical Institute.

References

Aggarwal, A., Liu, R., Chen, Y., Ralowicz, A. J., Bergerson, S. J., Tomaska, F., ... Podgorski, K. (2023, June). Glutamate indicators with improved activation kinetics and localization for imaging synaptic transmission. *Nature Methods*, *20*(6), 925–934. doi: 10.1038/s41592-023-01863-6

Amsalem, O., Eyal, G., Rogozinski, N., Gevaert, M., Kumbhar, P., Schürmann, F., & Segev, I. (2020, January). An efficient analytical reduction of detailed nonlinear neuron models. *Nature Communications*, *11*(1), 288. doi: 10.1038/s41467-019-13932-6

Aseyev, N., Ivanova, V., Balaban, P., & Nikitin, E. (2023, June). Current Practice in Using Voltage Imaging to Record Fast Neuronal Activity: Successful Examples from Invertebrate to Mammalian Studies. *Biosensors*, *13*(6), 648. doi: 10.3390/bios13060648

Bianco, S., Cadene, R., Celona, L., & Napoletano, P. (2018). Benchmark Analysis of Representative Deep Neural Network Architectures. *IEEE Access*, *6*, 64270–64277. doi: 10.1109/ACCESS.2018.2877890

Carnevale, N. T., & Hines, M. L. (2006). *The NEURON Book*. Cambridge: Cambridge University Press. doi: 10.1017/CBO9780511541612

Chen, R. T. Q., Rubanova, Y., Bettencourt, J., & Duvenaud, D. (2019, December). *Neural Ordinary Differential Equations.* arXiv. doi: 10.48550/arXiv.1806.07366

Ellis-Davies, G. C. R. (2011, April). Two-Photon Microscopy for Chemical Neuroscience. *ACS Chemical Neuroscience*, *2*(4), 185–197. doi: 10.1021/cn100111a

Fino, E., Araya, R., Peterka, D., Salierno, M., Etchenique, R., & Yuste, R. (2009). RuBi-Glutamate: Two-photon and visible-light photoactivation of neurons and dendritic spines. *Frontiers in Neural Circuits*, *3*.

Gegenfurtner, K. R. (1992, December). PRAXIS: Brent's algorithm for function minimization. *Behavior Research Methods, Instruments, & Computers*, *24*(4), 560–564. doi: 10.3758/BF03203605

Gidon, A., Zolnik, T. A., Fidzinski, P., Bolduan, F., Papoutsi, A., Poirazi, P., ... Larkum, M. E. (2020, January). Dendritic action potentials and computation in human layer 2/3 cortical neurons. *Science*, *367*(6473), 83–87. doi: 10.1126/science.aax6239

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., ... Chintala, S. (2019, December). *PyTorch: An Imperative Style, High-Performance Deep Learning Library.* arXiv. doi: 10.48550/arXiv.1912.01703

Schutter, E. D. (2008, May). Why Are Computational Neuroscience and Systems Biology So Separate? *PLOS Computational Biology*, *4*(5), e1000078. doi: 10.1371/journal.pcbi.1000078