MOSTIPS: Single-cell model optimisation using stimuli to isolate parameters

Felix B. Kern, Thomas Nowotny, György Kemenes

Sussex Neuroscience, University of Sussex, Brighton, UK — fbk21@sussex.ac.uk

Project support
UoS School of Life Sciences Doctoral Studentship grant
EPSRC (EP/J019690/1, EP/P006094/1)
HFSP (RGP0053/2015)
BBSRC (BB/P00766X/1)

PRINTED BY SCIENCEPOSTERS .co.uk

Introduction

Fitting computational, Hodgkin-Huxley-like neuron models to specific cells is a resource intensive and time consuming task and is usually done by matching the model against a voltage trace or derived features. This typically requires hours of offline processing, thus making it impossible to validate the resulting model in the original cell.

Non-linear interaction between parameters are a major impediment to fast, accurate model optimisation. Here, we propose a method that uses evolutionary algorithms to **design voltage clamp stimuli that reduce interactions between parameters**, emphasizing sensitivity to a target parameter. Using such stimuli, we can **reduce fitting time to mere minutes**, allowing immediate validation and use of the fitted model.

Convergence

Standard deviation of parameter vectors within fitting populations, pooled across stimulus sets generated with both cluster and bubble methods. Fitting methods, top: GA with 30% crossover probability; bottom: DE.

Using the information about which parameters are influential to a given observation has clear benefits, as shown by the greater convergence with graded (DE) and target-only (GA and DE) mutation schemes.

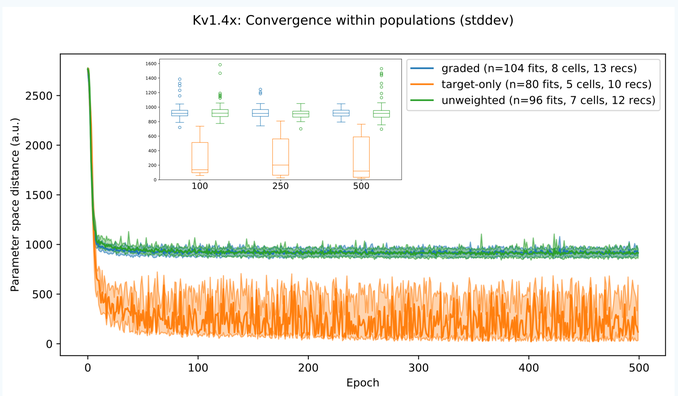

Model fitting

To fit models to live or simulated data, stimuli are applied in random order while a genetic algorithm (*GA*) or a differential evolution algorithm (*DE*) minimises the current residual within the corresponding observation windows. In each epoch, parameters are mutated according to one of three schemes (*unweighted*: all parameters are mutated equally; *graded*: mutation rate is proportional to parameter influence J_p in the next observation; *target-only*: only the parameter targeted by the next observation is mutated). By biasing exploration towards sensitive parameters, the graded and target-only mutation schemes lead to swift convergence without compromising fit quality.

As a proof of concept, we have attempted to fit a model of the *Lymnaea stagnalis* B1 motoneuron to synthetic data polluted with white noise, using as reference either the parameter set used during stimulus generation, or a randomly perturbed version.

Satisfied with the performance, we then turned to ion channels ($rK_V1.4$ and $rK_V2.1$) ectopically expressed in *Xenopus* oocytes, where the assumptions of single-compartment models apply. Each channel is modelled with two components. To obtain reference values for each parameter, we used classical voltage-clamp step protocols to measure (passive parameters) or least-squares fit (active parameters) the "true" parameter values.

Stimulus generation

The sensitivity of a model's output (here, the clamp current $I(t)$) to the value of a given target parameter changes dynamically with the model's internal state and any applied stimulus. To quantify this sensitivity, we approximate the Jacobian $\mathbf{J}(t)$ with respect to all parameters of interest, and define sensitivity to a parameter p as $J_p(t) = \partial/\partial p \; I(t) \; / \; |\mathbf{J}(t)|$. We judge stimulus/observation pairs to be highly capable of isolating p when they maximise J_p either in isolation (e.g. blue "*bubble*" for J_{gNa}), or while maintaining a consistent distribution of relative sensitivities across all parameters (e.g. purple "*cluster*" for J_{gK}). We then artificially evolve such stimulus/observation pairs, selecting for high isolation capability and screening for robustness to noise and parameter variation.

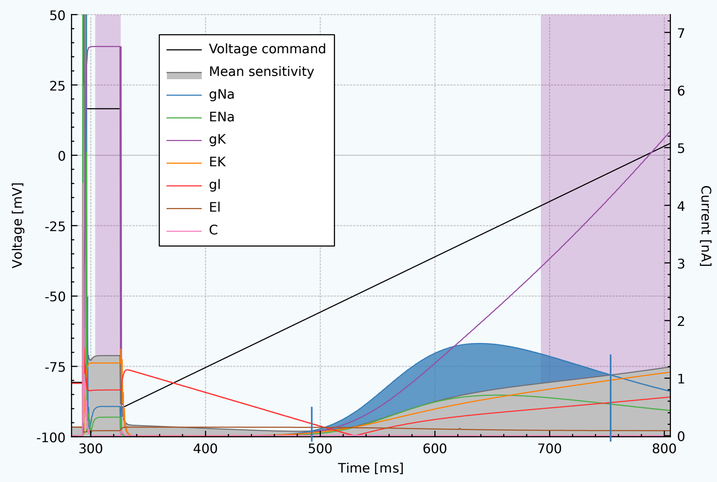

Proof of concept

Aggregate results fitting against **synthetic data**, using a model of an invertebrate neuron (Vehovszky, Szabo, Elliott, BMC Neurosci, 2005) are shown below. Stimuli were by the "bubble" method, and fitting was done with a target-only GA. Traces represent the average and worst case deviations from the respective target parameter of the best model at each epoch across fits, and are smoothed to improve legibility. Note the different scales.

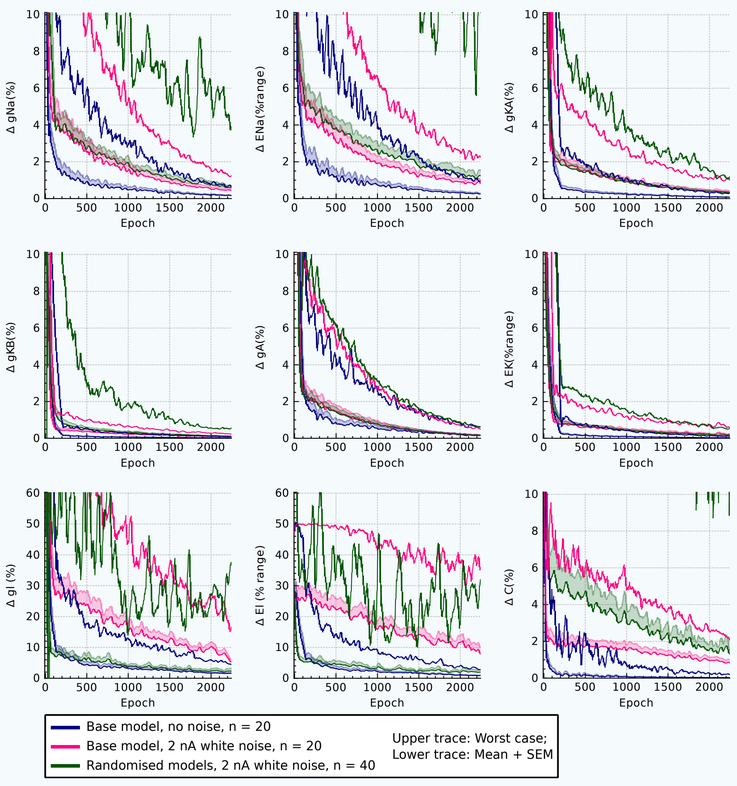

Similarity to classical fit

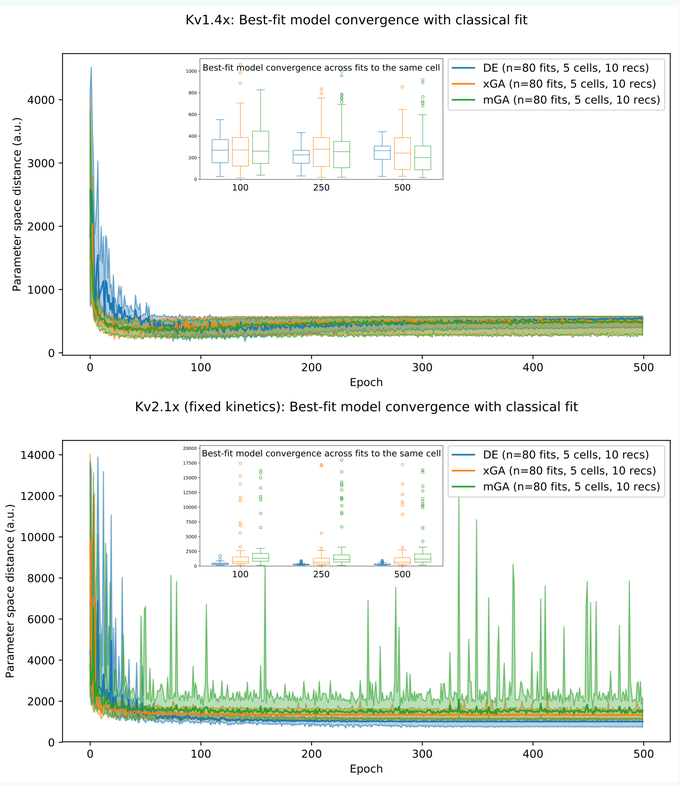

Shown on the left are parameter-space distances between models with the lowest fitting cost in each fit and the reference parameter values derived from classical voltage clamp protocol fitting, which includes direct measurement (of capacitance and leak current, at strongly hyperpolarised membrane potentials) and least-squares fitting against voltage step families to probe both activation and deactivation. Results pooled across stimulus sets generated with both bubble and cluster methods. Mutation scheme, top: target-only; bottom: graded. xGA and mGA refer to GA with (30%) or without crossover, respectively.

All tested methods arrive at their final parameter values within <100 epochs (~100,000 model evaluations). Although there is considerable divergence in the parameter values found between methods, results within a given method are largely stable across fits (see insets), differing less among each other than from the classical reference values. As the data here are pooled across several stimulus sets, this suggests that the MOSTIPS and classical fitting methods extract different information from the data provided by their respective voltage clamp protocols. Although agreement between different methods would be desirable, we consider cross-validation, i.e. generalisation against novel data, more important, see below.

Cross-validation

Cross-validation was performed by simulating the parameter sets found by the MOSTIPS and classical method against all stimuli used with the respective cell. Here, we report the root mean squared error between the recorded and simulated clamp current traces. The insets show the logarithm of the error ratio between the MOSTIPS fits and their respective classical reference fit; values below zero indicate that the MOSTIPS values provide a better overall fit. Mutation scheme, top: graded; fitting method, bottom: DE.

The cross-validation error, though well above noise level (estimated at 84 ± 172 nA across the recordings used here), is comparable between MOSTIPS and classical fits, with MOSTIPS outperforming the classical fit in many cases. The rather large absolute error is attributable to two causes: Firstly, we have used single, unfiltered recordings without leak subtraction. The leak current drifts over the course of the experiment, but since the MOSTIPS fits only have access to a short slice of data, we did not adjust the corresponding parameter value between the various recordings of each cell. Secondly, despite our best efforts in modelling the single channels used here with two components each (and freely adjustable kinetics, top), some model mismatch remains.

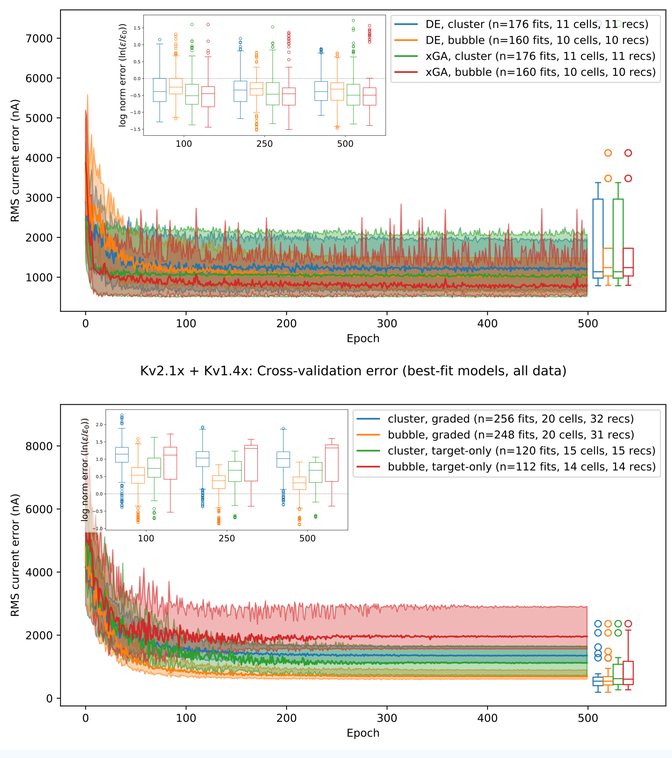

Conclusions

By shifting much of the time- and resource-intensive work from fitting towards the preparation of stimuli, the MOSTIPS method allows fast, reliable model optimisation with small amounts of data. The fits demonstrated here are based on just one second of stimulation and recording per parameter (6-8, except for the Kv2.1x kinetics fit with 28 parameters). Coupled with a simulation strategy that makes heavy use of parallel processing on a GPU, we arrive at fully optimised models in under ten minutes, which allows subsequent use of the model in the exact cell it was derived from, e.g. for dynamic clamp experiments that require a well-parametrised model.