Model-based optimal interpolation and filtering for noisy, intermittent biophysical recordings

Quentin JM Huys and Liam Paninski

Gatsby Computational Neuroscience Unit, UCL and Department of Statistics, Columbia University

qhuys@gatsby.ucl.ac.uk, liam@stat.columbia.edu, www.gatsby.ucl.ac.uk/~qhuys

Introduction

Imaging methods give us unprecedented access to detailed cellular data.

Djurisic et al. 2004

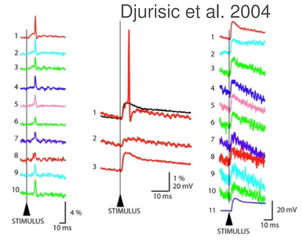

However, these data often suffer from

a) small signal-to-noise ratio (voltage-sensitive dyes)
b) confounds by unobserved variables (calcium for voltage dyes, voltage for calcium dyes)

Here we propose to address both issues in a probabilistic framework. We will use knowledge about the kinetics of the cell under investigation to smooth the noisy data, and to infer unobserved processes.

Method

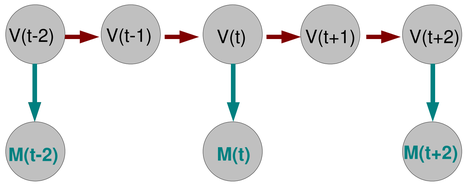

- Measure M at various points in time
- Want to infer underlying variable V over same time interval
- Only infer some probability distribution p(V|M) or even, just some aspects of that distribution, such as maximum, or mean.

$$\frac{dV}{dt} = \frac{1}{C}\sum_c \bar{g}_c o_c(t)(E_c - V(t)) + \sigma_V \eta_t = \phi(V, o, \bar{g}) + \sigma_V \eta_t$$
$$\frac{dm}{dt} = (1-m)\alpha_m(V) - m\beta_m(V)$$

Given V(t) we can evaluate V(t+1)

$$V_{t+1} = V_t + \phi(V_t, o_t, \bar{g})dt + \eta_t \qquad p(V_{t+1}|V_t, o_t, \bar{g}) = \mathcal{N}(\phi_t, \sigma)$$

Also know how measurement relates to voltage

$$M_t = V_t + \sigma_O \xi_t \qquad p(M_t|V_t) = \mathcal{N}(V_t, \sigma_O)$$

Together, these distributions define the likelihood for hidden voltage traces given observations

$$p(M_{1:T}|V_{1:t}) = p(V_1)\prod_{t=1}^{T-1} p(M_t|V_t)p(V_{t+1}|V_t)$$

Use this knowledge to find voltage traces with high likelihood.

Or rather, find particular voltage traces that allow us to evaluate quantities of interest, such as the mean:

$$< V_{1:T} > = \int dV_{1:T} p(V_{1:T}|M_{1:t}) = \int dV_{1:T} q(V_{1:T}) \underbrace{\frac{p(V_{1:T}|M_{1:t})}{q(V_{1:T})}}_{w} \approx \sum_i V^i_{1:T} w_i$$

So we draw samples from the approximate q(V(1:T))

$$V^i_{1:T} \sim q(V_{1:T}) \qquad w_i = \frac{p(V^i_{1:T}|M_{1:t})}{q(V^i_{1:T})}$$

The approximate q(V) is given recursively by

$$q(V_t) = p(V_t|V_{1:t-1}, M_{1:t}) \propto p(V_t|V_{t-1})p(V_{t-1}|V_{1:t-2}, M_{1:t-2})p(M_t|V_t)$$

Which means we simply run lots of neurones at once, and weigh them by the likelihood they give to the data.

$$w^{*i}_t = w^i_{t-1} p(M_t|V^i_t) \quad \tilde{w}^i_t = w^{*i}_t / (\textstyle\sum_j w^{*j}_t)$$

Results

If we have access to the true dynamics of the neurone under investigation, we can recover the voltage trace and the open fractions very efficiently from very noisy data.

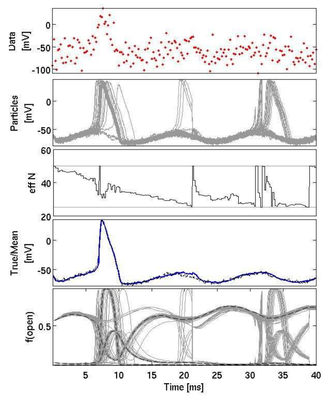

It also allows filtering – inference of voltage and open fractions at times at which no (noisy) samples were observed.

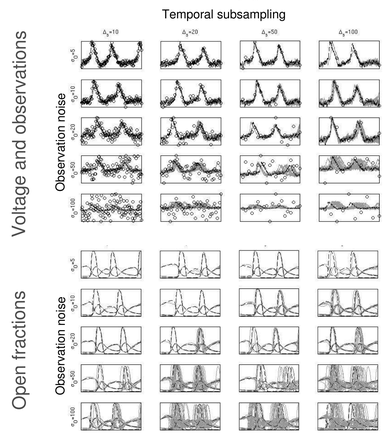

As long as we have a kinetic model, it allows us to infer other hidden variables too, in particular [Ca⁺⁺] can be inferred from Voltage signals, and conversely voltage signals can be inferred from [Ca⁺⁺] recordings

Inferring [Ca⁺⁺] from voltage data

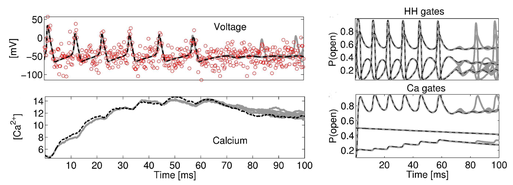

Inferring voltage from [Ca⁺⁺] data

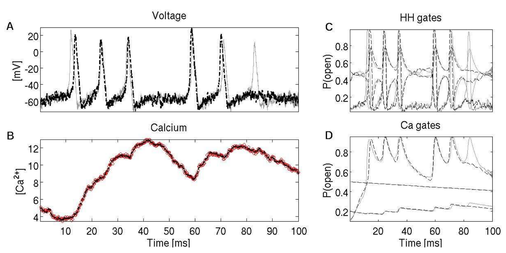

Channel densities

So far we have assumed exact knowledge of both channel kinetics and channel concentrations. Here we relax these assumptions and also infer channel concentrations.

EM

The EM algorithm iterates between finding the hidden variables, and updating the parameters that depend on them. Given channel densities g, we have above found a set of probable underlying voltages $\{V^i_{0:T}\}$. Given the true voltage, the densities are easily found by setting:

$$\hat{\mathbf{g}} = \arg\max_{\mathbf{g}} ||\dot{\mathbf{V}} - \mathbf{J}\mathbf{g}||^2 \qquad [\mathbf{J}\mathbf{g}]_t = \sum_c \bar{g}_c \underbrace{o_c(t)(E - V(t))}_{J_{ct}}$$

However, we have a set of samples rather than the true V, so we set instead

$$\hat{\mathbf{g}} = \arg\max_{\mathbf{g}} \left\langle ||\dot{\mathbf{V}} - \mathbf{J}\mathbf{g}||^2 \right\rangle_{q(V_{1:T})}$$

which depends only on three expected sufficient statistics

$$\langle J_{ct} J_{c't} \rangle \quad \langle J_{ct} V_t \rangle \qquad\qquad \langle J_{ct} V_{t+1} \rangle$$

where averages are over marginal distributions $p(V_t|M_{1:T})$ and $p(V_t, V_{t+1}|M_{1:T})$

To obtain marginals given all data, we add a backward update to the weights:

$$w^i_t = \tilde{w}^i_t \left(\sum_j \frac{w^j_{t+1} p(V^j_{t+1}|V^i_t)}{\sum_k \tilde{w}^k_t p(V^j_{t+1}|V^k_t)} \right) \qquad w^{ji}_t = \tilde{w}^i_t \left(\frac{w^j_{t+1} p(V^j_{t+1}|V^i_t)}{\sum_k \tilde{w}^k_t p(V^j_{t+1}|V^k_t)} \right)$$

Inferring channel densities from noisy data

Here we started from a non-spiking channel configuration and found a set of channel densities that leads to spiking and matches the data well. Note that only few data points make out the action potential.

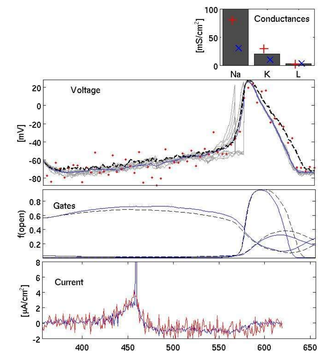

There may be many channel density combinations that lead to acceptable reconstructions.

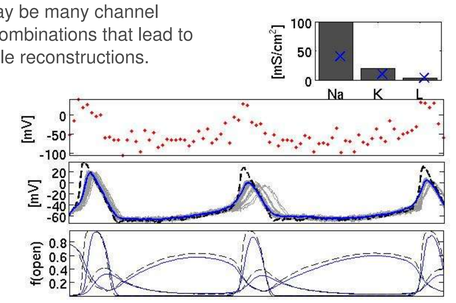

Discussion

Knowledge of kinetics and channel densities can be used to

smooth noisy, subsampled data
infer unobserved variables (voltage, open fractions, calcium)
infer even fast variables from slow (but not very noisily observed) variables – voltage from calcium

Knowledge of kinetics can be used to infer channel densities from very noisy data.

But:

EM is plagued by local minima – need good initial conditions (eg a spiking neurone)
Particle filters are just importance samplers – in high dimensions, rely on accurate sampling distribution.

References

Djurisic et al. 2004: Voltage imaging from dendrites of mitral cells: EPSP attenuation and spike trigger zones. *J. Neurosci.* 24 (30):6703-14; **Huys et al. 2006:** Efficient estimation of detailed single-neurone models. *J. Neurophysiol.* 96:872-96