Efficient computation of the most likely path in integrate-and-fire and more general state-space models

Shinsuke Koyama[1], Rob Kass[1], and Liam Paninski[2]

[1]Carnegie Mellon University, [2]Columbia University

A number of important models in neuroscience may be described in "state-space" form with point-process observations: a hidden state variable evolves according to some continuous Markovian dynamics, and the rate of the observed spike trains is some function of this underlying hidden state. Examples include the integrate-and-fire model [4] and models used in a number of spike train decoding applications [1, 6].

It is of interest to compute the most likely (ML) path that the hidden state variable traversed on a given trial, given the observed spike trains. The point-process filter algorithm introduced in [1] computes an approximation to this ML path, but this approximation is rigorously accurate only in the so-called "high-information" limit (when we have a large number of spikes per unit time, or when the hidden state dynamics are nearly deterministic). More recently, [5] discussed methods for computing the ML path exactly (and in more generality than the state-space setting), but the computational complexity of these general methods scales like $O(T^3)$, where T is the trial length, and are therefore inapplicable for long trials.

Here we develop $O(T)$ methods for computing the exact ML path. In the case of linear Gaussian state space dynamics, the ML path in continuous time satisfies a second-order nonlinear ordinary differential equation. This equation may be solved numerically via standard $O(T)$ relaxation methods. More generally, if the dynamics are linear and driven by innovations with a log-concave density, and the point-process observations satisfy certain standard assumptions, then the optimization problem is strictly concave (guaranteeing a unique ML path), and the Newton-Raphson algorithm may be applied; each Newton update here requires just $O(T)$ time, because the Hessian of the loglikelihood is block tridiagonal. In the case of the integrate-and-fire neuron with a hard voltage threshold, barrier methods may be applied to solve the resulting constrained optimization problem [3], again in $O(T)$ time.

We describe several applications of the resulting methods. First, we may compute the (marginal) likelihood of the observed data via Laplace's method [2]; thus we may fit the model parameters via directly maximizing this marginal likelihood, or alternatively via the expectation maximization algorithm. Second, we may apply similar Laplace approximation methods to develop more efficient proposal densities for sequential Monte Carlo ("particle filtering") applications. Finally, the ML path can be used to provide a good initialization for iterative algorithms (Monte Carlo or expectation propagation) for computing conditional expectations in these models.

References

[1] E. Brown, L. Frank, D. Tang, M. Quirk, and M. Wilson. *Journal of Neuroscience*, 18:7411–7425, 1998.
[2] R. Davis and G. Rodriguez-Yam. *Statistica Sinica*, 15:381–406, 2005.
[3] L. Paninski. *Journal of Computational Neuroscience*, 21:71–87, 2006.
[4] L. Paninski, S. Iyengar, R. Kass, and E. Brown. In *Stochastic Methods in Neuroscience*. Oxford University Press, 2008.
[5] J. Pillow and L. Paninski. *Under review, Neural Computation*, 2007.
[6] W. Truccolo, U. Eden, M. Fellows, J. Donoghue, and E. Brown. *Journal of Neurophysiology*, 93:1074–1089, 2005.