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

Quentin JM Huys[1] and Liam Paninski[2]
[1] Gatsby Computational Neuroscience Unit, University College London,
[2] Department of Statistics, Columbia University
qhuys@gatsby.ucl.ac.uk, liam@stat.columbia.edu

February 7, 2006

## Introduction

Noise is an omnipresent issue that is often handled suboptimally. For example, noise is an issue in voltage-sensitive imaging – even the best dyes achieve signal-to-noise ratios of no more than $\sim 1-6\%$ (Djurisic and Zecevic, 2005). Averaging noise out is not always possible and sometimes not even desirable. Missing data are often burdensome too: in voltage dye experiments, the laser has to be moved between sites of interest and thus the data is not acquired simultaneously, leading to gaps in the data. Despite advances (Bullen et al., 1997), this problem becomes more prominent the more sites one attempts to record from. More generally, we might even be interested in a variable that has not been observed directly at all, such as the voltage in a $Ca^{++}$imaging experiment. Principled methods to filter out noise, to interpolate between data points and to infer unobserved variables could substantially complement advances in data acquisition methodology.

Here we show how, when time series recordings of a dynamical system (*eg* the voltage of a cell) are made, knowledge of the dynamical system can be used to both filter and interpolate between the measurements, providing a principled alternative to heuristics such as temporal smoothing or low-pass filtering. Neural dynamics are usually specified as Markov chains. If these dynamics are hidden (due to noisy or indirect measurements), the task of recovering the distribution over the true underlying state evolution $\mathbf{s}_{0:T}$ of the neurone over time $p(\mathbf{s}_{0:T}|V_{0:T})$ is equivalent to inference in nonlinear state space models. These models, together with their discrete analogues such as intermittent Kalman filters and Hidden Markov Models have been analysed extensively and are very well understood. In particular, if the hidden variables do indeed evolve in a Markovian manner (as is often the case), a number of algorithms from the machine learning literature allow efficient sampling from $p(\mathbf{s}_{0:T}|V_{0:T})$ (cf. Doucet et al., 2001; Godsill et al., 2004), despite the huge size of this state space ($\mathcal{O}(N^T)$ for a state space of size $N$). In particular, we find that the combination of a nonlinear Gaussian state space model with Gaussian observation noise and the forward-backward particle filter, a simple algorithm frequently applied to Hidden Markov Models allows us to recover the true voltage of a Fitzhugh-Nagumo (FHN) spiking model very well.

## Methods

To illustrate the performance of the method, we run it on a noisy Fitzhugh-Nagumo neurone by writing

$$\begin{aligned} dV &= (-V(V-a)(V-1) - W + I)dt/\tau + \sqrt{dt}\sigma_V dN(t) \\ dW &= (V - gW)dt + \sqrt{dt}\sigma_W dN(t) \end{aligned}$$

where $V$ is the voltage and $W$ the recovery variable ($\mathbf{s} = \{V, W\}$), $I$ the constant current input, $dN(t)$ is iid Gaussian white noise, $\sigma^2_{v,w}$ are the variances of the current and hidden state noise respectively, which defines a Markov process

$$\begin{aligned} V(t+dt) &\sim \mathcal{N}(V(t) + f(V,W,I)dt/\tau, \sqrt{dt}\sigma_V) \\ W(t+dt) &\sim \mathcal{N}(W(t) + g(V,W)dt/\tau, \sqrt{dt}\sigma_W) \end{aligned}$$

The voltage $V(t)$ is sampled at sampling intervals $\Delta$ and corrupted with iid Gaussian white noise $\eta_t$ to obtain $\{\tilde{V}_t\}_{t=0}^T$. Thus noise is included both in the evolution equations themselves, and in the observation process.

The filtering task is now to recover the mean of $p(V_{0:T}|\tilde{V}_{0:T}, \theta)$ (or some other function involving the integral over this distribution) with $\theta$ parametrising our model for the same times as data points were observed, whereas

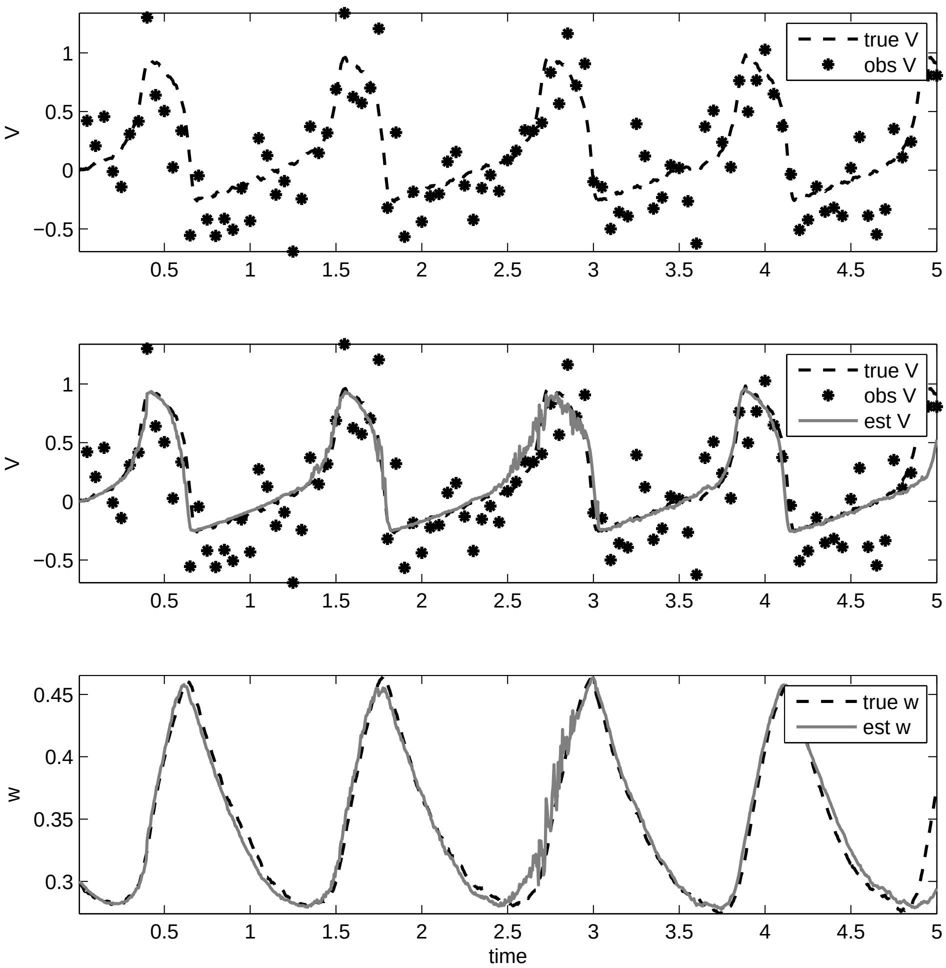


Figure 1: An application of the forward-backward particle filter to noisy, subsampled Fitzhugh-Nagumo model data. **Top panel**: True voltage path (dashed trace) and noisy, subsampled observations (dots; x10 subsampling ratio). We emphasize that the dotted trace is the only observed data here; the dashed trace is not observed directly. **Middle panel**: Voltage estimated by forward-backward particle filtering algorithm (gray solid trace) overlaid with true and observed voltage for comparison. **Bottom panel**: True and estimated $W$ variable. Again, the true $W$ trace is not observed directly; nonetheless, this trace may be accurately reconstructed given noisy, subsampled observations from $V(t)$. $N = 50$ particles used here; the noisy oscillatory behavior visible near time $t = 2.5$ disappears as more particles are employed.

the interpolation task is to do this at intervening times too. Finding the distribution over the entire path is feasible due to the Markovian nature of $V(t)$, and we find a *representation* of the distribution in terms of weighted samples from it (Minka, 1998). Quantities of interest include the mean and variance of $p(V_{0:T}|\tilde{V}_{0:T}, \theta)$, both of which are easily calculated in terms of the samples: Omitting $\theta$ from now on, for each time $t$ we draw $K$ samples $V_t^k \sim p(V_t|\tilde{V}_{1:t}, V_{1:t-1}^k)$ (and equally for $W$) and concatenate them: $V_{1:t}^k = (V_{1:t-1}^k, V_t^k)$. Averages over $p(V_{1:t})$ can now be approximated by a weighted sum over the samples

$$\int dV_{1:t}\, f(V_{1:t}) p(V_{1:t}|\tilde{V}_{1:t}) \approx \frac{1}{K} \sum_k f(V_{1:t}^k) \frac{b_t^k}{\sum_{k'} b_t^{k'}}$$

where the weights $b_t^k = b_{t-1}^k p(V_t|\tilde{V}_{1:t}, V_{1:t-1}^k)$ can be evaluated analytically and are only dependent on $V$ and not on $W$ (because the $W$'s are not observed). For interpolation times at which there is no observation, the sampling distributions for $V$ and $W$ are simply given by equations 1 and the weights remain unchanged (Doucet et al., 2001). As we march from $t = 0$ to $t = T$ sequentially, the weights only incorporate the effect of the past samples, and are thus not exactly representing the joint distribution $p(V_{0:T}, \tilde{V}_{0:T})$. To achieve this, the future sample path has to be taken into account as well, which is done by adding a reweighting of all samples, starting from time $T$, according to the transitions $p(V_{t-1}|V_t)$. Although the present exposition is couched in terms of observation times, this need not be the case.

Figure 1 illustrates the performance of the technique. It should also be noted that this processing is not overly computationally burdensome; these reconstructions took seconds on a laptop computer running Matlab.

## Discussion

We have here shown that a simple forward-backward particle filter can accurately recover multidimensional dynamics given intermittent samples from just one of the dynamical variables, at low signal-to-noise levels that qualitatively match those encountered in voltage imaging experiments. It is possible to formulate dynamical models of other, entirely unobserved variables (such as the voltage in a $Ca^{++}$imaging experiment) and apply the same techniques. Finally, the probabilistic form of this approach also naturally allows combination of measurements from different sources, such as voltage and $Ca^{++}$imaging, done simultaneously.

The present is a very simple algorithm and more sophisticated ones are available, such sequential importance sampling (Doucet et al., 2000). These are likely to be more efficient especially in more high-dimensional problems, especially when combined with the forward-backward algorithm. The Markovian nature reduces the complexity of sampling from exponential $\mathcal{O}(N^T)$ to linear $\mathcal{O}(N^2T)$ in the length of the recording $T$. Other methods have been applied in the past (*eg* extended Kalman filters (Voss et al., 2004)). The main advantage of the present approach lies in the addition of the backward recursion, which results in the use of samples from $p(V_t|\tilde{V}_{0:T})$ which are better informed by the data than samples from $p(V_t|\tilde{V}_{1:t})$.

We have here assumed knowledge of the dynamical system. For example, given spatiotemporal voltage data from a dendritic tree and some knowledge of the channel types present in the cell, it is possible to estimate a detailed compartmental model of the cell under investigation (Ahrens et al., 2006; Huys et al., 2006). Given such a model, the present method can be used to recover the underlying voltage trace. However, the method in Huys et al. (2006) is only efficient if the voltage was observed noiselessly in the first place. Although this is approximately true for path-clamp recordings, it does not hold of imaging techniques. In the presence of noisy measurements, it should possible to alternate between estimating the most likely (hidden) voltage path as described here, and estimating the parameters of the model (as described by (Huys et al., 2006)) in a general scheme termed expectation-maximisation (Roweis and Ghahramani, 1998), and thus generalise this model-based filtering and interpolation to approximately model-free but still optimal filtering and interpolation.

## References


Ahrens, M., Huys, Q. J. M., and Paninski, L. (2006). Large-scale biophysical parameter estimation in single neurons via contrained linear regression. In *Advances in Neural Information Processing Systems 18, In Press*, Cambridge, MA. MIT Press.

Bullen, A., Patel, S. S., and Saggau, P. (1997). High-speed, random-access fluorescence microscopy: I. High-resolution optical recording with voltage-sensitive dyes and ion indicators. *Biophys. J.*, 73(1):477–91.

Djurisic, M. and Zecevic, D. (2005). Imaging of spiking and subthreshold activity of mitral cells with voltage-sensitive dyes. *Ann N Y Acad Sci*, 1048:92–102.

Doucet, A., de Freitas, N., and Gordon, N., editors (2001). *Sequential Monte Carlo in Practice*. Springer.

Doucet, A., Godsill, S., and Andrieu, C. (2000). On sequential Monte Carlo sampling methods for Bayesian filtering. *Stat. Comput.*, 10:197–208.

Godsill, S., Doucet, A., and West, M. (2004). Monte Carlo smoothing for non-linear time series. *Journal of the American Statistical Association*, 99:156–168.

Huys, Q. J. M., Ahrens, M., and Paninski, L. (2006). Estimating detailed single-neuron models. *Submitted*.

Minka, T. P. (1998). From hidden markov models to linear dynamical systems. Technical report, University of Toronto.

Roweis, S. and Ghahramani, Z. (1998). Learning nonlinear dynamical systems using an EM algorithm. In *Advances in Neural Information Processing Systems (NIPS) 11*, Cambridge, MA. MIT Press.

Voss, H. U., Timmer, J., and Kurths, J. (2004). Nonlinear dynamical system identification from uncertain and indirect measurements. *Int. J. Bifurc. Chaos*, 14(6):1905–33.