

# A sequential Monte Carlo approach to estimate biophysical neural models from spikes

Liang Meng<sup>1</sup>, Mark A Kramer and Uri T Eden

Department of Mathematics and Statistics, Boston University, Boston, MA, 02215, USA

E-mail: [ljmeng@math.bu.edu](mailto:ljmeng@math.bu.edu)

Received 29 January 2011

Accepted for publication 8 June 2011

Published 4 November 2011

Online at [stacks.iop.org/JNE/8/065006](http://stacks.iop.org/JNE/8/065006)

## Abstract

Realistic computational models of neuronal activity typically involve many variables and parameters, most of which remain unknown or poorly constrained. Moreover, experimental observations of the neuronal system are typically limited to the times of action potentials, or spikes. One important component of developing a computational model is the optimal incorporation of these sparse experimental data. Here, we use point process statistical theory to develop a procedure for estimating parameters and hidden variables in neuronal computational models given only the observed spike times. We discuss the implementation of a sequential Monte Carlo method for this procedure and apply it to three simulated examples of neuronal spiking activity. We also address the issues of model identification and misspecification, and show that accurate estimates of model parameters and hidden variables are possible given only spike time data.

(Some figures in this article are in colour only in the electronic version)

## 1. Introduction

Developing biophysical models of neuronal systems is a challenging task. Before even attempting to solve a model, one must first confront many decisions, including the spatial scale of the model (e.g., mean field [4] versus single neuron [17]) and level of detail (e.g., multicompartiment conductance-based models [39] versus abstract mathematical models [20]). Typically, the choice of model scale and detail depends upon the available neurophysiological data. Perhaps the most common model choice for single neuron recordings *in vivo* or *in vitro* is a Hodgkin–Huxley-type neuron model [17]. This model class typically consists of four or more variables (representing the dynamics of the neuronal voltage and ionic currents) coupled nonlinearly and produces the characteristic action potential or ‘spike’ of neuronal voltage activity.

Even relatively simple neuronal models (such as of the Hodgkin–Huxley type) present significant challenges. First, complicated mathematical mechanisms govern the model

dynamics and understanding these mechanisms remains an active area of research [3, 21, 22, 34]. Second, most neuronal models possess a large number of parameters which typically remain experimentally unconstrained. For example, the original Hodgkin–Huxley model possesses at least seven parameters (e.g., capacitance, reversal potentials and maximal conductances for the sodium, potassium and leak currents) with wide ranges of possible values [17]. In computational neuroscience, a common procedure for estimating these parameters is ‘hand-tuning’ to produce simulated model dynamics that matches qualitatively the desired neuronal activity [30]. Hand-tuning approaches usually require a great deal of time and expertise [27, 38, 43]. Once a set of suitable parameters is found, it is often unclear whether the solution is unique or whether other model formulations compatible with the data exist.

Recently, researchers have attempted to replace this hand-tuning procedure with more rigorous approaches. These include brute-force approaches that conduct extensive searches of model dynamics over broad ranges of parameter sets [1, 30]. Other approaches directly utilize neuronal voltage activity to

<sup>1</sup> Author to whom any correspondence should be addressed.

estimate model parameters, in simulations [5, 16, 18, 19, 41, 45] and in experimental recordings [8, 42]. Similar estimation procedures have also been applied to estimate parameters in mean field models [32, 35]. In these studies, detailed knowledge of the voltage dynamics (e.g., the intracellular voltage trace) is required to estimate the model parameters.

Yet, often in experiments the intracellular voltage of an individual neuron remains unobserved. Instead, the only observable is a neuron's extracellular spike train response. Given only this limited knowledge—of when a neuron generates spikes, not the subthreshold neuronal voltage activity—parameter estimation is still possible in abstract neuronal models, such as the leaky integrate-and-fire (LIF) neuron [23, 26, 29]. In this work, we pursue a similar goal of quantitative parameter estimation from neuronal spike train responses. In doing so, we employ models with greater biophysical realism than the LIF model and show that, given only the neuronal spike train data, we may successfully estimate some of the biophysical model parameters.

We develop a parameter estimation framework that combines conductance-based neural modeling with point process statistical theory to estimate model components directly from a set of observed spike times. The estimation algorithm is constructed using sequential Monte Carlo (particle filter) [7, 9] methods that combine future and past spiking information to update a collection of model formulations that are consistent with the observed spiking data. When there exists a single set of parameters consistent with the observed spiking, we find that the algorithm rapidly converges to the correct values. When multiple parameter values consistent with the data exist, the algorithm produces a collection of estimates that tightly cover the space of suitable values. Finally, when no parameters associated with the selected model class are consistent with the data, the algorithm fails to converge to adequate parameter values, and the resulting dynamics do not explain the observed spike data.

The organization of this manuscript is as follows. In section 2, we introduce the parameter estimation procedure and model systems, which include the FitzHugh–Nagumo and two Hodgkin–Huxley-type models. In section 3, we simulate the models to generate spike train data, and then use these data to estimate the model parameters. We discuss the model identifiability problem for the standard Hodgkin–Huxley model, and suggest an experimental protocol to refine parameter estimates. To illustrate the issue of model misspecification, we consider two Hodgkin–Huxley-type model neurons—one with a slow intrinsic current, and one without—and show that the correct model choice (motivated by a biophysical understanding of the slow intrinsic current) is critical. We conclude in section 4 with a discussion of these results and proposals for future work.

## 2. Methods

### 2.1. The estimation procedure

**2.1.1. The state-space framework.** Biophysical models of a neuron's spiking activity are often expressed mathematically

using systems of several first-order differential equations. The Hodgkin–Huxley equations are a classic example of this sort of mathematical model [17]. In general, these models characterize the evolution of dynamic state variables representing, for example, the membrane voltage or various voltage-gated ion channels.

We express these neural dynamics in discrete time using a state-space model of the general form

$$\mathbf{S}_t - \mathbf{S}_{t-1} = \mathbf{F}(\mathbf{S}_{t-1}, \Theta) \cdot \Delta t + \epsilon_t, \quad (1)$$

where  $\mathbf{S}_t$  is a state vector including a set of unobserved state variables at time point  $t$  and  $\mathbf{F}$  represents a vector function that characterizes the expected evolution of the state variables. The function  $\mathbf{F}$  also requires the specification of a set of parameters  $\Theta$  that govern the dynamics of the state variables.  $\Delta t$  is a fixed discretization interval, and  $\epsilon_t$  is an additive white-noise term that leads to variability in the dynamic activity.

For any realization of the state process from the start time to time  $T$ ,  $\mathbf{S}_{1:T}$ , the spike times associated with that realization are determined exactly by the large amplitude depolarizations of the membrane voltage. If we assume that  $\mathbf{S}_t$  represents the true dynamics of the neuron whose spiking activity we observe, we could estimate the parameters  $\Theta$  and hidden variables in  $\mathbf{S}_t$  by computing the probability distribution of spikes occurring at the observed times based only on the dynamics specified by (1). However, this distribution is difficult to compute and estimation algorithms based on this approach may be computationally intractable. Therefore, instead of identifying  $\mathbf{S}_t$  as the exact dynamics of the observed neuron, we consider the state process to be a possible surrogate for the unobserved neural dynamics, which provides information about the probability of observing a spike at any moment in time rather than the exact time of a spike. In other words, we construct a model for the probability of observing a spike at any time step based on the realization  $\mathbf{S}_t$ . This probability of a spike at any time takes the form of a conditional intensity function  $\lambda_t = \lambda(\mathbf{S}_{1:t+k}, \mathbf{H}_t, \Theta) = \lim_{\Delta t \rightarrow 0} \frac{p(\Delta N_t = 1 | \mathbf{S}_{1:t+k}, \mathbf{H}_t, \Theta)}{\Delta t}$ , where  $\Delta N_t$  is the number of spikes occurring at time step  $t$  and  $\mathbf{H}_t$  is the past history of spiking up to time  $t$ . The intensity function might depend on both past ( $\mathbf{S}_{1:t}$ ) and future states ( $\mathbf{S}_{t+1:t+k}$ ). Given a specific form of this conditional intensity function, the probability of observing  $\Delta N_t$  spikes at time  $t$  is [37]

$$p(\Delta N_t | \mathbf{S}_{1:t+k}, \mathbf{H}_t, \Theta) = e^{\Delta N_t \log(\lambda_t \Delta t) - \lambda_t \Delta t}. \quad (2)$$

Here, we assume that the time step is small enough so that there is at most one spike in any single time step. Together, (1) and (2) form a state-space model with spike observations ( $\Delta N_t$ ), which will be the foundation for estimating the dynamic variables ( $\mathbf{S}$ ) and making inferences about the model parameters ( $\Theta$ ).

**2.1.2. Estimation algorithm.** In the previous subsection, we formalized models for both the hidden state variables (e.g., the membrane voltage and ionic gates) and the observable variable (the spike process,  $\Delta N_t$ ). Given the observable variable, we would now like to use these models to estimate values for the state variables and the model parameters. To do

so, we construct a sequential Monte Carlo algorithm (SMC) [7, 9], also known as a particle filter, to estimate the posterior probability distribution of these unknown quantities, given the observed spike data. Particle filters are so named because they represent the distribution of an unknown state process using a collection of weighted samples, or particles. At any time step, each particle represents possible values for the unknown variables and parameters, and the weighting function represents the probability associated with these values. As the number of particles becomes very large, this SMC characterization becomes more and more accurate.

In order to develop the details of the algorithm, let  $\mathbf{U}_t^i = \{\mathbf{S}_t^i, \Theta_t^i\}$  denote joint states containing both the state variables and parameters of the  $i$ th particle at time point  $t$ . Although the parameters are typically static, for the purpose of this algorithm, we allow the parameters to vary with time. The details of this treatment of model parameters is discussed in [24]. Let  $w_t^i$  be the weight associated with particle  $\mathbf{U}_t^i$ . When  $w_t^i$  is high, it suggests that values near  $\mathbf{U}_t^i$  are likely. There are multiple approaches to computing the values and weights of each particle at any time. In this case, we construct a bootstrap particle filter [9], where the initial values for the particles at time  $t$  are sampled from the particles at the previous time step. The values of each particle are then updated by simulating (1). The weights of each particle are updated by multiplying by the likelihood of the observed spiking data at time  $t$  given by (2),  $w_t^i = w_{t-1}^i p(\Delta N_t | \mathbf{U}_t^i)$ . Intuitively, each particle from the previous time step undergoes one step of the model dynamics. If the resulting state values are consistent with the newly observed data, the weight is increased. If the data are unlikely given the state values for a particle, its weight is decreased.

A common problem with particle filters is the degeneracy phenomenon, where after a few iterations, all but one particle will have negligible weight [7]. It has been shown that the variance of the weights can only increase over time, and thus, it is impossible to avoid the degeneracy phenomenon [6]. To reduce the effect of degeneracy, we use a resampling scheme. The basic idea of resampling is to eliminate particles that have small weights and to concentrate on particles with large weights. Here, we use a residual resampling scheme [24] whereby particles with large weights are replicated based on their weight and particles with small weights have some probability of surviving and some probability of being eliminated. Let  $n$  be the number of particles used. We retain  $M_i = \lfloor nw_t^i \rfloor$  copies of  $\mathbf{U}_t^i$ , where  $\lfloor \cdot \rfloor$  indicates rounding down to the nearest integer, and then obtain  $n - \sum_i M_i$  i.i.d. draws from  $\mathbf{U}_t^i$  with probabilities proportional to  $nw_t^i - M_i$ ,  $i = 1, \dots, n$ . After resampling, the weights of each particle are reset to  $1/n$ .

At every time step the algorithm produces a collection of particles, each containing proposed values for the unknown state variables and parameters. We construct estimates for the unknown quantities by computing their sample means over all the particles, and construct approximate 95% confidence intervals by computing the 2.5% and 97.5% percentiles of the particle values. For convenience, we provide a pseudo-code description of the algorithm in the [appendix](#).

## 2.2. Biophysical models

We employ three types of biophysical models. The first is a FitzHugh–Nagumo-type model, a simplified model of neuronal spiking activity [10, 28]. The model consists of two equations:

$$\begin{aligned}\dot{V} &= V(a - V)(V - 1) - w + I \\ \dot{w} &= bV - cw.\end{aligned}$$

Here,  $V$  mimics the membrane voltage and ‘recovery’ variable  $w$  mimics activation of an outward current. Parameter  $I$  mimics the resting current. Parameter  $a$  describes the shape of the cubic parabola  $V(a - V)(V - 1)$ , and parameters  $b > 0$  and  $c \geq 0$  describe the kinetics of the recovery variable  $w$ .

In order to apply the aforementioned (1) and (2), we need to discretize the continuous model equations. We do so in the standard way:

$$\begin{aligned}V_t &= V_{t-1} + [V_{t-1}(a - V_{t-1})(V_{t-1} - 1) - w_{t-1} + I] \\ &\quad \cdot \Delta t + \epsilon_t\end{aligned}\quad (3)$$

$$w_t = w_{t-1} + [bV_{t-1} - cw_{t-1}] \cdot \Delta t, \quad (4)$$

where  $\{\epsilon_t\}$  represents a noisy current term defined to be Gaussian white noise with variance  $\sigma^2 \Delta t$ .  $\sigma$  describes the magnitude of the noise term. These same discretization process and noisy current input are employed for all models throughout the rest of the paper. We use the model for two purposes—to generate the ‘experimental’ spike times defined as the local large amplitude peaks of the voltage process,  $V_t$ , and to estimate the model parameters and hidden variables—as discussed in the next section.

The second model we consider is the standard Hodgkin–Huxley model [17]. This model consists of four variables that describe the voltage  $V$  dynamics and three intrinsic currents—potassium current  $I_K$ , sodium current  $I_{Na}$  and leak current  $I_L$ . The equations for the Hodgkin–Huxley model are

$$\begin{aligned}C\dot{V} &= I - \overbrace{g_K n^4 (V - E_K)}^{I_K} - \overbrace{g_{Na} m^3 h (V - E_{Na})}^{I_{Na}} \\ &\quad - \overbrace{g_L (V - E_L)}^{I_L} \\ \dot{n} &= \alpha_n(V)(1 - n) - \beta_n(V)n \\ \dot{m} &= \alpha_m(V)(1 - m) - \beta_m(V)m \\ \dot{h} &= \alpha_h(V)(1 - h) - \beta_h(V)h\end{aligned}\quad (5)$$

where

$$\begin{aligned}\alpha_n(V) &= \alpha_0 \frac{10 - V}{\exp\left(\frac{10 - V}{10}\right) - 1} \\ \beta_n(V) &= \beta_0 \exp\left(\frac{-V}{80}\right) \\ \alpha_m(V) &= 0.1 \frac{25 - V}{\exp\left(\frac{25 - V}{10}\right) - 1} \\ \beta_m(V) &= 4 \exp\left(\frac{-V}{18}\right) \\ \alpha_h(V) &= 0.07 \exp\left(\frac{-V}{20}\right) \\ \beta_h(V) &= \frac{1}{\exp\left(\frac{30 - V}{10}\right) + 1}.\end{aligned}\quad (6)$$

These equations, provided in the original paper [17], correspond to the membrane potential shifted by approximately 65 mV, so that the resting potential is at  $V \approx 0$  mV. The shifted Nernst equilibrium potentials are

$$E_K = -12 \text{ mV}, \quad E_{Na} = 120 \text{ mV}, \quad E_L = 10.6 \text{ mV}.$$

Typical values of maximal conductances are

$$g_K = 36 \text{ mS cm}^{-2}, \quad g_{Na} = 120 \text{ mS cm}^{-2}, \\ g_L = 0.3 \text{ mS cm}^{-2},$$

and  $C = 1 \mu\text{F cm}^{-2}$  is the membrane capacitance [17]. The second example in the next section is based on these parameter values. The functions  $\alpha(V)$  and  $\beta(V)$  describe the transition rates between open and closed states of the channels, respectively. We fix  $\alpha_0 = 0.01$  and  $\beta_0 = 0.125$  and estimate these parameters in figure 7. We again discretize these continuous equations to be of the form of (1) and (2). We use this model both to generate ‘experimental’ spike times, defined as the local large amplitude peaks in the voltage process, and as the state equations for the particle filter algorithm.

The last model we consider is a standard Hodgkin–Huxley model modified in three ways. First, the sodium activation variable ( $m$  in (5)) is replaced with its equilibrium value  $m_0 = \alpha_m(V)/(\alpha_m(V) + \beta_m(V))$ . Second, the potassium activation variable ( $n$  in (5)) and sodium inactivation variable ( $h$  in (5)) are replaced with a single variable  $w$ . These first two modifications are standard techniques to simplify the dynamics of the Hodgkin–Huxley model [13]. Third, an additional potassium current with slow dynamics (e.g., a muscarinic receptor suppressed potassium current or M-current [33]) is added. The modified model consists of three variables—voltage  $V$ , a fast current  $w$  and a slow current  $B$ —with dynamic equations:

$$C \dot{V} = I - g_K w^4 (V - E_K) - g_{Na} m_0^3 (1 - w) (V - E_{Na}) \\ - g_B B (V - E_K) - g_L (V - E_L) \\ \dot{w} = \alpha_w(V)(1 - w) - \beta_w(V)w \\ \dot{B} = \alpha_B(V)(1 - B) - \beta_B(V)B$$

where

$$\alpha_w(V) = 0.4\alpha_n(V + 65) \\ \beta_w(V) = 0.4\beta_n(V + 65) \\ \alpha_B(V) = \frac{0.0008}{1 + \exp((-V - 20)/5)} \\ \beta_B(V) = 0.0004 \exp((-V - 43)/18).$$

In this model we have scaled the voltage to match current neuroscience conventions (i.e. so that the neuronal rest voltage is near  $-70$  mV). The parameters are set so that  $E_K = -95$  mV,  $E_{Na} = 50$  mV,  $E_L = -70$  mV,  $g_{Na} = 100 \text{ mS cm}^{-2}$  and  $g_L = 0.25 \text{ mS cm}^{-2}$ . We vary the parameters  $I$ ,  $g_K$  and  $g_B$  to generate simulated data and estimate hidden variables and parameters in the model as described below. We again discretize these continuous equations to be of the form of (1) and (2).

### 2.3. Conditional spike intensity function

At each time step, the weight of each particle is updated based on its probability distribution of generating the observed spiking activity at that time. There are multiple approaches for constructing the probability model for the spiking activity. A simple approach of assigning a constant, non-zero weight only to those particles for which spiking occurs at exactly the observed spike time will lead to most particles being eliminated after each spike event. This approach would therefore require a huge number of particles to accurately estimate the model parameters. A much more computationally intensive approach would be to compute the probability of spiking based on the set of all possible realizations of the state process for each particle. Here, we seek a middle ground between these two approaches by modeling the conditional intensity of spiking as a smooth function of the membrane voltage trajectory  $V_t$ , which is one component of state vector  $\mathbf{S}_t$ . This serves to smooth out the likelihood function, allowing particles with values that could have plausibly led to the observed spiking activity to locate the peaks of the likelihood surface. Here, we chose the form of the conditional intensity function as follows:

$$\lambda[t|V_{1:t-1}, V_{t:t+k}] = \sum_{\tau=1}^{t+k} g(V_\tau) \cdot f(\tau - t) \quad (7)$$

$$g(x) = \eta \cdot \frac{e^{\nu(x-V_{th})}}{1 + e^{\nu(x-V_{th})}} \quad (8)$$

$$f(x) = \begin{cases} p^{-x} & x \leq 0 \\ q^x & x > 0. \end{cases} \quad (9)$$

The sigmoid function  $g(\cdot)$  reflects the effect of increasing voltage values to the spike intensity, in which parameter  $\nu$  controls the growth rate of the function and parameter  $\eta$  sets the upper limit;  $V_{th}$  is a voltage threshold, above which the probability of spiking rapidly increases. The function  $f(\cdot)$  is a weighting function, which measures the effect of the voltage at times away from  $t$ ; both  $p$  and  $q$  are constants and less than 1. Adjusting  $p$  and  $q$  tunes the weights of past and future function  $g(\cdot)$  values, respectively. This conditional intensity model defines a spiking probability function that increases smoothly when  $V$  passes the specified threshold. It serves to smooth the shape of the likelihood function upon which the weights of each particle are based. The parameter values control the degree of smoothing. Presumably, if there is little smoothing, then the likelihood will have narrow isolated peaks that require a huge number of particles to find, while if there is too much smoothing, very few particles are eliminated and the algorithm will converge much more slowly. In a separate simulation study (not shown), we found that the parameter estimates were robust across a wide range of choices of  $\eta$ ,  $\nu$ ,  $V_{th}$ ,  $p$  and  $q$ .

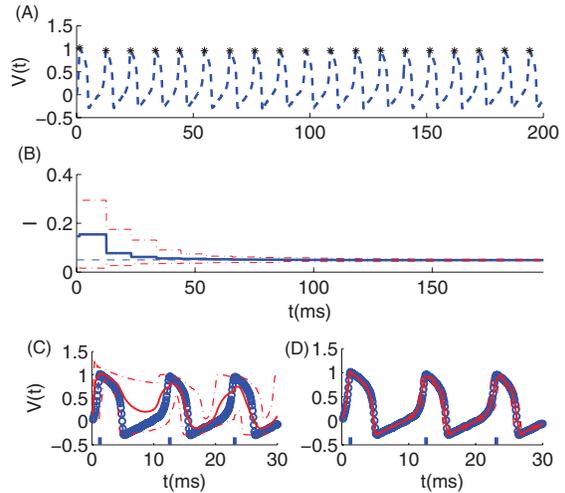
Note that the conditional intensity at time  $t$  depends on the value of the membrane voltage process,  $V$  through time  $t + k$ . This allows particles that are likely to produce spikes either just before or just after an observed spike to be retained. In practice, this means that we simulate the state process for each particle up to time  $t + k$  and compare the realizations to the spiking activity at time  $t$ .

### 3. Results

In this section we apply the estimation procedure to three types of biophysical models. In each case, we first simulate the model with a fixed set of parameter values and noise level to generate the spike time data. Then, given only the times of spike occurrence and assuming limited knowledge of some model variables or parameters, we attempt to estimate these variables and parameters. We evaluate the quality of the estimation algorithm in two ways. First, we compare the resulting parameter estimates to those used to simulate the ‘experimental’ data. Second, we compare the voltage traces that result from the model estimates to the (hidden) true voltage trace generated by the original simulation.

We begin by illustrating the application of the point process particle filter to the problem of estimating a single parameter in a simplified neuronal model. Specifically, we estimate the unknown resting current in the FitzHugh–Nagumo model. We simulate the model with parameter values:  $a = 0.1$ ,  $b = 0.01$ ,  $c = 0.02$  and noise level  $\sigma = 0.005$ . The conditional intensity function for the estimation procedure uses the following parameters:  $\Delta t = 0.1$  ms,  $\eta = 0.00329$ ,  $\nu = 30$ ,  $p = q = 0.9$ ,  $V_{th} = 0.8$ ,  $\rho = 0.96$ . A thousand particles were used in the sequential Monte Carlo algorithm. Figure 1 shows the result of the estimation procedure. The FitzHugh–Nagumo dynamics were simulated according to (3) and (4) with a constant resting current  $I = 0.05$ . A realization of the model voltage dynamics is plotted in figure 1(A), where the black asterisks denote the spike times. Our goal is to infer the (unknown) resting current  $I$  given only the simulated spike data, when all the other model parameters are known. We initialize the estimation algorithm with the prior distribution  $I \sim U(0, 0.3)$  and recursively generate the Monte Carlo approximation of the posterior distribution of  $I$ ,  $p(I|\Delta N_{1:T})$ . Then the expected current with respect to this posterior is used to estimate the true current value  $I = 0.05$ . Here, the state model is constructed so that the unknown value of the resting current is assumed to be constant. The estimation result is shown in figure 1(B). The expected resting current  $\hat{I}$  converges to the true resting current  $I$  quickly, and the 95% confidence bounds indicate increasing certainty about the estimate as time progresses. At the end of estimation procedure,  $\hat{I}$  approaches 0.049 and the confidence bounds approach (0.0459, 0.0526).

In addition to estimating the resting current  $I$ , we also approximate the time-varying unobserved state variables (in this case the voltage and recovery variable). In figure 1(C), we show the early tracking performance of the initial distribution of estimates. With only a limited number of spikes having occurred in the first 30 ms, the estimator behaves poorly and the confidence intervals are wide, indicating substantial uncertainty about the estimate. However, after spiking information over 200 ms is incorporated into the particle filter, the estimates converged to a very narrow range of parameter values (figure 1(B)). We then restarted the sequential Monte Carlo algorithm, fixing the parameter values to the converged estimates at the end of the estimation process, and tracking only the state variables. We show the tracking result in

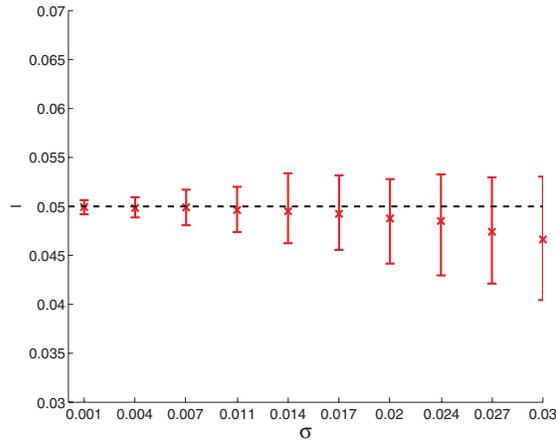


**Figure 1.** Estimating the resting current in the FitzHugh–Nagumo model from spike observations. (A) Simulated data from the FitzHugh–Nagumo model with a fixed resting current. Here we plot the continuous voltage trace (dashed blue curve) and indicate the spike times by black asterisks. (B) Given only the spike times in (A) we estimate the resting current (solid blue line). The two dot-dashed red lines indicate the 95% confidence interval of the estimation. All particles converge quickly and the confidence interval approaches a very narrow bound around the true resting current (dashed blue line) after only six spikes. (C) The initial estimates of voltages (solid red curve) and their 95% confidence interval (dot-dashed red curves) show deviation from the true voltages (blue circles) and large uncertainty during the first 30 ms. (D) The converged parameter estimates produce voltage traces with a mean (solid red curve) and 95% confidence interval (dot-dashed red curves) that approximates the true voltage trajectory (blue circles) well.

figure 1(D). The mean of the converged estimates matches the true voltage with high accuracy and the narrow confidence intervals indicate high confidence in these estimates. Thus, as long as our model describes the action potentials correctly, we can recover the full voltage information with high confidence by using only the times of the spikes. In the same way, the recovery variable ( $w$ ) is also tracked accurately (not shown).

We expect that increasing noise levels in the model dynamics will disrupt the parameter estimates. To explore this in the FitzHugh–Nagumo model, we illustrate in figure 2 the effect of increasing the noise level ( $\sigma$  is the standard deviation of the noise term,  $\epsilon_t$  in (3), for one time step) on the parameter estimates. When  $\sigma = 0.001$ , we have nearly perfectly regular spiking, while when  $\sigma = 0.03$ , spiking becomes highly irregular. As  $\sigma$  increases, the size of the confidence bounds about the parameter value tends to increase, but the bounds still tend to contain the true value (figure 2). These results demonstrate, for the FitzHugh–Nagumo model, the robustness of the estimation algorithm to increased noise.

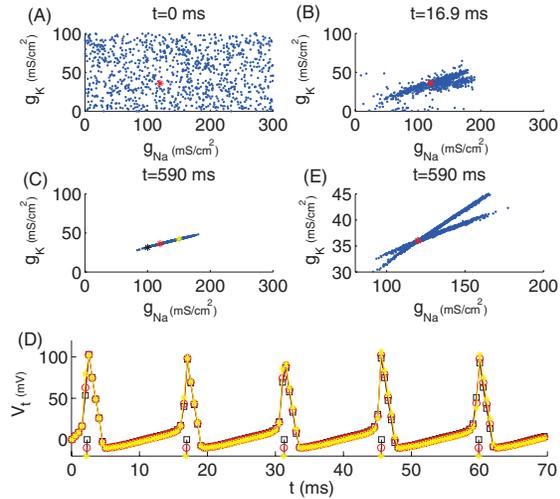
In the second example, we consider the problem of inferring multiple parameters simultaneously in a more advanced and physiologically realistic neuronal model.



**Figure 2.** Robustness of estimation to noise level. (A) Given simulated spike trains of the same length from the FitzHugh–Nagumo model with different levels of noise  $\sigma$ , we estimate the input current  $I$  whose true value is 0.05 (indicated by the dashed black line). The red stars represent the estimates of  $I$  and red bars are their corresponding 95% confidence intervals. The true current values lie within the confidence bounds of every estimate.

Specifically, we estimate the sodium and potassium conductance in the standard Hodgkin–Huxley equations. We begin by simulating the Hodgkin–Huxley model for a fixed value of input current (parameter  $I = 10 \mu\text{A cm}^{-2}$ ) and record the resulting spike times. Then, given only these spike times, we estimate two conductances  $g_K$  (true value  $36 \text{ mS cm}^{-2}$ ) and  $g_{Na}$  (true value  $120 \text{ mS cm}^{-2}$ ) with all other parameters fixed. The conditional intensity function for the estimation procedure uses the following parameters:  $\Delta t = 0.05 \text{ ms}$ ,  $\eta = 1.622$ ,  $\nu = 0.1$ ,  $p = q = 0.9$ ,  $V_{th} = 80 \text{ mV}$ . The sequential Monte Carlo algorithm used 10 000 particles.

Figures 3(A)–(C) show the temporal evolution of the distribution of parameter estimates from the point process particle filter. Initially, the particles are uniformly distributed in the region  $g_K \in (0, 100) \text{ mS cm}^{-2}$ ,  $g_{Na} \in (0, 300) \text{ mS cm}^{-2}$  (figure 3(A)). The distribution of estimates evolves in time, and by the second observed spike, the distribution of parameter estimates has narrowed to cover a much smaller region of the parameter space (figure 3(B)). After 590 ms and approximately 40 spikes, the distribution of the parameter estimates has not converged to a single point, but has stabilized to a narrow line segment in the parameter space that includes the true parameter values (figure 3(C)). These results suggest that multiple combinations of conductance values for  $g_K$  and  $g_{Na}$  can produce dynamics that are consistent with the observed spiking activity. To illustrate this, we simulate the Hodgkin–Huxley model for three different sets of parameter values that are contained in the linear subspace to which the estimates converged. The spike times produced by these different parameter sets are nearly identical and consistent with those produced by the true parameter values (figure 3(D)). It is not surprising that multiple sets of parameter values can produce similar neuronal dynamics [11, 15, 25, 30, 31, 36]. In this case,



**Figure 3.** Estimating conductance parameters in a Hodgkin–Huxley model. (A)–(C) Sequential parameter estimates for  $g_K$  and  $g_{Na}$ . The red asterisk denotes the true values of  $g_K$  and  $g_{Na}$ . The blue dots denote the parameter values for all of the particles. (A) The initial particle estimates are uniformly distributed in the two-dimensional parameter space. (B) Distribution of particles after the second observed spike. The parameter values of the particles begin to concentrate in a region that contains the true values of  $g_K$  and  $g_{Na}$ . (C) Distribution of particles after 40 spikes. The parameter estimates have converged to a narrow linear subspace of parameter values that are consistent with the spike data. The three asterisks indicate parameter values used in (D). (D) Three voltage traces (blue, red and yellow) corresponding to the three parameter choices (blue, red and yellow, respectively) indicated in (C). Although the parameter values differ, the three voltage traces and resulting spike times (plotted as colored symbols along the horizontal axis) are nearly indistinguishable. (E) By collecting data from a second experiment with an altered resting current, we obtain a new set of estimates that intersect with the first set around a single point, allowing us to compute a single accurate estimate of the true parameters.

the linear relationship between the parameters is consistent with our biophysical understanding of the neuron; the increase of outward current due to higher potassium conductance is approximately balanced by the increase of inward current with larger sodium conductance. The particle filter identifies the full space of parameter values that could have produced the observed data. Intuitively, this result suggests that the likelihood surface contains a flat ridge to which the estimates congregate.

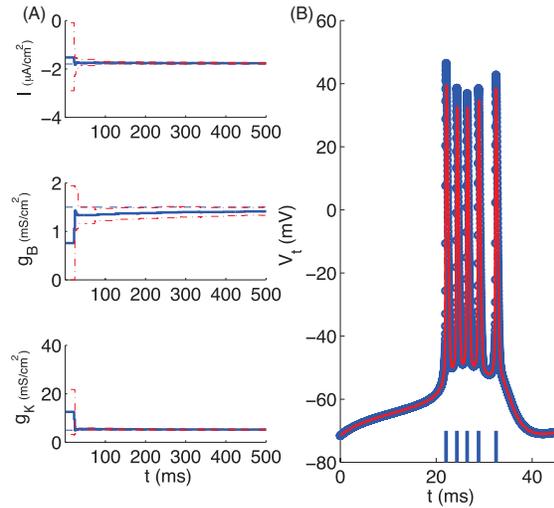
In this case, the true model parameters are not identifiable, because the recorded data cannot distinguish between certain parameter values. However, consider an experiment with the goal of determining a neuron’s true conductance values by injecting a fixed current and measuring the spike activity. Although we would not be able to uniquely identify the true conductance values from these data, this estimation framework suggests experimental manipulations that could help more accurately identify these true values. For example, figure 3(E) illustrates how the distribution of parameter estimates would change if we increased the resting current to  $I = 30 \mu\text{A cm}^{-2}$ . To determine these estimates we generate

new spike times from the Hodgkin–Huxley model with the adjusted resting current. From these new spike times, we estimate the conductances  $g_K$  and  $g_{Na}$  and again find values that lie along a line in the two-dimensional parameter space. However, the slope of this line depends upon the choice of resting current  $I$ . Visual inspection reveals that the two sets of parameter estimates, based on the data from two distinct resting currents, intersect near the true parameter values (at  $g_K = 36 \text{ mS cm}^{-2}$  and  $g_{Na} = 120 \text{ mS cm}^{-2}$ ). These results suggest a more sophisticated procedure for identifying specific parameters; by adjusting parameters that can be controlled, an estimation procedure combining multiple parameter values could be developed so that the likelihood has unique, easily identifiable peaks. For example, *in vitro* recordings from single neurons typically allow the experimenter to control the resting current [14]. Therefore, injecting different current values is experimentally feasible and may provide a protocol for accurate estimation of conductance values from *in vitro* neuronal spike train data.

The presence or absence of specific ionic currents may lead to different patterns of neuronal activity such as regular spiking or bursting [21, 30]. During regular spiking, a neuron generates action potentials at a constant (or nearly constant) rate, while during bursting only brief intervals of rapid spiking appear. Bursting activity typically requires a separation of time scales, in which a ‘slow subsystem’ represents an intrinsic current that, for example, gradually increases during the active phase of a burst. Eventually, the slow current increases enough to end the active phase of the burst and prevent additional spikes until this current depletes. Therefore, the slow subsystem is critical for setting the timescale and intensity of the bursts.

To simulate bursting activity here, we utilize a (modified) Hodgkin–Huxley model that includes a slow intrinsic current (see section 2). Without this slow current the model generates regular spiking activity, and the effect of the slow current is to modulate this regular spiking activity. We apply the estimation procedure to spike time data generated from this model and show that we can simultaneously approximate three model parameters—the resting current (denoted by  $I$ ), the potassium conductance (denoted by  $g_K$ ) and the conductance of the slow current (denoted by  $g_B$ )—during both regular spiking and bursting activity. We then conclude with an example of model misspecification to show that an inappropriate model choice causes the estimation procedure to fail.

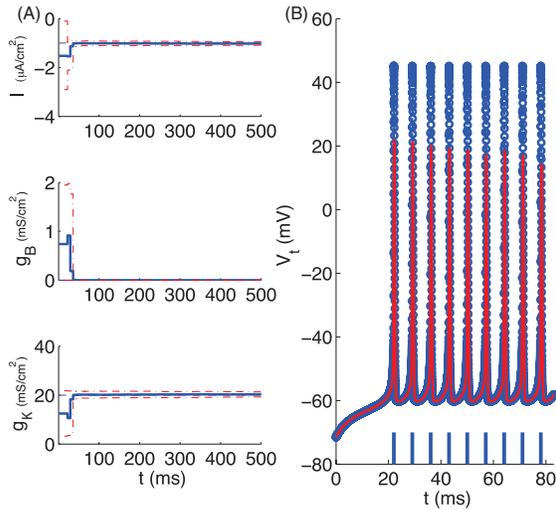
We simulated both regular spiking and bursting under this model and used the resulting spike times as the ‘experimental’ data for the point process particle filter. We initialized the algorithm with a prior distribution of estimates that was uninformative about which mode of spiking (regular or bursting) the neuron would produce. The initial particles were drawn with probability 0.5 from a distribution consistent with regular firing ( $I \sim U(-2, 0) \mu\text{A cm}^{-2}$ ,  $g_B = 0 \text{ mS cm}^{-2}$  (since there is no slow current in the regular spiking model),  $g_{Na} \sim U(18, 22) \text{ mS cm}^{-2}$ ) and with probability 0.5 from a distribution consistent with bursting ( $I \sim U(-3, -1) \mu\text{A cm}^{-2}$ ,  $g_B \sim U(1, 2) \text{ mS cm}^{-2}$ ,  $g_{Na} \sim U(3, 7) \text{ mS cm}^{-2}$ ). These prior distributions were



**Figure 4.** Estimating three parameters in a bursting model. (A) Simultaneous estimates of three model parameters,  $I$ ,  $g_B$  and  $g_K$ , over time. The dashed blue lines indicate the underlying parameter values ( $I = -1.8 \mu\text{A cm}^{-2}$ ,  $g_B = 1.5 \text{ mS cm}^{-2}$  and  $g_K = 5 \text{ mS cm}^{-2}$ ). The solid blue lines are the parameter estimates. The red dot-dashed lines represent 95% confidence intervals about the estimates. All three estimates converge to their true values quickly and consistently. (B) Based on the converged particles taken at the end of the estimation process, the expected voltage trace and estimated spike times (red curve) match the underlying voltage trace (thick blue circles) and spike times (blue bars). The conditional intensity function for the estimation procedure uses the following parameters:  $\Delta t = 0.01 \text{ ms}$ ,  $\eta = 13.8$ ,  $\nu = 0.1$ ,  $p = q = 0.9$ ,  $V_{th} = 20 \text{ mV}$ ,  $\rho = 0.96$ . The sequential Monte Carlo algorithm used 1000 particles.

chosen based on an existing understanding of the model dynamics, which helps limit the range of possible particle values. Exploring a broader prior range with the same density of coverage requires more particles, and therefore more computation.

With these general prior distributions capable of producing both regular spiking and bursting activity, we apply the estimation procedure to spike time data simulated from the modified Hodgkin–Huxley model. Figures 4(A) and 5(A) illustrate the estimated expected parameter values in comparison with the true parameter values for the bursting and regular spiking models, respectively. We find that, in both cases, all three estimates converge to the corresponding true parameter values quickly, with narrow confidence intervals, suggesting high certainty in the estimates. The particle filter successfully isolates the correct model class, and moreover accurately identifies the true parameter values in each case. To illustrate the accuracy of these estimates, we recompute the expected voltages and show the tracking result based on the converged particles. For both model classes, the underlying voltages are tracked accurately as are the spike times (figures 4(B) and 5(B)). We note that the spike heights in the mean estimated voltage traces remain below the true spike heights; this is expected because the observed data

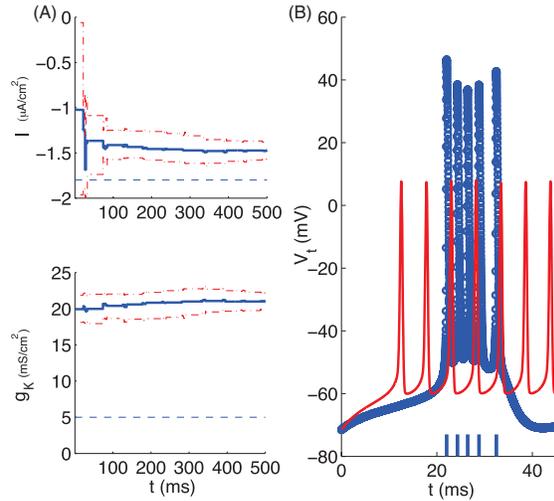


**Figure 5.** Estimating three parameters in a regular spiking model. (A) Simultaneous estimates of three model parameters  $I$ ,  $g_B$  and  $g_K$  over time. Line colors and symbols match those in figure 4(A). The underlying parameter values are  $I = -1 \mu\text{A cm}^{-2}$ ,  $g_B = 0 \text{ mS cm}^{-2}$  and  $g_K = 20 \text{ mS cm}^{-2}$ . All three estimates converge to the underlying values quickly and consistently, as in the case of the bursting model. (B) Expected voltage trace (red curve) based on the converged parameter estimates. The underlying voltage trace (thick blue circles) and spike times are tracked with high accuracy. The conditional intensity function for the estimation procedure uses the following parameters:  $\Delta t = 0.01 \text{ ms}$ ,  $\eta = 4.5$ ,  $\nu = 0.1$ ,  $p = q = 0.9$ ,  $V_{\text{th}} = 20 \text{ mV}$ ,  $\rho = 0.96$ . The sequential Monte Carlo algorithm used 1000 particles.

only contains information about the spike times, not the spike shapes nor the stochastic nature of the model. All of the gating variables are estimated accurately as well (not shown). We note that the estimate of parameter  $g_B$  goes to zero for the regular spiking data (figure 5(A)). The estimation procedure therefore reveals that this current is not necessary to capture the dynamics of the observed spike times. This is consistent with our biophysical intuition: setting  $g_B = 0 \text{ mS cm}^{-2}$  turns off the slow current, preventing bursting activity and promoting regular spiking activity.

The parameter estimates for both the bursting and regular spiking models approach unique points in the three-dimensional parameter space. For these estimates, the model identification issue seen in the previous example (figure 3) does not arise. In the previous example, the similarity of the spike times generated for multiple parameter combinations results from the complementary relationship between the sodium  $g_{\text{Na}}$  and potassium  $g_K$  currents. The same kind of similarity of the spike times cannot be induced by adjusting parameters in the modified Hodgkin–Huxley model considered here.

One possible concern about the particle filter approach is that, since the algorithm uses the observed spiking data, the resulting estimates will always produce activity consistent with the data. To explore this issue, we consider an example of model misspecification. First, we simulate the bursting model described above to generate spike time data. Then, we



**Figure 6.** Model misspecification results in an inaccurate representation of the data. (A) Estimates of parameters  $I$  and  $g_K$  under the regular firing model using bursting data. The parameter estimates (solid blue lines) fail to converge to the underlying parameter values (dashed blue line). (B) The expected voltage trace (red line) based on the converged particles does not match the underlying trace (dashed blue line).

attempt to estimate parameters and hidden variables for the regular spiking model (i.e. the model with no slow current component) given the bursting spike time data. In this case, we have intentionally chosen the wrong model class for the data. Since the slow current conductance is zero (i.e.  $g_B = 0 \text{ mS cm}^{-2}$ ) in the regular spiking model, only two parameters are estimated, the resting current  $I$  and potassium conductance  $g_K$ . Figure 6(A) presents the estimates of these two parameters. The estimates do not converge to a single set of values. Instead, the population estimates continue to have a large variability as illustrated by the size of the confidence intervals. One might suspect that this result is due to a model identifiability issue, but figure 6(B) shows that the estimated voltage trace deviates substantially from the true voltage trace. Clearly the estimates are very poor. This result suggests that if we start with the wrong model class, it is impossible to find a set of parameters that accurately explains the observed spike data. The modeler must therefore choose an appropriate model class (e.g., a model with the appropriate intrinsic currents) to produce an accurate representation of the data.

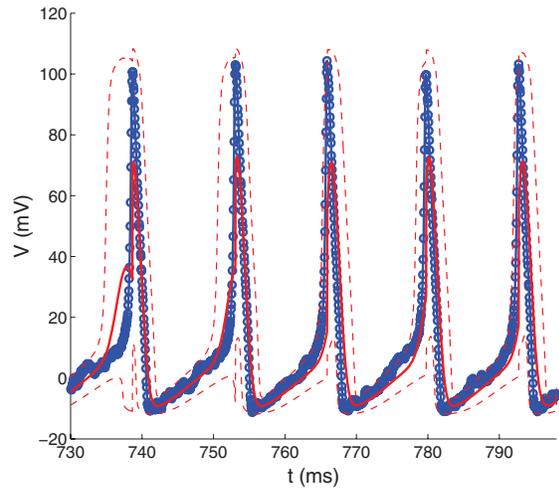
#### 4. Discussion

In neuroscience, biophysical models of neural spiking are often fitted using hand-tuning procedures to identify model parameters that are qualitatively consistent with particular patterns of spiking activity. While this approach has been useful for constructing models that capture general features of spiking data, it is more difficult to identify parameters that could produce a specific set of electrophysiological recordings. Recent work has focused on developing methods to estimate

biophysical model parameters from recorded data [29, 41], but these approaches typically require detailed knowledge of the voltage dynamics (e.g., the intracellular voltage trace), which are often unobserved. In this paper, instead of full information of the voltage dynamics, we utilize only spike times to estimate parameters and hidden variables (including the sub-threshold voltage) in biophysical neuronal models. The proposed estimation method is an adapted sequential Monte Carlo method which incorporates point process theory to deal with the discrete event spike data. This sequential Monte Carlo approach efficiently identifies parameter values that are most compatible with the observed spiking activity, by allowing each particle to explore its local likelihood surface and only replicating those particles where the likelihood remains high. Similarly, the algorithm is robust to noise in the state model because particles whose dynamic variables start to deviate from the true realization eventually are eliminated and replaced by others that follow the true dynamics more closely.

In general, it is unclear whether spike time data provide sufficient information to estimate accurately any set of parameters in a biophysical model. For some simple models, such as the FitzHugh–Nagumo model and the Hodgkin–Huxley model with slow current, the estimates we developed rapidly converged to the correct values. However, for the problem of simultaneously estimating the sodium and potassium conductances in the standard Hodgkin–Huxley model, we encountered a model identification issue: instead of identifying the true parameter values, our approach provided a group of parameter estimates all consistent with the observed spike time data. When model identification issues arise we expect that the variance of the resulting estimates for individual parameters may remain large, and the covariance between parameters may be far from 0. Therefore, the structure of the estimated posterior distribution should be investigated for all parameters simultaneously. In terms of understanding the model dynamics, this result is more informative than the traditional hand-tuning approach that typically yields only a single set of parameter values (e.g., the first set encountered that are consistent with the data) and ignores other parameter sets also compatible with the data. Therefore, in terms of data fitting, the method proposed here can obtain more general results beyond a single set of model parameters. However, this does not mean that the true parameter values are unobtainable. For the standard Hodgkin–Huxley model, the estimation procedure suggested an experimental manipulation that allowed us to estimate the true conductance values by injecting different levels of fixed resting current. In the examples considered here, we estimated a constant, unknown resting current. Often, we are interested in exploring models with time-varying currents. If these currents are known (i.e. controlled by the experimenter), then we can easily incorporate them into the state model. If we have a biophysical model of the time-varying inputs, we can similarly incorporate this model into the state model, although the complexity of the model and estimation procedure will increase.

There are several limitations of the approach proposed here. First of all, our approach requires pre-knowledge of the underlying equations for the neuronal model. Of course, this



**Figure 7.** Estimation of six parameters from the Hodgkin–Huxley model. The expected voltage trace and its 95% confidence interval when estimating six parameters ( $g_{Na}$ ,  $g_K$ ,  $g_L$ ,  $\alpha_0$ ,  $\beta_0$ ,  $I$ ) simultaneously in the Hodgkin–Huxley model with noise  $\sigma = 10$ . The red solid and dashed lines represent voltage estimates and their confidence bounds respectively over time. The blue circle line indicates the underlying true voltage trace.

works well for the simulated data considered here in which we know exactly the underlying dynamic model. But for real spiking data recorded in experiment, we do not know exactly (or usually even approximately) the latent model. One possible solution is to apply the estimation procedure to multiple model types and then choose the best one as the underlying model based on maximum likelihood theory. How well these techniques apply to experimental observations of spike times, and whether these techniques outperform existing methods [14], will be the subject of future work. Secondly, in this paper we have not discussed any of the multiple quantitative goodness-of-fit methods to characterize model misspecification that are present in the literature [2, 12, 44]. Point process theory provides a number of approaches for measuring goodness-of-fit of a spiking model, including methods based on penalized likelihoods [40] and methods based on time rescaling [2]; we will explore these procedures in future work.

Thirdly, the model estimation problem we considered is simplified in that only a small number of parameters were estimated (along with the set of dynamic variables). Detailed biophysical models typically require many parameters, and as the number of parameters to estimate increases, so does the dimensionality of the state space and the computational burden for the estimation algorithm. The question of how these methods will scale to large, biophysically realistic, estimation problems is an important area for future work. To begin to address this issue here, we examined another simulation for a Hodgkin–Huxley model, now with six unknown parameters, including the leak current conductance ( $g_L$  in (5)) and two parameters controlling the potassium gating variable dynamics ( $\alpha_0$  and  $\beta_0$  in (6)). Figure 7 illustrates the resulting fit for the

expected voltage trace for this model. As with the three-parameter model estimation problem, the final distribution of particles converges to a narrow subspace of parameter values (not shown), all of which lead to voltage dynamics that are consistent with the observed spike data. As the number of parameters to be estimated increases, issues related to model identification are more likely to arise, and computational techniques implementing parallelization will probably become critical.

Lastly, our proposed conditional intensity model is not derived purely from biophysics. Instead, we consider the state process to be a possible surrogate for the unobserved neural dynamics, which provides information about the probability of observing a spike at any moment in time rather than the exact time of a spike. However, we can develop our method directly from biophysical models. To do so requires treating  $\mathbf{S}_t$  as the exact dynamics of the observed neuron and computing the probability distribution of simulated spikes based on model simulations, which is computationally expensive. For the procedure proposed here, we sought a balance between computational time and accurate estimates.

To summarize, provided only the spike times of neuronal activity we constructed an adapted sequential Monte Carlo procedure which incorporates these sparse experimental data with biophysical models. Unlike the past hand-tuning approach and previous methods requiring full knowledge of the voltage dynamics, our method estimates model parameters directly from the spike time data.

## Acknowledgments

The authors would like to thank Hillel Chiel and Peter Thomas for organizing the special issue on Applied Dynamics: From Neural Dynamics to Neural Engineering. MAK holds a Career Award at the Scientific Interface from the Burroughs Wellcome Fund. This work was partially supported by grant IIS.0643993 from the National Science Foundation to UTE, and by Award Number R01NS072023 from the National Institute Of Neurological Disorders And Stroke to UTE and MAK.

## Appendix. The pseudo-code description of estimation algorithm

Select an initial probability distribution for  $\mathbf{U}$  denoted by  $p(\mathbf{U}_0)$ . We assume that the estimation procedure starts with the neuron at rest, and initialize the state variables to their resting values (steady-state values assuming a resting current of zero). If no such resting period is available, we might choose to initialize the algorithm at a time immediately following a spike, where the state variables typically follow some known stereotypical trajectories.

### (1) Initialization:

Set  $t = 0$  and for  $i = 1, \dots, n$  particles, draw the initial states and parameters from the  $p(\mathbf{U}_0)$  and set  $w_0^i = n^{-1}$  for all  $i$ . Set  $t = 1$ .

### (2) Importance sampling:

Using the prior distribution  $p(\mathbf{S}_t | \mathbf{S}_{t-1}, \Theta_t)$  as the proposal distribution  $\pi_t$ . Update all of the states based on (1) sample  $\mathbf{S}_t$  from  $\pi_t$ . Evaluate the importance weight of the  $i$ th particle

$$w_t^i = w_{t-1}^i \cdot p(\Delta N_t | \mathbf{S}_{1:t+k}^i),$$

where  $p(\Delta N_t | \mathbf{S}_{1:t+k}^i)$  is computed by (2). Then normalize the importance weights

$$\tilde{w}_t^i = \frac{w_t^i}{\sum_j w_t^j}.$$

### (3) Resampling:

Resampling can be performed at any fixed interval. In this case, we resample after each spike observation. Draw  $n$  particles  $\{\tilde{\mathbf{S}}_t^i : i = 1, \dots, n\}$  from  $\{\mathbf{S}_t^i : i = 1, \dots, n\}$  using the residual resampling approach (see text). Reset the weights to  $\tilde{w}_t^i = n^{-1}$  to obtain the Monte Carlo estimate of the probability density

$$p(\tilde{\mathbf{S}}_t | \Delta N_{1:t}) \approx n^{-1} \sum_{i=1}^n \delta(\tilde{\mathbf{S}}_t - \tilde{\mathbf{S}}_t^i),$$

where  $\delta(\cdot)$  is the Dirac delta function, indicating a point mass at 0.

### (4) Sample parameters:

Draw a new parameter vector  $\tilde{\Theta}_t^i$  from the  $i$ th normal component of the kernel density, namely

$$\tilde{\Theta}_t^i \sim N(\cdot | m_{t-1}^i, h^2 \Sigma_{t-1}),$$

where  $m_{t-1}^i = \rho \cdot \Theta_{t-1}^i + (1 - \rho) \cdot \bar{\Theta}_{t-1}$  and  $\Sigma_{t-1} = \hat{\text{Cov}}(\{\Theta_t^i\}_i)$ ,  $h^2 = 1 - \rho^2$ .  $\rho$  is a constant, called discount factor and  $\bar{\Theta}_{t-1}$  denotes the weighted average of all the parameters over particles.  $\rho$  is often selected to be between 0.960 and 0.999 based on standard practice [24]. Sampling the parameter values in this manner ensures that the variance of the estimate cannot increase as we observe more data. Thus, it will increase the convergence rate in the sense of the second moment.

### (5) Compute any statistics $g_t(\tilde{\mathbf{U}}_t)$ of interest based on approximated posterior distribution and repeat steps 2–5

$$E(g_t(\tilde{\mathbf{U}}_t)) \approx \sum_{i=1}^n \tilde{w}_t^i \cdot g_t(\tilde{\mathbf{U}}_t^i).$$

## References

- [1] Bhalla U S and Bower J M 1993 Exploring parameter space in detailed single neuron models: simulations of the mitral and granule cells of the olfactory bulb *J. Neurophysiol.* **69** 1948–65
- [2] Brown E N, Barbieri R, Ventura V, Kass R E and Frank L M 2002 The time-rescaling theorem and its application to neural spike train data analysis *Neural Comput.* **14** 325–46
- [3] Channell P, Cymbalyuk G and Shilnikov A 2007 Origin of bursting through homoclinic spike adding in a neuron model *Phys. Rev. Lett.* **98** 134101
- [4] Deco G, Jirsa V, Robinson P, Breakspear M and Friston K 2008 The dynamic brain: from spiking neurons to neural masses and cortical fields *PLoS Comput. Biol.* **4** e1000092