


Stochastic representations of ion channel kinetics and exact stochastic simulation of neuronal dynamics

David F. Anderson, Bard Ermentrout & Peter J. Thomas

Journal of Computational Neuroscience

ISSN 0929-5313
Volume 38
Number 1

J Comput Neurosci (2015) 38:67-82
DOI 10.1007/s10827-014-0528-2



Your article is protected by copyright and all rights are held exclusively by Springer Science +Business Media New York. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".

Stochastic representations of ion channel kinetics and exact stochastic simulation of neuronal dynamics

David F. Anderson · Bard Ermentrout · Peter J. Thomas

Received: 11 February 2014 / Revised: 18 August 2014 / Accepted: 3 September 2014 / Published online: 19 November 2014
© Springer Science+Business Media New York 2014

Abstract In this paper we provide two representations for stochastic ion channel kinetics, and compare the performance of exact simulation with a commonly used numerical approximation strategy. The first representation we present is a random time change representation, popularized by Thomas Kurtz, with the second being analogous to a “Gillespie” representation. Exact stochastic algorithms are provided for the different representations, which are preferable to either (a) fixed time step or (b) piecewise constant propensity algorithms, which still appear in the literature. As examples, we provide versions of the exact algorithms for the Morris-Lecar conductance based model, and detail the error induced, both in a weak and a strong sense, by the use of approximate algorithms on this model. We include ready-to-use implementations of the random time change algorithm in both XPP and Matlab. Finally, through the consideration of parametric sensitivity analysis, we show

how the representations presented here are useful in the development of further computational methods. The general representations and simulation strategies provided here are known in other parts of the sciences, but less so in the present setting.

Keywords Markov process · Conductance based model · Exact stochastic simulation · Morris-Lecar model

1 Introduction

Fluctuations in membrane potential arise in part due to stochastic switching in voltage-gated ion channel populations (Dorval Jr. and White 2005; Laing and Lord 2010; White et al. 2000). We consider a stochastic modeling, i.e. master equation, framework (Anderson and Kurtz 2011; Colquhoun and Hawkes 1983; Earnshaw and Keener 2010; Lee and Othmer 2010; Wilkinson 2011) for neural dynamics, with noise arising through the molecular fluctuations of ion channel states. We consider model nerve cells that may be represented by a single isopotential volume surrounded by a membrane with capacitance $C > 0$. Mathematically, these are hybrid stochastic models which include components, for example the voltage, that are continuous and piecewise differentiable and components, for example the number of open potassium channels, that make discrete transitions or jumps. These components are coupled because the parameters of the ODE for the voltage depend upon the number of open channels, and the propensity for the opening and closing of the channels depends explicitly upon the time-varying voltage.

These hybrid stochastic models are typically described in the neuroscience literature by providing an ODE governing the absolutely continuous portion of the system, which

Action Editor: David Terman

Electronic supplementary material The online version of this article (doi:10.1007/s10827-014-0528-2) contains supplementary material, which is available to authorized users.

D. F. Anderson
Department of Mathematics, University of Wisconsin,
Madison, WI, USA
e-mail: anderson@math.wisc.edu

B. Ermentrout
Department of Mathematics, University of Pittsburgh,
Pittsburgh, PA, USA
e-mail: bard@pitt.edu

P. J. Thomas (✉)
Department of Mathematics, Applied Mathematics, and Statistics,
Case Western Reserve University, Cleveland, OH, USA
e-mail: pjthomas@case.edu

is valid between jumps of the discrete components, and a chemical master equation providing the dynamics of the probability distribution of the jump portion, which itself depends upon the solution to the ODE. These models are piecewise-deterministic Markov processes (PDMP) and one can therefore also characterize them by providing (i) the ODE for the absolutely continuous portion of the system and (ii) both a rate function that determines when the next jump of the process occurs and a transition measure determining which type of jump occurs at that time (Davis 1984). Recent works making use of the PDMP formalism has led to limit theorems (Pakdaman et al. 2010; 2012), dimension reduction schemes (Wainrib et al. 2012), and extensions of the models to the spatial domain (Buckwar and Riedler 2011; Riedler et al. 2012).

In this paper, we take a different perspective. We will introduce here two pathwise stochastic representations for these models that are similar to Itô SDEs or Langevin models. The difference between the models presented here and Langevin models is that here the noise arises via stochastic counting processes as opposed to Brownian motions. These representations give a different type of insight into the models than master equation representations do, and, in particular, they imply different exact simulation strategies. These strategies are well known in some parts of the sciences, but less well known in the current context (Anderson and Kurtz 2011).¹

From a computational standpoint the change in perspective from the master equation to pathwise representations is useful for a number of reasons. First, the different representations naturally imply different exact simulation strategies. Second, and perhaps more importantly, the different representations themselves can be utilized to develop new, highly efficient, computational methods such as finite difference methods for the approximation of parametric sensitivities, and multi-level Monte Carlo methods for the approximation of expectations (Anderson 2012; Anderson and Higham 2012; Anderson et al. 2014). Third, the representations can be utilized for the rigorous analysis of different computational strategies and for the algorithmic reduction of models with multiple scales (Anderson et al. 2011; Anderson and Koyama 2014; Ball et al. 2006; Kang and Kurtz 2013).

¹For an early statement of an exact algorithm for the hybrid case in a neuroscience context see ((Clay and DeFelice 1983), Equations (2)–(3)). Strassberg and DeFelice further investigated circumstances under which it is possible for random microscopic events (single ion channel state transitions) to generate random macroscopic events (action potentials) (Strassberg and DeFelice 1993) using an exact simulation algorithm. Bressloff, Keener, and Newby used an exact algorithm in a recent study of channel noise dependent action potential generation in the Morris-Lecar model (Keener and Newby 2011; Newby et al. 2013). For a detailed introduction to stochastic ion channel models, see (Groff et al. 2009; Smith and Keizer 2002).

We note that the types of representations and simulation strategies highlighted here are known in other branches of the sciences, especially operations research and queueing theory (Glynn 1989; Haas 2002), and stochastically modeled biochemical processes (Anderson 2007; Anderson and Kurtz 2011). See also (Riedler and Notarangelo 2013) for a treatment of such stochastic hybrid systems and some corresponding approximation algorithms. However, the representations are not as well known in the context of computational neuroscience and as a result approximate methods for the simulation of sample paths including (a) fixed time step methods, or (b) piecewise constant propensity algorithms, are still utilized in the literature in situations where there is no need to make such approximations. Thus, the main contributions of this paper are: (i) the formulation of the two pathwise representations for the specific models under consideration, (ii) a presentation of the corresponding exact simulation strategies for the different representations, and (iii) a comparison of the error induced by utilizing an approximate simulation strategy on the Morris-Lecar model. Moreover, we show how to utilize the different representations in the development of methods for parametric sensitivity analysis.

The outline of the paper is as follows. In Section 2 we heuristically develop two distinct representations and provide the corresponding numerical methods. In Section 3, we present the random time change representation as introduced in Section 2 for a particular conductance based model, the planar Morris-Lecar model, with a single ion channel treated stochastically. Here, we illustrate the corresponding numerical strategy on this example and provide in the appendix both XPP and Matlab code for its implementation. In Section 4, we present an example of a conductance based model with more than one stochastic gating variable, namely the Morris-Lecar model with both ion channel types (calcium channels and potassium channels) treated stochastically. To illustrate both the strong and weak divergence between the exact and approximate algorithms, in Section 5 we compare trajectories and histograms generated by the exact algorithms and the piecewise constant approximate algorithm. In Section 6, we show how to utilize the different representations presented here in the development of methods for parametric sensitivity analysis, which is a powerful tool for determining parameters to which a system output is most responsive. In Section 7, we provide conclusions and discuss avenues for future research.

2 Two stochastic representations

The purpose of this section is to heuristically present two pathwise representations for the relevant models. In

Section 2.1 we present the random time change representation. In Section 2.2 we present a ‘‘Gillespie’’ representation, which is analogous to the PDMP formalism discussed in the introduction. In each subsection, we provide the corresponding numerical simulation strategy. See (Anderson and Kurtz 2011) for a development of the representations in the biochemical setting and see (Eithier and Kurtz 1986; Kurtz 1980, 1981) for a rigorous mathematical development.

2.1 Random time change representations

We begin with an example. Consider a model of a system that can be in one of two states, A or B , which represent a ‘‘closed’’ and an ‘‘open’’ ion channel, respectively. We model the dynamics of the system by assuming that the dwell times in states A and B are determined by independent exponential random variables with parameters $\alpha > 0$ and $\beta > 0$, respectively. A graphical representation for this model is



The usual formulation of the stochastic version of this model proceeds by assuming that the probability that a closed channel opens in the next increment of time Δs is $\alpha\Delta s + o(\Delta s)$, whereas the probability that an open channel closes is $\beta\Delta s + o(\Delta s)$. This type of stochastic model is often described mathematically by providing the ‘‘chemical master equation,’’ which for Eq. (1) is simply

$$\begin{aligned} \frac{d}{dt} p_{x_0}(A, t) &= -\alpha p_{x_0}(A, t) + \beta p_{x_0}(B, t) \\ \frac{d}{dt} p_{x_0}(B, t) &= -\beta p_{x_0}(B, t) + \alpha p_{x_0}(A, t), \end{aligned}$$


where $p_{x_0}(x, t)$ is the probability of being in state $x \in \{A, B\}$ at time t given an initial condition of state x_0 . Note that the chemical master equation is a linear ODE governing the dynamical behavior of the *probability distribution* of the model, and does not provide a stochastic representation for a particular *realization* of the process.

In order to construct such a pathwise representation, let $R_1(t)$ be the number of times the transition $A \rightarrow B$ has taken place by time t and, similarly, let $R_2(t)$ be the number of times the transition $B \rightarrow A$ has taken place by time t . We let $X_1(t) \in \{0, 1\}$ be one if the channel is closed at time t , and zero otherwise, and let $X_2(t) = 1 - X_1(t)$ take the value one if and only if the channel is open at time t . Then, denoting $X(t) = (X_1(t), X_2(t))^T$, we have

$$X(t) = X(0) + R_1(t) \begin{pmatrix} -1 \\ 1 \end{pmatrix} + R_2(t) \begin{pmatrix} 1 \\ -1 \end{pmatrix}.$$

We now consider how to represent the counting processes R_1, R_2 in a useful fashion, and we do so with unit-rate Poisson processes as our mathematical building blocks. We

recall that a unit-rate Poisson process can be constructed in the following manner (Anderson and Kurtz 2011). Let $\{e_i\}_{i=1}^\infty$ be independent exponential random variables with a parameter of one. Then, let $\tau_1 = e_1, \tau_2 = \tau_1 + e_2, \dots, \tau_n = \tau_{n-1} + e_n$, etc. The associated unit-rate Poisson process, $Y(s)$, is simply the counting process determined by the number of points $\{\tau_i\}_{i=1}^\infty$, that come before $s \geq 0$. For example, if we let ‘‘x’’ denote the points τ_n in the image below



then $Y(s) = 6$.

Let $\lambda : [0, \infty) \rightarrow \mathbb{R}_{\geq 0}$. If instead of moving at constant rate, s , along the horizontal axis, we move instead at rate $\lambda(s)$, then the number of points observed by time s is $Y(\int_0^s \lambda(r) dr)$. Further, from basic properties of exponential random variables, whenever $\lambda(s) > 0$ the probability of seeing a jump within the next small increment of time Δs is

$$P\left(Y\left(\int_0^{s+\Delta s} \lambda(r) dr\right) - Y\left(\int_0^s \lambda(r) dr\right) \geq 1\right) \approx \lambda(s)\Delta s.$$

Thus, the *propensity* for seeing another jump is precisely $\lambda(s)$.

Returning to the discussion directly following (1), and noting that $X_1(s) + X_2(s) = 1$ for all time, we note that the propensities of reactions 1 and 2 are

$$\lambda_1(X(s)) = \alpha X_1(s), \quad \lambda_2(X(s)) = \beta X_2(s).$$

Combining all of the above implies that we can represent R_1 and R_2 via

$$R_1(t) = Y_1\left(\int_0^t \alpha X_1(s) ds\right), \quad R_2(t) = Y_2\left(\int_0^t \beta X_2(s) ds\right),$$

and so a pathwise representation for the stochastic model (1) is

$$\begin{aligned} X(t) &= X_0 + Y_1\left(\int_0^t \alpha X_1(s) ds\right) \begin{pmatrix} -1 \\ 1 \end{pmatrix} \\ &\quad + Y_2\left(\int_0^t \beta X_2(s) ds\right) \begin{pmatrix} 1 \\ -1 \end{pmatrix}, \end{aligned} \tag{2}$$

where Y_1 and Y_2 are independent, unit-rate Poisson processes. Suppose now that $X_1(0) + X_2(0) = N \geq 1$. For example, perhaps we are now modeling the *number* of open and closed ion channels out of a total of N , as opposed to simply considering a single such channel. Suppose further that the propensity, or rate, at which ion channels are opening can be modeled as

$$\lambda_1(t, X(t)) = \alpha(t)X_1(t)$$

and the rate at which they are closing has propensity

$$\lambda_2(t, X(t)) = \beta(t)X_2(t),$$

where $\alpha(t), \beta(t)$ are non-negative functions of time, perhaps being voltage dependent. That is, suppose that for each $i \in \{1, 2\}$, the conditional probability of seeing the counting process R_i increase in the interval $[t, t+h)$ is $\lambda_i(t, X(t))h + o(h)$. The expression analogous to Eq. (2) is now

$$X(t) = X_0 + Y_1 \left(\int_0^t \alpha(s) X_1(s) ds \right) \begin{pmatrix} -1 \\ 1 \end{pmatrix} + Y_2 \left(\int_0^t \beta(s) X_2(s) ds \right) \begin{pmatrix} 1 \\ -1 \end{pmatrix}. \quad (3)$$

Having motivated the time dependent representation with the simple model above, we turn to the more general context. We now assume a jump model consisting of d chemical constituents (or ion channel states) undergoing transitions determined via $M > 0$ different reaction channels. For example, in the toy model above, the chemical constituents were $\{A, B\}$, and so $d = 2$, and the reactions were $A \rightarrow B$ and $B \rightarrow A$, giving $M = 2$. We suppose that $X_i(t)$ determines the value of the i th constituent at time t , so that $X(t) \in \mathbb{Z}^d$, and that the propensity function of the k th reaction is $\lambda_k(t, X(t))$. We further suppose that if the k th reaction channel takes place at time t , then the system is updated according to addition of the reaction vector $\zeta_k \in \mathbb{Z}^d$,

$$X(t) = X(t-) + \zeta_k.$$

The associated pathwise stochastic representation for this model is

$$X(t) = X_0 + \sum_k Y_k \left(\int_0^t \lambda_k(s, X(s)) ds \right) \zeta_k, \quad (4)$$

where the Y_k are independent unit-rate Poisson processes. The chemical master equation for this general model is

$$\frac{d}{dt} P_{X_0}(x, t) = \sum_{k=1}^M P_{X_0}(x - \zeta_k, t) \lambda_k(t, x - \zeta_k) - P_{X_0}(x, t) \sum_{k=1}^M \lambda_k(t, x),$$

where $P_{X_0}(x, t)$ is the probability of being in state $x \in \mathbb{Z}_{\geq 0}^d$ at time $t \geq 0$ given an initial condition of X_0 .

When the variable $X \in \mathbb{Z}^d$ represents the randomly fluctuating state of an ion channel in a single compartment conductance based neuronal model, we include the membrane potential $V \in \mathbb{R}$ as an additional dynamical variable. In contrast with neuronal models incorporating Gaussian noise processes, here we consider the voltage to evolve deterministically, conditional on the states of one or more ion channels. For illustration, suppose we have a single ion channel type with state variable X . Then, we supplement the path-

wise representation (4) with the solution of a differential equation obtained from Kirchoff's current conservation law:

$$C \frac{dV}{dt} = I_{\text{app}}(t) - I_V(V(t)) - \left(\sum_{i=1}^d g_i^o X_i(t) \right) (V(t) - V_X) \quad (5)$$

Here, g_i^o is the conductance of an individual channel when it is in the i^{th} state, for $1 \leq i \leq d$. The sum gives the total conductance associated with the channel represented by the vector X ; the reversal potential for this channel is the constant V_X . The term $I_V(V)$ captures any deterministic voltage-dependent currents due to other channels besides channel type X , and I_{app} represents a time-varying, deterministic applied current. In this case the propensity function will explicitly be a function of the voltage and we may replace $\lambda_k(s, X(s))$ in Eq. (4) with $\lambda_k(V(s), X(s))$. If multiple ion channel types are included in the model then, provided there are a finite number of types each with a finite number of individual channels, the vector $X \in \mathbb{Z}^d$ represents the aggregated channel state. For specific examples of handling a single or multiple ion channel types, see Sections 3 and 4, respectively.

2.1.1 Simulation of the representation Eqs. (4)–(5)

The construction of the representation (4) provided above implies a simulation strategy in which each point of the Poisson processes Y_k , denoted τ_n above, is generated sequentially and as needed. The time until the next reaction that occurs past time T is simply

$$\Delta = \min_k \left\{ \Delta_k : \int_0^{T+\Delta_k} \lambda_k(s, X(s)) ds = \tau_T^k \right\},$$

where τ_T^k is the first point associated with Y_k coming after $\int_0^T \lambda_k(s, X(s)) ds$:

$$\tau_T^k \equiv \inf \left\{ r > \int_0^T \lambda_k(s, X(s)) ds : Y_k(r) - Y_k \left(\int_0^T \lambda_k(s, X(s)) ds \right) = 1 \right\}.$$

The reaction that took place is indicated by the index at which the minimum is achieved. See Anderson (2007) for more discussion on this topic, including the algorithm provided below in which T_k will denote the value of the integrated intensity function $\int_0^t \lambda_k(s, X(s)) ds$ and τ_k will denote the first point associated with Y_k located after T_k .

All random numbers generated in the algorithm below are assumed to be independent.

Note that with a probability of one, the index determined in step 4 of Algorithm 1 is unique at each step. Note also that the algorithm relies on us being able to calculate a hitting time for each of the $T_k(t) = \int_0^t \lambda_k(s, X(s)) ds$ exactly.

Algorithm 1 (For the simulation of the representation (4)–(5))

1. Initialize: set the initial number of molecules of each species, X . Set the initial voltage value V . Set $t = 0$. For each k , set $\tau_k = 0$ and $T_k = 0$.
2. Generate M independent, uniform(0,1) random numbers $\{r_k\}_{k=1}^M$. For each $k \in \{1, \dots, M\}$ set $\tau_k = \ln(1/r_k)$.
3. Numerically integrate (5) forward in time until one of the following equalities hold:

$$\int_t^{t+\Delta} \lambda_k(V(s), X(s)) ds = \tau_k - T_k. \quad (6)$$
4. Let μ be the index of the reaction channel where the equality (6) holds.
5. For each k , set

$$T_k = T_k + \int_t^{t+\Delta} \lambda_k(V(s), X(s)) ds,$$
 where Δ is determined in Step (3).
6. Set $t = t + \Delta$ and $X \leftarrow X + \zeta_\mu$.
7. Let r be uniform(0,1) and set $\tau_\mu = \tau_\mu + \ln(1/r)$.
8. Return to step 3 or quit.

Of course, in general this is not possible. However, making use of any reliable integration software will almost always be sufficient. If the equations for the voltage and/or the intensity functions can be analytically solved for, as can happen in the Morris-Lecar model detailed below, then such numerical integration is unnecessary and efficiencies can be gained.

2.2 Gillespie representation

There are multiple alternative representations for the general stochastic process constructed in the previous section, with a ‘‘Gillespie’’ representation probably being the most useful in the current context. Following (Anderson and Kurtz 2011), we let Y be a unit rate Poisson process and let $\{\xi_i, i = 0, 1, 2, \dots\}$ be independent, uniform (0, 1) random variables that are independent of Y . Set

$$\lambda_0(V(s), X(s)) \equiv \sum_{k=1}^M \lambda_k(V(s), X(s)),$$

$$q_0 = 0 \text{ and for } k \in \{1, \dots, M\}$$

$$q_k(s) = \lambda_0(V(s), X(s))^{-1} \sum_{\ell=1}^k \lambda_\ell(V(s), X(s)),$$

where X and V satisfy

$$R_0(t) = Y \left(\int_0^t \lambda_0(V(s), X(s)) ds \right) \quad (7)$$

$$X(t) = X(0) + \sum_{k=1}^M \zeta_k \int_0^t 1_{\{\xi_{R_0(s-)} \in [q_{k-1}(s-), q_k(s-)]\}} dR_0(s) \quad (8)$$

$$C \frac{dV}{dt} = I_{\text{app}}(t) - I_V(V(t)) - \left(\sum_{i=1}^d g_i^o X_i(t) \right) (V(t) - V_X). \quad (9)$$

Then the stochastic process (X, V) defined via Eqs. (7)–(9) is a Markov process that is equivalent to Eq. (4)–(5), see (Anderson and Kurtz 2011). An intuitive way to understand the above is by noting that $R_0(t)$ simply determines the holding time in each state, whereas Eq. (8) simulates the embedded discrete time Markov chain (sometimes referred to as the *skeletal chain*) in the usual manner. Thus, this is the representation for the ‘‘Gillespie algorithm’’ (Gillespie 1977) with time dependent propensity functions (Anderson 2007). Note also that this representation is analogous to the PDMP formalism discussed in the introduction with λ_0 playing the role of the rate function that determines when the next jump takes place, and Eq. 8 implementing the transitions.

2.2.1 Simulation of the representation Eqs. (7)–(9)

Simulation of Eqs. (7)–(9) is the analog of using Gillespie’s algorithm in the time-homogeneous case. All random numbers generated in the algorithm below are assumed to be independent.

Algorithm 2 (For the simulation of the representation (7)–(9))

1. Initialize: set the initial number of molecules of each species, X . Set the initial voltage value V . Set $t = 0$.
2. Let r be uniform(0,1) and numerically integrate (9) forward in time until

$$\int_t^{t+\Delta} \lambda_0(V(s), X(s)) ds = \ln(1/r).$$
3. Let ξ be uniform(0,1) and find $k \in \{1, \dots, M\}$ for which

$$\xi \in [q_{k-1}((t + \Delta)-), q_k((t + \Delta)-)].$$
4. Set $t = t + \Delta$ and $X \leftarrow X + \zeta_k$.
5. Return to step 2 or quit.

Note again that the above algorithm relies on us being able to calculate a hitting time. If the relevant equations can

be analytically solved for, then such numerical integration is unnecessary and efficiencies can be gained.

3 Morris-Lecar

As a concrete illustration of the exact stochastic simulation algorithms, we will consider the well known Morris-Lecar system (Rinzel and Ermentrout 1989), developed as a model for oscillations observed in barnacle muscle fibers (Morris and Lecar 1981). The deterministic equations, which correspond to an appropriate scaling limit of the system, constitute a planar model for the evolution of the membrane potential $v(t)$ and the fraction of potassium gates, $n \in [0, 1]$, that are in the open or conducting state. In addition to a hyperpolarizing current carried by the potassium gates, there is a depolarizing calcium current gated by a rapidly equilibrating variable $m \in [0, 1]$. While a fully stochastic treatment of the Morris-Lecar system would include fluctuations in this calcium conductance, for simplicity in this section we will treat m as a fast, deterministic variable in the same manner as in the standard fast/slow decomposition, which we will refer to here as the planar Morris-Lecar model (Ermentrout and Terman 2010). See Section 4 for a treatment of the Morris-Lecar system with both the potassium and calcium gates represented as discrete stochastic processes.

The deterministic or mean field equations for the planar Morris-Lecar model are:

$$\frac{dv}{dt} = f(v, n) = \frac{1}{C} (I_{app} - g_{Ca} m_{\infty}(v)(v - v_{Ca}) - g_L(v - v_L) - g_K n(v - v_K)) \quad (10)$$

$$\frac{dn}{dt} = g(v, n) = \alpha(v)(1 - n) - \beta(v)n = (n_{\infty}(v) - n)/\tau(v) \quad (11)$$

The kinetics of the potassium channel may be specified either by the instantaneous time constant and asymptotic target, τ and n_{∞} , or equivalently by the *per capita* transition rates α and β . The terms m_{∞} , α , β , n_{∞} and τ satisfy

$$m_{\infty} = \frac{1}{2} \left(1 + \tanh \left(\frac{v - v_a}{v_b} \right) \right) \quad (12)$$

$$\alpha(v) = \frac{\phi \cosh(\xi/2)}{1 + e^{2\xi}} \quad (13)$$

$$\beta(v) = \frac{\phi \cosh(\xi/2)}{1 + e^{-2\xi}} \quad (14)$$

$$n_{\infty}(v) = \alpha(v)/(\alpha(v) + \beta(v)) = (1 + \tanh \xi)/2 \quad (15)$$

$$\tau(v) = 1/(\alpha(v) + \beta(v)) = 1/(\phi \cosh(\xi/2)) \quad (16)$$

where for convenience we define $\xi = (v - v_c)/v_d$. For definiteness, we adopt values of the parameters

$$v_K = -84, v_L = -60, v_{Ca} = 120 \quad (17)$$

$$I_{app} = 100, g_K = 8, g_L = 2, C = 20 \quad (18)$$

$$v_a = -1.2, v_b = 18 \quad (19)$$

$$v_c = 2, v_d = 30, \phi = 0.04, g_{Ca} = 4.4 \quad (20)$$

for which the deterministic system has a stable limit cycle. For smaller values of the applied current (e.g. $I_{app} = 75$) the system has a stable fixed point, that loses stability through a subcritical Hopf bifurcation as I_{app} increases ((Ermentrout and Terman 2010), Section 3).

In order to incorporate the effects of random ion channel gating, we will introduce a finite number of potassium channels, N_{tot} , and treat the number of channels in the open state as a discrete random process, $0 \leq N(t) \leq N_{tot}$. In this simple model, each potassium channel switches between two states – closed or open – independently of the others, with voltage-dependent *per capita* transition rates α and β , respectively. The entire population conductance ranges from 0 to $g_K^o N_{tot}$, where $g_K^o = g_K/N_{tot}$ is the single channel conductance, and g_K is the maximal whole cell conductance. For purposes of illustration and simulation we will typically use $N_{tot} = 40$ individual channels.²

Our random variables will therefore be the voltage, $V \in (-\infty, +\infty)$, and the number of open potassium channels, $N \in \{0, 1, 2, \dots, N_{tot}\}$. The number $N_{tot} \in \mathbb{N}$ is taken to be a fixed parameter. We follow the usual capital/lowercase convention in the stochastic processes literature: $N(t)$ and $V(t)$ are the random processes and n and v are values they might take. In the random time change representation of Section 2.1, the opening and closing of the potassium channels are driven by two independent, unit rate Poisson processes, $Y_{open}(t)$ and $Y_{close}(t)$.

The evolution of V and N is closely linked. Conditioned on N having a specific value, say $N = n$, the evolution of V obeys a *deterministic* differential equation,

$$\left. \frac{dV}{dt} \right|_{N=n} = f(V, n). \quad (21)$$

Although its conditional evolution is deterministic, V is nevertheless a random variable. Meanwhile, N evolves as a jump process, *i.e.* $N(t)$ is piecewise-constant, with transitions $N \rightarrow N \pm 1$ occurring with intensities that depend on

²Morris and Lecar used a value of $g_K = 8 \text{ mmho/cm}^2$ for the specific potassium conductance, corresponding to 80 pSiemens (pS) per square micron (Morris and Lecar 1981). The single channel conductance is determined by the structure of the potassium channel, and varies somewhat from species to species. However, conductances around 20 pS are typical (Shingai and Quandt 1986), which would give a density estimate of roughly 40 channels for a 10 square micron patch of cell membrane.

the voltage V . Conditioned on $V = v$, the transition rates for N are

$$N \rightarrow N + 1 \text{ with per capita rate } \alpha(v) \quad (22)$$

(i.e. with net rate $\alpha(v) \cdot (N_{\text{tot}} - N)$),

$$N \rightarrow N - 1 \text{ with per capita rate } \beta(v) \quad (23)$$

(i.e. with net rate $\beta(v) \cdot N$).

Graphically, we may visualize the state space for N in the following manner:

$$0 \xrightleftharpoons[\beta]{\alpha \cdot N_{\text{tot}}} 1 \xrightleftharpoons[2\beta]{\alpha \cdot (N_{\text{tot}}-1)} 2 \cdots (k-1) \xrightleftharpoons[k\beta]{\alpha \cdot (N_{\text{tot}}-k+1)} k \xrightleftharpoons[(k+1)\beta]{\alpha \cdot (N_{\text{tot}}-k)} (k+1) \cdots (N_{\text{tot}}-1) \xrightleftharpoons[N_{\text{tot}}\beta]{\alpha} N_{\text{tot}},$$

with the nodes of the above graph being the possible states for the process N , and the transition intensities located above and below the transition arrows.

Adopting the random time change representation of Section 2.1 we write our stochastic Morris-Lecar system as follows (cf. Eqs. (10)–(11)):

$$\frac{dV}{dt} = f(V(t), N(t)) = \quad (24)$$

$$= \frac{1}{C} (I_{\text{app}} - g_{Ca}m_{\infty}(V(t))(V(t) - V_{Ca}) - g_L(V - V_L) - g_K^o N(t)(V(t) - V_K)) \quad (25)$$

$$N(t) = N(0) - Y_{\text{close}} \left(\int_0^t \beta(V(s))N(s) ds \right) + Y_{\text{open}} \left(\int_0^t \alpha(V(s))(N_{\text{tot}} - N(s)) ds \right).$$

Supplementary Materials provide sample implementations of Algorithm 1 for the planar Morris-Lecar equations in XPP and Matlab, respectively. Figure 1 illustrates the results of the Matlab implementation.

4 Models with more than one channel type

We present here an example of a conductance based model with more than one stochastic gating variable. Specifically, we consider the original Morris-Lecar model, which has a three dimensional phase space, and we take both the calcium and potassium channels to be discrete. We include code in **Supplementary Materials**, for XPP and Matlab, respectively, for the implementation of Algorithm 1 for this model.

4.1 Random Time Change Representation for the Morris-Lecar Model with Two Stochastic Channel Types

In Morris and Lecar's original treatment of voltage oscillations in barnacle muscle fiber (Morris and Lecar 1981) the calcium gating variable m is included as a dynamical variable. The full (deterministic) equations have the form:

$$\frac{dv}{dt} = F(v, n, m) = \frac{1}{C} (I_{\text{app}} - g_L(v - v_L) - g_{Ca}m(v - v_{Ca}) - g_Kn(v - v_K)) \quad (26)$$

$$\frac{dn}{dt} = G(v, n, m) = \alpha_n(v)(1 - n) - \beta_n(v)n = (n_{\infty}(v) - n)/\tau_n(v) \quad (27)$$

$$\frac{dm}{dt} = H(v, n, m) = \alpha_m(v)(1 - m) - \beta_m(v)m = (m_{\infty}(v) - m)/\tau_m(v) \quad (28)$$

Here, rather than setting m to its asymptotic value $m_{\infty} = \alpha_m/(\alpha_m + \beta_m)$, we allow the number of calcium gates to evolve according to Eq. (27). The planar form (Eqs. (10)–(11)) is obtained by observing that m approaches equilibrium significantly more quickly than n and v . Using standard arguments from singular perturbation theory (Rinzel and Ermentrout 1989; Terman and Rubin 2002), one may




Fig. 1 Trajectory generated by Algorithm 1 (the random time change algorithm, (4)–(5)) for the planar Morris-Lecar model. We set $N_{\text{tot}} = 40$ potassium channels and used a driving current $I_{\text{app}} = 100$, which is above the Hopf bifurcation threshold for the parameters given. *Top Panel:* Number of open potassium channels (N), as a function of time. *Second Panel:* Voltage (V), as a function of time. *Bottom Panel:* Trajectory plotted in the (V, N) plane. Voltage varies along a continuum while open channel number remains discrete. Red curve: v -nullcline of the underlying deterministic system, obtained by setting the RHS of Eq. (10) equal to zero. Green curve: n -nullcline, obtained by setting the RHS of Eq. (11) equal to zero

approximate certain aspects of the full system (28)–(27) by setting m to $m_\infty(v)$, and replacing $F(v, n, m)$ and $G(v, n, m)$ with $f(v, n) = F(v, n, m_\infty(v))$ and $g(v, n) = G(v, n, m_\infty(v))$, respectively. This reduction to the slow dynamics leads to the planar model (10)–(11).

To specify the full 3D equations, we introduce $\xi_m = (v - v_a)/v_b$ in addition to $\xi_n = (v - v_c)/v_d$ already introduced for the 2D model. The variable ξ_x represents where the voltage falls along the activation curve for channel type x , relative to its half-activation point (v_a for calcium and v_c for potassium) and its slope (reciprocals of v_b for calcium and v_d for potassium). The per capita opening rates α_m, α_n and closing rates β_m, β_n for each channel type are given by

$$\alpha_m(v) = \frac{\phi_m \cosh(\xi_m/2)}{1 + e^{2\xi_m}}, \quad \beta_m(v) = \frac{\phi_m \cosh(\xi_m/2)}{1 + e^{-2\xi_m}} \quad (29)$$

$$\alpha_n(v) = \frac{\phi_n \cosh(\xi_n/2)}{1 + e^{2\xi_n}}, \quad \beta_n(v) = \frac{\phi_n \cosh(\xi_n/2)}{1 + e^{-2\xi_n}} \quad (30)$$

with parameters $v_a = -1.2, v_b = 18, v_c = 2, v_d = 30, \phi_m = 0.4, \phi_n = 0.04$. The asymptotic open probabilities for calcium and potassium are given, respectively, by the terms m_∞, n_∞ , and the time constants by τ_m and τ_n . These terms satisfy the relations

$$m_\infty(v) = \alpha_m(v)/(\alpha_m(v) + \beta_m(v)) = (1 + \tanh \xi_m)/2 \quad (31)$$

$$n_\infty(v) = \alpha_n(v)/(\alpha_n(v) + \beta_n(v)) = (1 + \tanh \xi_n)/2 \quad (32)$$

$$\tau_m(v) = 1/(\phi \cosh(\xi_m/2)) \quad (33)$$

$$\tau_n(v) = 1/(\phi \cosh(\xi_n/2)). \quad (34)$$

Assuming a finite population of M_{tot} calcium gates and N_{tot} potassium gates, we have a stochastic hybrid system with one continuous variable, $V(t)$, and two discrete variables, $M(t)$ and $N(t)$. The voltage evolves according to the sum of the applied, leak, calcium, and potassium currents:

$$\begin{aligned} \frac{dV}{dt} &= F(V(t), N(t), M(t)) \\ &= \frac{1}{C} \left(I_{\text{app}} - g_L(V(t) - v_L) - g_{Ca} \frac{M(t)}{M_{\text{tot}}}(V(t) - v_{Ca}) \right. \\ &\quad \left. - g_K \frac{N(t)}{N_{\text{tot}}}(V(t) - v_K) \right). \end{aligned} \quad (35)$$

In contrast, the number of open M gates and the number of open N gates change only by unit increases and decreases, while remaining constant between such changes. The discrete channel states $M(t)$ and $N(t)$ evolve according to

$$\begin{aligned} M(t) &= M(0) - Y_{\text{close}}^M \left(\int_0^t \beta_m(V(s))M(s) ds \right) + \\ &\quad Y_{\text{open}}^M \left(\int_0^t \alpha_m(V(s))(M_{\text{tot}} - M(s)) ds \right) \end{aligned} \quad (36)$$

$$\begin{aligned} N(t) &= N(0) - Y_{\text{close}}^N \left(\int_0^t \beta_n(V(s))N(s) ds \right) + \\ &\quad Y_{\text{open}}^N \left(\int_0^t \alpha_n(V(s))(N_{\text{tot}} - N(s)) ds \right). \end{aligned} \quad (37)$$

Figure 2 shows the results of the Matlab implementation for the 3D Morris-Lecar system, with both the potassium and calcium channel treated discretely, using Algorithm 1 (the random time change algorithm, (4)–(5)). Here $M_{\text{tot}} = N_{\text{tot}} = 40$ channels, and the applied current $I_{\text{app}} = 100$ puts the deterministic system at a stable limit cycle close to a Hopf bifurcation.

5 Comparison of the Exact Algorithm with a Piecewise Constant Propensity Approximation.

Exact versions of the stochastic simulation algorithm for hybrid ion channel models have been known since at least the 1980s (Clay and DeFelice 1983). Nevertheless, the implementation one finds most often used in the literature is an approximate method in which the per capita reaction propensities are held fixed between channel events. That is, in step 3 of Algorithm 1 the integral

$$\int_t^{t+\Delta_k} \lambda_k(V(s), X(s)) ds$$

is replaced with

$$\Delta_k \lambda_k(V(t), X(t))$$

leaving the remainder of the algorithm unchanged. Put another way, one generates the sequence of channel state jumps using the propensity immediately following the most recent jump, rather than taking into account the time dependence of the reaction propensities due to the continuously changing voltage. This piecewise constant propensity approximation is analogous, in a sense, to the forward Euler method for the numerical solution of ordinary differential equations.

Figure 3 shows a direct comparison of pathwise numerical solutions obtained by the exact method provided in Algorithm 1 and the approximate piecewise constant method detailed above. In general, the solution of a stochastic differential equation with a given initial condition is a map from the sample space Ω to a space of trajectories. In the present context, the underlying sample space consists of one independent unit rate Poisson process per reaction channel. For the planar Morris-Lecar model a point in Ω amounts to fixing two Poisson processes, Y_{open} and Y_{closed} , to drive the transitions of the potassium channel. For the full 3D Morris-Lecar model we have four processes, $Y_1 \equiv Y_{\text{Ca,open}}, Y_2 \equiv Y_{\text{Ca,closed}}, Y_3 \equiv Y_{\text{K,open}}$




Fig. 2 Trajectory generated by Algorithm 1 (the random time change algorithm, (4)–(5)) for the full three-dimensional Morris-Lecar model (Eqs. (27)–(28)). We set $N_{\text{tot}} = 40$ potassium channels and $M_{\text{tot}} = 40$ calcium channels, and used a driving current $I_{\text{app}} = 100$, a value above the Hopf bifurcation threshold of the mean field equations for the parameters given. *Top Left Panel:* Number of open calcium channels (M), as a function of time. *Second Left Panel:* Number of open potassium channels (N), as a function of time. *Third Left Panel:* Voltage

(V), as a function of time. *Right Panel:* Trajectory plotted in the (V, M, N) phase space. Voltage varies along a continuum while the joint channel state remains discrete. Note that the number of open calcium channels makes frequent excursions between $M = 0$ and $M = 40$, which demonstrates that neither a Langevin approximation nor an approximate algorithm such as τ -leaping (Euler's method) would provide a good approximation to the dynamics of the system

and $Y_4 \equiv Y_{K,\text{closed}}$. In this case the exact algorithm provides a numerical solution of the map from $\{Y_k\}_{k=1}^4 \in \Omega$ and initial conditions (M_0, N_0, V_0) to the trajectory $(M(t), N(t), V(t))$. The approximate piecewise constant algorithm gives a map from the same domain to a different trajectory, $(\tilde{M}(t), \tilde{N}(t), \tilde{V}(t))$. To make a pathwise comparison for the full Morris-Lecar model, therefore, we fix both

the initial conditions and the four Poisson processes, and compare the resulting trajectories.




Fig. 3 Comparison of the exact algorithm with the piecewise constant propensity approximation. Blue solid lines denote the solution $(M(t), N(t), V(t))$ obtained using the Algorithm 1. Red dashed lines denote the solution $(\tilde{M}(t), \tilde{N}(t), \tilde{V}(t))$ obtained using the piecewise constant approximation. Both algorithms were begun with identical initial conditions (M_0, N_0, V_0) , and driven by the same four Poisson process streams Y_1, \dots, Y_4 . Note the gradual divergence of the trajectories as differences between the exact and the forward approximate algorithms accumulate, demonstrating “strong” (or pathwise) divergence of the two methods. The exact and approximate trajectories diverge as time increases, even though they are driven by identical noise sources

Several features are evident in Fig. 3. Both algorithms produce a sequence of noise-dependent voltage spikes, with similar firing rates. The trajectories (M, N, V) and $(\tilde{M}, \tilde{N}, \tilde{V})$ initially remain close together, and the timing of the first spike (taken, e.g., as an upcrossing of V from negative to positive voltage) is similar for both algorithms. Over time, however, discrepancies between the trajectories accumulate. The timing of the second and third spikes is noticeably different, and before ten spikes have accumulated the spike trains have become effectively uncorrelated.

Even though trajectories generated by the exact and approximate algorithms diverge when driven by identical Poisson processes, the two processes could still generate sample paths with similar time-dependent or stationary distributions. That is, even though the two algorithms show strong divergence, they could still be close in a weak sense.

Given M_{tot} calcium and N_{tot} potassium channels, the density for the hybrid Markov process may be written

$$\rho_{m,n}(v, t) = \frac{1}{dv} \Pr \{M(t) = m, N(t) = n, V \in [v, v + dv)\}, \quad (38)$$

and obeys a master equation

$$\begin{aligned} \frac{\partial \rho_{m,n}(v, t)}{\partial t} = & - \frac{\partial (F(v, n, m) \rho_{m,n}(v, t))}{\partial v} \\ & - (\alpha_m(v)(M_{\text{tot}} - m) + \beta_m(v)m + \alpha_n(v) \\ & \times (N_{\text{tot}} - n) + \beta_n(v)n) \rho_{m,n}(v, t) \\ & + (M_{\text{tot}} - m + 1) \alpha_m(v) \rho_{m-1,n}(v, t) \\ & + (m + 1) \beta_m(v) \rho_{m+1,n}(v, t) \\ & + (N_{\text{tot}} - n + 1) \alpha_n(v) \rho_{m,n-1}(v, t) \\ & + (n + 1) \beta_n(v) \rho_{m,n+1}(v, t), \end{aligned} \quad (39)$$

with initial condition $\rho_{m,n}(v, 0) \geq 0$ given by any (integrable) density such that $\int_{v \in \mathbb{R}} \sum_{m,n} \rho_{m,n}(v, 0) dv \equiv 1$, and boundary conditions $\rho \rightarrow 0$ as $|v| \rightarrow \infty$ and $\rho \equiv 0$ for either $m, n < 0$ or $m > M_{\text{tot}}$ or $n > N_{\text{tot}}$.

In contrast, the approximate algorithm with piecewise constant propensities does not generate a Markov process, since the transition probabilities depend on the past rather than the present values of the voltage component. Consequently they do not satisfy a master equation. Nevertheless it is plausible that they may have a unique stationary distribution.

Figure 4 shows pseudocolor plots of the histograms viewed in the (v, n) plane, i.e. with entries summed over

m , for Algorithm 1 (“Exact”) and the approximate piecewise constant (“Approx”) algorithm, with $M_{\text{tot}} = N_{\text{tot}} = k$ channels for $k = 1, 2, 5, 10, 20$ and 40 . The two algorithms were run with independent random number streams in the limit cycle regime (with $I_{\text{app}} = 100$) for $t_{\text{max}} \approx 200,000$ time units, sampled every 10 time units, which generated $\geq 17,000$ data points per histogram. For $k < 5$ the difference in the histograms is obvious at a glance. For $k \geq 10$ the histograms appear increasingly similar.

Figure 5 shows bar plots of the histograms with 2,000,000 sample points projected on the voltage axis, i.e. with entries summed over m and n , for the same data as in Fig. 4, with $M_{\text{tot}} = N_{\text{tot}} = k$ channels ranging from

Fig. 4 Normalized histograms of voltage and N for exact and approximate piecewise constant algorithms. The V -axis was partitioned into 100 equal width bins for each pair of histograms. The N -axis takes on $N_{\text{tot}} + 1$ discrete values. Color scale indicates relative frequency with which a bin was occupied (dark = infrequent, lighter = more frequent)


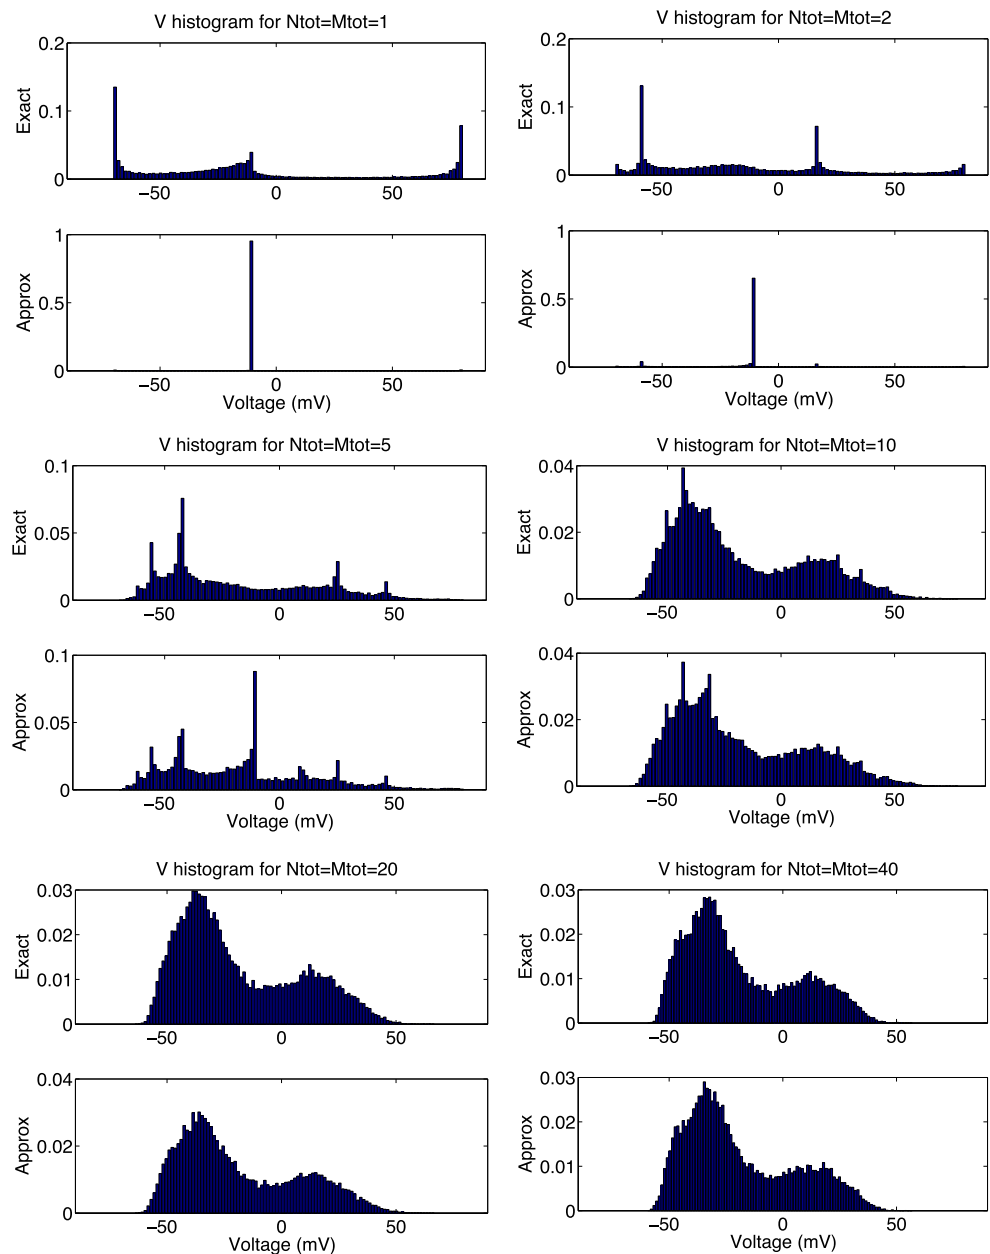


Fig. 5 Histograms of voltage for exact and approximate piecewise constant algorithms. The V -axis was partitioned into 100 equal width bins for each pair of histograms



$k = 1$ to $k = 100$. Again, for $k \leq 5$, the two algorithms generate histograms that are clearly distinct. For $k \geq 20$ they appear similar, while $k = 10$ appears to be a borderline case.

To quantify the similarity of the histograms we calculated the empirical L_1 difference between the histograms in two ways: first for the full (v, n, m) histograms, and then for the histograms collapsed to the voltage axis. Let $\rho_{m,n}(v)$ and $\tilde{\rho}_{m,n}(v)$ denote the stationary distributions for the exact and the approximate algorithms, respectively (assuming these exist). To compare the two distributions we approximate the L_1 distance between them, i.e.

$$d(\rho, \tilde{\rho}) = \int_{v_{\min}}^{v_{\max}} \left(\sum_{m=0}^{M_{\text{tot}}} \sum_{n=0}^{N_{\text{tot}}} |\rho_{m,n}(v) - \tilde{\rho}_{m,n}(v)| \right) dv$$

where v_{\min} and v_{\max} are chosen so that for all m, n we have $F(v_{\min}, n, m) > 0$ and $F(v_{\max}, n, m) < 0$. It is easy to see that such values of v_{\min} and v_{\max} must exist, since $F(v, n, m)$ is a linear and monotonically decreasing function of v for any fixed pair (n, m) , cf. Eq. (28). Therefore, for any exact simulation algorithm, once the voltage component of a trajectory falls in the interval $v_{\min} \leq v \leq v_{\max}$, it remains in this interval for all time. We approximate the voltage integral by summing over 100 evenly spaced bins along the voltage axis. Figure 6, top panel, shows empirical




Fig. 6 L_1 differences between empirical histograms generated by exact and approximate algorithms. $H_j(v, n, m)$ is the number of samples for which $V \in [v, v + \Delta v)$, $N = n$ and $M = m$, where Δv is a discretization of the voltage axis into 100 equal width bins; $j = 1$ represents an empirical histogram obtained by an exact algorithm (Algorithm 1), and $j = 2$ represents an empirical histogram obtained by a piecewise-constant propensity approximation. Note the normalized L_1 difference of two histograms can range from 0 to 2. Based on 2,000,000 samples for each algorithm. **Top:** Difference for the full histograms, $L_1(H_1 - H_2) \equiv \sum_{v,n,m} |H_1(v, n, m) - H_2(v, n, m)| / \#samples$. Here v refers to binned voltage values. The total number of bins is $100 M_{tot} N_{tot}$. **Bottom:** Difference for the corresponding voltage histograms, $L_1(VH_1 - VH_2) \equiv \sum_v |VH_1(v) - VH_2(v)| / \#samples$, where $VH_j(v) = \sum_{n,m} H_j(v, n, m)$. The total number of bins is 100

estimates of $d(\rho, \tilde{\rho})$ for values of $M_{tot} = N_{tot}$ ranging from 1 to 100. The bottom panel shows the histogram for voltage considered alone.

6 Coupling, variance reduction, and parametric sensitivities

It is relatively straightforward to derive the exact simulation strategies from the different representations provided here. Less obvious, and perhaps even more useful, is the fact that the random time change formalism also lends itself to the development of new computational methods that potentially allow for significantly greater computational speed, with no loss in accuracy. They do so by providing ways to *couple* two processes in order to reduce the variance, and thereby increase the speed, of different natural estimators.

For an example of a common computational problem where such a benefit can be had, consider the computation of *parametric sensitivities*, which is a powerful tool in determining parameters to which a system output is most responsive. Specifically, suppose that the

intensity/propensity functions are dependent on some vector of parameters θ . For instance, θ may represent a subset of the system’s mass action kinetics constants, the cell capacitance, or the underlying number of channels of each type. We may wish to know how sensitively a quantity such as the average firing rate or interspike interval variance depends on the parameters θ . We then consider a family of models (V^θ, X^θ) , parameterized by θ , with stochastic equations

$$\begin{aligned} \frac{d}{dt} V^\theta(t) &= f(\theta, V^\theta(t), X^\theta(t)) \\ X^\theta(t) &= X_0^\theta + \sum_k Y_k \left(\int_0^t \lambda_k^\theta(V^\theta(s), X^\theta(s)) ds \right) \zeta_k, \end{aligned} \quad (40)$$

where f is some function, and all other notation is as before. Some quantities arising in neural models depend on the entire sample path, such as the mean firing rate and the interspike interval variance. Let $g(\theta, V^\theta, X^\theta)$ be a path functional capturing the quantity of interest. In order to efficiently and accurately evaluate the relative shift in the expectation of g due to a perturbation of the parameter vector, we would estimate

$$\begin{aligned} \tilde{s} &= \epsilon^{-1} \mathbb{E} \left[g(\theta', V^{\theta'}, X^{\theta'}) - g(\theta, V^\theta, X^\theta) \right] \\ &\approx \frac{1}{\epsilon N} \sum_{i=1}^N \left[g(\theta', V_{[i]}^{\theta'}, X_{[i]}^{\theta'}) - g(\theta, V_{[i]}^\theta, X_{[i]}^\theta) \right] \end{aligned} \quad (41)$$

where $\epsilon = \|\theta - \theta'\|$ and where $(V_{[i]}^\theta, X_{[i]}^\theta)$ represents the i th path generated with a parameter choice of θ , and N is the number of sample paths computed for the estimation. That is, we would use a finite difference and Monte Carlo sampling to approximate the change in the expectation.

If the paths $(V_{[i]}^{\theta'}, X_{[i]}^{\theta'})$ and $(V_{[i]}^\theta, X_{[i]}^\theta)$ are generated independently, then the variance of the estimator (41) for \tilde{s} is $O(N^{-1}\epsilon^{-2})$, and the standard deviation of the estimator is $O(N^{-1/2}\epsilon^{-1})$. This implies that in order to reduce the confidence interval of the estimator to a target level of $\rho > 0$, we require

$$N^{-1/2}\epsilon^{-1} \lesssim \rho \implies N \gtrsim \epsilon^{-2}\rho^{-2},$$

which can be prohibitive. Reducing the variance of the estimator can be achieved by *coupling* the processes $(V_{[i]}^{\theta'}, X_{[i]}^{\theta'})$ and $(V_{[i]}^\theta, X_{[i]}^\theta)$, so that they are correlated, by constructing them on the same probability space. The representations we present here lead rather directly to schemes for reducing the variance of the estimator. We discuss different possibilities.

1. **The common random number (CRN) method.** The CRN method simulates both processes according to the Gillespie representation (7)–(9) with the same Poisson process Y and the same stream of uniform random

variables $\{\xi_i\}$. In terms of implementation, one simply reuses the same two streams of uniform random variables in Algorithm 2 for both processes.

2. **The common reaction path method (CRP).** The CRP method simulates both processes according to the random time change representation with the same Poisson processes Y_k . In terms of implementation, one simply makes one stream of uniform random variables for each reaction channel, and then uses these streams for the simulation of both processes. See (Rathinam et al. 2010), where this method was introduced in the context of biochemical models.
3. **The coupled finite difference method (CFD).** The CFD method utilizes a “split coupling” introduced in (Anderson 2012). Specifically, it splits the counting process for each of the reaction channels into three pieces: one counting process that is shared by X^θ and $X^{\theta'}$ (and has propensity equal to the minimum of their respective intensities for that reaction channel), one that only accounts for the jumps of X^θ , and one that only accounts for the jumps of $X^{\theta'}$. Letting $a \wedge b = \min\{a, b\}$, the precise coupling is

$$\begin{aligned}
 X^{\theta'}(t) &= X_0^{\theta'} + \sum_k Y_{k,1} \left(\int_0^t m_k(\theta, \theta', s) ds \right) \zeta_k \\
 &\quad + \sum_k Y_{k,2} \left(\int_0^t \lambda_k^{\theta'}(V^{\theta'}(s), X^{\theta'}(s)) - m_k(\theta, \theta', s) ds \right) \zeta_k \\
 X^\theta(t) &= X^\theta(t) + \sum_k Y_{k,1} \left(\int_0^t m_k(\theta, \theta', s) ds \right) \zeta_k \\
 &\quad + \sum_k Y_{k,3} \left(\int_0^t \lambda_k^\theta(V^\theta(s), X^\theta(s)) - m_k(\theta, \theta', s) ds \right) \zeta_k,
 \end{aligned}
 \tag{42}$$

where

$$m_k(\theta, \theta', s) \equiv \lambda_k^\theta(V^\theta(s), X^\theta(s)) \wedge \lambda_k^{\theta'}(V^{\theta'}(s), X^{\theta'}(s)),$$

and where $\{Y_{k,1}, Y_{k,2}, Y_{k,3}\}$ are independent unit-rate Poisson processes. Implementation is then carried out by Algorithm 1 in the obvious manner.

While the different representations provided in this paper imply different exact simulation strategies (i.e. Algorithms 1 and 2), those strategies still produce statistically equivalent paths. This is *not* the case for the methods for parametric sensitivities provided above. To be precise, each of the methods constructs a coupled pair of processes $((V^\theta, X^\theta), (V^{\theta'}, X^{\theta'}))$, and the marginal processes (V^θ, X^θ) and $(V^{\theta'}, X^{\theta'})$ are all statistically equivalent no matter the method used. However, the covariance $\text{Cov}((V^\theta(t), X^\theta(t)), (V^{\theta'}(t), X^{\theta'}(t)))$ can be drastically different. This is important, since it is variance reduction we are after, and for any component X_j ,

$$\text{Var}(X_j^\theta(t) - X_j^{\theta'}(t)) = \text{Var}(X_j^\theta(t)) + \text{Var}(X_j^{\theta'}(t)) - 2\text{Cov}(X_j^\theta(t), X_j^{\theta'}(t)).$$

Thus, minimizing variance is equivalent to maximizing covariance. Typically, the CRN method does the worst job of maximizing the covariance, even though it is the most widely used method (Srivastava et al. 2013). The CFD method with the split coupling procedure typically does the best job of maximizing the covariance, though examples exist in which the CRP method is the most efficient (Anderson 2012; Anderson and Koyama 2014; Srivastava et al. 2013).

7 Discussion

We have provided two general representations for stochastic ion channel kinetics, one based on the random time change formalism, and one extending Gillespie’s algorithm to the case of ion channels driven by time-varying membrane potential. These representations are known in other branches of science, but less so in the neuroscience community. We believe that the random time change representation (Algorithm 1, (4)–(5)) will be particularly useful to the computational neuroscience community as it allows for generalizations of computational methods developed in the context of biochemistry, in which the propensities depend upon the state of the jump process only. For example, variance reduction strategies for the efficient computation of first and second order sensitivities (Anderson 2012; Anderson and Wolf 2012; Rathinam et al. 2010), as discussed in Section 6, and for the efficient computation of expectations using multi-level Monte Carlo (Anderson and Higham 2012; Giles 2008) now become feasible.

The random time change approach avoids several approximations that are commonly found in the literature. In simulation algorithms based on a fixed time step chemical Langevin approach, it is necessary to assume that the increments in channel state are approximately Gaussian distributed over an appropriate time interval (Fox and Yan-nan 1994; Goldwyn et al. 2011; Goldwyn and Shea-Brown 2011; Mino et al. 2002; White et al. 1998). However, in exact simulations with small membrane patches corresponding to $M_{\text{tot}} = 40$ calcium channels, the exact algorithm routinely visits the states $M(t) = 0$ and $M(t) = 40$, for which the Gaussian increment approximation is invalid regardless of time step size. The main alternative algorithm typically found in the literature is the piecewise constant propensity or approximate forward algorithm (Fisch et al. 2012; Kispersky and White 2008; Schwalger et al. 2010). However, this algorithm ignores changes to membrane potential during the intervals between channel state changes. As the sensitivity of ion channel opening and closing to voltage is fundamental to the neurophysiology of cellular excitability, these algorithms are not appropriate unless the time between openings

and closing is especially small. The exact algorithm (Clay and DeFelice 1983; Keener and Newby 2011; Newby et al. 2013) is straightforward to implement and avoids these approximations and pitfalls.

Our study restricts attention to the suprathreshold regime of the Morris-Lecar model, in which the applied current (here, $I_{\text{app}} = 100$) puts the deterministic system well above the Hopf bifurcation marking the onset of oscillations. In this regime, spiking is not the result of noise-facilitated release, as it is in the subthreshold or excitable regime. Bressloff, Keener and Newby used eigenfunction expansion methods, path integrals, and the theory of large deviations to study spike initiation as a noise facilitated escape problem in the excitable regime, as well as to incorporate synaptic currents into a stochastic network model; they use versions of the exact algorithm presented in this paper (Bressloff and Newby 2014b; Newby et al. 2013). Interestingly, these authors find that the usual separation-of-time scales picture breaks down. That is, the firing rate obtained by 1D Kramers rate theory when the slow recovery variable (fraction of open potassium gates) is taken to be fixed does not match that obtained by direct simulation with the exact algorithm. Moreover, by considering a maximum likelihood approach to the 2D escape problem, the authors show that, counterintuitively, spontaneous closure of potassium channels contributes more significantly to noise induced escape than spontaneous opening of sodium channels, as might naively have been expected.

The algorithms presented here are broadly applicable beyond the effects of channel noise on the regularity of action potential firing in a single compartment neuron model. Exact simulation of hybrid stochastic models have been used, for instance, to study spontaneous dendritic NMDA spikes (Bressloff and Newby 2014a), intracellular growth of the T7 bacteriophage (Alfonsi et al. 2005), and hybrid stochastic network models taking into account piecewise deterministic synaptic currents (Bressloff and Newby 2013). This latter work represents a significant extension of the neural master equation approach to stochastic population models (Bressloff 2009; Buice et al. 2010).

For any simulation algorithm, it is reasonable to ask about the growth of complexity of the algorithm as the underlying stochastic model is enriched. For example, the natural jump Markov interpretation of Hodgkin and Huxley's model for the sodium channel comprises eight distinct states with twenty different state-to-state transitions, each driven by an independent Poisson process in the random time change representation. Recent investigations of sodium channel kinetics have led neurophysiologists to formulate models with as many as 26 distinct states connected by 72 transitions (Milescu et al. 2010). While the random time change representation extends naturally to such scenarios, it may also be fruitful to combine it with complexity

reduction methods such as the stochastic shielding algorithm introduced by Schmandt and Galán (Schmandt and Galán 2012), and analyzed by Schmidt and Thomas (Schmidt and Thomas 2014). For example, of the twenty independent Poisson processes driving a discrete Markov model of the Hodgkin-Huxley sodium channel, only four of the processes directly affect the conductance of the channel; fluctuations associated with the remaining sixteen Poisson processes may be ignored with negligible loss of accuracy in the first and second moments of the channel occupancy. Similarly, for the 26 state sodium channel model proposed in (Milescu et al. 2010), all but 12 of the 72 Poisson processes representing state-to-state transitions can be replaced by their expected values. Analysis of algorithms combining stochastic shielding and the random time change framework are a promising direction for future research.

Acknowledgments Anderson was supported by NSF grant DMS-1318832 and Army Research Office grant W911NF-14-1-0401. Ermentrout was supported by NSF grant DMS-1219754. Thomas was supported by NSF grants EF-1038677, DMS-1010434, and DMS-1413770, by a grant from the Simons Foundation (#259837), and by the Council for the International Exchange of Scholars (CIES). We gratefully acknowledge the Mathematical Biosciences Institute (MBI, supported by NSF grant DMS 0931642) at The Ohio State University for hosting a workshop at which this research was initiated. The authors thank David Friel and Casey Bennett for helpful discussions and testing of the algorithms.

Conflict of interest The authors declare that they have no conflict of interest.

References

- Alfonsi, A., Cancès, E., Turinici, G., Di Ventura, B., Huisinga, W. (2005). Adaptive simulation of hybrid stochastic and deterministic models for biochemical systems. In *ESAIM: Proceedings* (Vol. 14, pp. 1–13). EDP Sciences.
- Anderson, D.F. (2007). A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *Journal of Chemical Physics*, 127(21), 214107.
- Anderson, D.F. (2012). An efficient finite difference method for parameter sensitivities of continuous time Markov chains. *SIAM Journal on Numerical Analysis*, 50(5), 2237–2258.
- Anderson, D.F., Ganguly, A., Kurtz, T.G. (2011). Error analysis of tau-leap simulation methods. *Annals of Applied Probability*, 21(6), 2226–2262.
- Anderson, D.F., & Higham, D.J. (2012). Multi-level Monte Carlo for continuous time Markov chains, with applications in biochemical kinetics. *SIAM: Multiscale Modeling and Simulation*, 10(1), 146–179.
- Anderson, D.F., Higham, D.J., Sun, Y. (2014). *Complexity analysis of multilevel monte carlo tau-leaping*. Submitted.
- Anderson, D.F., & Koyama, M. (2014). An asymptotic relationship between coupling methods for stochastically modeled population processes. Accepted for publication to *IMA Journal of Numerical Analysis*.

- Anderson, D.F., & Kurtz, T.G. (2011). *Design and analysis of biomolecular circuits, chapter 1. Continuous Time Markov chain models for chemical reaction networks*. Springer.
- Anderson, D.F., & Wolf, E.S. (2012). A finite difference method for estimating second order parameter sensitivities of discrete stochastic chemical reaction networks. *Journal of Chemical Physics*, 137(22), 224112.
- Ball, K., Kurtz, T.G., Popovic, L., Rempala, G. (2006). Asymptotic analysis of multiscale approximations to reaction networks. *Annals of Applied Probability*, 16(4), 1925–1961.
- Bressloff, P.C. (2009). Stochastic neural field theory and the system-size expansion. *SIAM Journal on Applied Mathematics*, 70(5), 1488–1521.
- Bressloff, P.C., & Newby, J.M. (2013). Metastability in a stochastic neural network modeled as a velocity jump Markov process. *SIAM Journal on Applied Dynamical Systems*, 12(3), 1394–1435.
- Bressloff, P.C., & Newby, J.M. (2014a). Stochastic hybrid model of spontaneous dendritic NMDA spikes. *Physical Biology*, 11(1), 016006.
- Bressloff, P.C., & Newby, J.M. (2014b). Path integrals and large deviations in stochastic hybrid systems. *Physical Review E*, 89(4), 042701.
- Buckwar, E., & Riedler, M.G. (2011). An exact stochastic hybrid model of excitable membranes including spatio-temporal evolution. *Journal of Mathematical Biology*, 63, 1051–1093.
- Buice, M.A., Cowan, J.D., Chow, C.C. (2010). Systematic fluctuation expansion for neural network activity equations. *Neural Computation*, 22(2), 377–426.
- Cao, Y., Gillespie, D.T., Petzold, L.R. (2006). Efficient step size selection for the tau-leaping simulation method. *Journal of Chemical Physics*, 124(4), 044109.
- Clay, J.R., & DeFelice, L.J. (1983). Relationship between membrane excitability and single channel open-close kinetics. *Biophysical Journal*, 42(2), 151–7.
- Colquhoun, D., & Hawkes, A.G. (1983). *Single-channel recording, chapter the principles of the stochastic interpretation of ion-channel mechanisms*. New York: Plenum Press.
- Davis, M.H.A. (1984). Piecewise-deterministic markov processes: a general class of non-diffusion stochastic models. *Journal of the Royal Statistical Society. Series B*, 46(3), 353–388.
- Dorval Jr., A.D., & White, J.A. (2005). Channel noise is essential for perithreshold oscillations in entorhinal stellate neurons. *The Journal of Neuroscience*, 25(43), 10025–10028.
- Earnshaw, B.A., & Keener, J.P. (2010). Invariant manifolds of binomial-like nonautonomous master equations. *SIAM Journal Applied Dynamical Systems*, 9(2), 568–588.
- Ermentrout, G.B., & Terman, D.H. (2010). *Foundations of mathematical neuroscience*. Springer.
- Ethier, S.N., & Kurtz, T.G. (1986). *Markov processes: characterization and convergence*. New York: John Wiley.
- Fisch, K., Schwalger, T., Lindner, B., Herz, A.V.M., Benda, J. (2012). Channel noise from both slow adaptation currents and fast currents is required to explain spike-response variability in a sensory neuron. *Journal of Neuroscience*, 32(48), 17332–44.
- Fox, R.F., & Yan-nan, L. (1994). Emergent collective behavior in large numbers of globally coupled independently stochastic ion channels. *Physical Review E Statistical Physics Plasmas Fluids Related Interdisciplinary Topics*, 49(4), 3421–3431.
- Giles, M.B. (2008). Multilevel Monte Carlo path simulation. *Operations Research*, 56, 607–617.
- Gillespie, D.T. (1977). Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81, 2340–2361.
- Gillespie, D.T. (2007). Stochastic simulation of chemical kinetics. *Annual Review of Physical Chemistry*, 58, 35–55.
- Glynn, P.W. (1989). A GSMP formalism for discrete event systems. *Proceedings of the IEEE*, 77(1), 14–23.
- Goldwyn, J.H., Imennov, N.S., Famulare, M., Shea-Brown, E. (2011). Stochastic differential equation models for ion channel noise in Hodgkin-Huxley neurons. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, 83(4 Pt 1), 041908.
- Goldwyn, J.H., & Shea-Brown, E. (2011). The what and where of adding channel noise to the Hodgkin-Huxley equations. *PLoS Computational Biology*, 7(11), e1002247.
- Groff, J.R., DeRemigio, H., Smith, G.D. (2009) Markov chain models of ion channels and calcium release sites. In *Stochastic methods in neuroscience* (pp. 29–64).
- Haas, P.J. (2002). *Stochastic petri nets: modelling stability, simulation*, 1st edn. New York: Springer.
- Hodgkin, A.L., & Huxley, A.F. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *Journal of Physiology*, 117, 500–544.
- Kang, H.W., & Kurtz, T.G. (2013). Separation of time-scales and model reduction for stochastic reaction models. *Annals of Applied Probability*, 23, 529–583.
- Keener, J.P., & Newby, J.M. (2011). Perturbation analysis of spontaneous action potential initiation by stochastic ion channels. *Physical Review E, Statistical, Nonlinear, and Soft Matter Physics*, 84(1-1), 011918.
- Kispersky, T., & White, J.A. (2008). Stochastic models of ion channel gating. *Scholarpedia*, 3(1), 1327.
- Kurtz, T.G. (1980). Representations of markov processes as multiparameter time changes. *Annals of Probability*, 8(4), 682–715.
- Kurtz, T.G. (1981). Approximation of population processes, CBMS-NSF Reg. Conf. Series in Appl. Math.: 36, SIAM.
- Laing, C., & Lord, G.J. (Eds.) (2010). *Stochastic methods in neuroscience*. Oxford University Press.
- Lee, C., & Othmer, H. (2010). A multi-time-scale analysis of chemical reaction networks: I. deterministic systems. *Journal of Mathematical Biology*, 60, 387–450. doi:10.1007/s00285-009-0269-4.
- Milescu, L.S., Yamanishi, T., Ptak, K., Smith, J.C. (2010). Kinetic properties and functional dynamics of sodium channels during repetitive spiking in a slow pacemaker neuron. *Journal of Neuroscience*, 30(36), 12113–27.
- Mino, H., Rubinstein, J.T., White, J.A. (2002). Comparison of algorithms for the simulation of action potentials with stochastic sodium channels. *Annals of Biomedical Engineering*, 30(4), 578–87.
- Morris, C., & Lecar, H. (1981). Voltage oscillations in the barnacle giant muscle fiber. *Biophysical Journal*, 35(1), 193–213.
- Newby, J.M., Bressloff, P.C., Keener, J.P. (2013). Breakdown of fast-slow analysis in an excitable system with channel noise. *Physical Review Letters*, 111(12), 128101.
- Pakdaman, K., Thieullen, M., Wainrib, G. (2010). Fluid limit theorems for stochastic hybrid systems with application to neuron models. *Advances in Applied Probability*, 42(3), 761–794.
- Pakdaman, K., Thieullen, M., Wainrib, G. (2012). Asymptotic expansion and central limit theorem for multiscale piecewise-deterministic Markov processes. *Stochastic Proceedings of Applied*, 122, 2292–2318.
- Rathinam, M., Sheppard, P.W., Khammash, M. (2010). Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks. *Journal of Chemical Physics*, 132, 034103.

- Riedler, M., & Notarangelo, G. (2013). Strong Error Analysis for the Θ -Method for Stochastic Hybrid Systems arXiv preprint arXiv:1310.0392.
- Riedler, M.G., Thieullen, M., Wainrib, G. (2012). Limit theorems for infinite-dimensional piecewise deterministic Markov processes. Applications to stochastic excitable membrane models. *Electronic Journal of Probability*, 17(55), 1–48.
- Rinzel, J., & Ermentrout, G.B. (1989). Analysis of neural excitability and oscillations. In C. Koch, & I. Segev (Eds.), *Methods in neuronal modeling*, 2nd edn. MIT Press.
- Terman, D., & Rubin, J. (2002). Geometric singular perturbation analysis of neuronal dynamics. In B. Fiedler (Ed.), *Handbook of dynamical systems, vol. 2: towards applications* (pp. 93–146). Elsevier.
- Schmandt, N.T., & Galán, R.F. (2012). Stochastic-shielding approximation of Markov chains and its application to efficiently simulate random ion-channel gating. *Physical Review Letters*, 109(11), 118101.
- Schmidt, D.R., & Thomas, P.J. (2014). Measuring edge importance: a quantitative analysis of the stochastic shielding approximation for random processes on graphs. *The Journal of Mathematical Neuroscience*, 4(1), 6.
- Schwalger, T., Fisch, K., Benda, J.an., Lindner, B. (2010). How noisy adaptation of neurons shapes interspike interval histograms and correlations. *PLoS Computers in Biology*, 6(12), e1001026.
- Shingai, R., & Quandt, F.N. (1986). Single inward rectifier channels in horizontal cells. *Brain Research*, 369(1-2), 65–74.
- Skaugen, E., & Walløe, L. (1979). Firing behaviour in a stochastic nerve membrane model based upon the Hodgkin-Huxley equations. *Acta Physiologica Scandinavica*, 107(4), 343–63.
- Smith, G.D., & Keizer, J. (2002). Modeling the stochastic gating of ion channels. In *Computational cell biology* (pp. 285–319). New York: Springer.
- Srivastava, R., Anderson, D.F., Rawlings, J.B. (2013). Comparison of finite difference based methods to obtain sensitivities of stochastic chemical kinetic models. *Journal of Chemical Physics*, 138(7), 074110.
- Strassberg, A.F., & DeFelice, L.J. (1993). Limitations of the Hodgkin-Huxley formalism: Effects of single channel kinetics on transmembrane voltage dynamics. *Neural Computation*, 5, 843–855.
- Wainrib, G., Thieullen, M., Pakdaman, K. (2012). Reduction of stochastic conductance-based neuron models with time-scales separation. *Journal of Computational Neuroscience*, 32, 327–346.
- White, J.A., Chow, C.C., Ritt, J., Soto-Trevino, C., Kopell, N. (1998). Synchronization and oscillatory dynamics in heterogeneous, mutually inhibited neurons. *Journal of Computational Neuroscience*, 5, 5–16.
- White, J.A., Rubinstein, J.T., Kay, A.R. (2000). Channel noise in neurons. *Trends in Neurosciences*, 23, 131–137.
- Wilkinson, D.J. (2011). *Stochastic modelling for systems biology*. Chapman & Hall/CRC.