J Comput Neurosci (2006) 21:71–87
DOI 10.1007/s10827-006-7200-4

The most likely voltage path and large deviations approximations for integrate-and-fire neurons

Liam Paninski

Received: 29 October 2004 / Revised: 13 December 2005 / Accepted: 13 January 2006 / Published online: 22 April 2006

Abstract We develop theory and numerical methods for computing the most likely subthreshold voltage path of a noisy integrate-and-fire (IF) neuron, given observations of the neuron's superthreshold spiking activity. This optimal voltage path satisfies a second-order ordinary differential (Euler-Lagrange) equation which may be solved analytically in a number of special cases, and which may be solved numerically in general via a simple "shooting" algorithm. Our results are applicable for both linear and nonlinear subthreshold dynamics, and in certain cases may be extended to correlated subthreshold noise sources. We also show how this optimal voltage may be used to obtain approximations to (1) the likelihood that an IF cell with a given set of parameters was responsible for the observed spike train; and (2) the instantaneous firing rate and interspike interval distribution of a given noisy IF cell. The latter probability approximations are based on the classical Freidlin-Wentzell theory of large deviations principles for stochastic differential equations. We close by comparing this most likely voltage path to the true observed subthreshold voltage trace in a case when intracellular voltage recordings are available *in vitro*.

Keywords Likelihood · Stochastic dynamics · Freidlin-Wentzell · Calculus of variations · Intracellular recordings

Action Editor: Peter Latham

L. Paninski (✉)
Department of Statistics, Columbia University
e-mail: liam@stat.columbia.edu;
http://www.stat.columbia.edu/∼liam

1. Introduction

A classic and recurring problem in theoretical neuroscience is to estimate the probability that a noisy integrate-and-fire-type (IF-type) neuronal model, which has fired at time $t = 0$, will not fire again until time $t = T$. This problem appears in a number of contexts, including firing rate computations (Brunel and Latham, 2003), statistical model fitting (Paninski et al., 2004b), and decoding (Brown et al., 2004).

According to the theory of large deviations (Freidlin and Wentzell, 1984; Dembo and Zeitouni, 1993), as the noise level σ tends to zero, this probability may be asymptotically approximated (on a logarithmic scale) by simply considering the likelihood of the "most likely" (ML) voltage path $V(t)$ satisfying the constraints that (1) the voltage at time 0 begins at a given reset value ($V(0) = V_{\text{reset}}$); (2) the voltage at time T reaches a threshold value ($V(T) = V_{\text{th}}$); and (3) the voltage does not cross threshold before time T ($V(t) \leq V_{\text{th}}, 0 < t < T$; note, however, that the voltage is allowed to touch threshold, $V(t) = V_{\text{th}}$, for times between 0 and T). Thus by computing the likelihood of the ML path we may obtain a useful approximation for the firing probability of interest.

Of course, this ML path is of independent interest. The likelihood of different voltage paths is of intrinsic importance for understanding the subthreshold encoding of information, for "looking under the membrane" given the spike train (Fig. 1). The ML path may also be used as an estimate of the subthreshold (unobserved) voltage $V(t)$ given the spike train.

We focus on obtaining this most likely (ML) path here. Despite the huge literature on large deviations principles for stochastic physical and biological systems (Freidlin

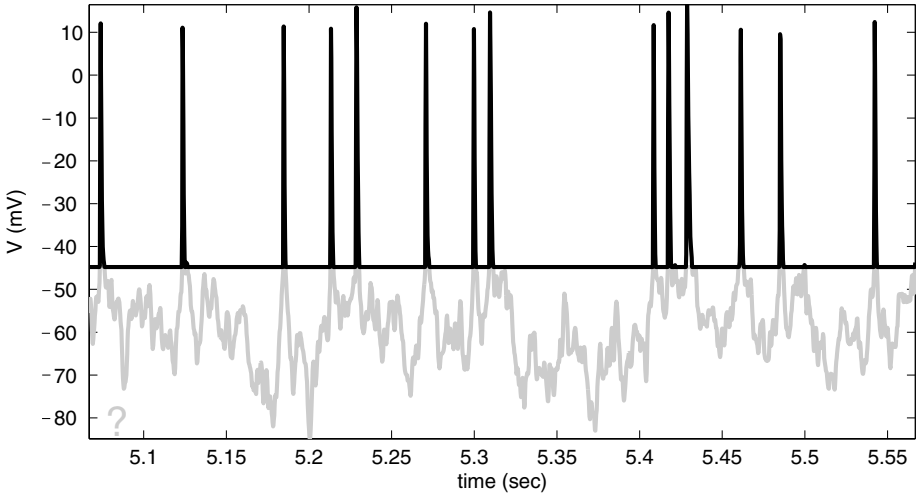

Fig. 1 Illustration of the problem: a spike train (black) is given as data, and our task is to infer the underlying subthreshold (unobserved) voltage fluctuations (gray). C.f. Fig. 10, below, illustrating an *in vitro* application of the same idea

and Wentzell, 1984; Dembo and Zeitouni, 1993), and the growing body of work on the effects of noise on neural activity and on likelihood-based approaches for fitting statistical models to neural data, to our knowledge this problem has not previously been investigated in depth, with the exception of the independent recent work of Badel et al. (2005), where the spike-triggered average of a linear integrate-and-fire-type cell was studied in the low-noise limit. Thus we give full details of the derivations of the ML path here, despite the fact that the techniques (Kuhn-Tucker conditions, Euler-Lagrange equation) are fairly standard from a mathematical point of view.

In Section 2 we give our main results on the identity of the ML path, including a key differential equation that determines this path. We apply these results to a variety of models (linear and nonlinear membrane dynamics, white and correlated noise, constant and time-varying conductance and rest potential) in Section 3. We discuss the corresponding probability approximations in Section 4 We close with a brief application to predicting subthreshold physiological voltage traces (recorded *in vitro*) from spike train data in Section 5. An appendix gives details of computational methods for sampling from, computing conditional densities of, and maximizing likelihoods for the hidden state of a hidden Markov model (HMM), of which the integrate-and-fire model may be considered a special case.

2. Main results

The noisy IF-type neuron behaves according to the stochastic differential equation

$$dV(t) = f(V(t), t)dt + \sigma N_t, \tag{1}$$

with $f(V(t), t)$ denoting the (possibly nonlinear and inhomogeneous) deterministic dynamics of the model, and N_t denoting a stochastic noise source multiplied by the scale factor σ. Note that, for simplicity, we have restricted our attention to dynamics which are one-dimensional in V, at the cost of limiting the dynamical repertoire of the model. (Generalizations of the results below to the case of multidimensional subthreshold dynamics are possible, but are not pursued here.) We assume throughout that the function f is smooth; in the following, we will use $V_0(t)$ to denote the (unique) deterministic ($\sigma = 0$) solution to the above equation given the initial condition $V(0)$.

As discussed above, we will consider the conditional dynamics of a noisy IF cell which has been observed to spike at time $t = 0$ and $t = T$. For simplicity, we set the reset and threshold potential to 0 and 1, respectively; this entails no loss of generality, by the usual rescaling of the V axis. Thus we will constrain voltage paths from Eq. (1) to satisfy $V(0) = 0, V(T) = 1, V(t) \leq 1, 0 < t < T$.

Under standard white Gaussian noise N_t, we may formally write the likelihood of a given path $V(t)$ as the likelihood of the noise trace which generated $V(t)$, that is,

$$\text{“likelihood}\,(\{V(t)\}) = \frac{1}{Z} e^{-\frac{1}{2}\int_0^T N(t)^2 dt} = \frac{1}{Z} e^{-\frac{1}{2\sigma^2}\int_0^T \left(\dot{V}(t) - f(V(t),t)\right)^2 dt},''$$

with $\dot{V}(t)$ denoting the formal time derivative of $V(t)$ (recall from Eq. (1) that $\dot{V}(t)$ differs from $f(V(t), t)$ by exactly $\sigma\, N(t)$). We have put the above expression in quotes to emphasize its formal nature; $\dot{V}$ exists only in the sense of generalized functions (Hida, 1980), (that is, $\dot{V}$ exists as a regular function with probability zero). However, this formula may be justified either by taking continuous limits of the corresponding formula in discrete time, or by an appeal to Girsanov's formula (Karatzas and Shreve,

1997), and therefore it is reasonable to define the most likely path $V_{\mathrm{ML}}(t)$ as the path which minimizes the rescaled cost

$$L_0(\{V(t)\}) \equiv \int_0^T (\dot{V}(t) - f(V(t),t))^2 dt.$$

To enforce the constraints $V(0) = 0$ and $V(T) = 1$, we will solve $V(t)$ minimizing

$$L(\{V(t)\}) = \int_0^T \left(\dot{V}(t) - f(V(t),t)\right)^2 dt + \lambda_0 V(0)^2 + \lambda_T (V(T)-1)^2,$$

with λ_i appropriate multipliers. It might be helpful to think of the above problem in the language of optimal control theory: we want to find some input $N_{\mathrm{ML}}(t)$, such that if we define $V_{\mathrm{ML}}(t)$ to be the solution of the perturbed (noiseless) ODE

$$\frac{\partial V_{\mathrm{ML}}(t)}{\partial t} = f(V_{\mathrm{ML}}(t),t) + N_{\mathrm{ML}}(t),$$

then $V_{\mathrm{ML}}(t)$ satisfies the threshold constraints and $N_{\mathrm{ML}}(t)$ minimizes the cost function

$$\|N\|_2^2 = \int_0^T N(t)^2 dt \sim L_0(\{V\}).$$

When f is affine in V (that is, the linear leaky integrate-and-fire case, with possibly time-varying leak and rest potential, $f(V,t) = -g(t)\,V(t) + I(t)$), L will be strictly convex in $N(t)$, and therefore also strictly convex in paths $V(t)$ with fixed endpoints $V(0) = 0$ and $V(T) = 1$; since the constraint space of allowed voltage paths $V(t)$ is convex, the minimum will be unique in the class of constraint functions $V(t)$. This uniqueness does not need to hold in general; thus the condition we derive below will be necessary but not in general sufficient for a minimum to hold.[1]

The above minimization problem is a standard problem in the calculus of variations, which we will review here briefly for those readers unfamiliar with the theory. The usual Kuhn-Tucker (KT) conditions in this case are that the functional gradient of $L(\{V\})$ with respect to $V(t)$ should either be: zero, in the case that none of the inequality constraints are met exactly at the minimizer; or, more generally, contained in the vector cone spanned by the negative of the vectors defining the inequality constraints that are met exactly. Here a cone spanned by a set S of vectors $\{\vec{x}(t)\}_{t\in S}$ is defined simply as the set of all nonnegative combinations of these vectors,

$$\left\{\vec{y} : \vec{y} = \sum_{t\in S} \eta(t)\vec{x}(t);\ \eta(t) \geq 0\ \forall t\right\};$$

the vectors $\vec{x}(t)$ here are the negatives of the vectors defining the inequality constraints $V(t) < 1, 0 < t < T$, and thus have the simple form

$$\vec{x}(t) = -\delta_t,$$

with δ_t denoting the delta function centered at time t. The set S, again, is the set of all times t where the constraints are met exactly; if this set S is empty, this cone contains just the point $\vec{0}$ (i.e., the functional gradient of $L(\{V\})$ is zero). These KT conditions are illustrated with a simple two-dimensional example in Fig. 2.

Rewriting the above sum $\sum_t \eta(t)\vec{x}(t)$ as an integral and expressing the gradient in terms of the directional derivative

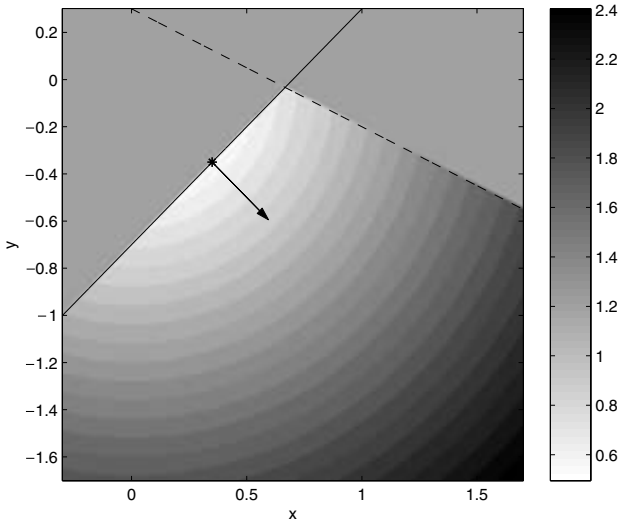

Fig. 2 Illustration of the KT conditions in two dimensions. The function we are trying to minimize is the two-norm, $(x^2 + y^2)^{1/2}$, under two linear inequality constraints, as indicated with the two diagonal black lines. The gray area indicates the disallowed region (on which the inequality constraints are not satisfied). In this case, only one constraint (indicated by a solid line) is relevant; since at the minimizer of the objective function on the allowed space (indicated by the black asterisk) the upper right inequality (dashed trace) holds strictly (whereas the inequality on the upper left becomes a strict equality) the gradient (black arrow) must lie within the cone spanned by the negative of the vector defining the upper left inequality. In this case, since only one constraint is met exactly, this cone is just the one-dimensional ray orthogonal to the constraint surface, as shown

[1] Existence of this minimizer, on the other hand, may be guaranteed under Lipschitz continuity of f and its space and time derivatives (c.f. Eq. (3), via standard compactness and lower semicontinuity arguments from optimal control (Vinter, 2000).

in the direction of an arbitrary vector $U(t)$, this means that the optimal path $V_{\mathrm{ML}}(t)$ must satisfy

$$\begin{aligned} L(\{V_{\mathrm{ML}} + \epsilon U\}) &= L(\{V_{\mathrm{ML}}\}) + \epsilon \left(\sum_{t \in S} \eta_{V_{\mathrm{ML}}}(t)(-\delta_t) \right) \\ &\quad \times U + o(\epsilon),\ \epsilon \to 0 \\ &= L(\{V_{\mathrm{ML}}\}) - \epsilon \int_0^T U(t) d\eta_{V_{\mathrm{ML}}}(t) \\ &\quad + o(\epsilon),\ \epsilon \to 0 \end{aligned} \tag{2}$$

for all nice functions $U(t)$, where η_{VML} is a positive measure with support on S, the set of points t where $V_{\mathrm{ML}}(t) = 1$, i.e., those points where the constraints $V_{\mathrm{ML}}(t) \leq 1$ are met exactly.

We write $L(\{V + \epsilon U\})$ to first order as

$$\begin{aligned} L(\{V + \epsilon U\}) &= \int_0^T (\dot{V}(t) + \epsilon \dot{U}(t) - f(V(t), t) \\ &\quad -\epsilon f_V(V(t), t) U(t))^2 dt \\ &\quad +\lambda_0 (V(0) + \epsilon U(0))^2 + \lambda_T (V(T) \\ &\quad +\epsilon U(T) - 1)^2 + o(\epsilon), \end{aligned}$$

with f_V denoting $\frac{\partial f}{\partial V}$; collecting terms of order 2ϵ, we have

$$\begin{aligned} L(\{V + \epsilon U\}) &= L(\{V\}) \\ &\quad +2\epsilon \int_0^T (\dot{V}(t)\dot{U}(t) - f(V(t), t)\dot{U}(t) \\ &\quad -f_V(V(t), t)\dot{V}(t)U(t) \\ &\quad +f(V(t), t) f_V(V(t), t) U(t)) dt \\ &\quad +2\epsilon(\lambda_0 V(0) U(0) + \lambda_T U(T)(V(T) - 1)) \\ &\quad +o(\epsilon). \end{aligned}$$

Integrating the second line by parts gives

$$\begin{aligned} L(\{V + \epsilon U\}) &= L(\{V\}) \\ &\quad +2\epsilon \int_0^T U(t)(-\ddot{V}(t) + f_t(V(t), t) \\ &\quad +f(V(t), t) f_V(V(t), t)) dt \\ &\quad +2\epsilon [U(0)(f(V(0), 0) - \dot{V}(0) + \lambda_0 V(0)) \\ &\quad +U(T)(\dot{V}(T) - f(V(T), T) \\ &\quad +\lambda_T (V(T) - 1))] \\ &\quad +o(\epsilon). \end{aligned}$$

Now for the KT conditions to hold for all λ and nice $U(t)$, the ϵ terms in the above expression must have the special dot-product form given in the $O(\epsilon)$ term in Eq. (2). This, in turn, implies that

$$\begin{aligned} \eta_{VML} \equiv -2(&-\ddot{V}_{\mathrm{ML}}(t) + f_t(V_{\mathrm{ML}}(t), t) \\ &+f(V_{\mathrm{ML}}(t), t) f_V(V_{\mathrm{ML}}(t), t)) \end{aligned}$$

is a positive measure with support on the points t where $V_{\mathrm{ML}}(t) = 1$. In other words, $V_{\mathrm{ML}}(t)$ must satisfy the Euler-Lagrange differential equation

$$-\ddot{V}_{\mathrm{ML}}(t) + f_t(V_{\mathrm{ML}}(t), t) + f(V_{\mathrm{ML}}(t), t) f_V(V_{\mathrm{ML}}(t), t) \begin{cases} =0 & V_{\mathrm{ML}}(t) \leq 1 \\ \leq 0 & V_{\mathrm{ML}}(t) = 1 \end{cases}, \tag{3}$$

with initial condition $V_{\mathrm{ML}}(0) = 0$ and end condition $V_{\mathrm{ML}}(T) = 1$ (Fig. 3); this equation will be key to the rest of our development. See also Freidlin and Wentzell(1984) for further discussion of similar variational problems related to first-passage times.

Some discussion of this equation is in order. First, note that if $V_0(T) = 1$ and $V_0(t) \leq 1$ for $t < T$, with $V_0(t)$ the solution to the original IF-type differential Eq. (1) under

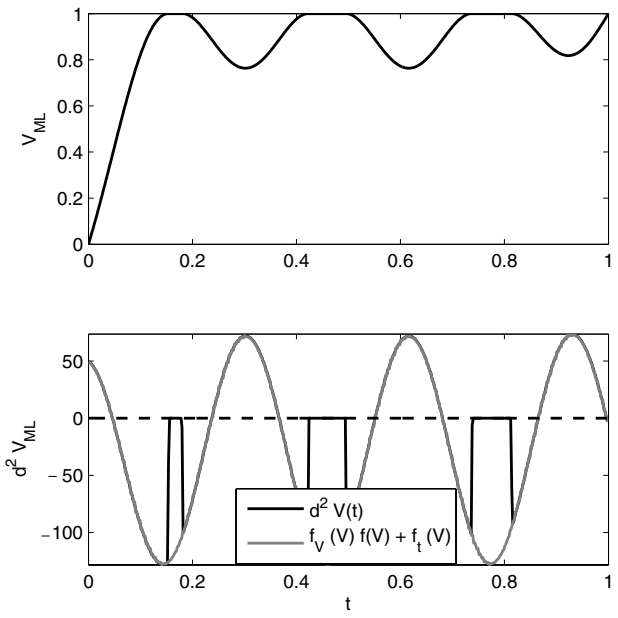

Fig. 3 Illustration of solution $V_{\mathrm{ML}}(t)$ to the differential Eq. (3). Top: Solution for $V_{\mathrm{ML}}(t)$, given a linear, constant-conductance neuron ($f(V, t) = -gV + I(t)$) and sinusoidally varying input current $I(t)$; recall that this solution is unique here, by the strict convexity of $L(\{V\})$ under linear dynamics $f(V, t)$. Note that the solution meets the voltage threshold $V = 1$ several times; in each case, close inspection reveals that $\dot{V}_{\mathrm{ML}}(t) = 0$ for all t such that $V_{\mathrm{ML}}(t) = 1$. Bottom: Comparison of the terms on the left-hand side of Eq. (3): $\ddot{V}_{\mathrm{ML}}(t)$ vs. $f_V(V_{\mathrm{ML}}(t)) f(V_{\mathrm{ML}}(t)) + f_t(V_{\mathrm{ML}}(t)) = g^2 V_{\mathrm{ML}}(t) - gI(t) + \dot{I}(t)$. Note that, as predicted, these two terms are equal except for t such that $V_{\mathrm{ML}}(t) = 1$, in which case $\ddot{V}_{\mathrm{ML}}(t)$ jumps to zero and the inequality in (3) holds strictly

zero noise with initial condition $V(0) = 0$, then $V_0(t)$ automatically solves Eq. (3), and is clearly the ML voltage path. This can be seen by simply setting σ to zero in equation (1) and differentiating both sides with respect to time t; of course, this is as we would expect, since $L(\{V_0\}) = 0$ and is therefore uniquely minimized in this case. It is also worth emphasizing that the ML solution $V_{\rm ML}(t)$ is independent of σ (since σ appears in neither (3) nor the definition of $L(\{V\})$), despite the fact that conditional average voltages given spike times (e.g., spike-triggered averages) do depend on σ (see Badel et al., 2005; Paninski, 2005 and Fig. 5 ahead).

Second, the solutions to Eq. (3) have some interesting features. For example, assuming $f(V, t)$ is continuous, if the solution $V_{\rm ML}(t)$ reaches the threshold $V = 1$ at some time $t < T$, then $V_{\rm ML}(t)$ must reach the threshold smoothly, with $\dot{V}_{\rm ML}(t) = 0$. Otherwise, $\ddot{V}_{\rm ML}(t)$ would have a negative singularity (corresponding to a downward kink in $V_{\rm ML}(t)$, with $\dot{V}_{\rm ML}(t)$ jumping discontinuously from a positive value to zero), contradicting the inequality in (3). Similarly, $\dot{V}_{\rm ML}(t)$ is constrained to be zero when $V_{\rm ML}(t)$ is leaving threshold for times $t < T$. See Fig. 3 for an illustration of these effects.

These features place certain constraints on the solutions to Eq. (3) which are useful when developing algorithms to solve this equation. These algorithms have a "shooting" flavor: as usual with a second-order ODE, we are free to choose the first derivative at time $t = 0$, $\dot{V}(0)$. (Recall that $V_{\rm ML}(0)$ is constrained to be zero, by assumption; in the case of an unconstrained second-order ODE, fixing $V(0)$ and $\dot{V}(0)$ would uniquely specify the solution, but in the present case $V_{\rm ML}(t)$ must be chosen to satisfy the threshold constraints $V_{\rm ML}(t) < 1$ as well.) Given these initial conditions, the solution evolves autonomously as time t increases, until either time runs out, $t = T$, or threshold is reached, $V(t_0) = 1$ for some time $t_0 < T$. In the first case, we can re-run the voltage path with a larger choice for $V(0)$ (to get closer to the end-threshold condition, $V(T) = 1$). In the second case, either $V(t_0) = 0$ or $\dot{V}(t_0) > 0$. The second case is not allowed, by the preceding discussion (thus we need to reduce the original $\dot{V}(0)$; therefore, assume $\dot{V}(t_0) = 0$. Now $\ddot{V}(t_0^+) = 0$ as well, so we need to check that the inequality in (3) is satisfied. If not, we again need to go back and reduce $\dot{V}(0)$; if the inequality is satisfied, on the other hand, then we need to choose the time t_1, with $t_0 < t_1 < t_2$, that V will leave $V = 1$ again. (Here t_2 is defined as the first time after t_0 that the inequality in (3) fails, that is, $f_t(V(t), t) + f(V(t), t) f_v(V(t), t) > 0$; recall that $\ddot{V}(t) = 0$ for all $t_0 < t < t_1$.) We should emphasize that $\dot{V}(t_1) = 0$, that is, $V(t)$ again evolves autonomously (with no spare degrees of freedom) once t_1 is chosen, and the process may be iterated unti $t = T$.

3. Examples

In this section we provide a number of applications of the above results, to obtain exact solutions to the optimal path $V_{\rm ML}(t)$ under a variety of subthreshold dynamics $f(V, t)$. We also describe some alternate methods for computing $V_{\rm ML}(t)$; these methods can lead to much simpler solutions in certain special cases.

3.1. Linear dynamics

The most important special case is the linear homogeneous case (the usual linear leaky integrate-and-fire cell),

$$f(V(t), t) = -gV.$$

Equation (3) gives

$$\begin{aligned}\ddot{V}_{\rm ML}(t) &= f_t(V_{\rm ML}(t), t) + f(V_{\rm ML}(t), t) f_V(V_{\rm ML}(t), t)\\ &= 0 + (-gV_{\rm ML}(t))(-g)\\ &= g^2 V_{\rm ML}(t),\end{aligned}$$

a linear homogeneous second-order ODE which is solved in general by $ae^{gt} + be^{-gt}$, for a, b arbitrary constants; applying the endpoint conditions gives

$$V_{\rm ML}(t) = \frac{1}{e^{gT} - e^{-gT}}(e^{gt} - e^{-gt}).$$

In the zero-leak case $g \to 0$, this reduces to the simple linear ramp solution,

$$V_{\rm ML}(t) = \frac{t}{T};$$

see Fig. 4 for a few examples.

Alternatively, this linear case may be solved via a maximum principle. First, writing

$$V(t) = e^{-gt} * N(t)$$

as the usual (unique) solution to the linear equation $\dot{V} = -gV + N(t)$, note that our original problem may be restated in this case as: find $N(t)$ to minimize $\|N\|_2$ under the linear constraint that $\langle k, N\rangle = 1$, with the linear convolution functional $\langle k, N\rangle$ defined as

$$\begin{aligned}\langle k, N\rangle &\equiv e^{-gt} * N(t)\Big|_{t=T}\\ &= \int_0^T e^{-g(T-t)} N(t) dt\\ &= e^{-gT} \int_0^T e^{gt} N(t) dt.\end{aligned}$$

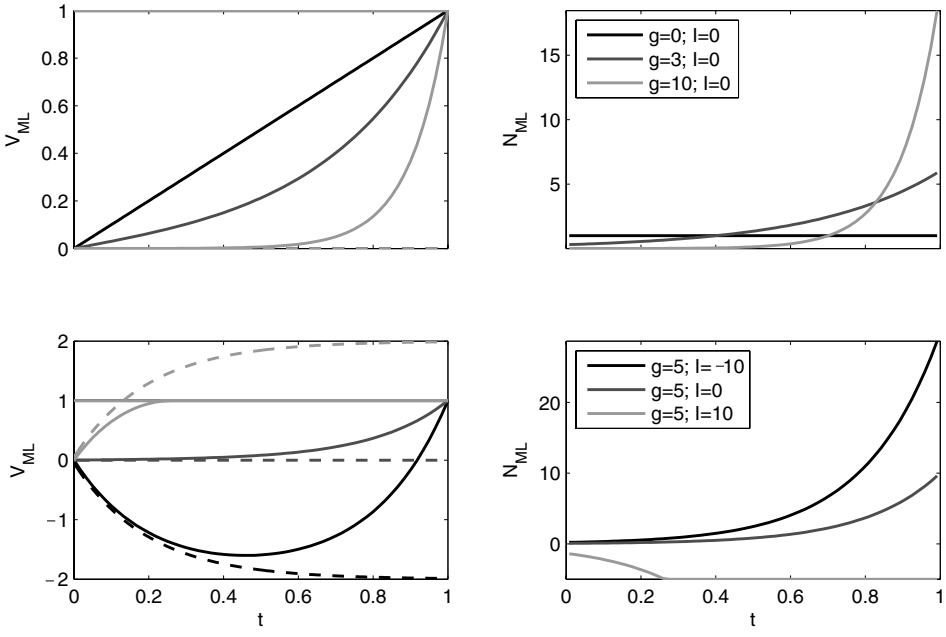

Fig. 4 Illustration of $V_{\mathrm{ML}}(t)$ for linear cell with constant input and constant conductance, $f(V, t) = -gV + I$. Top: Solutions $V_{\mathrm{ML}}(t)$ for $I = 0$ and g varying; $V_{\mathrm{ML}}(t)$ becomes more sharply convex and concentrated at $t \approx T$ as g increases. The right panel shows the corresponding optimal $N_{ML}(t) = \dot{V}_{ML}(t) - f(V_{\mathrm{ML}}(t))$; for $g = 0$, N_{ML} is constant, while for positive values of g, N_{ML} grows exponentially with time. Bottom: Solutions $V_{\mathrm{ML}}(t)$ for g fixed and varying I; dashed lines show corresponding unconditioned noiseless traces $V_0(t)$. $V_{\mathrm{ML}}(t)$ becomes more convex or concave depending on whether I is large negative or positive, respectively. In particular, for I sufficiently large and positive, $V_{\mathrm{ML}}(t)$ meets the threshold at $t_0 < T$ and remains at $V = 1$ for the remaining time $t_0 \le t \le T$; note that in this case $V_{\mathrm{ML}}(t)$ grows more slowly than $V_0(t)$, in order to satisfy $\dot{V}_{\mathrm{ML}}(t_0) = 0$

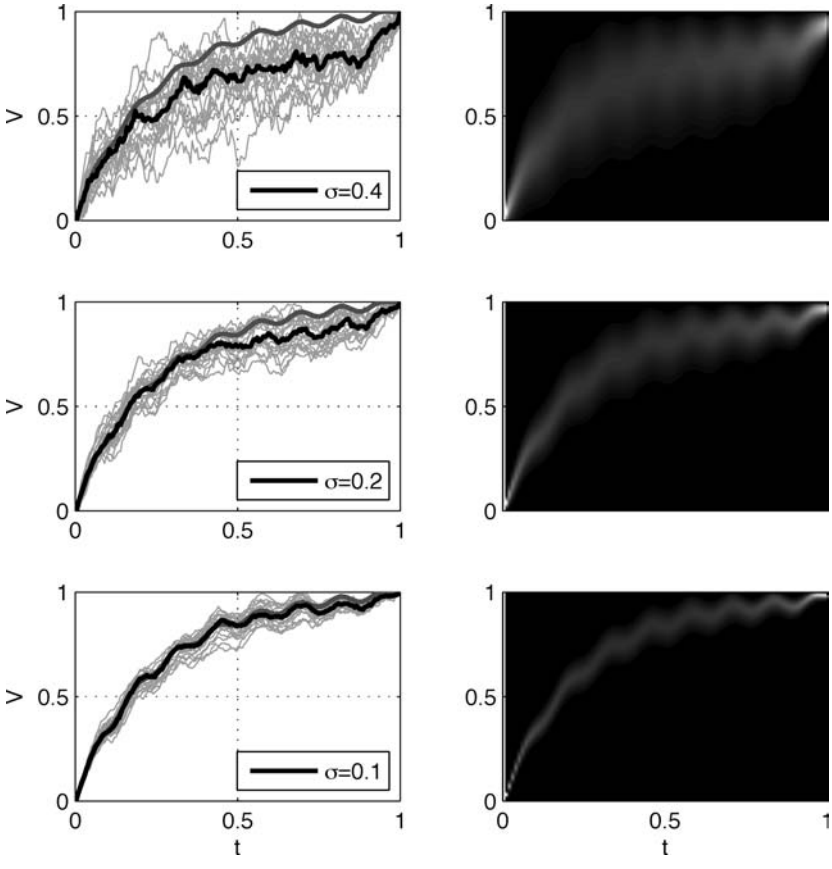

Fig. 5 Samples from the conditional distributions $p(V(t)|\text{spike at } 0, 1)$, for different values of σ. Left: Samples from conditional distribution, for the indicated values of σ. Model cell was linear with fixed conductance g, driven by a sinusoidal (8 Hz) input current $I(t)$. Bold traces indicate $V_{\mathrm{ML}}(t)$ (gray) and the conditional empirical median path (black). Note that the ML and conditional mean paths converge as $\sigma \to 0$, as expected (c.f. section 4.2. Right: Corresponding exact conditional distributions $p(V(t)|\text{spikeat } 0,1)$ (see appendix for details on the computational method for obtaining these distributions). The conditional distributions mirror the sample distributions shown on the left

As usual for this kind of linearly-constrained L_2-norm minimization problem, the $N_{\mathrm{ML}}(t)$ solving this condition is exactly proportional to the vector defining the constraint, in this case $N_{\mathrm{ML}}(t) = ck(t) = ce^{gt}$, for c chosen such that $\langle k, N_{\mathrm{ML}}\rangle = \langle k, ck\rangle = c\|k\|_2^2 = 1$ (the solution is again unique, by the strict convexity of the norm $\|\cdot\|_2$ and by the convexity of linear functionals); plugging into the convolution formula for $V(t)$ above we arrive at our original solution. (Note that the inequality constraints $V_{\mathrm{ML}}(t) \leq 1, 0 < t < T$, are satisfied automatically here.) For the no-leak case $g \to 0$, we obtain the common-sense solution that the optimal input noise current $N_{\mathrm{ML}}(t)$ is constant; for $g > 0$, $N_{\mathrm{ML}}(t)$ grows monotonically in t, in a sense "saving its effort" since $N(t)$ is lost to leak for t small. Interestingly, for g negative (modeling active subthreshold dynamics, with an unstable equilibrium at the rest potential $V = 0$) we get the same solution for $V_{\mathrm{ML}}(t)$, but the solution for $N_{\mathrm{ML}}(t)$ is time-reversed; the interpretation here is that $N_{\mathrm{ML}}(t)$ expends most of its effort to get $V_{\mathrm{ML}}(t)$ moving away from the equlibrium point $V = 0$, and then lets $V_{\mathrm{ML}}(t)$ glide towards threshold as the driving force $f(V(t)) = -gV(t)$ becomes larger.

Time-varying conductance

The case of time-varying $g(t)$ (for example, in cases where the conductance might be rapidly changing in the wake of an action potential (Stevens and Zador, 1998), $f(V, t) = -g(t)V(t)$, can be handled similarly. In this case, we have that

$$V(t) = \int_0^t e^{-\int_s^t g(u)du} N(s)ds;$$

thus, ignoring the inequality constraint $V(t) \leq 1$ for the moment, the optimal $N_{\mathrm{ML}}(t)$ is proportional to $e^{\int_T^t g(u)du}$, and the optimal $V_{\mathrm{ML}}(t)$ is obtained by plugging this $N_{\mathrm{ML}}(t)$ back into the above formula:

$$\begin{aligned} V_{\mathrm{ML}}(t) &= c\int_0^t e^{-\int_s^t g(u)du} e^{\int_T^s g(u)du} ds \\ &= ce^{-\int^t g(u)du} \int_0^t e^{2\int^s g(u)du} ds, \end{aligned}$$

with c chosen to satisfy $V_{\mathrm{ML}}(T) = 1$ ($V_{\mathrm{ML}}(0) = 0$ is satisfied automatically).

So far we have proceeded as if the equality in (3) holds strictly, and that the corresponding solutions to the differential equation satisfy the constraint that $V_{\mathrm{ML}}(t) \leq 1$ for all $0 < t < T$, given the endpoint constraints $V_{\mathrm{ML}}(0) = 0$ and $V_{\mathrm{ML}}(T) = 1$; for example, in the linear homogeneous case considered above, $V_{\mathrm{ML}}(t)$ increases monotonically up to the threshold, thus automatically satisfying the threshold inequality. However, it is not immediately clear that the above solution for arbitrary $g(t)$ necesarily remains subthreshold for $0 < t < T$, and in general we might have to consider the inequality constraints as well. We can establish that the equality constraints are satisfied strictly in at least two cases: for example, it is sufficient that either $\dot{V}_{\mathrm{ML}}(t) > 0$ or $\ddot{V}_{\mathrm{ML}}(t) \geq 0$ for all $V_{\mathrm{ML}}(t) > V^*$, for some $V^* < 1$, since either of these conditions prevents $V_{\mathrm{ML}}(t)$ from either "kissing" the threshold $V = 1$ (recall the discussion following Eq. (3)) or curving back down from above after a threshold crossing. Applying the latter condition to Eq. (3) can give general conditions for the analytic solution to be in the constraint set; we explore these conditions further in Section 3.2.

In the present case, we have that equality holds in (3) for all t if $g(t)^2 - \dot{g}(t)$ is nonnegative, since in this case $\ddot{V}_{\mathrm{ML}}(t) \geq 0$ whenever $V_{\mathrm{ML}}(t) > 0$. Second, computing the time derivative of the above formula for $V_{\mathrm{ML}}(t)$, we have

$$\dot{V}_{\mathrm{ML}}(t) = ce^{-\int^t g(u)du}\left(e^{2\int^t g(u)du} - g(t)\int_0^t e^{2\int^s g(u)du} ds\right);$$

from this, it is also clear that strict equality holds in (3) whenever $g(t) \leq 0$ for all t. Again, however, it is not altogether clear that the equality holds strictly for arbitrary $g(t)$, and in general the above expression for $V_{\mathrm{ML}}(t)$ might only apply to the strictly subthreshold portion of the optimal path.

Nonzero rest potential

Now the case $f(V(t), t) = -gV + I(t)$, in which the neuron is driven by some input current $I(t)$, is fairly straightforward. We again have a simple exponential formula for the subthreshold solution, $V(t) = e^{-gt} * [I(t) + N(t)]$; again, we want to choose N_{ML} as small as possible (that is, minimize $||N||_2$) while ensuring that $V_{\mathrm{ML}}(T) = 1$. As above, if the boundary conditions may be ignored, we can accomplish this by simply choosing $N_{\mathrm{ML}}(t) = ce^{gt}$, with c chosen to enforce the end constraint. The boundary conditions, in turn, may be ignored whenever $I(t)$ is such that $V_{\mathrm{ML}}(t)$ remains subthreshold until $t = T$; for example, it is sufficient (though not necessary) that $H(t) \equiv I(t) * e^{-gt}$ be monotonically increasing, with $H(T) \leq 1$.

The simplest example for which the boundary conditions come into play is as follows: let $I(t)$ be constant: $I(t) = a$. For simplicity, let $g = T = 1$. Then condition (3) reads

$$V(t) - a \begin{cases} = \ddot{V}_{\mathrm{ML}}(t) & V_{\mathrm{ML}}(t) < 1 \\ \leq \ddot{V}_{\mathrm{ML}}(t) = 0 & V_{\mathrm{ML}}(t) = 1. \end{cases}$$

Solutions to the equality above with end conditions $V_{\mathrm{ML}}(0) = 0$ and $V_{\mathrm{ML}}(T) = 1$ are of the form

$$V_{\mathrm{ML}}(t) = ce^t + de^{-t} + a,$$

with c and d satisfying the two-by-two system of linear equations

$$c + d = -a$$
$$ec + e^{-1}d = 1 - a.$$

For a sufficiently small, c and d may be chosen so that $V_{\mathrm{ML}}(t)$ satisfies $V_{\mathrm{ML}}(t) \leq 1$ for $t < T$, and is therefore a solution to our maximization problem. For larger a, however, $V_{\mathrm{ML}}(t)$ becomes strongly concave, implying that $V_{\mathrm{ML}}(t)$ of the above form with $V_{\mathrm{ML}}(T) = 1$ must approach 1 from above at T, and in turn that $V_{\mathrm{ML}}(t)$ must hit threshold for some $t_1 < T$ (with $\dot{V}_{\mathrm{ML}}(t_1) = 0$, as discussed above). For such a voltage path, $V_{\mathrm{ML}}(t) = 1$ for $t_1 \leq t \leq T$; see Fig. 4.

Time-correlated noise

The maximum-principle approach can be generalized easily to the case in which $N(t)$ is non-white, that is, temporally correlated. However, a different approach leads to a simpler solution in this case. We note first that in the absence of any constraints,[2] the random solutions $V(t)$ of Eq. (1) under linear dynamics, $f(V, t) = -g(t)\,V(t) + I(t)$, form a Gaussian process (Karlin and Taylor, 1981), with mean

$$V_0(t) = \int_0^t e^{-\int_s^t g(u)du} I(s)ds$$

—recall that this is the deterministic ($\sigma = 0$) solution to Eq. (1) for linear $f(V, t)$—and covariance

$$C = \sigma^2 C_N C_{\mathrm{OU}} C_N,$$

where C_N is the covariance function of the noise N_t (we assume N_t to have zero mean, with no loss of generality), and we have denoted the Ornstein-Uhlenbeck (OU) covariance

$$C_{\mathrm{OU}}(t_1, t_2) = e^{-\int_{t_1}^{t_2} g(u)du} \int_0^{t_1} e^{-2\int_s^{t_1} g(u)du} ds$$

(assuming $t_1 < t_2$, and with t_2 replacing t_1 in the opposite case). This covariance C may often be computed analytically, particularly in the case of stationary input noise, when $C_N(t_1, t_2) = c_N(|t_1 - t_2|)$ for some suitable real function c_N.

Now we apply the usual Gaussian identities: the conditional mean, given $V(T) = 1$, is

$$V(t) = V_0(t) + \frac{C(t, T)}{C(T, T)}(1 - V_0(T));$$

since the mean of a Gaussian is also the ML solution, the above formula automatically gives $V_{\mathrm{ML}}(t)$ if $V(t)$ satisfies $V(t) \leq 1$ for $0 < t < T$. Otherwise, as usual, the KT conditions must be applied.

General linear case: Quadratic programming

Of course, in all of the above cases, we can also numerically optimize $L(\{V\})$ (or a suitably modified version thereof in the case of correlated noise $N(t)$); for example, lsqnonneg.m in matlab is designed to solve exactly the kind of linearly-constrained quadratic optimization (quadratic programming) problem considered here. The projection of the analytic solution $V(t)$ onto the convex constraint set may be used as a starting point for such a numerical optimization; if this $V(t)$ is itself within the constraint set, then no further optimization is necessary.

3.2. Nonlinear dynamics

In the case of nonlinear dynamics $f(V)$ (for example, in cases where the conductance g depends on voltage), we must rely on Eq. (3). We review the method for solving second-order nonlinear differential equations here. We need to solve $\dot{V}_{\mathrm{ML}}(t) = h(V_{\mathrm{ML}}(t))$, for $h(V) = f_V(V)f(V)$ (for simplicity, we will restrict our attention below to the time-homogeneous case $f(V, t) = f(V)$. Let $W = dV/dt$ then as usual we have a matrix equality,

$$W = dV/dt$$
$$dW/dt = h(V).$$

Multiplying these equalities gives

$$\frac{dV}{dt}h(V) = \frac{dW}{dt}W;$$

multiplying both sides by dt and integrating gives

$$\int h(V)dV = \frac{1}{2}W^2 + c_1.$$

Denote the left-hand side as $H(V)$. Then

$$\frac{dV}{dt} = W = (2(H(V) - c_1))^{1/2}.$$

[2] In the following, we will assume that the noise is "reset" by the spike, that is, V_{t_1} is independent of V_{t_2} whenever a spike has been observed in the interval (t_1, t_2) (and hence the corresponding noise N_t is independent in this sense as well). The case in which N_t is not reset by the spike proceeds similarly, except we have to condition on $V(t)$ at all spike times simultaneously—not just the spikes at $t = 0$ and T—since past spike times will influence the current conditional noise distribution. See Paninski et al. (2004b), Section 6.4, and Fig. 11 ahead for two different illustrations of a similar coupled-noise effect.

Separating variables, we have that $V_{\mathrm{ML}}(t) = G^{-1}(t)$, where G^{-1} is the inverse of

$$t = G(V) \equiv c_2 + \int \frac{dV}{(2(H(V) - c_1))^{1/2}}.$$

By choosing the free variables c_1 and c_2 correctly, we may satisfy the end conditions $V_{\mathrm{ML}}(0) = 0$ and $V_{\mathrm{ML}}(T) = 1$. It turns out, for many important cases, that the inequality conditions $V_{\mathrm{ML}}(t) \leq 1,\ 0 < t < T$, are guaranteed automatically as well. For example, it is sufficient that the cell's dynamics are "active" in the sense that $f(V) f_V(V) \geq 0$ for $V > 1$, since this implies, with Eq. (3), that $\dot{V}$ must be nonnegative as well in this range, preventing any solution to the above equation from curving back down to $V = 1$ from above.

Exponential cell

We apply the above general method to a special case here. Let $f(V) = e^{aV}$, with $a > 0$. This is a simple model of the nonlinear subthreshold dynamics of a repetitively firing cell (see, e.g.,(Fourcaud and Brunel, 2002), though note that the model discussed there behaves differently for strongly hyperpolarized voltages). Then $h(V) = ae^{2aV}$, $H(V) = \frac{1}{2}e^{2aV}$, and we have (by 2.315 of Jeffrey and Zwillinger, (2000) that:

$$\begin{aligned} G(V) &= \int \frac{dV}{(e^{2aV} - 2c_1)^{1/2}} + c_2 \\ &= \frac{1}{2a\sqrt{-2c_1}} \log \frac{\sqrt{e^{2aV} - 2c_1} - \sqrt{-2c_1}}{\sqrt{e^{2aV} - 2c_1} + \sqrt{-2c_1}} + c_2; \\ &\qquad\qquad c_1 < 0 \\ &= \frac{1}{a\sqrt{2c_1}} \operatorname{atan} \frac{\sqrt{e^{2aV} - 2c_1}}{\sqrt{2c_1}} + c_2; \quad c_1 > 0. \end{aligned}$$

This implies

$$V_{\mathrm{ML}}(t) = \frac{1}{2a} \log \left[-2c_1 \left(\frac{e^{2a\sqrt{-2c_1}(t-c_2)} + 1}{e^{2a\sqrt{-2c_1}(t-c_2)} - 1} \right)^2 + 2c_1 \right], \quad c_1 < 0,$$

or

$$V_{\mathrm{ML}}(t) = \frac{1}{2a} \log \left[2c_1 \left(\tan \left(a\sqrt{2c_1}\,(t - c_2) \right) \right)^2 + 2c_1 \right], \quad c_1 > 0;$$

the constants c_1 and c_2, as before, are fixed by the end conditions $V_{\mathrm{ML}}(0) = 0$ and $V_{\mathrm{ML}}(T) = 1$ (Fig. 6). Since this model satisfies the "active" constraint that $f_V(V)\, f(V) > 0$, the solution automatically remains subthreshold until $t = T$.

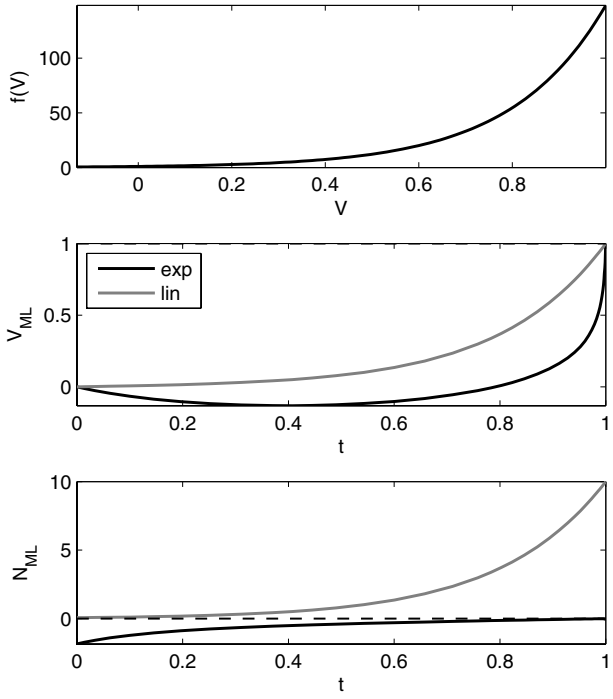

Fig. 6 Optimal path for a cell with exponential subthreshold dynamics. Top: $f(V) = e^{aV}$, with $a = 5$. Middle: Optimal voltage trace $V_{\mathrm{ML}}(t)$, with corresponding $V_{\mathrm{ML}}(t)$ under linear dynamics $f(V) = -aV$ for comparison. Note that the voltage trace under exponential dynamics climbs to threshold much more sharply. Bottom: Optimal noise traces $N_{ML}(t) = \dot{V}_{ML}(t) - f(V_{ML}(t), t)$. Note that exponential and linear noise traces are opposite, in a sense: the linear noise path, as discussed above, grows monotonically as t increases and the voltage climbs away from the rest potential at $V = 0$; on the other hand, the exponential path expends its effort slowing the climb towards threshold for small t (where the driving force $f(V(t))$ is smallest), and then relaxes towards zero as the dynamics carry the voltage towards threshold. Also note that the total energy expended, $\int_o^t |\, N_{\mathrm{ML}}(t)\, |^{\,2}\, dt$, is much smaller in the exponential case, as expected given the positivity of $f(V)$ in the exponential case and the negativity of $f(V)$ in the linear case on the relevant range $0 < V < 1$

Quadratic cell

Another interesting example is the quadratic integrate-and-fire cell (Ermentrout and Kopell, 1986; Hansel and Mato, 2003; Brunel and Latham, 2003),

$$f(V, t) = a(V - b)^2 + c, \quad a > 0.$$

For $c < 0$, this model has two fixed points, V_1 and V_2, with $V_1 < V_2$, V_1 stable, and V_2 unstable. When the unstable fixed point is subthreshold, $V_2 < 1$, the optimal noise path $N_{\mathrm{ML}}(t)$ effectively only has to push V_{ML} up to V_2, and then passively follow the dynamics, which will automatically push $V(t)$ up to threshold. This effect is illustrated in Fig. 7 in this case, the analytic approach outlined above proved to be unenlightening in our hands, and so we simply solved Eq. (3) numerically using the "shooting" method described in Section 2. Since

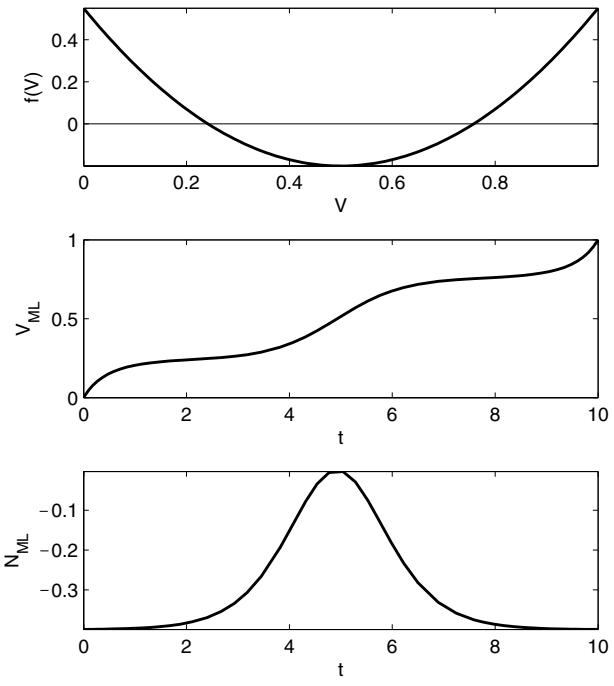

Fig. 7 Optimal path for a cell with quadratic subthreshold dynamics. Top: $f(V) = a(V-b)^2 + c$, with $a = 3, b = .5$, and $c = -.2$. Middle: $V_{\mathrm{ML}}(t)$. Bottom: Optimal noise trace $N_{ML}(t) = \dot{V}_{ML}(t) - f(V_{ML}(t), t)$. Note that $V_{\mathrm{ML}}(t)$ follows the $f(V)$ dynamics near $t = 0$ and $t = T$ (where the dynamics are "going the right way," that is, $f(V(t)) > 0$), but has to expend some energy, in the form of a peak in $N_{\mathrm{ML}}(t)$, to push through the region of V-space for which $f(V) < 0$, i.e., the dynamics are pushing V back towards the stable fixed point at $V \approx 0.25$

this model has "active" subthreshold dynamics ($f(V) f_{\mathrm{V}}(V) > 0$ for $V > 0.8$), as described above, the condition $V_{\mathrm{ML}}(t) < 1$ is automatically enforced given the end conditions $V_{\mathrm{ML}}(0) = 0$ and $V_{\mathrm{ML}}(T) = 1$.

4. Large deviations approximations

In the following sections we discuss some applications of the most likely voltage path $V_{\mathrm{ML}}(t)$ obtained above. Our first application is based on the well-known large deviations principle (Freidlin and Wentzell, 1984; Dembo and Zeitouni, 1993) for stochastic differential equations of the form (1). For applications (Paninski et al., 2004b) we are often interested in computing probability integrals of the form

$$P_\sigma(C) = \int_{V \in C} dP_\sigma(V),$$

that is, the probability P_σ that voltage paths generated from model (1) under a given noise scale σ will lie within some constraint set C (the example we have in mind, of course, is the constraint set of all voltage paths that begin at $V(0) = 0$ and end at $V(T) \in [1-\epsilon, 1]$, with $V(t) < 1$ for $0 < t < T$ and $\epsilon \to 0$).

The problem of approximating this type of integral has received extensive attention in the probability and physics literature. In particular, for the spike-time constraint set C considered above and under standard white Gaussian noise N_t, it is well-known that the following (Freidlin-Wentzell) exponential approximation holds (Freidlin and Wentzell, 1984; Dembo and Zeitouni, 1993) for $\sigma \to 0$:

$$-2\sigma^2 \log P_\sigma(C) = L_0(\{V_{\mathrm{ML}}\}) + o(1) = \|N_{\mathrm{ML}}\|_2^2 + o(1);$$

this may be rewritten in the possibly more intuitive form

$$\begin{aligned} P_\sigma(C) &= \exp\left[-\frac{1}{2\sigma^2} L_0(\{V_{\mathrm{ML}}\}) + o\left(\frac{1}{\sigma^2}\right)\right] \\ &= \exp\left[-\frac{1}{2\sigma^2}\int_0^T (\dot{V}_{\mathrm{ML}}(t) - f(V_{\mathrm{ML}}, t))^2 dt \right. \\ &\quad \left. + o\left(\frac{1}{\sigma^2}\right)\right], \sigma \to 0. \end{aligned}$$

In other words, the probability that any voltage path generated from model (1) will lie within the constraint set C may be approximated, for σ small enough and on an exponential scale, simply by the likelihood of the most likely voltage path in C, $V_{\mathrm{ML}}(t)$. Clearly, once $V_{\mathrm{ML}}(t)$ (or, equivalently, $N_{\mathrm{ML}}(t)$) is known, this approximation may be easily computed. In addition, it is fairly straightforward to take the gradient of this approximation with respect to the relevant parameters (σ, g, $I(t)$, etc.) for settings in which we wish to maximize the likelihood of the observed spikes under these parameters (Paninski et al., 2004b). The large-deviations approximation, which may be computed stably and easily for arbitrarily small σ (and, by extension, arbitrarily small $P_\sigma(C)$), is quite useful in this numerical optimization setting (Paninski et al., 2005), where we typically deal with log-likelihoods, and small errors in the computation of very small likelihoods can be disastrous.

4.1. Inter-spike interval density approximations

Let us apply this idea to approximate the interspike interval probability density, that is, the probability density that a cell which has spiked at time $t = 0$ will fire next at $t = T$. According to the above, we need to compute, for each desired value of T, the most likely voltage path $V_{\mathrm{ML}}^{\mathrm{T}}(t)$ which starts at $V = 0$ and crosses threshold for the first time at T. Once this optimal path is obtained, we need only compute the corresponding "energy," $L(\{V_{\mathrm{ML}}^T\})$, multiply

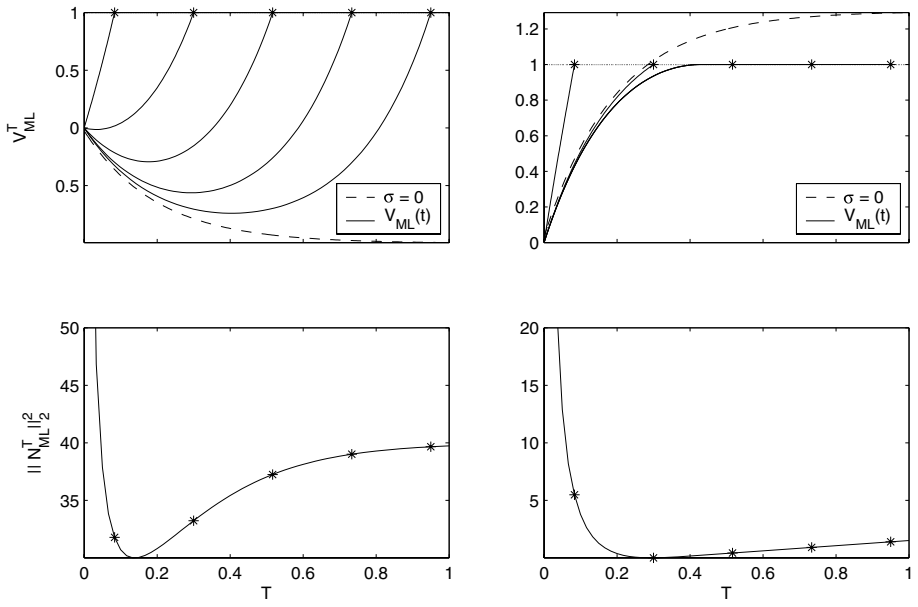

Fig. 8 Illustration of computations underlying large deviations approximation for interspike interval distribution. Left: Optimal paths(top) and corresponding energies $L_0(\{V_{\mathrm{ML}}^T\}) = ||N_{\mathrm{ML}}^T||_2^2$ (bottom) for a linear homogeneous cell driven by a constant negative current I, $f(V,t) = -gV(t) + I$. Note that optimal voltage paths are all convex, and for longer times match noiseless path, $V_0(t)$ (dashed trace), more accurately. Energy minimum occurs for fairly small t, after which energy climbs to a finite asymptote (for large enough T, the optimal path matches $V_0(t)$—incurring no $L(\{V\})$ cost—until t sufficiently close to T, at which time $V_{\mathrm{ML}}(t)$ climbs away from $V_0(t)$ up towards the threshold $V = 1$). Also note that no zero-energy path exists in this case, since $V_0(t)$ does not reach threshold. Right: Corresponding plots for same neuron driven by positive fixed input current I. In this case, the optimal voltage paths are concave and meet threshold strictly before T for T large enough. Also, $V_0(t)$ meets threshold at a finite time t_0 in this case, implying that the energy reaches a unique minimum of 0 at t_0; for $t > t_0$, the energy asymptotically grows linearly, corresponding to the continuing cost of keeping $V_{\mathrm{ML}}(t)$ below threshold

by $-1/2\ \sigma^2$, and exponentiate to obtain our probability approximation.

The most straightforward example of these computations, as usual, is the leaky integrate-and-fire cell, $f(V,t) = -gV + I$, with I fixed at 0 for simplicity. Recall from Section 3.1 that

$$V_{ML}^T(t) = c_T(e^{gt} - e^{-gt}),$$

with $c_T = (e^{gT} - e^{-gT})^{-1}$. We have that

$$\begin{aligned} L_0\left(\{V_{\mathrm{ML}}^T\}\right) &= \int_0^T \left(\dot{V}_{\mathrm{ML}}^T(t) - \left(-gV_{\mathrm{ML}}^T(t)\right)\right)^2 dt \\ &= \int_0^T c_T(ge^{gt} + ge^{-gt}) \\ &\quad + c_T(ge^{gt} - ge^{-gt}))^2 dt \\ &= 4c_T^2 \int_0^T g^2 e^{2gt} dt \\ &= \frac{2g(e^{2gT} - 1)}{(e^{gT} - e^{-gT})^2}. \end{aligned}$$

In the no-leak ($g = 0$) case, this reduces to

$$L_0\left(\{V_{\mathrm{ML}}^T\}\right) = 1/T,$$

implying that

$$p(T) \sim e^{-1/(2\sigma^2 T)},$$

with $p(T)$ the probability density of observing the first spike since reset at time T. In this case the exact formula is known (Hida, 1980) and can be compared:

$$p(T) = \frac{1}{\sigma(2\pi T^3)^{1/2}} e^{-1/(2\sigma^2 T)};$$

the approximation agrees on a log scale, as promised (though note the key difference that the approximation is nonintegrable). Further examples of these ISI approximations appear in Figs. 8 and 9.

4.2. Conditional probability approximations

It is worth pointing out that the large deviations principle used above is actually much more general; in fact, for any sufficiently regular set A, we have (Freidlin and Wentzell, 1984; Dembo and Zeitouni, 1993)

$$\begin{aligned} -2\sigma^2 \log P_\sigma(A) &= \inf_{V \in A} L_0(\{V\}) + o(1) \\ &= \inf_{V \in A} ||N_V||_2^2 + o(1), \quad \sigma \to 0; \end{aligned}$$

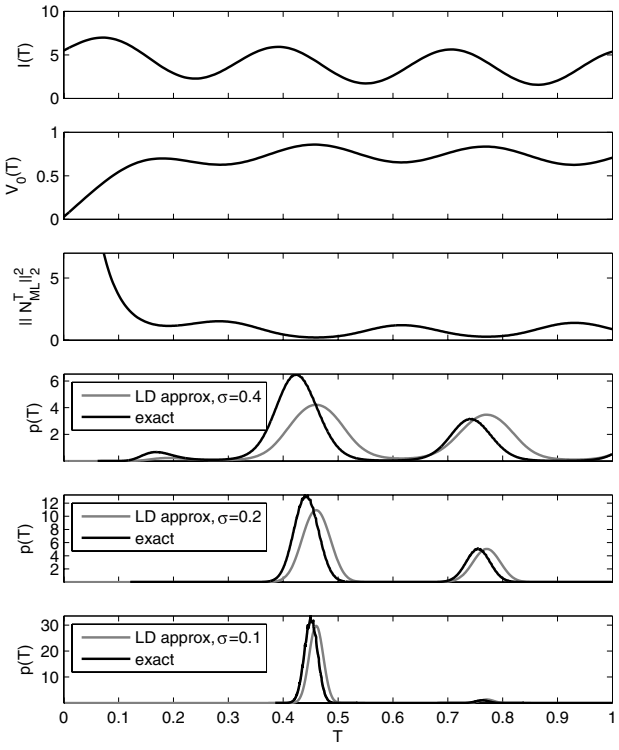

Fig. 9 Illustration of large deviations approximation for interspike interval distribution. Cell shown here is linear, with fixed g and sinusoidally-varying input current $I(t)$. Top: Input current $I(t)$. Second panel: Noiseless voltage path $V_0(t)$. Third: Energy $L_0(\{V_{\mathrm{ML}}^T\}) = ||N_{\mathrm{ML}}^T||_2^2$ as a function of time T, as computed in Fig. 9. Energy minima roughly correspond to peaks in $V_0(t)$ as shown in second panel. Lower panels: Comparison of exact and approximate interspike interval densities for varying values of σ; exact density computed by usual Fokker-Planck methods (see appendix; Gerstner and Kistler, 2002; Paninski et al., 2004b). Note that approximation becomes more accurate as σ or $t \to 0$, as predicted, and that approximation captures multimodal nature of density even for larger values of σ. Also note that approximation fails systematically for larger σ: the true density consistently shows a phase-shift to the left, corresponding to noise pushing the voltage over threshold sooner than predicted by the simple energy approximation

here we have used the notation N_{V} for the noise path $N(t) = \dot{V}(t) - f(V(t), t)$ corresponding to a given observed voltage trace $V(t)$. Thus, as $\sigma \to 0$, $P_\sigma(A)$ decays exponentially, with the rate of decay given by $-2\sigma^2 \inf_{V\in A} L_0(\{V\})$, the loglikelihood of the most likely path $V(t)$ in A; if the noiseless path $V_0(t)$ is contained in A, on the other hand, this exponential decay does not occur (since in this case $\inf_{V\in A} L_0(\{V\}) = L_0(\{V_0\}) = 0$). Clearly the approximation for $P_\sigma\ (C)$ in the preceding section is a special case of this general result, with $\inf_{V\in C} L_0(\{V\}) = L_0(\{V_{\mathrm{ML}}\})$.

Now it is possible to show, using the above general large deviations principle, along with the strict convexity of $L_0(\{V\})$ in V and the convexity of C, that a similar conditional probability approximation holds (Cover and Thomas, 1991; Dembo and Zeitouni, 1993): given the observed spike train (and therefore that the subthreshold $V(t)$ is contained in the set C), the conditional probability $P_\sigma(V \in A \mid V \in C)$ of a given set $A \subset C$ obeys

$$-2\sigma^2 \log P_\sigma(V \in A \mid V \in C)$$
$$= \inf_{V\in A} L_0(\{V\}) - L_0(\{V_{\mathrm{ML}}\}) + o(1);$$

thus the conditional probability of any voltage path outside of an $\|N\|_2$-ball of the optimal path $V_{\mathrm{ML}}(t)$ decays exponentially in σ^{-2}. This, in turn, implies that the mean conditional path closely approximates the ML path for σ sufficiently small, as expected. This effect is visible in Fig. 5 Note that the conditional mean consistently sags below the ML path, with the sag (caused by the fact that superthreshold paths are "killed" by the fact that we have conditioned on subthreshold paths, thus skewing the mean downwards) roughly proportional to σ; see Badel et al. (2005) and Paninski (2005) for more details.

5. Demonstration on *in vitro* data

In this last section, we provide a simple demonstration of how the above ideas may be applied to physiological data. We analyzed *in vitro* recordings described in Paninski et al. (2003): briefly, dual-electrode whole-cell somatic recordings were made from pyramidal cells from layers III and V in slices from sensorimotor cortex of rats aged P14–24; non-repeating Gaussian noise current stimuli were delivered through one electrode, while voltage responses were recorded through the other electrode.

From these paired intracellular voltage and input-current recordings, we fit the following integrate-and-fire-based model:

$$dV(t) = -gV + I(t) + \sigma N_t;$$
$$I(t) = I_{DC} + k * I_{\mathrm{in}}(t) + h(t - t_{i-1}),$$

with $k * I_{\mathrm{in}}(t))$ denoting the original injected current $I_{\mathrm{in}}(t)$ filtered by a simple five-delay linear filter $\vec{k} = \{k_0, k_{-1}, \ldots k_{-4}\}$, and $h(t - t_i)$ an after-spike current whose value depends only on $(t - t_{i-1})$, the time since the last observed spike t_{i-1}. The model cell emits a spike whenever a threshold, V_{th}, is crossed. Thus the full model parameters are $\theta = \{g, \sigma, I_{DC}, \vec{k}, h(\tau), V_{\mathrm{th}}\}$; see Paninski et al. (2004a) for details on the (maximum likelihood) estimation of these parameters (the fitting procedure effectively comes down to a linear regression, as discussed in Stevens and Zador (1998) and Jolivet et al. (2004). While the detailed shape of the action potential may also be qualitatively described using a similar model (Jolivet et al., 2004; Paninski et al., 2004a), for simplicity we did not attempt to model the

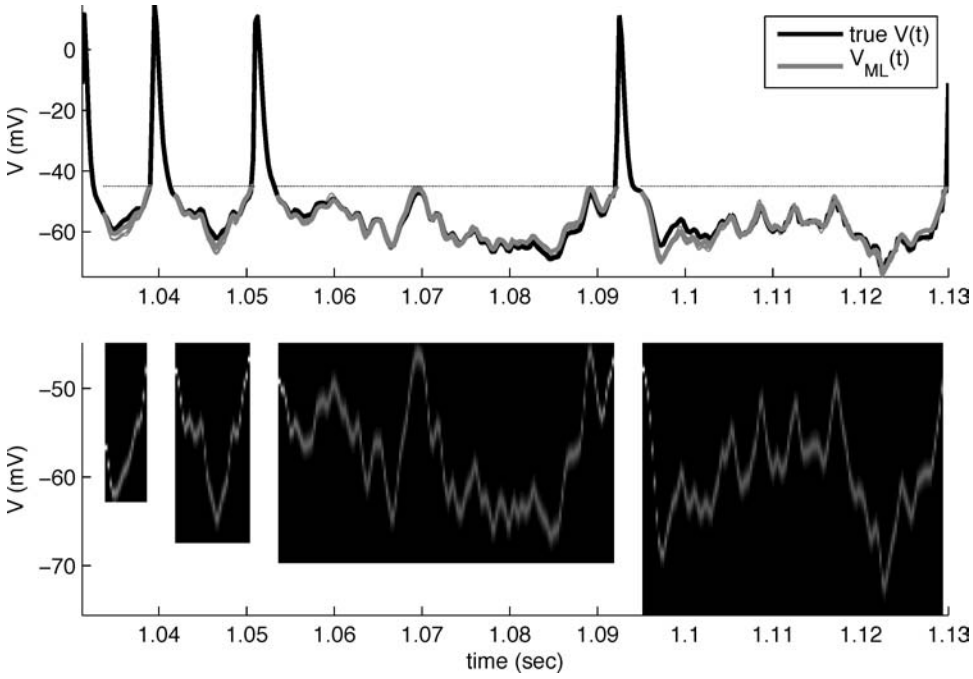

Fig. 10 Computing the most likely voltage path given *in vitro* physiological data. Top: Comparison of true voltage path (bold black trace) with computed $V_{\mathrm{ML}}(t)$ (bold gray trace) and samples from conditional distribution $p(V(t) \mid spikes, \{I(t)\}_{0\le t\le T}, \theta)$ (thin traces, partially obscured; recall θ denotes the estimated IF model parameters). Trace shown is a randomly chosen segment of a 25-second long white noise experiment; dashed trace indicates estimated threshold. Bottom: Conditional distributions $p(V(t) \mid spikes, \{I(t)\}_{0\le t\le T}, \theta)$. White space indicates gaps in time where voltage was superthreshold, and thus not modeled. Note small scale of estimated noise σ (Mainen and Sejnowski, 1995)

superthreshold behavior of the recorded neuron, but merely chose the initial condition of the above differential equation model after each spike as $V(t_{i-1}+\Delta) = V_{\mathrm{obs}}(t_{i-1}+\Delta))$, where $V_{\mathrm{obs}}(t)$ is the true observed voltage, and $\Delta = 2.5$ ms is the empirically chosen spike width of the cell. Thus, in short, when the voltage V crosses threshold V_{th}, a spike is emitted (whose precise superthreshold shape we ignore here), and then the voltage V is reset to the true voltage $V_{\mathrm{obs}}(t_{i-1}+\Delta)$, as observed at the time of the end of the spike, $(t_{i-1}+\Delta)$.

Our model thus fixed, we could then apply the ideas discussed above to examine the maximum-likelihood voltage path under this model, given a novel (that is, cross-validated) set of spike data $\{t_i\}$ and corresponding stimulus $\{I(t)\}$, and to compare the resulting $V_{\mathrm{ML}}(t)$ with the true observed voltage trace, $V_{\mathrm{obs}}(t)$. The results are shown in Fig. 10. $V_{\mathrm{ML}}(t)$ is seen to match the true subthreshold $V_{\mathrm{obs}}(t)$ quite well.

6. Conclusion

In this work we have developed analytical and numerical methods for the investigation of $V_{\mathrm{ML}}(t)$, the most likely subthreshold voltage path given an observed spike train. This optimal path turns out to satisfy a simple second-order ordinary differential equation, (3), which can be solved exactly in a number of simple but physiologically relevant cases; analysis of this equation in turn leads to further insight into the general structure of the optimal path. We also adapted existing methods for sampling from and computing conditional densities of the latent state of a hidden Markov process ($V(t)$, here), given the observable data (the superthreshold spike train). We made use of this optimal path $V_{\mathrm{ML}}(t)$ to apply large deviations principles for stochastic differential equations to obtain approximations for the likelihoods and interspike interval densities associated with a given set of model parameters for the subthreshold dynamics. Finally, we demonstrated that the ML method can provide an accurate reconstruction of the true subthreshold somatic voltage as a function of time, given accurate methods for fitting integrate-and-fire-based models to these subthreshold dynamics (Stevens and Zador, 1998; Jolivet et al., 2004; Paninski et al., 2004a).

We conclude by listing a few open directions for future research.

Nonhomogeneous noise level σ

One obvious extension to the simple model (1) considered here would be to let σ depend on V and t:

$$dV(t) = f(V(t),t)dt + \sigma(V,t)dW_t.$$

We can recreate our analysis of Section 2 in this case, by examining the functional derivative of

$$\int_0^T \frac{[\dot{V}(t) - f(V(t))]^2}{\sigma^2(V(t),t)}dt;$$

unfortunately, however, the resulting differential equations turn out to be less tractable in this more general case (see Freidlin and Wentzell, 1984 for further discussion). However, some special cases might be worth examining in further

detail. For example, in the case of linear f and for $\sigma(V, t) = \sigma(t)$ (i.e., σ constant in V), the underlying distributions remain Gaussian, implying that the standard linear maximum principle and Gaussian identities used above can again be employed here.

Incorporating model uncertainty

The second way our model assumptions might be relaxed would be to place less confidence in how well we have estimated our model $\{f(V, t), \sigma, V_{\rm th}, V_{\rm reset}\}$. This uncertainty in the model parameters could take several plausible forms; the classical Bayesian approach would be integrate over some prior distribution on the model parameters. For example, it would be worthwhile (and fairly straightforward) to model the observed variability in the voltage reset after a spike (c.f. Fig. 10). It might also be useful to relax our assumptions about the existence of a sharp voltage threshold, or at least model noise in the threshold process. One tractable generalization of the deterministic threshold model used here would be to assume spikes are generated via a Poisson process with rate $r(V(t))$, with $r(V)$ a positive, increasing function with a sharp increase at some threshold value $V_{\rm th}$ (Stevens and Zador, 1996; Plesser and Gerstner, 2000; Gerstner and Kistler, 2002; Paninski, 2004). In this case, computing the optimal voltage path would require the maximization of the loglikelihood

$$\log(r(V(T))) - \int_0^T r(V(t))dt - \int_0^T \frac{[\dot{V}(t) - f(V,t)]^2}{2\sigma^2};$$

interestingly, the solution here will depend in general on the level of the subthreshold noise σ, which was not the case in the deterministic threshold case analyzed above. It is also worth noting that in the case of linear subthreshold dynamics $f(V, t) = -g(t)V(t) + I(t)$ and a firing rate function $r(V)$ which is convex and whose logarithm is concave, the loglikelihood above is strictly concave, with no non-global local maxima, and therefore has a unique maximum which may be computed easily via simple gradient ascent techniques.

Modeling multiple simultaneous spike trains

Assume k spike trains are observed simultaneously: can we infer the optimal voltage paths if there exists a correlated noise source linking these cells together? A first step in this direction is to assume a model of the form:

$$dV_i(t) = f_i(V_i(t), t)dt + dW_{i,t}, \quad 1 \le i \le k$$

with the noise sources $dW_{i,t}$ Gaussian, temporally white, and correlated with some covariance matrix C (more formally, the Brownian motions W_i have time-independent increments with between-cell covariance $|t_2 - t_1|C$).

For this model we have the likelihood

$$L(\{\vec{V}(t)\}) = -\frac{1}{2}\int_0^T \vec{n}(t)^t C^{-1}\vec{n}(t)dt,$$

with

$$\vec{n}(t) = (\dot{V}_1(t) - f_1(V_1(t), t), \dot{V}_2(t) - f_2(V_2(t), t), \ldots, \dot{V}_k(t) - f_k(V_k(t), t))^t.$$

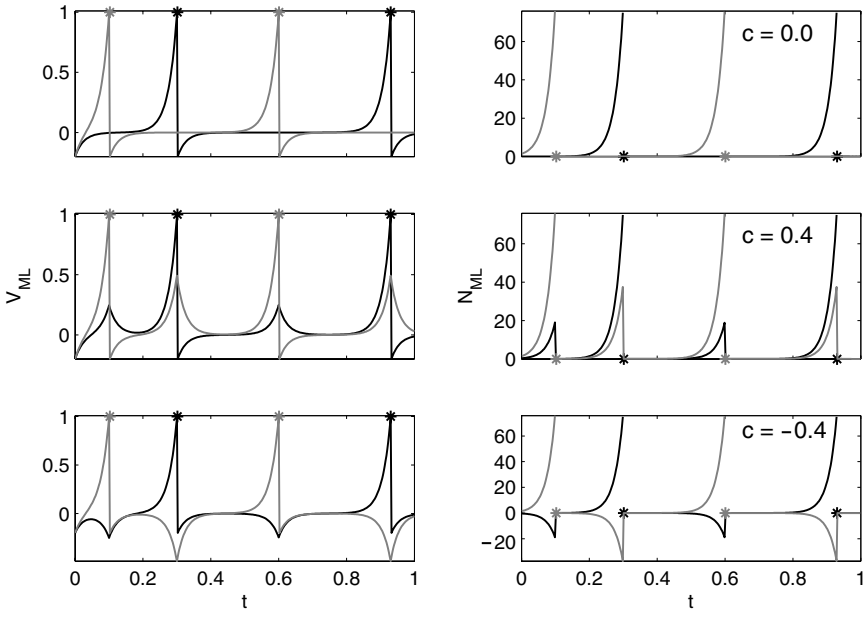

Fig. 11 Computing the most likely voltage path under correlated subthreshold noise. Two identical linear homogeneous cells, $f(V, t) = -gV + I$, with $g = 40$, $V_{\rm reset} = -0.2$, $V_{\rm th} = 1$, $I = 0$; the only difference between the two cells is the noise level $\sigma_1^2 = 1$ (the voltage of cell 1 is indicated by the black traces) and $\sigma_1^2 = 2$ (gray). Spike times indicated by asterisks. Top: Uncorrelated noise case, $c = 0$. Middle: positively correlated noise, $c = 0.4$; note that voltage of cell 1 tends to move upward with that of cell 2, and vice versa. Middle: negatively correlated noise, $c = -0.4$

Clearly in the case that C is diagonal (uncorrelated noise), the joint likelihood decomposes into a sum of k independent integrals, implying that maximizing the full likelihood is equivalent to individually maximizing the independent likelihoods

$$L(\{\vec{V}(t)\}) = -\frac{1}{2\sigma_i^2}\int_0^T \left(\dot{V}_i(t) - f_i(V_i(t),t)\right)^2 \vec{dt},$$

bringing us back to our original problem.

The corresponding maximum likelihood problem in the nondiagonal C case may still be solved easily for linear dynamics $f(V, t)$ using quadratic programming techniques; see Fig. 11 for an example. The general nonlinear case may again be attacked via the Euler-Lagrange variational approach.

Appendix A: Hidden Markov model techniques for exact conditional sampling, density computation, and likelihood maximization

As mentioned above, the integrate-and-fire cell is a special case of a continuous-time hidden Markov model (HMM): V acts as the "hidden" variable, which evolves according to Markovian dynamics, and the presence or absence of a spike in time bin t is the observed variable, which is dependent (in this case deterministically) only on $V(t)$ at the single time point t. In this appendix we briefly describe how to adapt the existing methods for computing (and sampling from) the conditional density of the hidden variable of an HMM, conditioned on its beginning- and end-states, to this special IF model case (c.f. Fig. 5 above). See, e.g., Rabiner, (1989) for further detail.

Computing conditional densities

We want to compute p ($V(t) \mid s([0,T])$), with $s([0,T])$ denoting the observed spike data on the interval $[0, T]$. We have

$$\begin{aligned} & p(V(t)|s([0,T])) \\ &= \frac{p(s([0,T]) \mid V(t))p(V(t))}{p(s([0,T]))} \\ &= \frac{p(s([0,t])|V(t))p(s((t,T])|V(t))p(V(t))}{p(s([0,T]))} \\ &= \frac{p(V(t) \mid s([0,t]))p(s([0,t]))p(s((t,T]) \mid V(t))}{p(s([0,T]))}, \end{aligned}$$

where the second equality reflects the conditional independence of $s([0, t])$ and $s((t, T))$ given $V(t)$. The denominator and first two terms in the numerator may be computed by standard density propagation methods, that is, recursively iterating the Chapman-Kolmogorov equations forward (Karlin and Taylor, 1981), or in the case of diffusion processes as considered here, solving the corresponding Fokker-Planck (forward) equation (Burkitt and Clark, 1999; Gerstner and Kistler, 2002; Paninski et al., 2004b). More simply, of course, the ratio $p(s([0,t]))/p(s([0,T]))$, which is constant in V, may just be taken as a normalization factor that ensures the conditional V-probability integrates to one.

The last term in the numerator above may be computed by forward density propagation starting from each possible point $V(t)$; however, this simple approach will be highly inefficient if we want this function for all times $0 < t < T$. The following recursive "backwards" approach is more efficient. We initialize at $t = T - dt$

$$\begin{aligned} & p(s([T-dt,T])|V(T-dt)) \\ &= H(V_{\text{th}} - V(T-dt)) \int_{V_{\text{th}}}^{\infty} \mathcal{N}_{\mu,\sigma^2 dt}(V)dV, \end{aligned}$$

with $H(\cdot)$ the Heaviside function and $\mathcal{N}_{\mu,\sigma^2}(V)$ the normal density on V with mean μ and variance σ^2. Here we have made the abbreviation

$$\mu = V(T-dt) + f(V(T-dt), T-dt)dt.$$

This complicated-looking initialization is just the result of integrating out the probability that the noise N_{T-dt} will push V over threshold in the next time step dt.

Now use the backward recursive formula

$$\begin{aligned} & p(s([t,T])|V(t)) \\ &= H(V_{\text{th}} - V(t)) \int p(s([t+dt,T])|V(t+dt)) \\ &\times p(V(t+dt) \mid V(t)) \\ &dV(t+dt); \end{aligned}$$

again, this is just the Chapman-Kolmogorov equation (in the case of the diffusion processes considered here, we can take limits as $dt \to 0$ and obtain the Kolmogorov backward partial differential equation (Karlin and Taylor, 1981), but this is beyond the scope of this brief appendix; see e.g. Paninski (2005) for an application of this idea). In the case of stationary dynamics $f(V,t) = f(V)$, recursing this formula corresponds to a simple iterated matrix multiplication.

So in total computing $p(V(t) \mid s([0,T]))$ requires just one forward density propagation and one backward propagation; see Fig. 12 for an illustration.

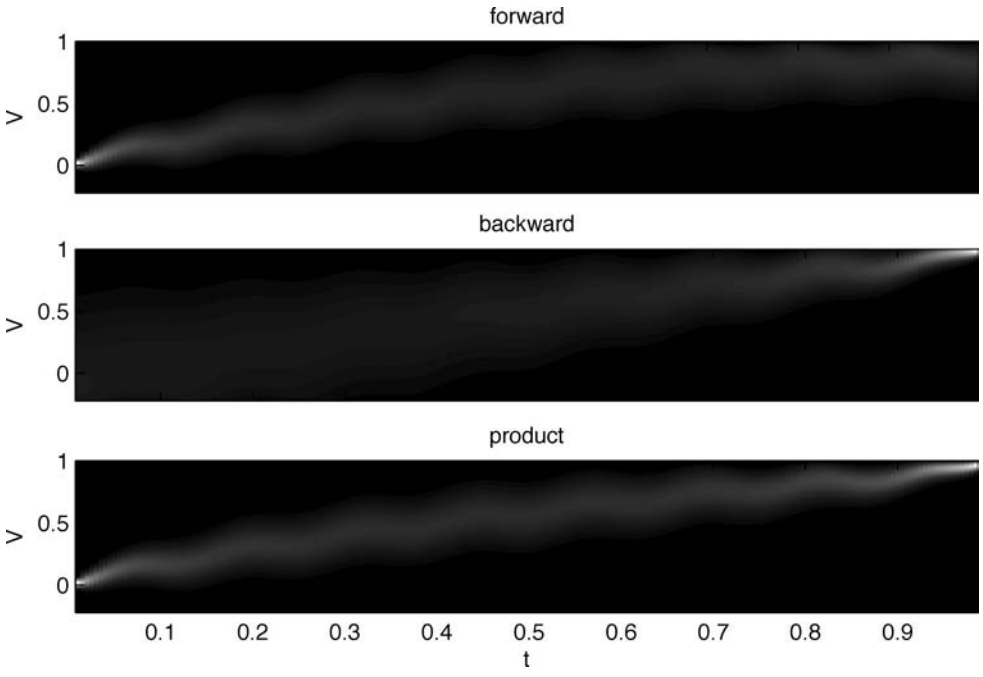

Fig. 12. Illustration of forward and backward propagation steps for computing conditional densities $p(V(t) \mid s([0, T]))$. Top: "Forward" density $p(V(t) \mid s([0, t])$. Neuron is linear, $f(V, t) = -gV(t) + I(t)$, with fixed g; input current $I(t)$ varies sinusoidally (8 Hz). Middle: "Backward" conditional firing rate $p(s([t, T]) \mid V(t))$. Bottom: Normalized product, $p(V(t) \mid s([0, T]))$

Sampling

The exact sampling method used here, a variant of the forward-backward algorithm described above for computing conditional densities, was suggested by Ghahramani; see also, e.g., Neal et al. (2003). An identical procedure is used in Paninski, (2005). We initialize $V(T) = 1$. (This initial condition is due to the data that a spike occurred at time T, as above.)

Now, for $T > t > 0$, sample backwards:

$$\begin{aligned} V(t) &\sim p(V(t) \mid \{V(u)\}_{t<u<T}, s([0, T])) \\ &= p(V(t) \mid V(t + dt), s([0, t])) \\ &= \frac{1}{Z} p(V(t + dt), s([0, t])|V(t))p(V(t)) \\ &= \frac{1}{Z} p(V(t + dt) \mid V(t))p(s([0, t])|V(t))p(V(t)) \\ &= \frac{1}{Z} p(V(t + dt) \mid V(t))p(V(t)|s([0, t])). \end{aligned}$$

Thus sampling on each time step simply requires that we draw independently from a one-dimensional density, proportional to the product in the last line. Once this product has been computed, this sampling can be done using standard methods (namely, the inverse cumulative probability transform (Press et al., 1992). The term on the right may be precomputed for all t via a single density propagation forward; we emphasize that this only has to be done once, no matter how many samples are required. The second term is computed directly from the Gaussian stochastic dynamics (1), given each $V(t + dt)$. Putting the samples together, for $0 < t < T$, clearly gives a sample from $p(\{V(t)\}_{0<t<T}|s([0, T])$, as desired.

Optimizing

For completeness, we present the dynamic programming (Viterbi) method for computing an optimal $V(t)$ path (Bellman, 1957). In our hands this technique did not lead to any further analytic insight, but it does provide a robust, efficient algorithm which works for arbitrary hidden Markov models (this is the main technique used in computational language processing applications, for example Rabiner (1989). The main idea is that we do not need to search over all possible paths for the optimal path; the complexity of this global search would grow exponentially with the number of time samples desired on $[0, T]$. Instead, we work recursively: given the most likely path up to some arbitrary $V(t - dt)$, for $0 < t < T$, we look for the path that maximizes the likelihood and ends at any given $V(t)$; the Markov nature of $V(t)$ implies that, given the path up to $V(t - dt)$, this involves just an optimization over single steps from $V(t - dt)$ to $V(t)$, and not an exponentially-large search over all paths up to $V(t)$. Thus, the algorithm proceeds as:

Initialize: $V(0) = 0$.

For $0 < t < T$: given optimal path up to $V(t - dt)$, for all possible values of $V(t - dt)$ (along with corresponding likelihood for each such path), choose optimal path up to $V(t)$, for each value of $V(t)$.

Finally, choose optimal path ending at $V(T) = 1$, given optimal paths (and corresponding likelihoods) up to $V(T - dt)$.

Acknowledgments Thanks to J. Pillow and E. Simoncelli for suggesting the problem of determining the most likely voltage path, to Z. Ghahramani for suggesting the forward-backward method for exact conditional sampling from hidden Markov models, to Q. Huys and P. Latham for many helpful comments, and to B. Lau and A. Reyes for their collaboration in collecting the data discussed in Fig. 10. This work was partially supported by funding from the Gatsby Charitable Trust and by a Royal Society International Research Fellowship to LP. A preliminary account of this work appeared in the COSYNE '05 (Computational and Systems Neuroscience Meeting 2005, Salt Lake City, Utah) Abstracts volume.

References

Badel L, Richardson M, Gerstner W (2005) Dependence of the spike-triggered average voltage on membrane response properties. Neurocomputing, In press.

Bellman R (1957) Dynamic Programming. Princeton University Press.

Brown E, Kass R, Mitra, P. (2004) Multiple neural spike train data analysis: state-of-the-art and future challenges. Nature Neurosci. 7: 456–461.

Brunel, N. Latham, P (2003) Firing rate of the noisy quadratic integrate-and-fire neuron. Neural Comput. 15: 2281–2306.

Burkitt A, Clark G (1999) Analysis of integrate-and-fire neurons: Synchronization of synaptic input and spike output. Neural Comput. 11: 871–901.

Cover T, Thomas J (1991). Elements of Information Theory. Wiley, New York.

Dembo A, Zeitouni O (1993). Large Deviations Techniques and Applications. Springer, New York.

Ermentrout G, Kopell N (1986) Parabolic bursting in an excitable system coupled with a slow oscillation. SIAM J. Appl. Math 2: 233–253.

Fourcaud N, Brunel N (2002) Dynamics of the firing probability of noisy integrate-and-fire neurons. Neural Comput. 14: 2057–2110.

Freidlin M, Wentzell A (1984) Random Perturbations of Dynamical Systems.Springer-Verlag.

Gerstner W, Kistler W (2002) Spiking Neuron Models: Single Neurons, Populations, Plasticity. Cambridge University Press.

Hansel D, Mato G (2003) Asynchronous states and the emergence of synchrony in large networks of interacting excitatory and inhibitory neurons. Neural Comput. 15: 1–56.

Hida T (1980) Brownian Motion. Springer-Verlag, New York.

Jeffrey A, Zwillinger D eds. (2000) Gradshteyn and Ryzhik's Table of Integrals, Series, and Products. 6th edn. Academic Press.

Jolivet R, Lewis T, Gerstner W (2004) Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy. J. Neurophysiology, 92: 959–976.

Karatzas I, Shreve S (1997) Brownian Motion and Stochastic Calculus. Springer.

Karlin S, Taylor H (1981) A Second Course in Stochastic Processes. Academic Press, New York.

Mainen Z, Sejnowski T (1995) Reliability of spike timing in neocortical neurons. Science, 268: 1503–1506.

Neal R, Beal M, Roweis S (2003) Inferring state sequences for non-linear systems with embedded hidden Markov models. NIPS, 16.

Paninski L (2004) Maximum likelihood estimation of cascade point-process neural encoding models. Network: Comput. Neural Sys. 15: 243–262.

Paninski L (2005) The spike-triggered average of the integrate-and-fire cell driven by gaussian white noise. In press.

Paninski L, Haith A, Pillow J, Williams C (2005) Improved numerical methods for computing likelihoods in the stochastic integrate-and-fire model. Comp. Sys. Neur. '05.

Paninski L, Lau B, Reyes A (2003) Noise-driven adaptation: in vitro and mathematical analysis. Neurocomputing 52: 877–883.

Paninski L, Pillow J, Simoncelli E (2004a) Comparing integrate-and-fire-like models estimated using intracellular and extracellular data. Neurocomputing, 65: 379–385.

Paninski L, Pillow J, Simoncelli E (2004b) Maximum likelihood estimation of a stochastic integrate-and-fire neural model. Neural Comput. 16: 2533–2561.

Plesser H, Gerstner W (2000) Noise in integrate-and-fire neurons: From stochastic input to escape rates. Neural Comput. 12: 367–384.

Press W, Teukolsky S, Vetterling W, Flannery B (1992). Numerical recipes in C. Cambridge University Press.

Rabiner, L (1989) A tutorial on hidden Markov models and selected applications in speech recognition. Proc. IEEE, 77: 257–286.

Stevens C, Zador A, (1996) When is an integrate-and-fire neuron like a Poisson neuron? NIPS 8: 103–109.

Stevens C, Zador A (1998) Novel integrate-and-fire-like model of repetitive firing in cortical neurons. Proc. 5th Joint Symp. Neural Comput. UCSD.

Vinter R (2000) Optimal Control. Birkhauser.