UNIVERSITY OF COPENHAGEN

Optimal Design for Estimation in Diffusion Processes from First Hitting Times

Iolov, Alexandre; Ditlevsen, Susanne; Longtin, André

Optimal Design for Estimation in Diffusion Processes from First Hitting Times[*]

Alexandre Iolov[†], Susanne Ditlevsen[‡], and André Longtin[§]

Abstract. We consider the optimal design problem for the Ornstein–Uhlenbeck process with fixed threshold, commonly used to describe a leaky, noisy integrate-and-fire neuron. We present a solution to the problem of devising the best external time-dependent perturbation to the process in order to facilitate the estimation of the characteristic time parameter for this process. The optimal design problem is constrained here by the fact that only the times between threshold crossings from below, known as hitting times, are observable. The optimal control is based on a maximization of the mutual information between the posterior of the unknown parameter given these observations and the distribution of the hitting times. Our approach is based on the adjoint method for computing the gradient of a functional of a solution to a Fokker–Planck partial differential equation with respect to an input function (i.e., to the control). Our method also enables the estimation of other parameters, in the case when more than one parameter is unknown.

Key words. design of experiments, mutual information, leaky integrate-and-fire neuronal models, Fokker–Planck equation, probability density function control method, Ornstein–Uhlenbeck process

AMS subject classifications. 62K05, 62L05, 93E20, 92C20, 92D25, 49L20, 49M25

DOI. 10.1137/16M1060376

1. Introduction. Techniques for the estimation of model parameters from observations of a stochastic process are an active area of research. Generally one hopes to have access to sufficient measurements of the state variables to guide the estimation of these parameters. However, in some cases the state variables are not measurable, and one has to work instead with times at which specific events occur, such as an earthquake or an epileptic seizure. Here we consider the problem of optimal design for estimation of parameters in a stochastic differential equation (SDE) given only observations of hitting times, i.e., of times at which the process crosses a fixed threshold. More specifically, since a primary motivation is the activity of neurons, we study the leaky integrate-and-fire (LIF) neuronal model, where the hitting times are identified with crossings from below of a threshold voltage. This corresponds to so-called extracellular potential measurements of spikes, which occur when the voltage inside the cell, evolving according to an Ornstein–Uhlenbeck process, crosses a fixed threshold value. While our motivation is based in neuroscience, the technique we develop is more generally

[†]Department of Mathematics and Statistics, University of Ottawa, Ottawa, ON K1N 6N5, Canada (aviolov@yahoo.com).
[‡]Department of Mathematical Sciences, University of Copenhagen, DK-2100 Copenhagen, Denmark (susanne@math.ku.dk).
[§]Department of Mathematics and Statistics and Department of Physics, University of Ottawa, Ottawa, ON K1N 6N5, Canada (alongtin@uottawa.ca).

applicable to dynamical models from which a point process can be extracted.

We assume that the controller has influence over some part of the SDE by causing time-varying perturbations of the process. In the case of a neuron, this would correspond to the case where one could influence the voltage evolution, e.g., through modulation of synaptic input, while one could only reliably record threshold crossing times. We further assume that the objective is to obtain crossing time observations, known as hitting times or spike times in the neural context, that are most informative in both formal and informal senses. That is, we choose the control as the solution to a well-posed optimization problem, where the ultimate goal is to obtain more accurate and more precise estimates of the unknown parameter(s).

We espouse a Bayesian perspective and proceed by assuming a prior distribution of the unknown parameters. Then we seek to maximize the mutual information (MI) between the posterior of the unknown parameters given the observations and the distribution of the hitting times. We note that, even for a fixed and known parameter set, the hitting time observations are still random due to the inherent stochasticity of the SDE.

Using MI for the selection of maximally informative experiments has been advocated by several recent lines of research, for example, in experimental psychology [5, 19], computational neuroscience [22, 21, 15], and quantum physics [9]. An alternative approach is to maximize the expected Fisher information of the experiment as in [10]. This is especially effective if the Fisher information is available as an analytical formula. That is the case when one observes trajectories from the diffusion process, as in [10], but not in the much more limited context of observations of hitting times.

The MI measures the expected reduction in uncertainty in the parameter value given an observation, under a fixed experimental stimulation. Thus, the stimulation that maximizes the MI should extract the maximum information about the model parameters, on average. Using MI for estimating statistical quantities is formally justified; see, e.g., [21], where it is proven under fairly weak assumptions that using a stimulus input that maximizes the MI between future observations and the parameters of interest leads to consistent and efficient parameter estimates. However, while theoretically sound, optimizing or even merely computing the MI may be computationally prohibitive. For our problem, the primary issue is the low level of information available given that one has access only to hitting times, rather than measurements of the underlying state variable.

We explicitly single out a paper [9] that describes a strategy to experimentally estimate the parameters of the Hamiltonian of a quantum system. The qubit state outcomes of zeroes or ones follow a Bernoulli distribution whose parameter is sought, and the time at which observations are made is under experimental control. They use a simple filtering technique originally suggested by [16] to maintain and update the prior distribution of the unknown parameters. In principle, their approach is similar to ours; the main difference is the nature of the observations, which in our case are the hitting times.

As posed, our optimization problem reduces to a slightly nonstandard partial differential equation (PDE) optimization problem, in particular a Fokker–Planck optimization problem. The field of PDE optimization is vast, and we just mention one of many good references, [4]. A series of papers on Fokker–Planck PDE optimization has also appeared; see, e.g., [1, 3]. None of these, however, deals with the boundary-based term that arises from the objective based on hitting times. This is done in our previous work [12], where we also solve a hitting time–

inspired optimization problem. The fundamental difference is that there we assume model parameters are known, whereas the uncertainty in the model parameters in the present work complicates the problem, as it requires the solution of a system of (loosely) coupled PDEs (the number depending on the form of the prior distribution), rather than just one as in [12]. That work, as well as the present paper, are part of growing efforts to control nerve cells (see, e.g., [25]).

The present paper is also part of the general literature on estimating parameters of diffusion processes from observations only of hitting times. Partly due to the applications of this problem in theoretical neuroscience, a lot has been written on the subject, for example in [23, 7, 13, 18], as well as in our own work [11], where we deal with the the more exotic case of a sinusoidal driving force. The literature on estimation from hitting times of the class of diffusion models we consider asserts that one parameter in particular, the characteristic time (i.e., the time constant of the LIF model), is the most difficult to estimate and has by far the widest confidence bounds. Thus, we focus on devising an input control that will provide the most accurate estimates for this parameter, assuming that the other parameters are known or at least nominally known in a sense that we explain below. We also shortly touch upon estimation of all parameters, showing that it is in fact possible to do.

This paper can be summarized as solving a PDE optimization arising from maximizing the MI between a prior of the characteristic time parameter and the random (future) hitting times. Its outline is as follows: in section 2 we give the mathematical formulation of the optimal design problem; in section 3 we describe the optimization procedure, in particular the derivation of the objective gradient and the numerical methods used; in section 4 we discuss the general shape of the optimal solution and its dependence on the various parameters in the model as well as the robustness of the optimization algorithm; in section 5 we show the performance of the optimally designed stimulation when many observations are first collected and only a single-pass, *batch* estimation of parameters is performed afterwards; in section 6 we turn to the *online* problem, where the estimation is performed after each observation and the stimulation is adjusted to reflect the newly updated prior of the unknown parameters; in section 7 we make concluding and forward-looking remarks.

2. Problem setup.

2.1. Mathematical model of a neuron.
Consider a noisy LIF model. This is a diffusion process that has an absorbing barrier at an upper boundary and is unbounded from below. In particular, for $t \geq 0$ and $k = 1, 2, \ldots$, it is governed by

$$dX(t) = \left(\alpha(t) + \frac{1}{\tau}(\mu - X(t)) \right) \mathrm{d}s + \sigma \, \mathrm{d}W(t), \quad X(0) = 0,$$

(1)
$$X(T_{\mathrm{sp}}) = x_{th} \implies \begin{cases} X(T_{\mathrm{sp}}^+) & = 0, \\ t_k & = T_{\mathrm{sp}}, \\ k & = k+1. \end{cases}$$

The random variables T_{sp} are the first hitting times of the process on the upper boundary, x_{th}, which is set to $x_{th} = 1$ without loss of generality. At each hitting time, the process is reset to 0 and then continues from there. Here, T_{sp}^+ denotes the right limit at T_{sp}. The control

$\alpha(t)$ is manipulated by us and can be altered at any time. The parameter set, $\theta = \{\mu, \tau, \sigma\}$, is assumed to be constant.

In the context of parameter estimation, we assume that the parameter set $\theta = \{\mu, \tau, \sigma\}$ or a subset thereof is unknown and only the hitting times, $\{t_k\}$, are observable. First we consider only estimating the characteristic time τ, assuming μ and σ are known. The literature on estimating parameters of LIF models from hitting times (see, e.g., [7, 18]) has at length discussed that the estimation of τ is clearly the hardest. In fact, the literature is not clear about whether τ is hard to estimate (practically unidentifiable, needing unrealistically large sample sizes) or whether it is structurally unidentifiable (hitting times contain no information on τ, which can therefore never be estimated, no matter the sample size). Since these studies were done in the context of no control, $\alpha = 0$, or at least α constant, we posit that a judicious choice for the shape of α can significantly improve the estimation of τ, so that it is in effect identifiable, also for realistic sample sizes.

The goal is to choose the control, $\alpha(t)$, such as to best estimate τ given that only the spike times $\{t_k\}$ are observed. Equivalently, setting $t_0 = 0$, we observe only the interspike intervals, $\{s_k\}$, defined as the differences between consecutive spike times, $s_k = t_k - t_{k-1}$.

The probability density of the length of the kth interval, which is a random variable denoted by S_k, conditional on an applied control $\alpha(\cdot)$, is denoted by

$$(2) \qquad g_k(s) := \frac{1}{dt} \mathbb{P}[S_k \in [s, s + dt] \,|\, \alpha(\cdot)].$$

We drop the index k when there is no confusion, and we sometimes write $g(s|\theta, \alpha(\cdot))$ to emphasize the dependence on the parameters and the control. The likelihood of the observed interspike intervals, $\{s_k\}_{k=1,\ldots,n}$, is then

$$(3) \qquad L(\{s_k\}_{k=1,\ldots,n}|\theta, \alpha(\cdot)) = \prod_{k=1}^{n} g(s_k|\theta, \alpha(\cdot)),$$

and the maximum likelihood estimator of θ is

$$(4) \qquad \hat{\theta} = \arg\max_{\theta} L(\{s_k\}_{k=1,\ldots,n}|\theta, \alpha(\cdot))$$

for a given applied control and resulting spike train.

The transition density for the state variable, X_t, $t \in [s_k, s_{k+1})$, is denoted by

$$(5) \qquad f(x,t) := \frac{1}{dx} \mathbb{P}\left[X_t \in [x, x + dx) \,|\, X_{s_k} = 0, X_s < 1 \text{ for } s_k < s < t\right].$$

The transition density satisfies a Fokker–Planck PDE over the domain, $x \in (-\infty, x_{th}]$, which we approximate with the domain $x \in (x_-, x_{th}]$, for some lower boundary x_- such that the probability that the process takes values smaller than this is negligible. For notational convenience, we denote the drift by $U(x,t) = -(\alpha(t) + (\mu - x)/\tau)$ and obtain

$$\begin{aligned} \partial_t f(x,t) &= \tfrac{1}{2}\sigma^2 \partial_x^2 f + \partial_x\Big(U(x,t)\,f\Big) \\ &= -\partial_x \phi(x,t), \end{aligned}$$

$$(6) \qquad \begin{cases} f(x,0) &= \delta(x), \\ \tfrac{1}{2}\sigma^2 \partial_x f + Uf\,|_{x=x_-} &= 0, \\ f\,|_{x=x_{th}} &= 0. \end{cases}$$

We have imposed *absorbing* boundary conditions at x_{th} and *reflecting* boundary conditions at x_-. The time-dependent probability flux at the threshold, $\phi(x_{th}, t)$, is important, as it is related to the density of hitting times, (2), via

$$(7) \qquad\qquad g(t) = \phi(x_{th}, t) = -\frac{1}{2}\sigma^2 \partial_x f|_{x=x_{th}},$$

since $Uf = 0$ at the absorbing boundary, x_{th}.

To see why (7) is true, note that the integral of $f(x, t)$ over the domain (x_-, x_{th}) equals the probability that X has not left the domain until time t. This is related to g through

$$\Pr[T_{\text{sp}} > t] = 1 - \int_0^t g(s)\, \mathrm{d}s = \int_{x_-}^{x_{th}} f(\xi, t)\, \mathrm{d}\xi.$$

Differentiating this expression with respect to t, we obtain

$$g(t) = -\partial_t \int_{x_-}^{x_{th}} f(\xi, t)\, \mathrm{d}\xi = \int_{x_-}^{x_{th}} \partial_\xi \phi(\xi, t)\, \mathrm{d}\xi = \phi(x, t)\,\Big|_{x_-}^{x_{th}}.$$

This gives (7), since $\phi(x_-, t) = 0$ is our *reflecting* boundary condition in (6); see [24] for details.

Intuitively, we wish to find a measure that quantifies the information the sampled spike train carries about the parameter as a function of the control. We take a Bayesian approach and assume some *prior distribution* of the random variable Θ over possible values of θ. We take a discrete prior distribution and denote it by

$$\rho(\theta) := \mathbb{P}(\Theta = \theta).$$

Given a single observation, T_{sp}, the parameter posterior distribution is

$$(8) \qquad\qquad p(\theta|T_{\text{sp}}; \alpha) = \frac{g(T_{\text{sp}}|\theta; \alpha)\rho(\theta)}{\int_\Theta g(T_{\text{sp}}|\theta; \alpha)\rho(\theta)\, \mathrm{d}\theta},$$

where $g(T_{\text{sp}}|\theta; \alpha)$ is the likelihood of the hitting time given in (2). We then choose the control, $\alpha(\cdot)$, that maximizes the MI between the two random variables, Θ and T_{sp}. Conditional on $\alpha(\cdot)$, the MI, I, is given by

$$(9) \qquad\qquad I[\alpha] = \int_\Theta \int_0^\infty g(t|\theta)\rho(\theta) \log\left(\frac{g(t|\theta)}{\int_\Theta g(t|\theta)\rho(\theta)\, \mathrm{d}\theta}\right) \mathrm{d}t\, \mathrm{d}\theta;$$

see Appendix A for the derivation. For different controls the MI will be different, since the hitting time density depends on the shape of α, whereas the prior, ρ, does not. Note that the Bayesian approach is only used to choose the control and thus obtain a sample spike train. The estimator is still the standard maximum likelihood estimator (4), and can take any value, including values which are not in the support of the prior nor in that of the posterior distribution. This will become important later, when a simple discrete prior distribution will be used and shown to be sufficient to find a control that improves the estimation over the

noncontrolled case. A discrete prior on only a few values of θ, in contrast to a continuous distribution over the entire parameter space, is chosen for computational reasons.

To summarize, we seek the control input $\alpha(\cdot)$ which maximizes I in (9), and we then verify that observations obtained under such optimal stimulation lead to parameter estimates that are more accurate and more precise than those obtained for observations when the system is perturbed by suboptimal stimulations, or no stimulation at all.

2.2. Experiment description. We consider two possible experimental setups. In the first, given a prior, ρ, a single stimulus control α^* is found by maximizing I; then that same waveform is applied N times, resetting it along with the voltage at each hitting time. We call this *batch estimation*.

In the second setup, we allow for updating α after each observed hitting time. The idea is that observations update the prior and that a different optimal $\alpha(\cdot)$ may correspond to this updated prior. We call this *online estimation*. Obviously, online estimation is computationally more demanding.

A subtle issue is that, even in the batch context, applying a *sequence* of not necessarily equal controls might result in more accurate estimation. However, there is no simple generalization of the MI-optimizing formalism that would produce a sequence of optimal controls, and so we leave this for future work.

3. Maximizing the mutual information. Consider the optimization problem of maximizing the MI in (9),

$$\alpha(\cdot) = \underset{\text{admissible } \alpha}{\arg\max} \ I[\alpha].$$

We consider as admissible controls those that are continuous and bounded, $\alpha(t) \in [\alpha_{\min}, \alpha_{\max}]$ for all t. First we calculate the gradient of I with respect to α and then use it in a standard gradient-based optimization procedure. There is the added complexity that the gradient here is an infinite-dimensional object, which raises some subtle functional analysis questions from a purely mathematical point of view; we refer the reader to the references [14, 4] for rigorous justification of our manipulations.

3.1. Calculation of the gradient using the adjoint method. A standard approach of deriving a gradient of a functional of a solution to a PDE with respect to an input is the *adjoint* method. We first rewrite (9) in terms of the transition density using (7),

$$(10) \qquad I[\alpha] = -\int_\Theta \int_0^\infty \partial_x f_\theta(x_{th}, t|\theta)\rho(\theta) \log\left(\frac{\partial_x f_\theta(x_{th}, t|\theta)}{\int_\Theta \partial_x f_\theta(x_{th}, t|\theta)\rho(\theta)\, \mathrm{d}\theta}\right) \mathrm{d}t\, \mathrm{d}\theta,$$

where the constant term $\sigma^2/2$ is dropped, as it is irrelevant for obtaining the maximizing value of α. The objective, I, is augmented with the dynamics of f by introducing the *adjoint* variable p:

$$I = -\int_\Theta \int_0^\infty \partial_x f_\theta(x_{th}, t|\theta)\rho(\theta) \log\left(\frac{\partial_x f_\theta(x_{th}, t|\theta)}{\int_\Theta \partial_x f_\theta(x_{th}, t|\theta)\rho(\theta)\, \mathrm{d}\theta}\right) \mathrm{d}t\, \mathrm{d}\theta$$

$$(11) \qquad + \int_\Theta \rho(\theta) \int_0^\infty \int_{x_-}^{x_{th}} p_\theta(x, t) \left(\partial_t f_\theta(x, t) + \partial_x \phi_\theta(x, t)\right) \mathrm{d}x\, \mathrm{d}t\, \mathrm{d}\theta,$$

where ϕ is the probability flux; see (6).

The added (11) is always zero for all choices of adjoint variable p_θ, since f satisfies the PDE in (6). However, by adding this null term and choosing p_θ in an appropriate way, we can remove the dependence of I on first-order perturbations of f. This allows one to compute the differential of I with respect to α, without needing to additionally compute the differential of f with respect to α as the chain rule would otherwise require. This is useful, since the differential of f with respect to α is a complicated mathematical object because any value of $f(x, t)$ depends on the entire history of $\alpha(\cdot)$, i.e., on all values $\alpha(s)$ for $s \leq t$.

The elimination of I's dependence on f is done by transferring the derivatives of f to p using repeated integration by parts. For completeness, we demonstrate in detail how this is done, dropping θ from the notation. First we introduce a finite (large) final time t_f where we can approximate $f(x, t_f) = 0$. Then

$$
\int_0^{t_f} \int_{x_-}^{x_{th}} p\,[\partial_t f + \partial_x \phi]\,\mathrm{d}x\,\mathrm{d}t = \int_{x_-}^{x_{th}} p\,f\Big|_0^{t_f}\,\mathrm{d}x - \int_0^{t_f} \int_{x_-}^{x_{th}} \partial_t p\,f\,\mathrm{d}x\,\mathrm{d}t
$$
$$
+ \int_0^{t_f} p\,\phi\Big|_{x_-}^{x_{th}}\,\mathrm{d}t - \int_0^{t_f} \int_{x_-}^{x_{th}} \partial_x p\,\phi\,\mathrm{d}x\,\mathrm{d}t.
$$

Using that $\phi = -Uf - \frac{1}{2}\sigma^2 \partial_x f$ yields

$$
-\int_0^{t_f} \int_{x_-}^{x_{th}} \partial_x p\,\phi\,\mathrm{d}x\,\mathrm{d}t = \int_0^{t_f} \int_{x_-}^{x_{th}} \partial_x p\,U f\,\mathrm{d}x\,\mathrm{d}t
$$
$$
+ \int_0^{t_f} \frac{1}{2}\sigma^2 \partial_x p\,f\Big|_{x_-}^{x_{th}}\,\mathrm{d}t - \int_0^{t_f} \int_{x_-}^{x_{th}} \frac{1}{2}\sigma^2 \partial_x^2 p f\,\mathrm{d}x\,\mathrm{d}t.
$$

Some simplifications can be made based on the terminal and boundary conditions for f. Since the initial conditions for f are fixed so that perturbations of f at $t = 0$ do not exist, perturbations of the control, α, do not perturb the initial conditions of f. The same is true for $t = t_f$. Thus, we disregard any terms involving $f(x, 0)$ or $f(x, t_f)$, since they are constant in α and therefore vanish when taking derivatives. Furthermore, $\phi(x_-, t) = 0$ and $f(x_{th}, t) = 0$. Let C be a constant; then the adjoint term simplifies to

$$
\int_0^{t_f} \int_{x_-}^{x_{th}} p\,[\partial_t f + \partial_x \phi]\,\mathrm{d}x\,\mathrm{d}t = C + \int_0^{t_f} \int_{x_-}^{x_{th}} \left[-\partial_t p - \frac{1}{2}\sigma^2 \partial_x^2 p + U \partial_x p\right] f\,\mathrm{d}x\,\mathrm{d}t
$$
$$
- \int_0^{t_f} \frac{1}{2}\sigma^2 p(x_{th}, t)\partial_x f(x_{th}, t)\,\mathrm{d}t
$$
$$
+ \int_0^{t_f} \frac{1}{2}\sigma^2 \partial_x p(x_-, t)f(x_-, t)\,\mathrm{d}t.
$$

The adjoint term thus breaks down into a constant term, a spatial term giving the (backwards) evolution for the adjoint variable, p, and two terms providing the boundary conditions for p.

Now replace the adjoint term in the original equation for I, (11),

$$I = \int_{\Theta} \rho(\theta) \left[- \int_0^{t_f} \partial_x f_\theta(x_{th}, t) \log \left(\frac{\partial_x f_\theta(x_{th}, t)}{\int_{\Theta} \partial_x f_\theta(x_{th}, t) \rho(\theta) \, d\theta} \right) dt \right.$$
$$+ \int_0^{t_f} \int_{x_-}^{x_{th}} \left[-\partial_t p_\theta - \frac{1}{2} \sigma^2 \partial_x^2 p_\theta + U \partial_x p_\theta \right] f_\theta \, dx \, dt$$
$$- \int_0^{t_f} \frac{1}{2} \sigma^2 p_\theta(x_{th}, t) \, \partial_x f_\theta(x_{th}, t) \, dt$$
$$\left. + \int_0^{t_f} \frac{1}{2} \sigma^2 \partial_x p_\theta(x_-, t) \, f_\theta(x_-, t) \, dt \right] d\theta.$$

Consider the effect of small perturbations, $\delta\alpha$, of the control on our objective,

$$\delta I_\epsilon = \frac{I(\alpha + \epsilon \delta\alpha) - I(\alpha)}{\epsilon}.$$

The natural assumption is that given $\alpha + \epsilon\delta\alpha$, there is a corresponding solution to the PDE, $f + \epsilon\delta f$. Taking the limit of $\epsilon \to 0$, we get

(12)
$$\delta I = \int_{\Theta} \rho(\theta) \int_0^{t_f} \left[- \partial_x \delta f_\theta(x_{th}, t) \log \left(\frac{\partial_x f_\theta(x_{th}, t)}{\int_{\Theta} \partial_x f_\theta(x_{th}, t) \rho(\theta) \, d\theta} \right) \right.$$
$$- \partial_x \delta f_\theta(x_{th}, t) + \frac{\partial_x f_\theta(x_{th}, t) \int_{\Theta} \rho(\theta) \partial_x \delta f_\theta(x_{th}, t) \, d\theta}{\int_{\Theta} \partial_x f_\theta(x_{th}, t) \rho(\theta) \, d\theta}$$
$$+ \int_{x_-}^{x_{th}} \left[-\partial_t p_\theta - \frac{1}{2} \sigma^2 \partial_x^2 p_\theta + U \partial_x p_\theta \right] \delta f_\theta \, dx - \int_{x_-}^{x_{th}} \partial_x p_\theta \, f_\theta \, \delta\alpha \, dx$$
$$\left. - \frac{1}{2} \sigma^2 p_\theta(x_{th}, t) \, \partial_x \delta f_\theta(x_{th}, t) + \frac{1}{2} \sigma^2 \partial_x p_\theta(x_-, t) \, \delta f_\theta(x_-, t) \right] dt \, d\theta.$$

Thus, δI depends on both $\delta\alpha$ and δf_θ. In fact, it only depends on $\delta\alpha$ through the third-to-last term and on δf_θ through all the others. However, we can select p_θ in a judicious matter, such that the dependence on δf vanishes. To eliminate δf on the lower boundary, x_-, we need that

$$\partial_x p(x_-, t) = 0.$$

To eliminate the spatial dependence on δf, we enforce that p evolves (backwards) as

$$\partial_t p + \frac{1}{2} \sigma^2 \partial_x^2 p - U \partial_x p = 0.$$

Note that, by interchanging the order of integration with respect to θ,

$$\int_{\Theta} \rho(\theta) \int_0^{t_f} \left[- \partial_x \delta f_\theta(x_{th}, t) + \frac{\partial_x f_\theta(x_{th}, t) \int_{\Theta} \rho(\theta) \partial_x \delta f_\theta(x_{th}, t) \, d\theta}{\int_{\Theta} \rho(\theta) \partial_x f_\theta(x_{th}, t) \, d\theta} \right] dt \, d\theta$$
$$= \int_{\Theta} \rho(\theta) \int_0^{t_f} \left[-1 + 1 \right] \partial_x \delta f_\theta(x_{th}, t) \, dt \, d\theta$$
$$= 0.$$

Finally, to eliminate the dependence of δI on $\partial_x \delta f_\theta$, we apply the simple boundary condition,

$$p_\theta(x_{th}, t) = -\frac{2}{\sigma^2} \log \left(\frac{\partial_x f_\theta(x_{th}, t)}{\int_\Theta \partial_x f_\theta(x_{th}, t) \rho(\theta) \, d\theta} \right).$$

In summary, the adjoint variables, p_θ, must satisfy the following adjoint PDE:

$$\partial_t p_\theta = -\frac{1}{2} \sigma^2 \partial_x^2 p_\theta + U(x, t; \theta) \partial_x p_\theta,$$

(13)
$$\begin{cases} p_\theta(x_{th}, t) & = -\frac{2}{\sigma^2} \log \left(\frac{\partial_x f_\theta(x_{th}, t)}{\int_\Theta \partial_x f_\theta(x_{th}, t) \rho(\theta) \, d\theta} \right), \\ \partial_x p_\theta \big|_{x = x_-} & = 0, \\ p_\theta(x, t_f) & = 0. \end{cases}$$

This finally provides us with a simple equation to calculate the objective gradient. Given p_θ and f_θ, the differential of I in (12) with respect to the control $\alpha(t)$ is

(14)
$$\delta I = \left[-\int_\Theta \rho(\theta) \left(\int_{x_-}^{x_{th}} \partial_x p_\theta \, f_\theta \, dx \right) d\theta \right] \delta\alpha.$$

The computational details on how to calculate δI are given in the next section.

3.2. Gradient ascent procedure. Maximization of the MI, (10), is approached as a gradient-based iterative method, where the infinite-dimensional gradient, δI in (14), is calculated by some finite-dimensional approximation. The process involves three basic stages:

1. Given the current control iterate, $\alpha_n(t)$, numerically solve for the corresponding f, p from their respective PDEs, (6) and (13).

2. Form the adjoint differential, $\delta I / \delta\alpha$, using (14).

3. Increment α_{n+1} in the direction of increasing $\delta I / \delta\alpha$,

(15)
$$\alpha_{n+1}(t) = \alpha_n(t) + \frac{\delta I}{\delta\alpha}(t) \, \Delta\alpha$$

for some step-size $\Delta\alpha$.

In practice, we need to align the various discretization grids, and the simplest thing to do is to discretize α, f, p at the same time points, t_j. The gradient evaluated at t_j is

$$\frac{\delta I}{\delta\alpha}(t_j) = -\int_\Theta \rho(\theta) \left(\int_{x_-}^{x_{th}} \partial_x p_\theta(t_j, x) f_\theta(t_j, x) \, dx \right) d\theta.$$

Equation (15) is the simplest gradient ascent scheme one can devise. The literature on gradient-based optimization [20] and the subset on PDE-based gradient optimization [4] offers more sophisticated schemes, but for simplicity we start with this basic one. An example of a more sophisticated scheme is the nonlinear conjugate-gradient ascent, as recommended in the literature on Fokker–Planck control [2]. A full description of the optimization algorithm is provided in the supplemental material.

In general, there is no *a priori* way to select the terminal time, t_f. For practical purposes we proceed as follows. Given a set of plausible parameters we set t_f large enough for $f \approx 0$,

and we further select a subinterval, $[0, t_{opt}] \subset [0, t_f]$, such that we only seek to optimize over $t \in [0, t_{opt}]$ and for $t > t_{opt}$ we let $\alpha = \alpha_{\max}$. This should ensure that

$$\int_{t > t_f} g(t) \, \mathrm{d}t \ll 1.$$

Finally, the calculation of the posterior distribution should be addressed. For computational reasons, we need to approximate the integral with respect to θ by a sum. For a given $\alpha(t)$, we solve for p, f for a few sampled values of θ from the updated prior distribution $\rho(\theta)$ and their corresponding probabilities. This is equivalent to assuming a discrete or a particle prior, i.e., a sum of weighted Dirac delta masses. In the simplest case, these are N values with equal probability,

(16) $$\rho(\tau) = \sum_i \tfrac{1}{N} \delta(\tau - \tau_i).$$

4. Properties of the optimal control. We now investigate the optimization algorithm and the properties of the optimal control. In the calculations below, we set $\mu = 0$, $\sigma = 1$, $t_f = 15$, $t_{opt} = 10$, $\alpha_{\min} = -2$, $\alpha_{\max} = 2$ unless otherwise stated. We use the same numerical scheme, the Crank–Nicolson method, for the integration of the forward and adjoint PDEs as in our previous work, [12], as this is a well-known classic technique for parabolic PDEs.

In the computational neuroscience literature this parameter set for μ, σ corresponds to what is known as the subthreshold and high-noise regime [11], as long as the control α equals zero. Increasing the control to $\alpha > 1$ will move it into the qualitatively different superthreshold regime, where the process can hit the threshold even in the absence of noise. In the subthreshold regime, a nonzero noise parameter is required if any spikes are to occur.

4.1. Optimization algorithm. A single control increment is shown in Figure 1, while an example of a full optimization ascent is given in Figure 2 using an assumed prior with two possible values of τ, namely 1/4 and 4.

Intuitively, the control tries to separate the hitting time distributions, such that an observed hitting time can more clearly be attributed to one of the two potential parameter values, τ. For the combination of known parameters μ, σ and prior ρ, used in Figure 2, the optimal control first suppresses firing for an initial time interval and then maximally stimulates firing in the remaining time. It is visually obvious that the result on $g(t|\tau)$ is to go from a case where the two hitting time distributions overlap to one where they are clearly delineated. Thus, given an observation from the optimally stimulated system, one can more confidently estimate what the underlying value of τ was that produced the observation.

For reference, a single optimization ascent takes on the order of 5 seconds on a dual-core i5 Intel laptop (2.67GHz). We should also point out that in general the computational time scales linearly with the number of points in the prior. However, for certain points in the prior a much finer discretization grid may be warranted to accurately solve the PDE.

4.2. Switch point sweep. Given the optimal solution obtained in Figure 2, we suspect that the optimal solution is approximately of a simple switching type, in the sense of first holding everything back and then exciting maximally ("inhibit-excite"). It is interesting to
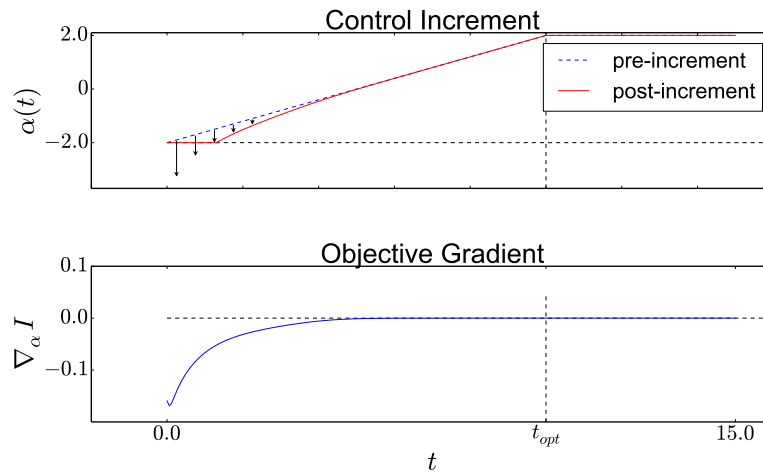
Figure 1. *Illustration of how the control is incremented given the gradient. Top panel: the initial control in dashed and the incremented control in solid with arrows indicating the incremental change. Bottom panel: the corresponding value of $\delta I/\delta \alpha(t)$. Note that we only optimize $\alpha(t)$ for $t \in [0, t_{opt}]$ and set $\alpha = \alpha_{max}$ for $t \in [t_{opt}, t_f]$. Note that we have applied the thresholding at α_{min} for the post-increment, so that $\alpha \in [\alpha_{min}, \alpha_{max}]$.*

check how sensitive the objective is to the exact value of the switching time, t_{switch}. If it is not very sensitive, it would be practically useful, as the exact optimization of the PDE-based objective will not be necessary, and any such switching-type solution will achieve satisfactory improvement in the estimation. We therefore sweep through different points of exactly when the switch from inhibition to excitation occurs, using a sharp sigmoid function, $\alpha(t) = \alpha_{max} \cdot \tanh(4(t - t_{switch}))$, as the applied control, and see how it impacts the resulting objective, I. The results are in Figure 3. It is clear that, at least for this parameter set, there is an improvement in the objective by delaying the switching time, with later switching times producing similar and even marginally better objective values than the objective corresponding to the solution of our iterative optimization. This suggests that the optimization procedure has trouble finding the exact maxima either due to initial conditions of the optimizer or other numeric artifacts. However, the differences in terms of objective value are marginal.

4.3. Sensitivity of the optimization method. In this subsection, we explore the effect of the number of points in the prior, the variance of the prior, the initial guess for the control $\alpha(\cdot)$, and the values of the known parameters μ and σ.

Number of points and shape of the prior. It is of interest to know whether the exact shape of the prior matters or whether the first two moments (mean and variance) are sufficient to determine the optimal control. We make a simple experiment where we compute the optimal control for two similar priors, the first using 2 points in the prior and the second using 11 points, such that both priors have the same log-mean and log-variance. Results are presented in Figure 4, where it is seen that the optimal control is essentially the same. This suggests that, for practical purposes, it suffices to choose a 2-point prior that matches the (log-) variance of the more detailed prior distribution. This would substantially reduce the computational burden.
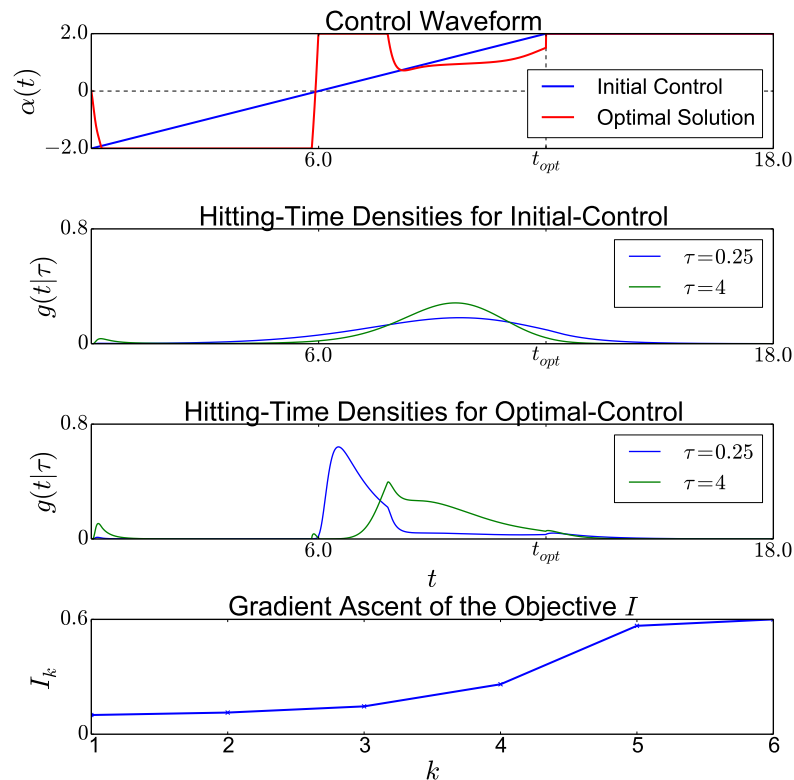
Figure 2. *Example of the gradient ascent algorithm for the optimization of I in (10) illustrating that the effect of the optimal control is to separate the hitting time distributions corresponding to the potential values of τ, here for two values of τ. Top panel: initial and optimal controls, $\alpha_0(t)$ and $\alpha_{opt}(t)$. Second panel: hitting time distributions $g(t|\tau)$ conditional on using the initial control, α_0. Third panel: $g(t|\tau)$ conditional on using the optimal control, α_{opt}. Bottom panel: progress of the mutual information, (9), as a function of the gradient ascent iterations (here the algorithm converged in nine iterations). The optimal control significantly improves the objective in comparison to the initial guess.*

The dispersion of the prior. It is also of interest to know how the dispersion of the prior, here quantified by the log-variance, changes the optimal control. Again, we take two 2-point priors, one with particles at $\tau \in \{1/4, 4\}$ and the other at $\tau \in \{3/4, 4/3\}$. Results are shown in Figure 5. It is clear that while the general shape of the optimal control is the same, regardless of the width of the prior, there is some difference. We further investigate the practical difference between the two priors in section 5.

The initial guess for the control α_0. Like all gradient-based optimization procedures, the algorithm is sensitive to the choice of initial guess for the independent variable, i.e., the choice of α_0. We illustrate this in Figure 6. For three different priors, namely a wide, a medium, and a narrow prior, we run the optimization scheme from four different initial guesses for α_0:

Mutual information as function of switching time
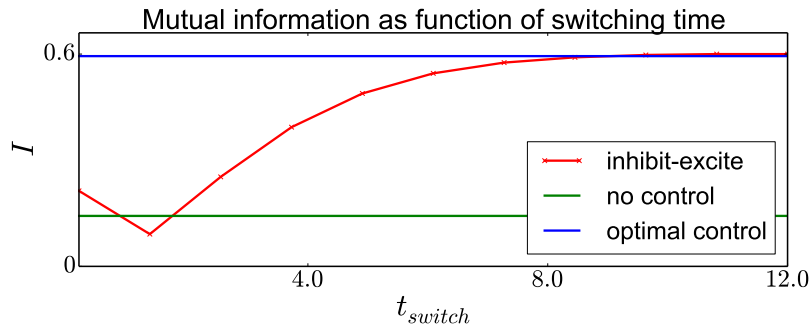
Figure 3. *Effect of switching time, t_s, in an inhibit-excite control on the MI. The applied control is a sharp sigmoid function, $\alpha(t) = \alpha_{max} \cdot \tanh(4(t - t_{switch}))$. The red line is the value of the MI as a function of the switching time, t_s, between maximally inhibitive to maximally excitatory control. The green line, plotted for reference, is the value of the mutual information for no control ($\alpha = 0$ for all t). And the blue line is the objective value achieved by the optimal control as Figure 2; all simulation parameters are also as in Figure 2.*
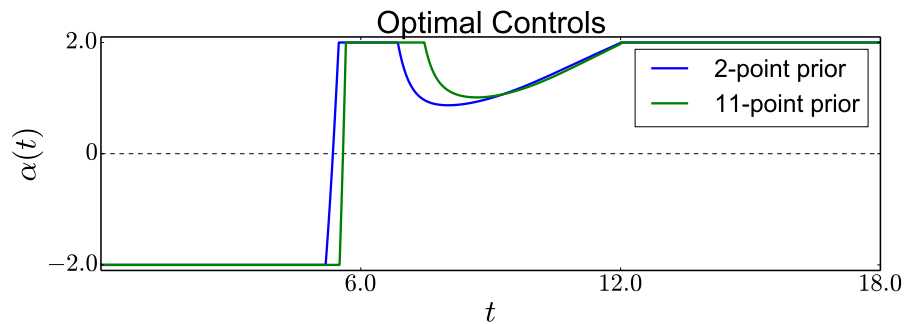
Optimal Controls

Figure 4. *Effect of the number of points in the prior on the optimal control. The blue curve was calculated with a 2-point prior and the green curve with an 11-point prior. The priors have the same log-mean and log-variance. For the 2-point prior, we use $\tau = [0.5, 2]$; for the 11-point prior, we set the weights at $\tau_n = \exp(n\Delta\tau)$, where n runs from -5 to 5 and $\Delta\tau$ is chosen to match the log-variance of the 2-point prior. All other parameters are as in Figure 2.*
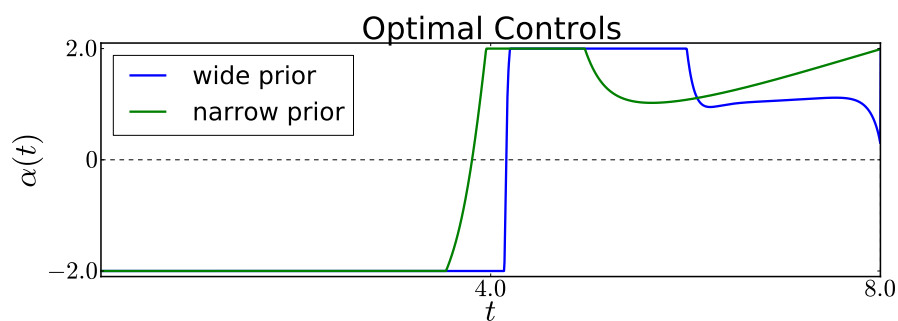
Optimal Controls

Figure 5. *Effect of the variance of a 2-point prior on the optimal control. The wide prior has probability 0.5 for both the values $\tau = 1/4$ and 4, while the narrow prior has points positioned much closer to each other, $\tau = 3/4$ or $4/3$.*

(a) Wide prior

(b) Medium prior
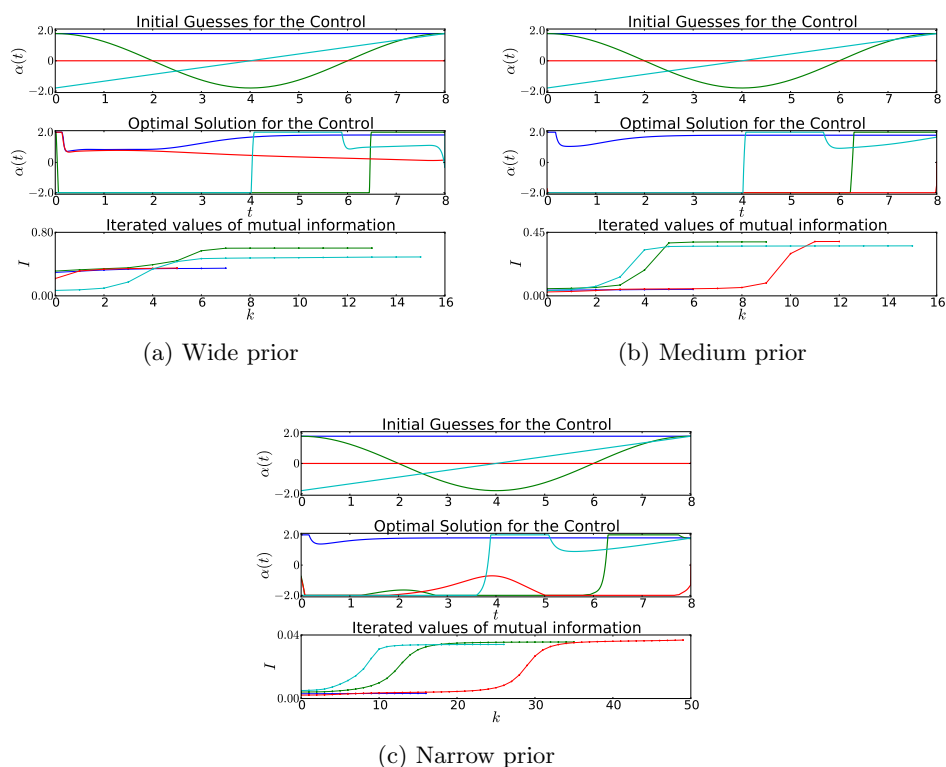
(c) Narrow prior

Figure 6. *Illustration of the dependence of the optimization scheme on the initial guess for the optimal control for* (a) *a wide prior,* (b) *a medium prior, and* (c) *a narrow prior. In the top panels the initial guesses are shown; in the middle panels the corresponding solutions obtained by the optimization routine are shown, and in the bottom panels the corresponding evolutions of the MI are shown. The various initial conditions considered are maximum stimulation for all t (blue); sinusoidal stimulation (green); no control at all, $\alpha = 0$ (red); and a linearly increasing control from min to max (cyan).*

1. zero for the entire (optimization) time interval;

2. linearly increasing from the lower to the upper control bound over the optimization interval;

3. a sinusoidal wave from maximum to minimum and again to maximum;

4. maximum for the entire optimization time interval.

In Figure 6, we see that for different initial guesses the optimization routine finds different optimal controls (see middle panels). An initial guess which is nonconstant, here linearly increasing or sinusoidal, is better than the constant initial guesses (the MI is larger; see lower panels). It is also seen that the shape of the prior does not have much influence on the optimal control found by the algorithm (the optimal solutions for the same initial guesses but different priors are close, as one can see by comparing the middle panels). The final value of the MI is larger for a wider prior (see lower panels), but this is a scaling issue, since of course the final parameter estimation will be the same if the control is the same. To summarize, the initial guess of the control is important, and different initial guesses should be tried to see whether

the objective can be improved, whereas the choice of prior distribution is less important.

Values of the known parameters μ, σ. So far we have restricted ourselves to one scenario of the assumed known parameters, $\mu = 0$ and $\sigma = 1$.

We also redid the basic analyses by varying these parameters, analyzing six qualitatively different sets of parameter values. In all cases the same 2-point prior was used, with probability one half for each of $\tau = 1/2$ and 2. The following parameter sets were investigated: $(\mu, \sigma) = (-0.5, 1), (0.1, 1), (1, 1), (0, 0.3), (0, 0.9)$, and $(0, 1.5)$. In all cases the algorithm succeeded in finding an optimal control (a different one for each pair) that significantly increased the MI relative to that of the initial control guess (results not shown). Although not an exhaustive study, this strongly suggests that our optimal design procedure is applicable and would be effective for a wide range of the values of μ, σ.

5. Batch estimation. We now proceed to generate a large sample of hitting times with several types of controlled stimulation and then estimate the unknown parameter after observing the whole sample. For each controlled stimulation, that same waveform is applied N times, resetting it along with the voltage at each hitting time. We call this batch estimation as opposed to online estimation. In the latter case, parameter estimates are updated after every single observation, and the control is updated according to an updated prior distribution. Online estimation is performed in section 6.

We assume that the true parameters governing (1) are as before: $\mu = 0$, $\tau = 1$, and $\sigma = 1$. We will assume σ, μ are known and only estimate the time constant τ, so we will maximize the MI between T_{sp} and τ. To obtain the optimal control, we take a discrete uniform prior on τ, using 10 points uniformly spaced in the interval $[1/4, 4]$; i.e., each point in the prior has probability $1/10$. Note that neither the mean nor the log-mean of the prior corresponds to the unknown true τ or $\log(\tau)$.

The obtained control is denoted by *opt*. In addition, we consider two optimal controls obtained using a wide and a narrow 2-point prior, as shown in Figure 5. We call these *opt-wide* for the prior with $\tau = [1/4, 4]$ and *opt-narrow* for the prior with $\tau = [3/4, 4/3]$. We also use the control *crit* where α is set to the fixed critical value $\alpha = \alpha_{crit} = x_{th}/\tau = 1$. The critical value defines the smallest value for α such that the process converges to the spiking threshold in the absence of noise (i.e., for $\sigma = 0$). Finally, we use a control *max* set to its upper bound, $\alpha = \alpha_{\max}$.

5.1. The estimation algorithm for the batch problem. In this section, we only aim to estimate one parameter, τ, which amounts to single-variable optimization. The negative log-likelihood of an observed hitting time set $\{t_n\}$ is

$$(17) \qquad l(\tau) = -\sum_n \log(g(t_n|\tau)) = -\sum_n \log\left(-\frac{\sigma^2}{2}\partial_x f(t_n|\tau)|_{x_{th}}\right).$$

We minimize (17) using the standard single-variable optimization routine in NumPy based on Brent's method.

5.2. Comparison of estimators. To summarize, we will stimulate the system using five stimulation waveforms:

    1. *opt*: the optimal control α_{opt}, based on a 10-point uniform prior between $[1/4, 4]$;

Table 1

Results for the estimates arising from simulations using various values of α (opt, crit, max). In each subtable there are N_b parameter estimates for each distinct α, with N_s hitting times used to form a τ-estimate. The true value of τ is $\tau = 1$. See also Figure 7.

Control type	mean($\hat{\tau}$)	std($\hat{\tau}$)	Control type	mean($\hat{\tau}$)	std($\hat{\tau}$)
opt	0.887	0.27	opt	1.011	0.08
opt-wide	0.766	0.32	opt-wide	0.946	0.20
opt-narrow	0.826	0.30	opt-narrow	1.005	0.07
crit	0.913	0.53	crit	0.797	0.07
max	0.872	0.37	max	0.787	0.06
(a) $N_b = 10^3, N_s = 10^2$			(b) $N_b = 10^2, N_s = 10^3$		

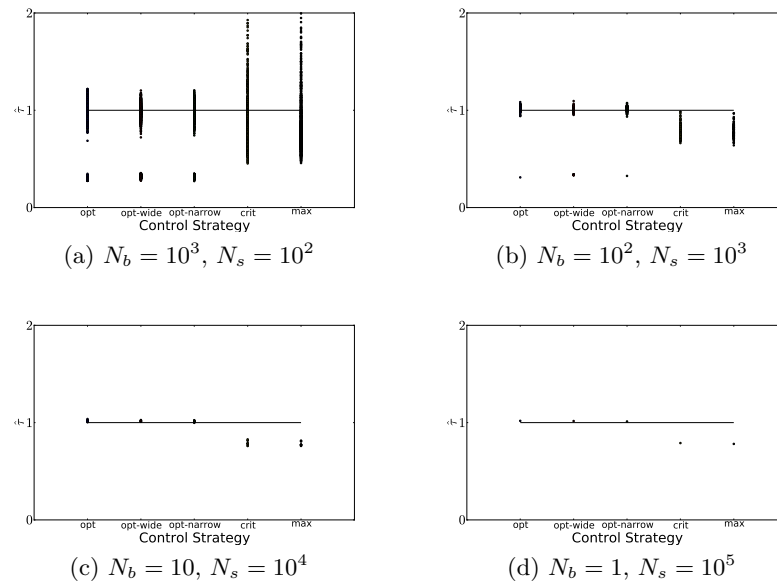| Control type	mean($\hat{\tau}$)	std($\hat{\tau}$)	Control type	mean($\hat{\tau}$)	std($\hat{\tau}$)
opt	1.018	0.01	opt	1.018	-
opt-wide	1.016	0.01	opt-wide	1.016	-
opt-narrow	1.012	0.01	opt-narrow	1.012	-
crit	0.791	0.02	crit	0.790	-
max	0.782	0.02	max	0.781	-
(c) $N_b = 10, N_s = 10^4$			(d) $N_b = 1, N_s = 10^5$		

2. *opt-wide*: the optimal control using a 2-point prior on $[1/4, 4]$;

3. *opt-narrow*: the optimal control using a 2-point prior on $[3/4, 4/3]$;

4. *crit*: the constant control α_{crit}, $(= x_{th}/\tau)$;

5. *max*: the max constant control, $\alpha_{\max}$ $(= 2)$.

We simulate N_b blocks of N_s hitting times each for each of the five α's above. That is, we apply each control N_s times to obtain N_s hitting times, and after each hitting time the control is reset. Thus, observations are independent and identically distributed. We then estimate τ over each set using maximum likelihood with the computed expression for the density, $g(t|\tau; \alpha(t))$. For a fair comparison between controls, we use the same Gaussian random draws per block of N_s hitting times for each control. The estimation results are tabulated in Table 1, and scatterplots of estimates can be found in Figure 7.

Together, Table 1 and Figure 7 confirm a clear advantage to using the optimal control, α_{opt}, over the simpler, constant controls. In particular, the bias of the estimates is significantly reduced, except for $N_s = 100$, where the bias is comparable. From Figure 7(a) it is seen that this is caused by the presence of a few outlier estimates for the optimal control, but for most data sets, the optimal control is clearly better. These outliers disappear for larger sample sizes.

Comparing the different optimal controls based on various priors, *opt* vs. *opt-wide* or *opt-narrow*, we see that they all perform approximately the same. This confirms the findings of the previous section that the number of points in the discrete prior is not that essential, and computational time can be considerably reduced by choosing priors with only a few points.

5.3. Estimating the entire parameter set. So far we have assumed that the drift parameter μ and the variance parameter σ are known. We now relax that assumption and attempt to also estimate these two parameters, in addition to the characteristic time τ. We will still use

(a) $N_b = 10^3$, $N_s = 10^2$

(b) $N_b = 10^2$, $N_s = 10^3$

(c) $N_b = 10$, $N_s = 10^4$

(d) $N_b = 1$, $N_s = 10^5$

Figure 7. *Visualization of the maximum likelihood estimates of τ for the different controls. Note that in (d) there is only one estimate per control; all simulated hitting times are used in that estimation. See also Table 1.*

the optimal control obtained assuming μ, σ are known. Why is this expected to work? First of all we found that the optimal control seems largely insensitive to the exact values of μ, σ as long as they are not too far from the nominal values ($\mu = 0$, $\sigma = 1$). This is in agreement with the findings of [8], where it is shown that the estimation of the time constant in the Ornstein–Uhlenbeck process from observations of the process itself can be hugely improved by perturbating the process, leaving the estimation of the other parameters largely unaffected. Moreover, as discussed in the introduction, the τ parameter is well known to be the hardest to estimate. Thus, a scheme designed to estimate it well is conjectured to help with estimation of the entire parameter set, since reducing fluctuations in one parameter estimate is expected to reduce (or leave unaffected) fluctuations in the others. Finally, the computational complexity quickly increases when optimizing the MI for three-dimensional priors, due to the potentially much larger number of PDEs required to be solved.

We use the same simulated hitting time data set as in Table 1 and Figure 7, using $N_s = 10^4$ hits and $N_b = 10$ separate samples; i.e., we make 10 estimates. The results in Figure 8 make it clear that the optimally stimulated samples yield much more accurate and precise estimates. In particular, all stimulation schemes come up with high-fidelity estimates for σ; it is in resolving the interplay between μ and τ that the optimal stimulation excels. The constant stimulations do not allow one to distinguish a high μ and a small τ from a low μ and a high τ; this can be most easily seen in the scatterplot on the right of Figure 8, where the constant stimulations tend to overestimate μ and greatly underestimate τ. On the other hand, the MI-optimizing dynamic stimulations do not have this erroneous correlation estimating both parameters fairly accurately, without a noticeable relation between the estimation errors.

(a) Individual estimates (b) μ vs. τ scatterplots
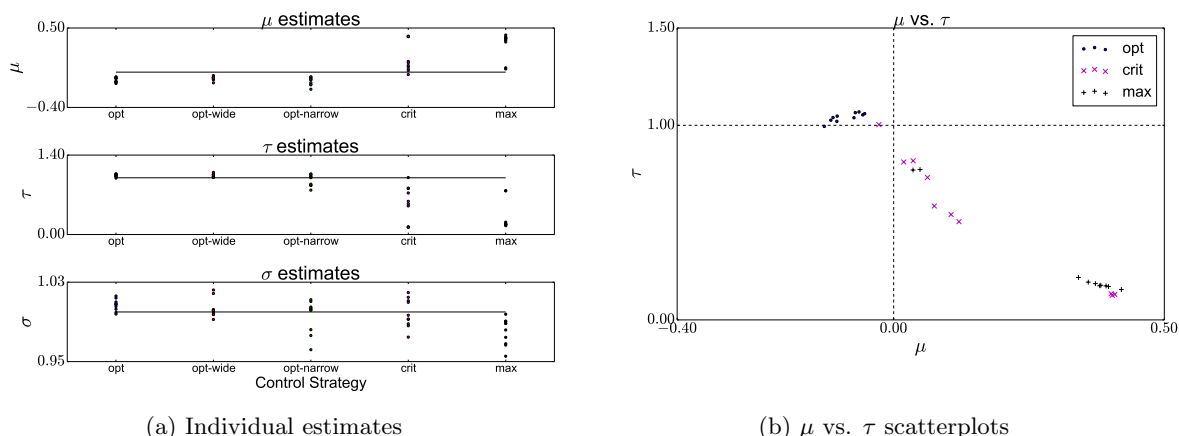
Figure 8. *Optimal stimulation significantly improves estimates when all model parameters are jointly estimated. In the left panel, we show $N_b = 10$ estimates of μ, τ, and σ, formed after observing $N_s = 10^4$ hitting times per estimate. There are three optimal stimulations (opt, opt-wide, and opt-narrow) and two constant ones (crit and max). The scatter dots are the individual estimates per stimulation-parameter pair. The horizontal line indicates the true value of the parameters, i.e., the value used to simulate the hitting times in the data samples. In the right panel, we focus on the interaction between the μ vs. τ estimates produced by the opt, crit, and max stimulations. The constant stimulations tend to consistently overestimate μ and underestimate τ.*

6. Online estimation. In the *online* optimization, we proceed as follows:

1. Find α_{opt} using the gradient ascent for the prior ρ.
2. Apply α_{opt}, and measure N_s hitting times t_k, $k = 1, 2, \ldots, N_s$.
3. Update the ρ into a posterior conditional on the observed $\{t_k\}$.
4. Recalibrate α_{opt} using the new ρ; i.e., go back to 1.

We progressively double N_s during the course of the simulation, starting from $N_s = 1$, so we recompute α_{opt} after the 1st spike and then after the 3rd, 7th, and so on. The heuristic is that at the beginning of the experiment the updated prior would be changing more rapidly than after already having observed several hundred spikes.

Computational efficiency considerations aside, we have all the tools to carry out points 1, 2, and 4 above. For point 3, updating the prior, we use the same particle filtering scheme as detailed in [9], in particular section 4 therein; the original reference is [16]. The algorithm is also provided in the online supplemental material, linked from the main article webpage.

6.1. Single hitting time illustration. In Figure 9, we illustrate one iteration of the update, that is, one hitting time given a stimulation from either the MI optimal control, or from a random control or from zero control; the random control just gives a randomly chosen constant stimulation for the computation of each hitting time. Lower values of τ imply a higher restoring force of the LIF process towards the origin and therefore longer time between hitting times. Since in this sample the observed hitting times were fairly long, especially for the MI-optimal stimulation, weights for smaller τ's grow larger, while weights for bigger τ's become smaller. However, it is immediately clear that the MI-optimal stimulation is more discerning, as it has almost entirely discarded (correctly) the possibility that $\tau > 2$, while the
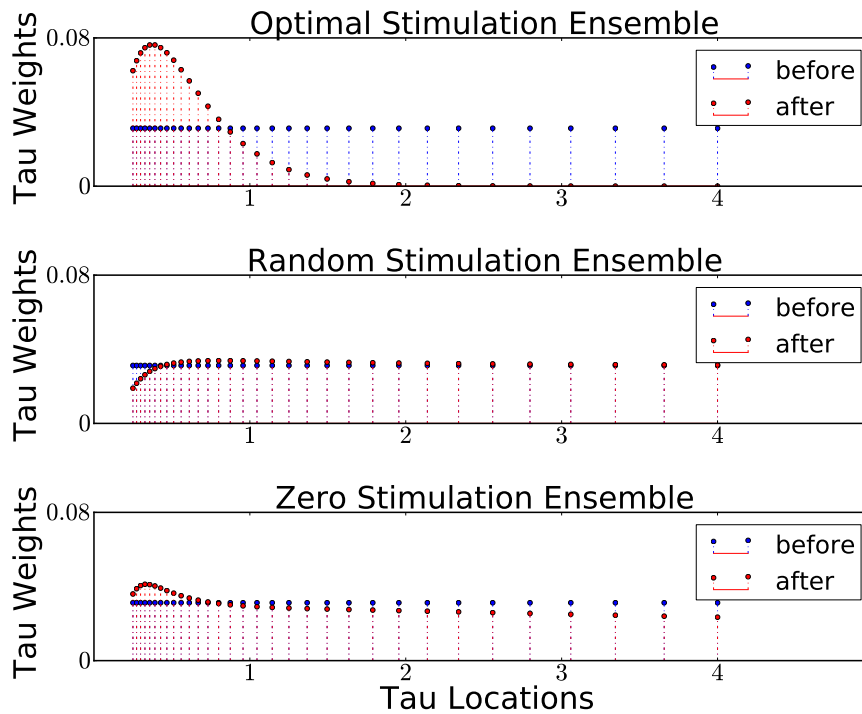
Figure 9. *Examples of a single iteration of the Online Stimulation-Estimation scheme. The actual hitting time observations were* 8.16 *(optimal stimulation),* 1.02 *(random stimulation),* 7.62 *(zero-stimulation).*

other two stimulations have resulted in only mild perturbations of the prior distribution.

6.2. Full multiple hitting time experiment. Finally, let us compute an entire estimation experiment. We stimulate the process and obtain a sequence of hitting times and online update our parameter prior distribution after every observation. We also online update our MI-optimal stimulation as the prior evolves after the 1st, 3rd, 7th, and so on observations.

We repeat and simulate 50 independent experiments of 500 hitting times each and then average the obtained ensembles of estimates. The aggregated evolutions for the updated prior distributions are visualized in Figure 10. We plot the quantiles of the running means of the updated priors for the N experiments. It is immediately clear that the MI-optimal control produces more accurate and more precise estimates and that the increase in precision is especially clear in the earlier parts of the experiment where less information is available. That is, using the MI-optimal control protocol reduces the uncertainty in the parameters much faster.

The fact that the base cases of constant controls are neither accurate nor particularly precise is consistent with the general results on the estimation of the characteristic time parameter. The novel finding is that a suitably tuned stimulation can result in much more accurate estimates.

It turns out that for the problem and the illustrative parameters, online upating the prior
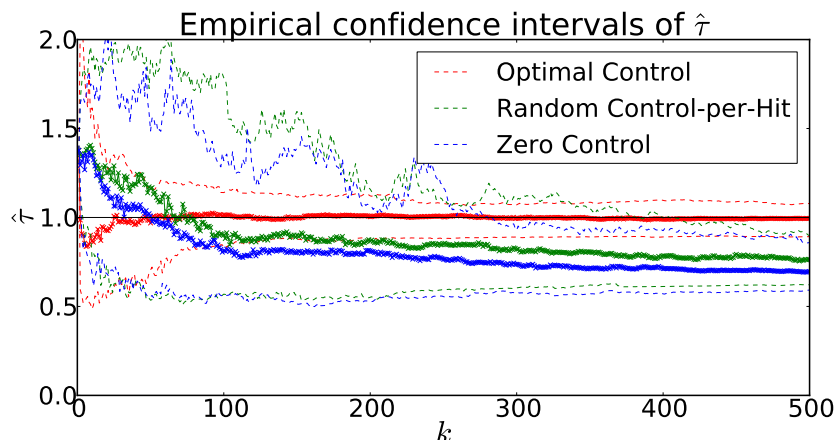
Figure 10. *Optimal-designed stimulation produces more precise and more accurate parameter estimates. Visualized are the median and 5/95th percentiles of the mean of the updated prior distribution over the $N = 50$ experiments. That is, for each experiment, j, and each hitting time, k, we compute the mean of the updated prior $\rho_{j,k}(\tau)$, and then we plot the median (crosses) of these N means as well as the third smallest and third largest of them (dashes). The black line indicates the true value of τ used to generate the synthetic observations. Again we see that the median and the extreme quantiles of the optimally stimulated estimates are much closer to the true value of the parameter, τ, than the estimates obtained from the naive stimulations.*

does not result in significant differences in the computed optimal control (results not shown). This is somewhat expected given the results in subsection 4.3. Therefore, if recomputing the optimal control is relatively time-consuming (as might be the case in neural experiments), one can just apply the batch control with comparable performance.

7. Discussion. We introduced a method for optimal design given hitting time observations. This contrasts with earlier work where the control is based on observations of the state variable of the SDE whose parameters are to be estimated. Our method is based on maximizing the MI between the observed hitting times and the posterior distribution of the parameters. The optimal control tends to separate the hitting time distributions associated with alternative values of the unknown parameter, thereby facilitating the identification of the parameter once an observation is made. The simulations show that the resulting estimates from the optimally stimulated system have higher precision and accuracy than sensible alternatives, such as random stimulation, critical stimulation at $\alpha = x_{th}/\tau$, or no stimulation at all.

Since the MI is expressed in terms of the hitting time density and the hitting time density can be related to the boundary term of a Fokker–Planck PDE, we approach the maximization of the MI as a PDE optimization problem. The standard use of an adjoint variable to obtain the objective gradient is made more complicated by the outer integration with respect to the unknown parameter prior. However, we are still able to derive the adjoint PDEs and thus to compute the objective gradient despite this added complication.

Our PDE-based adjoint optimization methods can be expensive and sensitive to the various parameters of the optimization. The main computational cost is the numerical solution to the forward and backward PDEs, which have to be solved numerous times during the op-

timization search. In principle, one has to solve as many PDEs as there are particles in the prior distribution. However, we have seen that in practice it is sufficient to use very coarse approximations to the prior. For example, as few as two or three particles which match the mean/variance of the more detailed higher-resolution prior seem to be sufficient to obtain the same optimal control as that obtained by using all the particles in the prior, at least for our problem.

In general, the optimally stimulated samples give (much) more accurate estimates than the ones that are naively stimulated. However, we see in Figure 7(b) that occasionally (once), the optimally stimulated sample can give very wrong estimates. We have taken a closer look at what happens in those situations. It seems that in the bad cases, there are enough extreme (i.e., very long) hitting times that the estimation procedure infers a very small characteristic time (i.e., the attractive force towards $\mu = 0$ is very strong). Generally in long samples such extreme hitting times are relatively rare, and the correct parameter value is inferred.

Our results point to a novel way to estimate the characteristic time of the LIF process. Parameter estimation techniques applied to the hitting time process have reported difficulties for this particular parameter. This is unfortunate given that it is a parameter that is deemed important to understand the integrative properties of single neurons. It is also a parameter that can vary based on the amount of inputs the neuron receives from other cells; more specifically, it reflects the ongoing conductance of the cell. It would be particularly interesting to see whether the batch or online implementations of the hitting time–based optimal design method could be implemented in an experimental setting and what new challenges arise in this context. Our method has the advantage of dealing with the intrinsic stochasticity of the underlying process for designing stimulation protocols that best reveal system parameters.

Doing a batch estimation, when the control is computed only once and then applied for all observations, is the computationally simpler approach. More interestingly, we show that the optimal control can be evolved online as new observations are assimilated into the parameter prior, and the optimal control is updated accordingly. However, we note that the optimal control is fairly similar for a wide range of priors, which somewhat limits the additional value of recomputing the optimal control as the observations from the experiment are recorded.

In line with the observation that the optimal control is fairly similar for various priors of τ, we note that the slight differences in the actually obtained stimulations do not make a big difference when it comes to estimating the parameters. Independent of the prior used, the resulting estimates in all cases are much better, more accurate and precise, than the estimates obtained using the base case of no stimulation. In fact, we empirically demonstrate that the optimal stimulation scheme also results in excellent estimation results for *all* system parameters, even though it is designed only to estimate τ.

In general, we note that the control shape that consists of an initial segment that inhibits spiking, $\alpha(t) \approx \alpha_{\min}$, followed by a segment that promotes spiking, where $\alpha(t) \approx \alpha_{\max}$, is obtained for a wide variety of priors or model parameters, μ, σ. In some contexts, however, we have observed a more complex optimal control, where for an initial region that is excitation, then inhibition, and then excitation again. It is very much an interesting open problem to determine how exactly these switching times depend on the various model parameters.

Future work could address the case where all the parameters are treated in a full Bayesian setting and the parameter prior is over all of them, not just over τ. In principle our derivations

carry over unchanged. The only potential complication is that the prior will need to contain more points to fully describe the uncertainty in all three parameters. It would also be of interest to test the estimation scheme for dynamical regimes other than the subthreshold high-noise one studied here. We anticipate that the estimation will improve in suprathreshold and/or lower-noise cases. Other minor tweaks involve using a more sophisticated gradient optimization method, such as a nonlinear conjugate gradient, which might be more computationally efficient.

Our formalism is not restricted to the linear SDE model presented here. In particular, the SDE could include nonlinear drift and/or diffusion terms and be treated in exactly the same manner by our optimal design method. The resulting PDEs (forward and adjoint) will remain linear, and the optimization calculations will be the same. More esoteric processes, where the SDE depends on its own density, can also be considered, although those result in nonlinear forward PDEs and the concomitant complications.

Our work can be viewed as an optimal design for parameter estimation of PDEs. It would be interesting to see whether in general the MI criterion can be applied successfully for the selection of perturbation in other PDE parameter estimation contexts.

Appendix A. Mutual information. The MI, I, between two random variables X and Y is defined by

$$(18) \qquad I(X, Y) = \int_{\mathcal{Y}} \int_{\mathcal{X}} p(x, y) \log \frac{p(x, y)}{p(x) p(y)} \, \mathrm{d}x \, \mathrm{d}y,$$

where $p(\cdot)$ are density functions of their arguments; see, e.g., [17, 6]. Intuitively, it measures the average reduction in uncertainty about one of the variables by the knowledge of the other variable, and it is zero if and only if the two variables are independent.

Here we show that (9) is the MI between the random variables T_{sp} and Θ. The marginal distribution of Θ is simply the prior, $p(\theta) = \rho(\theta)$. The joint distribution is $p(t, \theta) = g(t|\theta)\rho(\theta)$ by the Bayes formula. The marginal of T_{sp} is $p(t) = \int_{\Theta} g(t|\theta)\rho(\theta) \, \mathrm{d}\theta$. Plugging the three expressions into the definition in (18) yields

$$(19) \qquad I = \int_{\Theta} \int_0^{\infty} g(t|\theta)\rho(\theta) \log \left(\frac{g(t|\theta)\rho(\theta)}{\rho(\theta) \int_{\Theta} g(t|\theta)\rho(\theta) \, \mathrm{d}\theta} \right) \mathrm{d}t \, \mathrm{d}\theta.$$

After canceling $\rho(\theta)$ inside the log, we get (9).

REFERENCES

[1] M. ANNUNZIATO AND A. BORZÌ, *Optimal control of probability density functions of stochastic processes*, Math. Model. Anal., 15 (2010), pp. 393–407, https://doi.org/10.3846/1392-6292.2010.15.393-407.

[2] M. ANNUNZIATO AND A. BORZÌ, *A Fokker-Planck control framework for multidimensional stochastic processes*, J. Comput. Appl. Math., 237 (2013), pp. 487–507, https://doi.org/10.1016/j.cam.2012.06.019.

[3] M. ANNUNZIATO, A. BORZÌ, F. NOBILE, AND R. TEMPONE, *On the connection between the Hamilton-Jacobi-Bellman and the Fokker-Planck control frameworks*, Appl. Math., 5 (2014), pp. 2476–2484, https://doi.org/10.4236/am.2014.516239.

[4] A. BORZÌ AND V. SCHULZ, *Computational Optimization of Systems Governed by Partial Differential Equations*, SIAM, Philadelphia, 2012, https://doi.org/10.1137/1.9781611972054.

[5] D. R. CAVAGNARO, J. I. MYUNG, M. A. PITT, AND J. V. KUJALA, *Adaptive design optimization: A mutual information-based approach to model discrimination in cognitive science*, Neural Comput., 22 (2010), pp. 887–905.

[6] T. COVER AND J. THOMAS, *Elements of Information Theory*, 2nd ed., Wiley-Interscience, Hoboken, NJ, 2006.

[7] S. DITLEVSEN AND P. LÁNSKÝ, *Parameters of stochastic diffusion processes estimated from observations of first-hitting times: Application to the leaky integrate-and-fire neuronal model*, Phys. Rev. E (3), 76 (2007), 041906, https://doi.org/10.1103/PhysRevE.76.041906.

[8] S. DITLEVSEN AND P. LÁNSKÝ, *Only through perturbation can relaxation times be estimated*, Phys. Rev. E (3), 86 (2012), 050102(R), https://doi.org/10.1103/PhysRevE.86.050102.

[9] C. E. GRANADE, C. FERRIE, N. WIEBE, AND D. G. CORY, *Robust online Hamiltonian learning*, New J. Phys., 14 (2012), 103013, https://doi.org/10.1088/1367-2630/14/10/103013.

[10] G. HOOKER, K. K. LIN, AND B. ROGERS, *Control theory and experimental design in diffusion processes*, SIAM/ASA J. Uncertain. Quantif., 3 (2015), pp. 234–264, https://doi.org/10.1137/140962280.

[11] A. IOLOV, S. DITLEVSEN, AND A. LONGTIN, *Fokker-Planck and Fortet equation-based parameter estimation for a leaky integrate-and-fire model with sinusoidal and stochastic forcing*, J. Math. Neurosci., 4 (2014), pp. 1–34, https://doi.org/10.1186/2190-8567-4-4.

[12] A. IOLOV, S. DITLEVSEN, AND A. LONGTIN, *Stochastic optimal control of single neuron spike trains*, J. Neural Eng., 11 (2014), 046004, https://doi.org/10.1088/1741-2560/11/4/046004.

[13] P. LÁNSKÝ AND S. DITLEVSEN, *A review of the methods for signal estimation in stochastic diffusion leaky integrate-and-fire neuronal models*, Biol. Cybern., 99 (2008), pp. 253–262, https://doi.org/10.1007/s00422-008-0237-x.

[14] S. LENHART AND J. T. WORKMAN, *Optimal Control Applied to Biological Models*, Chapman Hall/CRC Math. Comput. Biol. Ser., Chapman and Hall/CRC, Boca Raton, FL, 2007.

[15] J. LEWI, R. BUTERA, AND L. PANINSKI, *Sequential optimal design of neurophysiology experiments*, Neural Comput., 21 (2009), pp. 619–687, https://doi.org/10.1162/neco.2008.08-07-594.

[16] J. LIU AND M. WEST, *Combined parameter and state estimation in simulation-based filtering*, in Sequential Monte Carlo Methods in Practice, Springer, New York, 2001, pp. 197–223, https://doi.org/10.1007/978-1-4757-3437-9_10.

[17] D. J. C. MACKAY, *Information Theory, Inference and Learning Algorithms*, Cambridge University Press, New York, 2003.

[18] P. MULLOWNEY AND S. IYENGAR, *Parameter estimation for a leaky integrate-and-fire neuronal model from ISI data*, J. Comput. Neurosci., 24 (2008), pp. 179–194, https://doi.org/10.1007/s10827-007-0047-5.

[19] J. I. MYUNG, D. R. CAVAGNARO, AND M. A. PITT, *A tutorial on adaptive design optimization*, J. Math. Psych., 57 (2013), pp. 53–67, https://doi.org/10.1016/j.jmp.2013.05.005.

[20] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer Ser. Oper. Res., Springer-Verlag, New York, 1999, https://doi.org/10.1007/b98874.

[21] L. PANINSKI, *Asymptotic theory of information-theoretic experimental design*, Neural Comput., 17 (2005), pp. 1480–1507, https://doi.org/10.1162/0899766053723032.

[22] L. PANINSKI, J. PILLOW, AND J. LEWI, *Statistical models for neural encoding, decoding, and optimal stimulus design*, Prog. Brain Res., 165 (2007), pp. 493–507, https://doi.org/10.1016/S0079-6123(06)65031-0.

[23] L. PANINSKI, J. PILLOW, AND E. SIMONCELLI, *Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model*, Neural Comput., 16 (2004), pp. 2533–2561, https://doi.org/10.1162/0899766042321797.

[24] H. TUCKWELL, *Introduction to Theoretical Neurobiology, Volume* 2: *Nonlinear and Stochastic Theories*, Cambridge University Press, Cambridge, UK, 1988.

[25] G. ULLAH AND S. J. SCHIFF, *Tracking and control of neuronal Hodgkin-Huxley dynamics*, Phys. Rev. E (3), 79 (2009), 040901, https://doi.org/10.1103/PhysRevE.79.040901.