Statistical models of spike trains*

Liam Paninski[1], Emery N. Brown[2], Satish Iyengar[3], and Robert E. Kass[4]

[1]Department of Statistics and Center for Theoretical Neuroscience
Columbia University
http://www.stat.columbia.edu/∼liam

[2]Department of Anesthesia and Critical Care
Massachusetts General Hospital, Harvard Medical School
Department of Brain and Cognitive Sciences
Harvard-MIT Division of Health Sciences and Technology
Massachusetts Institute of Technology
https://neurostat.mgh.harvard.edu/brown/emeryhomepage.htm

[3]Department of Statistics
University of Pittsburgh
http://www.stat.pitt.edu/si/

[4]Department of Statistics and Center for the Neural Basis of Cognition
Carnegie Mellon University
http://www.stat.cmu.edu/∼kass

April 16, 2008

*We thank J. Pillow, E. Simoncelli, E.J. Chichilnisky, and S. Koyama for their collaboration on the work described here. LP is partially supported by an NSF CAREER award and an Alfred P. Sloan Research Fellowship; EB is partially supported by NIH grants R01MH59733 and R01DA015644; SI is partially supported by NIMH grant MH041712-20; RK is partially supported by NIMH grant RO1-MH064537-04.

Spiking neurons make inviting targets for analytical methods based on stochastic processes: spike trains carry information in their temporal patterning, yet they are often highly irregular across time and across experimental replications. The bulk of this volume is devoted to mathematical and biophysical models useful in understanding neurophysiological processes. In this chapter we consider *statistical models* for analyzing spike train data.

Strictly speaking, what we would call a statistical model for spike trains is simply a probabilistic description of the sequence of spikes. But it is somewhat misleading to ignore the data-analytical context of these models. In particular, we want to make use of these probabilistic tools for the purpose of scientific inference.

The leap from simple descriptive uses of probability to inferential applications is worth emphasizing for two reasons. First, this leap was one of the great conceptual advances in science, taking roughly two hundred years. It was not until the late 1700s that there emerged any clear notion of inductive (or what we would now call *statistical*) reasoning; it was not until the first half of the twentieth century that modern methods began to be developed systematically; and it was only in the second half of the twentieth century that these methods became well understood in terms of both theory and practice. Second, the focus on inference changes the way one goes about the modeling process. It is this change in perspective we want to highlight here, and we will do so by discussing one of the most important models in neuroscience, the stochastic integrate-and-fire (IF) model for spike trains.

The stochastic IF model has a long history (Gerstein and Mandelbrot, 1964; Stein, 1965; Knight, 1972; Burkitt, 2006): it is the simplest dynamical model that captures the basic properties of neurons, including the temporal integration of noisy subthreshold inputs, all-or-none spiking, and refractoriness. Of course, the IF model is a caricature of true neural dynamics (see, e.g., (Ermentrout and Kopell, 1986; Brunel and Latham, 2003; Izhikevich, 2007) for more elaborate models) but, as demonstrated in this book and others (Ricciardi, 1977; Tuckwell, 1989; Gerstner and Kistler, 2002), it has provided much insight into the behavior of single neurons and neural populations.

In this chapter we explore some of the key *statistical* questions that arise when we use this model to perform inference with real neuronal spike train data. How can we efficiently fit the model to spike train data? Once we have estimated the model parameters, what can the model tell us about the encoding properties of the observed neuron? We also briefly consider some more general approaches to statistical modeling of spike train data.

We begin, in section 1, by discussing three distinct useful ways of approaching the IF model, via the language of stochastic (diffusion) processes, hidden Markov models, and point processes, respectively. Each of these viewpoints comes equipped with its own specialized analytical tools, and the power of the IF model is most evident when all of these tools may be brought to bear simultaneously. We discuss three applications of these methods in section 2, and then close in 3 by indicating the scope of the general point process framework of which the IF model is a part, and the possibilities for solving some key outstanding data-analytic problems in systems neuroscience.

1 The IF model from three different points of view

The leaky stochastic IF model is defined, in the simplest case, in terms of a one-dimensional stochastic voltage process $V(t)$ that follows the linear stochastic dynamics

$$dV(t) = \big(- g(t)V(t) + I(t)\big)dt + \sigma dB_t, \tag{1}$$

where the random term B_t denotes a standard Brownian motion[1]. We model the observed spike train as the passage times of the random process $V(t)$ through the threshold voltage V_{th}; after each threshold crossing, $V(t)$ is reset to $V_{reset} < V_{th}$.

Here $I(t)$ and $g(t)$ denote the input current and membrane conductance at time t, respectively. In the simplest case, we take the membrane conductance and input current to be constant, $g(t) = g, I(t) = I$. However, it is natural to consider the time-varying case as well. For example, in the popular "spike-response" model (Gerstner and Kistler, 2002; Paninski et al., 2004c), we add a time-varying current $h(t-t_i)$ into the cell after each spike time t_i, i.e., $I(t) = I + \sum_{t_i<t} h(t-t_i)$; this post-spike current models the lumped effects of all currents that enter a real neuron following an action potential. By changing the shape and magnitude of $h(.)$, we can model a variety of interspike interval behavior. For example, a negative but sharply time-limited $h(.)$ can extend the relative refractory period of the IF neuron, since it will take longer for the voltage V to climb back up to the threshold V_{th}; this in turn leads to firing rate saturation or spike-rate adaptation, since increasing the input current I will just increase the hyperpolarizing effect of the total spike history effect $\sum_{t_i<t} h(t-t_i)$. Similarly, a positive or biphasic $h(.)$ can induce burst effects in the spike train; see (Gerstner and Kistler, 2002; Paninski et al., 2004c) for further examples. It is also natural to consider similar models for the conductance $g(t)$ following a spike time (Stevens and Zador, 1998; Jolivet et al., 2004).

Just as importantly, we would like to model the effects of an external stimulus on the observed spike train. One natural way to include stimulus effects in the model is to let the input current $I(t)$ depend on the stimulus $\vec{x}(t)$. In the simplest case, this dependence could be linear,

$$I(t) = \sum_i a_i x_i(t), \tag{2}$$

but of course it is possible to include nonlinear effects as well (Paninski et al., 2004c).

When we turn to fitting models to data, and using them for scientific inference, there are various approaches we may consider. Some of the most powerful methods are based on the *likelihood function* which, in turn, is defined in terms of the joint probability density function (pdf) for the spike train. Let us write the pdf for the sequence of spike times $\{t_i\}$ as $p(\{t_i\}|\theta)$, where θ is a multidimensional parameter vector. In the setting described above, θ could include $\{V_{th}, V_{reset}, \sigma, \vec{a}, h(.), g\}$. The likelihood function reverses the dependence displayed in the pdf: the pdf is a function of $\{t_i\}$ for each fixed θ, while the likelihood function is a function of θ for each fixed $\{t_i\}$. In other words, the pdf fixes the parameter θ and assigns probability to the spike trains that might then occur, while the likelihood function fixes the spike train that was actually observed, and assigns relative likelihood to the parameter values that might have produced it. Because likelihood-based methods use the information in the data as efficiently as possible (see, e.g., (Brown et al., 2003; Kass et al., 2005) and the references therein), a great deal of effort has been devoted to developing such methods and for carrying out the relevant computations. In the next subsections we discuss three different conceptual approaches for performing statistical inference with the IF model via likelihood-based methods.

However, before we dive into these different approaches for computing and maximizing the likelihood, it is worth noting a surprising and important fact about the likelihood in this model: the logarithm of the likelihood turns out to be *concave* as a function of many of the

[1]Or in discrete time,

$$V(t+dt) - V(t) = \big(- g(t)V(t) + I(t)\big)dt + \sigma\sqrt{dt}\epsilon_t,$$

where ϵ_t is an i.i.d. standard normal random variable.

model parameter values θ (Paninski et al., 2004c; Paninski, 2005; Mullowney and Iyengar, 2007). This makes model fitting via maximum likelihood or Bayesian methods surprisingly tractable, since concave functions lack any suboptimal local maxima that could trap a numerical optimizer. (This concavity property is also extremely useful in the context of *decoding*, i.e., inferring the stimulus $\vec{x}$ given the observed spike train data $\{t_i\}$ (Pillow and Paninski, 2007; Ahmadian et al., 2008); we will discuss some simple decoding applications of the IF model below.) Indeed, without this concavity property the IF model would be much less attractive from a statistical point of view.

1.1 The IF model as a diffusion process

The most direct, and probably best-explored, view of the IF model is as a linear *diffusion process* (Ricciardi, 1977; Tuckwell, 1989; Burkitt, 2006). This connection allows us to apply powerful tools from stochastic calculus (Karatzas and Shreve, 1997) to understand the behavior of this model. For example, in the nonleaky case ($g(t) = 0$) with constant current input ($I(t) = I$), we may employ the "reflection principle" from the basic theory of Brownian motion to derive the so-called "inverse Gaussian" first passage time density (Gerstein and Mandelbrot, 1964; Seshadri, 1993; Iyengar and Liao, 1997; Brown, 2005); denoting $p(\tau|\theta)$ as the probability density that the next interspike interval will be of length τ, we may explicitly calculate

$$p(\tau|\theta) = \frac{V_{th} - V_{reset}}{\sqrt{2\pi\sigma^2\tau^3}} e^{-[(V_{th}-V_{reset})-I\tau]^2/2\sigma^2\tau}.$$

In this simplest case, each interspike interval is independent and identically distributed (this point process is a special case of a "renewal" process, since it resets itself after each spike), and therefore the likelihood of a full spike train is obtained by taking products over the observed interspike intervals $\tau_i = t_{i+1} - t_i$:

$$p(\{t_i\}|\theta) = \prod_i p(\tau_i|\theta).$$

This explicit analytical formula for the spike train likelihood allows us, for example, to derive closed-form expressions for the maximum likelihood estimate of the model parameters θ here (Seshadri, 1993); thus the classical stochastic process theory makes model fitting in this simple case quite straightforward.

In the more general case of nonconstant $I(t)$ and nonzero $g(t)$, things are somewhat more complicated; we have no explicit general formula for $p(\tau|\theta)$, for example. However, the stochastic calculus leads directly to partial differential equation (Risken, 1996; Paninski et al., 2004c) or integral equation (Siegert, 1951; Buoncore et al., 1987; Plesser and Tanaka, 1997; Burkitt and Clark, 1999; DiNardo et al., 2001; Paninski et al., 2007a; Mullowney and Iyengar, 2007) methods for computing the likelihood. For example, it is well-known (Tuckwell, 1989) that the probability density of the next spike satisfies

$$p(\tau|\theta) = -\frac{\partial}{\partial t}\left(\int P(v,t)dv\right)\Big|_{t=\tau},$$

where the density $P(v,t) = p(V(t) = v, \tau_i > t)$ satisfies the Kolmogorov "forward" (Fokker-Planck) PDE

$$\frac{\partial P(v,t)}{\partial t} = \frac{\sigma^2}{2}\frac{\partial^2 P(v,t)}{\partial v^2} + g(t)\frac{\partial\left[\big(v - I(t)/g(t)\big)P(v,t)\right]}{\partial v} \tag{3}$$

with boundary conditions

$$P(v, t_{i-1}^+) = \delta(v - V_{reset}),$$

corresponding to the fact that the voltage resets to V_{reset} following the last spike, and

$$P(V_{th}, t) = 0,$$

corresponding to the fact that the voltage $V(t)$ is never greater than the threshold V_{th}. In the constant-current, zero-leak case this PDE may be solved explicitly by the method of images (Daniels, 1982; Paninski, 2006b); more generally, we may solve the PDE numerically via efficient implicit techniques such as the Crank-Nicholson method (Press et al., 1992; Paninski et al., 2004c). These methods are mathematically closely related to techniques developed to understand the behavior of populations of neurons (Brunel and Hakim, 1999; Nykamp and Tranchina, 2000; Knight et al., 2000; Fourcaud and Brunel, 2002), though the application here is very different.

Alternately, $p(\tau|\theta)$ solves a number of integral equations of the form

$$f(t) = \int_0^t K(\tau, t) p(\tau|\theta) d\tau$$

(Siegert, 1951; Ricciardi, 1977; Plesser and Tanaka, 1997; Burkitt and Clark, 1999; DiNardo et al., 2001; Paninski et al., 2007a; Mullowney and Iyengar, 2007), where the kernel function $K(\tau, t)$ and the left-hand-side $f(t)$ may be computed as simple functions of the model parameters θ. In the case of constant current and conductance ($I(t) = I, g(t) = g$), this integral equation may be solved efficiently by Laplace transform techniques (Mullowney and Iyengar, 2007); more generally, the equation may be solved by numerical integration methods (Plesser and Tanaka, 1997; DiNardo et al., 2001) or by direct matrix inversion methods (Paninski et al., 2007a) (the latter methods also lead to simple formulas for the derivative of the likelihood with respect to the model parameters; this gradient information is useful for optimization of the model parameters, and for assessing the accuracy of the parameter estimates).

Thus, to summarize, stochastic process methods may be used to derive efficient schemes for computing the likelihood in this model. However, this approach leaves some questions open. For example, it is unclear how to use the methods described above to obtain various *conditional* quantities of interest given observed spike train data. One such important quantity is the probability $p(V(t)|\{t_i\}, \theta)$ that the voltage at time t is at a certain level, given the observed spike sequence $\{t_i\}$ and model parameters θ. To address these questions, we adopt a "state-space" approach to the IF model.

1.2 The IF model as a state-space model

A second way to approach the IF model is by considering a broader class of statistical models, known as "state space," or "hidden Markov" models. As the name implies, this model consists of two processes: an unobserved ("hidden") Markovian process, and an observed process which is related to the hidden process in a simple instantaneous manner (Rabiner, 1989).

Because this hidden Markov framework is natural in a wide variety of contexts, from speech recognition to genetic transcription, statisticians have developed a number of useful inference tools for models of this form. State-space models have recently been introduced in a number of neural contexts (Brown et al., 1998; Smith and Brown, 2003; Czanner et al., 2008;

Yu et al., 2006; Srinivasan et al., 2006; Kulkarni and Paninski, 2007; Huys and Paninski, 2007; Dombeck et al., 2007; Vogelstein and Paninski, 2007), with applications ranging from neural prosthetic design to signal processing of calcium imaging data.

The IF model may be cast in this framework easily. $V(t)$ is a hidden Markovian process which we observe only indirectly, through the spike times $\{t_i\}$, which may be considered a simple function of $V(t)$: the observed spike variable at time t is zero if $V(t)$ is below threshold, and one if $V(t) = V_{th}$. Now we may easily exploit standard methods from the state-space literature for computing the likelihood: for example, the standard "forward recursion" for the computation of the likelihood in hidden Markov models (Rabiner, 1989) corresponds exactly to the computation of the Kolmogorov forward density $P(v, t)$ discussed in the previous section (Paninski et al., 2004c; Nikitchenko and Paninski, 2007).

The standard method for computing the maximum likelihood estimate for the model parameter θ in a hidden Markov model is based on the so-called "Expectation-Maximization" (EM) algorithm (Dempster et al., 1977; Rabiner, 1989). This algorithm is guaranteed to increase the likelihood with each iteration (i.e., EM serves as a likelihood ascent algorithm — though interestingly, the basic algorithm does not require a direct computation of the likelihood itself). In a nutshell, the EM algorithm applied to the IF model iterates between two conceptually straightforward steps (Czanner and Iyengar, 2004; Nikitchenko and Paninski, 2007). In the Expectation step, we compute the *conditional expectation* $E(V(t)|\{t_i\}, \theta)$ of the voltage $V(t)$ at each time t, given all the observed spike data, via the standard "forward-backward" recursive algorithm from hidden Markov models (we will discuss an interesting application of this conditional expectation computation below, in section 2.1); in the Maximization step, we fit the parameters of model (1) directly, using the inferred conditional expectations $E(V(t)|\{t_i\}, \theta)$ as a substitute for the expectations (sufficient statistics) computed from the actual subthreshold voltages $V(t)$. This M step is straightforward because it is much easier to fit model (1) given direct voltage observations; in this case, the model reduces to a fairly standard autoregressive model, for which the estimation theory is quite well-established. See, for example, (Jolivet et al., 2004; Paninski et al., 2004b; Huys et al., 2006; Lansky et al., 2006) for some examples in which autoregressive models similar to our equation (1) have been fit directly to intracellularly recorded voltage data.

Thus the state-space setting provides a very useful framework for optimizing the IF model parameters θ, and for making inferences about the hidden subthreshold voltage $V(t)$ (see section 2.1 below). One additional major advantage of the state-space representation is its flexibility: it turns out to be fairly straightforward to generalize these techniques to include other important dynamical biophysical variables in our model beyond just the voltage $V(t)$. We will discuss some of these extensions in section 3 below.

1.3 Approximating the IF model via simpler point-process models

The diffusion- and state-space-based approaches above are conceptually elegant but computationally somewhat expensive. A third approach for performing inference with integrate-and-fire data is not to fit the IF model at all, but rather to use a simpler model that retains the IF model's key properties. (It is worth recalling here the adage that all models are wrong, but some are useful. In other words, when performing statistical analysis we are not bound to any one model in particular; instead, we should choose the model which best suits our needs.) A particularly simple framework for analysis of point process data involves likelihood methods based on the *conditional intensity function* $\lambda(t|H_t, \theta)$, the instantaneous firing rate

at time t given H_t, the complete stimulus and spiking history preceding time t (see (Brown et al., 2003; Brown, 2005; Kass et al., 2005) and references therein for further discussion). To be concrete, we define $P(\text{spike at time } t|H_t) = \lambda(t|H_t, \theta)dt$ or, more precisely,

$$\lambda(t|H_t, \theta) = \lim_{\Delta t \to 0} \frac{P(\text{spike in } (t, t + \Delta t)|H_t, \theta)}{\Delta t}.$$

The point-process likelihood function is then given by the simple and explicit pdf (Snyder and Miller, 1991; Brown et al., 2003; Paninski, 2004; Brown, 2005; Truccolo et al., 2005)

$$p(\{t_i\}|\theta) = e^{-\int_0^T \lambda(t|H_t,\theta)dt} \prod_i \lambda(t_i|H_{t_i}, \theta), \tag{4}$$

where $[0, T]$ indicates the time interval over which the spike train is observed. With this formula in hand it becomes natural to approach the fitting of spike-train data by trying to find tractable and useful forms for the conditional intensity function.

With this in mind, let us write the time of the spike preceding t as $s_*(t)$, and note that the usual stimulus-driven IF model (without effects $h(.)$ that involve spikes prior to $s_*(t)$) depends only on the experimental clock time t, relative to stimulus onset, and the elapsed time $t - s_*(t)$, i.e., it satisfies

$$\lambda(t|H_t, \theta) = \lambda(t, t - s_*(t)) \tag{5}$$

for a suitable nonnegative function $\lambda(., .)$. (More generally, this equation is a good approximation whenever the spike history effect $h(\tau)$ is negligible for τ larger than the typical interspike interval in the data.) Models of the general form of Equation (5) have been called "inhomogeneous Markov interval (IMI) models" by (Kass and Ventura, 2001).

We may consider three special cases of Equation (5):

- The multiplicative IMI (m-IMI) model:

$$\lambda(t, s_*(t)) = \lambda_0(t)r(t - s_*(t)). \tag{6}$$

  Here, $\lambda_0(t)$ modulates the firing rate only as a function of the experimental clock time, while $r(t-s_*(t))$ represents non-Poisson spiking behavior. See (Kass and Ventura, 2001) and references therein.

- The time-rescaled renewal process (TRRP) model:

$$\lambda(t, s_*(t)) = \lambda_0(t)r\big(\Lambda_0(t) - \Lambda_0(s_*(t))\big), \tag{7}$$

  where r is the hazard function of a renewal process (basically, a nonnegative function such that $\lim_{t\to\infty} \int^t r(s)ds = \infty$), and $\Lambda_0(t)$ has the time-rescaling form (Brown et al., 2002)

$$\Lambda_0(t) = \int_0^t \lambda_0(u)du. \tag{8}$$

- The additive IMI model:

$$\lambda(t, s_*(t)) = f[y(t) + h(t - s_*(t)], \tag{9}$$

  where $f(.)$ is some smooth nonnegative function, $y(t)$ is a scalar function of time, and the spike-history effect $h(.)$ only looks back at the last spike time $s_*(t)$, instead of over the full spike history as discussed in the models above.

Of course, Equation (5) may be generalized to include multi-spike history effects (see, e.g., (Kass and Ventura, 2001)).

Each of these models separate the dependence on t from the dependence on the time since the last spike $t - s_*(t)$, but they do so differently. Here, we may call λ_0 or $f[y(t)]$ the *excitability function* to indicate that this modulates the amplitude of the firing rate, and r the *recovery function*, since it affects the way the neuron recovers its ability to fire after generating an action potential. The fundamental difference between the models is the way the excitability interacts with the recovery function. In the m-IMI model, for example, the refractory period represented in the recovery function is not affected by excitability or firing rate variations, while in the TRRP model the refractory period is no longer fixed but is scaled by the firing rate (Reich et al., 1998; Barbieri et al., 2001).

(Koyama and Kass, 2008) investigated the relationship of m-IMI and TRRP models to the stimulus-driven IF model. Specifically, given a joint probability distribution of spike trains generated by an IF model, they asked how close this distribution is to each of the best-fitting m-IMI and TRRP models. Some of the main results in Koyama and Kass (2007) are provided by Figure 1. The authors considered the case that the input to the IF neuron was periodic, and numerically computed the probability that a test (specifically, a time-rescaled Kolmogorov-Smirnov test, as described in (Brown et al., 2002)) will reject the (false) hypothesis that the data were generated by a Poisson, m-IMI, or TRPP model, instead of the true IF model. Figure 1 displays the rate of rejecting the models as a function of the number of spikes. It takes very little data to reject the inhomogeneous Poisson process, while it takes much larger data sets to reject the m-IMI model and the TRRP. Because the recovery functions in the m-IMI and TRRP models allow us to mimic the interspike interval distributions of the IF cell, for many data-analytic applications the results using the m-IMI, TRRP, and IF models will be very similar to one another—but very different from those based on a Poisson model.

Clearly these three model classes overlap somewhat. For example, if the nonlinearity $f(.)$ in equation (9) is taken to be the exponential function, then the multiplicative and additive IMI models coincide. The additive formulation has become popular recently (Brillinger, 1988; Chornoboy et al., 1988; Brillinger, 1992; Paninski, 2004; Truccolo et al., 2005; Paninski et al., 2007b), largely because stimulus terms (e.g., if we set $y(t) = \sum a_i x_i(t)$, given the stimulus $\vec{x}(t)$) and spike history terms are treated in a unified, linear manner within the nonlinear function $f(.)$ (for this reason, the additive model is often described as a "generalized linear model"). It is worth making the connection between the IF and this additive model more explicit. Consider the inhomogeneous Poisson process with rate given by $f[V(t)]$, where $V(t)$ is the solution of the leaky-integrator differential equation $dV/dt = -gV(t) + I(t)$, starting at the initial value V_{reset} after every spike. This model is a simple version of the "escape-rate" approximation to the noisy integrate-and-fire (IF) model described in (Gerstner and Kistler, 2002). Since this differential equation is linear, $V(t)$ here may be written (up to boundary conditions) as the convolution of $I(t)$ with the exponential function e^{-gt}; that is, this soft-threshold IF model is just a version of the additive IMI model (with the input current redefined suitably; see (Paninski, 2004; Paninski et al., 2007b) for details), and therefore the additive parameters may be indirectly interpreted in biophysical terms. See (Paninski et al., 2007b) for an explicit comparison of the additive and IF models applied to data from a retinal ganglion neuron; again, as in Fig. 1, the additive model provides an excellent approximation to the behavior of the IF neuron.

Thus, to summarize, these point process models may be used to provide good approximations to the IF model. Some biophysical interpretability may be lost, because the "intracellu-

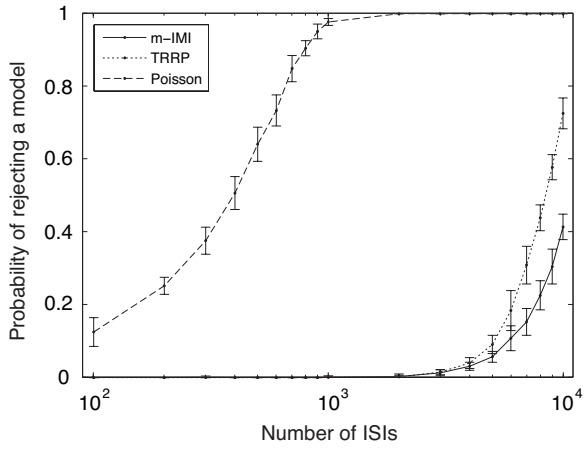

Figure 1: Probability of rejecting three statistical models as a function of the number of ISIs. The spike trains were generated from a IF neuron with periodic input. The solid line, dotted line and dashed represent the m-IMI model, the TRRP model and the inhomogeneous Poisson process. The mean and the standard deviation at each point were calculated with 10 repetitions. With relatively small sample sizes it can be easy to reject the hypothesis of Poisson firing, but impossible to distinguish either the m-IMI model or the TRRP model from the IF model. Even after 10,000 spikes the m-IMI model is, for practical purposes, statistically indistinguishable from the IF model. Modified from (Koyama and Kass, 2008).

lar" noise in the IF model (due to the diffusion term) is replaced by "point-process" variability (due to the randomness of the spike times). What we gain is a simpler, more tractable, and more general framework; in particular, the explicit likelihood formula (Eq. (4)) makes various inference tasks much more straightforward and computationally tractable (Paninski et al., 2007b).

2 Applications

2.1 Computing the spike-triggered average and the most likely voltage path

The simplicity of the IF model allows us to analytically compute certain quantities of great theoretical interest. A key question in neural coding is: what does a single spike "mean" (Rieke et al., 1997)? More generally, what does a pattern of spikes mean (de Ruyter van Steveninck and Bialek, 1988)? We may make this question quantitative by computing the conditional distribution of a behaviorally-relevant signal given the observation of a spike pattern; if we can characterize these conditional distributions for every conceivable spike pattern, we can reasonably say that we have answered the question, "what is the neural code"?

As is well-known, these questions may be answered explicitly in the case of simple "cascade" Poisson-based models (Chichilnisky, 2001; Schnitzer and Meister, 2003; Simoncelli et al., 2004). It turns out we can also solve this problem in the case of some of the more biophysical encoding models we have been working with so far. For a concrete example, let's look at the "doublet-triggered density" $P(V(t)|[t_1, t_2])$, the probability density on voltage of a nonleaky

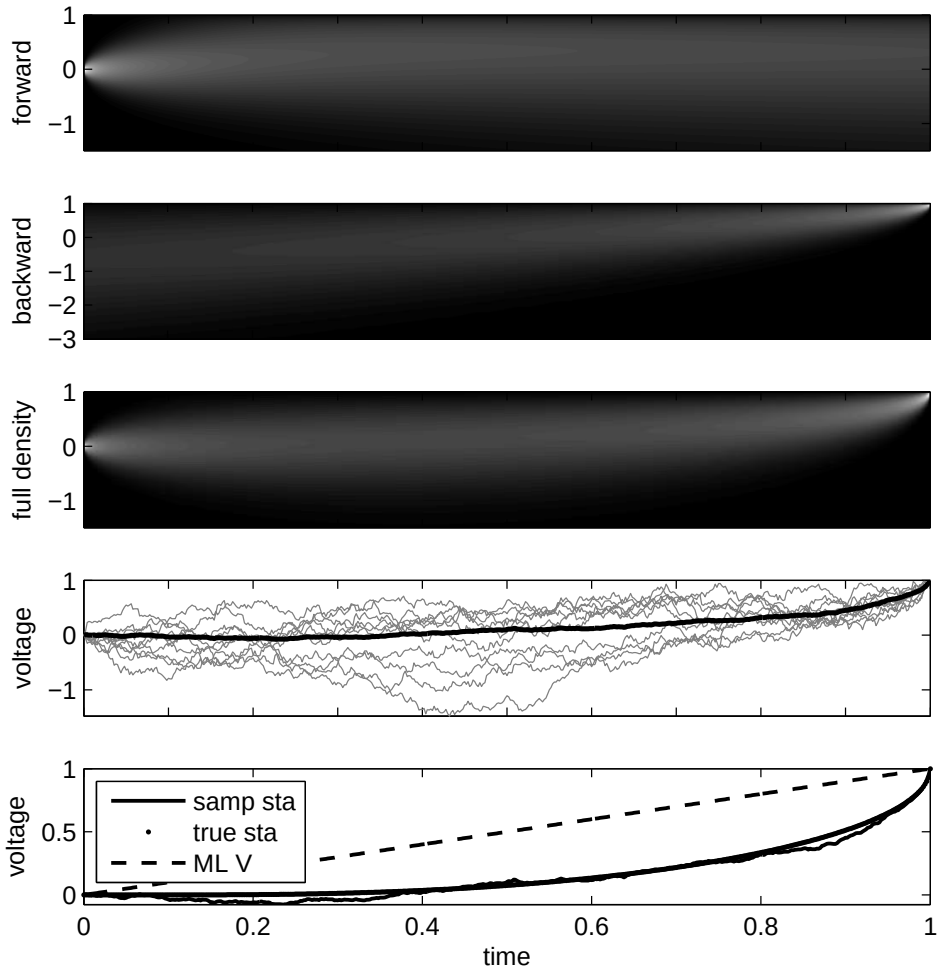

Figure 2: The doublet-triggered average of the nonleaky IF model (adapted from (Paninski, 2006b)). Panels 1-3 show evolution of densities $P_f(t), P_b(t)$, and $P(V(t)|[t_1, t_2])$, for $t \in [0, 1]$; grayscale indicates height of density. Panel 4 shows some samples from the conditional voltage path distribution given spikes at $t_1 = 0$ and $t_2 = 1$ (the standard forward-backward sampling method for hidden Markov models was employed here, after discretizing the state space $V(t)$), with the empirical mean conditional voltage given 100 samples shown in black. The bottom panel shows the most likely path (dotted trace), the analytical doublet-triggered average (dashed), and the empirical doublet-triggered average (solid). Parameters: $V_{th} = 1, V_{reset} = 0, \sigma^2 = 1, I = 1, t_1 = 0, t_2 = 1$. Note that the conditional mean voltage "sags" below V_{th}, since voltage paths that cross V_{th} are killed and are therefore subtracted out of the expectation $E(V(t)|[t_1, t_2])$; see (Paninski, 2006b) for further discussion of this effect.

integrate-and-fire cell that has been observed to fire at times t_1 and t_2 (with $t_2 > t_1$, and no spikes observed in the intervening period).

This density may be computed exactly (Paninski, 2006b) by applying simple stochastic calculus tools, similar to those we have discussed above in section 1.1. We condition the voltage to start at $V(t_1^+) = V_{reset}$ and end at $V(t_2^-) = V_{th}$, because we observed spikes at time

t_1 and t_2. The nonleaky IF model with constant current inputs may be described in terms of a Brownian motion with constant drift (Gerstein and Mandelbrot, 1964; Karatzas and Shreve, 1997). Because the Brownian motion is a Gaussian process, and our constraints are linear equality constraints, the conditioned process (known as a "Brownian bridge") is also Gaussian, with mean and variance which may be calculated by the usual Gaussian conditioning formulas. We denote the marginal density of the Brownian bridge at time t as

$$P(V(t)|V(t_1) = V_{reset}, V(t_2) = V_{th}) = \mathcal{N}\left(V_{reset} + m(t), s(t)\right),$$

where $\mathcal{N}(\mu, \sigma^2)$ denotes the Gaussian density with mean μ and variance σ^2, and we have made the abbreviations $m(t) = (t - t_1)(V_{th} - V_{reset})/(t_2 - t_1)$ and $s(t) = \sigma^2(t - t_1)(t_2 - t)/(t_2 - t_1)$.

Now we know not only that the cell has fired at two specific times t_1 and t_2, but also that the voltage is below threshold for all times $t_1 < t < t_2$. Thus we need to compute the conditional density of the Brownian bridge given that the process lies below threshold,

$$P(V(t)|[t_1, t_2]) = P(V(t)|V(t_1) = V_{reset}; V(t_2) = V_{th}; V(t) < V_{th}, t_1 < t < t_2).$$

Now we may apply the method of images, as in the unconditioned case, to obtain the remarkably simple solution

$$P(V(t)|[t_1, t_2]) \quad = \quad \frac{1}{Z}\left[\mathcal{N}\big(V_{reset} + m(t), s(t)\big) - \mathcal{N}\big(V_{th} + (V_{th} - V_{reset}) - m(t), s(t)\big)\right].$$

This solves our doublet-triggered density problem for the simple nonleaky, linear integrate-and-fire model. What about the leaky case, or more generally if the neuron obeys nonlinear and/or multidimensional subthreshold dynamics (Badel et al., 2005)? Here we may apply the state-space techniques discussed in section 1.2 above: we use a version of the forward-backward method to obtain (Paninski, 2006b):

$$P(V(t)|[t_1, t_2]) \quad = \quad \frac{1}{Z(t)} P_f(V, t) P_b(V, t)$$

where $P_f(V, t) = P(V(t), \tau_1 > t)$ satisfies the forward Kolmogorov PDE discussed in section 1.1, and $P_b(V, t) = P(\tau_1 = t_2 | V(t))$ satisfies a corresponding Kolmogorov "backward" PDE. See (Paninski, 2006b) for further details of the computation of these forward and backward probabilities; Figure 2 illustrates the components of this solution.

Once we have obtained the doublet-triggered average voltage $E(V(t)|[t_1, t_2])$, it turns out to be straightforward to apply renewal theory techniques (recalling that the spike trains of the IF model with constant coefficients may be described as a renewal model, as mentioned in section 1.1) to obtain the single-spike-triggered average $E(V(t)|[t_1])$; again, see (Paninski, 2006b) for details.

While obtaining the conditional average voltage $E(V(t)|\{t_i\})$ given some observed spike data $\{t_i\}$ is conceptually straightforward, it does require some numerical computation in the case of a nonconstant input current $I(t)$, or a nonzero membrane conductance $g > 0$. It turns out to be even easier to obtain the most likely conditional voltage path,

$$\mathbf{V}_{ML} = \arg\max_{\mathbf{V} \in C} p(\mathbf{V}|\{t_i\}),$$

than to compute the conditional average path. Here the optimization is taken over the constraint space

$$C = \{\mathbf{V} : V(t) \leq V_{th}\ \forall t;\ V(t_i^-) = V_{th};\ V(t_i^+) = V_{reset}\}$$

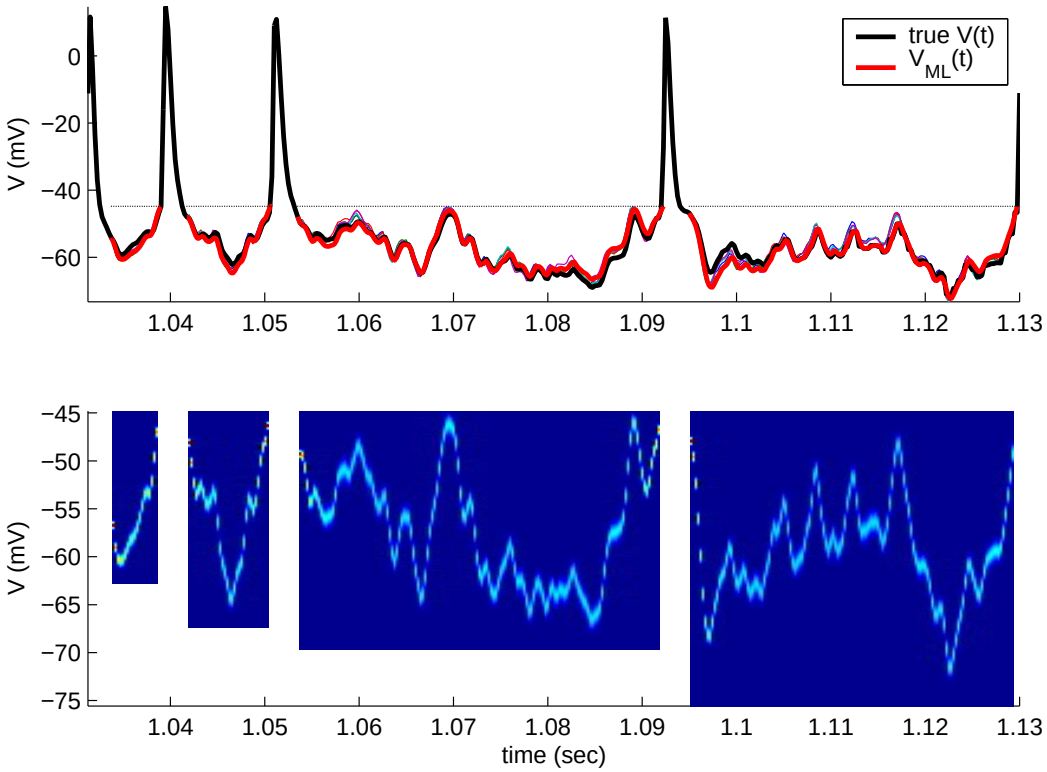

Figure 3: Computing the most likely voltage path given *in vitro* physiological data. **Top**: Comparison of true voltage path (bold black trace) with computed $V_{ML}(t)$ (bold gray trace) and samples from conditional distribution $p(V(t)|spikes, \{I(t)\}_{0\leq t\leq T}, \hat{\theta})$ (thin traces, partially obscured; here $\hat{\theta}$ denotes the IF model parameters estimated from a separate sequence of training data, and $I(t)$ is the observed fluctuating input current). Trace shown is a randomly chosen segment of a 25-second long white noise experiment; dashed trace indicates estimated threshold. **Bottom**: Conditional distributions $p(V(t)|spikes, \{I(t)\}_{0\leq t\leq T}, \hat{\theta})$. 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). Adapted from (Paninski, 2006a).

of all subthreshold voltage paths consistent with the observed spike train $\{t_i\}$. By using the state-space representation, it is straightforward to show that $p(\mathbf{V}|\{t_i\})$ is in fact just a linearly-constrained quadratic function of $\mathbf{V}$, and may therefore be optimized quite easily (Paninski, 2006a; Koyama et al., 2008). In fact, it turns out to be possible to obtain the most likely path analytically in greater generality than is possible for the conditional average path $E(V(t)|\{t_i\})$; further, one can prove that the conditional average and conditionally most likely voltage paths coincide in the case of small noise σ (Paninski, 2006a; Badel et al., 2005). See Fig. 3 for an application to real data, in which the most likely voltage path provides a very accurate prediction of the unobserved subthreshold voltage $V(t)$, given the superthreshold observed spike train $\{t_i\}$ and the input current $I(t)$ (Paninski, 2006a).

2.2 Tracking plasticity in hippocampal place fields

Neural receptive fields are frequently plastic: a neural response to a stimulus can change over time as a result of experience. (Frank et al., 2002) used a multiplicative IMI model to characterize spatial receptive fields of neurons from both the CA1 region of the hippocampus and the deep layers of the entorhinal cortex (EC) in awake, behaving rats. In their model, each neuronal spike train was described in terms of a conditional intensity function

$$\lambda(t|H_t, \theta) = \lambda^S(x(t))r(t - s_*), \tag{10}$$

which is of the form (6), except that the spatial intensity factor λ^S has replaced the temporal factor λ_0. In Equation (6), both the factors $\lambda^S(x(t))$ and $r(t - s_*)$ are dynamic, evolving simultaneously according to a state-space model specified by a point process adaptive filter (Brown et al., 2001; Brown et al., 2003). This allowed the authors to describe the way the receptive fields evolve across space and time. They found consistent patterns of plasticity in both CA1 and deep EC neurons, which were distinct: the spatial intensity functions of CA1 neurons showed a consistent increase over time, whereas those of deep EC neurons tended to decrease (Fig. 4). They also found that the ISI-modulating factor $r(t - s_*(t))$ of CA1 neurons increased only in the "theta" region (75-150 milliseconds), whereas those of deep EC neurons decreased in the region between 20 and 75 milliseconds. In addition, the minority of deep EC neurons whose spatial intensity functions increased in area over time fired in a significantly more spatially specific manner than non-increasing deep EC neurons. This led them to suggest that this subset of deep EC neurons may receive more direct input from CA1 and may be part of a neural circuit that transmits information about the animal's location to the neocortex.

2.3 Encoding and decoding spike train information in the retina

Our final application illustrates that the IF model may be tractably fit to data via maximum likelihood, and that the model allows us to: 1) predict the spike times of a sensory neuron at near-millisecond precision and 2) decode the responses of these sensory neurons with an accuracy greater than is possible with a Poisson model of the responses (Pillow et al., 2005). The basic results are illustrated in Figs. 5-6. The authors recorded the spike train responses of primate retinal ganglion cells which were stimulated with a temporally varying full-field light stimulus. Data from a long sequence of nonrepeated training stimuli were used to fit both the IF model and a simpler inhomogeneous Poisson model; Figure 5 compares the predictions of these two models given a novel repeated test stimulus. The IF model does a significantly better job of capturing the true neuron's time-varying mean response, variability, and fine spike timing details.

See also (Keat et al., 2001) for some related results using a similar model fit using non-likelihood-based methods, and (Iyengar and Liao, 1997) for an illustration of a generalized inverse Gaussian model's ability to capture the interspike interval properties of a goldfish retinal ganglion cell under constant illumination (i.e., no time-varying stimulus).

Figure 6 illustrates an important application of the likelihood-based approach. Once we have obtained an accurate approximation for $p(\{t_i\}|\vec{x})$, the probability of a spike response $\{t_i\}$ given the stimulus $\vec{x}$ (namely, the approximation is provided by our estimated model, $p(\{t_i\}|\vec{x}) \approx p(\{t_i\}|\vec{x}, \hat{\theta})$, with $\hat{\theta}$ denoting our estimate of the parameter), then we may invert

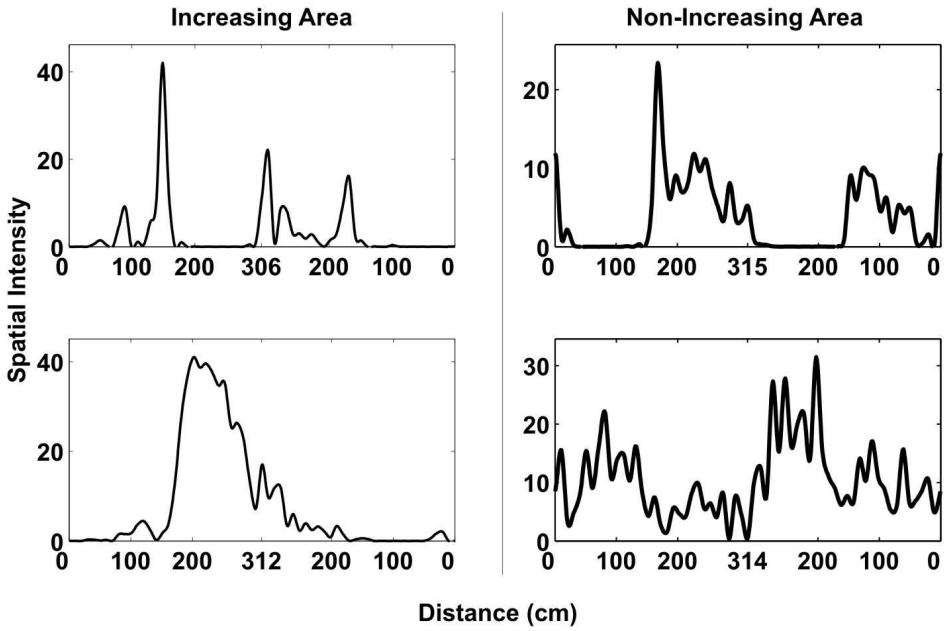

Figure 4: The instantaneous spatial intensity function for a CA1 neuron at the outset of the experiment (upper left) and at the end of the experiment (loer left) illustrating the increasing in the spatial firing pattern typical of these neurons in this experiment. The instantaneous spatial intensity function for a deep EC neuron at the outset of the experiment (upper right) and at the end of the experiment (lower right) illustrating the decreasing in the spatial firing pattern typical of these neurons in this experiment. Reprinted with permission of the Journal of Neursocience.

this likelihood via Bayes' rule to obtain the *decoding* posterior probability,

$$p(\vec{x}|\{t_i\}) \propto p(\{t_i\}|\vec{x})p(\vec{x}).$$

In many cases the prior probability $p(\vec{x})$ is known, allowing us to compute the posterior $p(\vec{x}|\{t_i\})$ explicitly, and then we may use this posterior distribution to decode $\vec{x}$ optimally, given the observed data $\{t_i\}$. This is what was done in the experiment illustrated in Fig. 6. The authors asked how well it was possible to discriminate two possible stimuli, given a short observed segment of the retinal spike train response. As expected, the IF model permits more accurate posterior decoding than does the inhomogeneous Poisson model, since (as we saw in Fig. 5) the IF model provides a better approximation to the true likelihood $p(\{t_i\}|\vec{x})$.

3 Conclusion

Integrate-and-fire models play a distinguished role in theoretical neuroscience, but only recently have they been taken seriously as statistical models for analyzing spike train data. We have placed the IF neuron within a more general point process modeling framework, and have described several useful variations on the general likelihood-based modeling theme. In particular, the three applications we reviewed show how this class of models may be used to describe stimulus-related effects: the first and third via IF models, the second via an explicit conditional-intensity model.

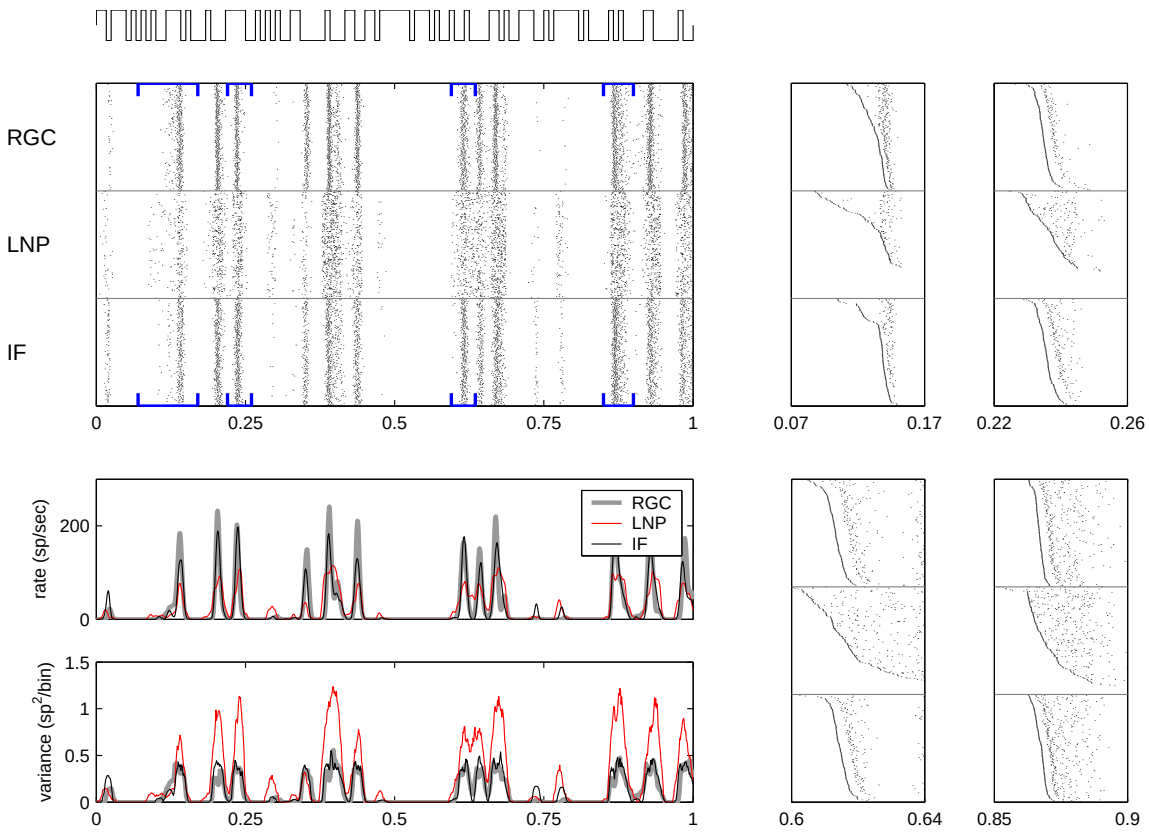

Figure 5: Responses of an OFF-type retinal ganglion cell (RGC) to a repeated stimulus. Top left: recorded responses of a single macaque RGC to repeated full-field temporally-varying light stimulus (top; inset trace indicates light intensity as a function of time), simulated inhomogeneous Poisson model (middle; LNP), and IF model (bottom) spike trains. (Both the IF and Poisson models were based on covariates including multiple time lags of the temporally-varying stimulus; the models were fit via ML to a nonoverlapping, nonrepeating segment of training data not shown here.) Each row corresponds to the response during a single stimulus repeat; 167 repeats are shown. Bottom left: peri-stimulus time histogram (mean observed spike rate) and variance for the RGC, Poisson model, and IF model. For this cell, the IF model accounts for 91% of the variance of the PSTH, whereas the Poisson model accounts for 75%. Right: magnified sections of rasters, with rows sorted in order of first spike time within the window. The four sections shown are indicated by black brackets in the upper left panels. Note that the IF model matches the true precision of the spike train much more accurately than does the Poisson model, as expected, since the Poisson model does not incorporate any spike-history effects. See (Pillow et al., 2005) for details. Reprinted with permission of the Journal of Neursocience.

Statistical models may have varying degrees of interpretability and biophysical plausibility. The primary goal of any statistical analysis of neural data is to develop an accurate and parsimonious description of the experimental data that is interpretable and consistent to the extent possible with known neurophysiological principles. The accuracy of the data description should be guided by formal goodness-of-fit assesments (Brown et al., 2002) and

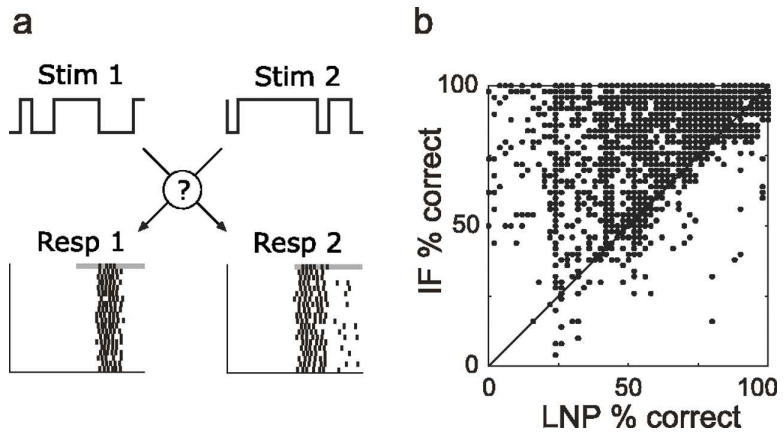

Figure 6: Decoding responses using model-derived likelihoods. a, Two stimulus (Stim) fragments and corresponding fragments of the RGC response (Resp) raster. Gray boxes highlight a 50 ms interval of the first row of each response raster. A two-alternative forced-choice (2AFC) discrimination task was performed on these response fragments, in which the task was to determine which stimulus gave rise to each response. The IF and Poisson (LNP) models were used to compute the likelihood of these responses given the "correct" and "incorrect" pairing of stimuli and responses, and the pairing with the higher likelihood was selected. This discrimination procedure was applied to each row of the response raster and used to obtain the percentage correct for the discrimination performance of each model. b, Discrimination performance of the IF and Poisson models. Each point corresponds to the percentage correct of a 2AFC discrimination task using two randomly selected 50 ms windows of the response. Although both models attain perfect performance (100%) for a majority of such randomly selected response windows, the scatter of points above the diagonal shows that when discrimination performance is imperfect, the IF model is far better at decoding the neural spike responses. See (Pillow et al., 2005) for full details. Reprinted with permission of the Journal of Neursocience.

the standard model selection approaches (Burnham and Anderson, 2002; Truccolo et al., 2005). From the point of view of analyzing a particular set of spike-train data, the decision whether to rely on IF models must consider interpretation, computational effort, and ability to account for apparent features of the data.

Several extensions of the IF models presented here are possible. In principle, the random term in equation (1) could be replaced with more realistic random term with positive temporal correlations, though this introduces substantial new computational issues (Haskell et al., 2001; Fourcaud and Brunel, 2002; Moreno et al., 2002). However, the state-space methods discussed above may still be applied in a straightforward manner. One interesting application is to model not just the voltage in a single somatic compartment, as we have discussed here, but also other dynamical variables (Badel et al., 2005; Vogelstein and Paninski, 2007) or the voltage in other unobserved compartments (Huys and Paninski, 2007). Along similar lines, state-space techniques may be used to include synaptic conductance noise effects in the IF model (Paninski, 2007), though the Gaussian noise assumption needs to be modified somewhat here, since conductances are nonnegative.

The general conditional-intensity framework allows inclusion of a wide variety of measured effects, including trial-to-trial variation terms, the local field potential (LFP), etc. In informal

schematic form, we would have

$$\text{log conditional intensity} = \text{stimulus} + \text{stimulus history} + \text{spiking history} + \text{trial} + \text{LFP}.$$

History effects may be of the IMI form discussed here, or they may reach back further in time, incoporating effects of many spikes (Kass and Ventura, 2001; Paninski, 2004; Kass et al., 2005; Truccolo et al., 2005; Paninski et al., 2007b), while trial-to-trial variation may be accommodated using slowly-varying, trial-dependent contributions to firing rate (Ventura et al., 2005; Czanner et al., 2008). The advantages of this sort of model become more apparent when one considers multiple simultaneously-recorded spike trains (Brown et al., 2004), where interactions among neurons may be modeled by inclusion of additional terms that define the conditional intensity (Chornoboy et al., 1988; Paninski et al., 2004a; Okatan et al., 2005; Truccolo et al., 2005; Kulkarni and Paninski, 2007; Pillow et al., 2008; Czanner et al., 2008). IF models have been successfully employed, for example, to explore the cross-correlation properties of pairs of simultaneously-recorded neurons (Iyengar, 1985; de la Rocha et al., 2007; Carandini et al., 2007). We expect the future of spike train data analysis to be driven by insightful extensions of these point process modeling techniques.

References

Ahmadian, Y., Pillow, J., and Paninski, L. (2008). Efficient Markov Chain Monte Carlo methods for decoding population spike trains. *Under review, Neural Computation*.

Badel, L., Richardson, M., and Gerstner, W. (2005). Dependence of the spike-triggered average voltage on membrane response properties. *Neurocomputing*, 69:1062–1065.

Barbieri, R., Quirk, M., Frank, L., Wilson, M., and Brown, E. (2001). Construction and analysis of non-poisson stimulus-response models of neural spiking activity. *Journal of Neuroscience Methods*, 105:25–37.

Brillinger, D. (1988). Maximum likelihood analysis of spike trains of interacting nerve cells. *Biological Cyberkinetics*, 59:189–200.

Brillinger, D. (1992). Nerve cell spike train data analysis: a progression of technique. *Journal of the American Statistical Association*, 87:260–271.

Brown, E. (2005). The theory of point processes for neural systems. In Chow, C., Gutkin, B., Hansel, D., Meunier, C., and Dalibard, J., editors, *Methods and Models in Neurophysics*, pages 691–726. Elsevier.

Brown, E., Barbieri, R., Eden, U., and Frank, L. (2003). Likelihood methods for neural data analysis. In Feng, J., editor, *Computational Neuroscience: a comprehensive approach*, pages 253–286. CRC Press.

Brown, E., Barbieri, R., Ventura, V., Kass, R., and Frank, L. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14:325–346.

Brown, E., Frank, L., Tang, D., Quirk, M., and Wilson, M. (1998). A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *Journal of Neuroscience*, 18:7411–7425.

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

Brown, E., Nguyen, D., Frank, L., Wilson, M., and Solo, V. (2001). An analysis of neural receptive field plasticity by point process adaptive filtering. *PNAS*, 98:12261–12266.

Brunel, N. and Hakim, V. (1999). Fast global oscillations in networks of integrate-and-fire neurons with low firing rates. *Neural Computation*, 11:1621–1671.

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

Buoncore, A., Nobile, A., and Ricciardi, L. (1987). A new integral equation for the evaluation of first-passage-time probability densities. *Adv Appl Prob*, 19:784–800.

Burkitt, A. (2006). A review of the integrate-and-fire neuron model. *Biological Cybernetics*, 95:1–19, 97–112.

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

Burnham, K. and Anderson, D. (2002). *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach*. Springer.

Carandini, M., Horton, J., and Sincich, L. (2007). Thalamic filtering of retinal spike trains by postsynaptic summation. *Journal of Vision*, 7:1–11.

Chichilnisky, E. (2001). A simple white noise analysis of neuronal light responses. *Network: Computation in Neural Systems*, 12:199–213.

Chornoboy, E., Schramm, L., and Karr, A. (1988). Maximum likelihood identification of neural point process systems. *Biological Cybernetics*, 59:265–275.

Czanner, G., Eden, U., Wirth, S., Yanike, M., Suzuki, W., and Brown, E. (2008). Analysis of between-trial and within-trial neural spiking dynamics. *Journal of Neurophysiology*, In press.

Czanner, G. and Iyengar, S. (2004). Expectation-Maximization (EM) estimation of an integrate-and-fire model with spike-frequency adaptation. *COSYNE*.

Daniels, H. (1982). Sequential tests contructed from images. *Annals of Statistics*, 10:394–400.

de la Rocha, J., Doiron, B., Shea-Brown, E., Josic, K., and Reyes, A. (2007). Correlation between neural spike trains increases with firing rate. *Nature*, 448:802–806.

de Ruyter van Steveninck, R. and Bialek, W. (1988). Real-time performance of a movement-senstivive neuron in the blowfly visual system: coding and information transmission in short spike sequences. *Proc. R. Soc. Lond. B*, 234:379–414.

Dempster, A., Laird, N., and Rubin, D. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal Royal Stat. Soc., Series B*, 39:1–38.

DiNardo, E., Nobile, A., Pirozzi, E., and Ricciardi, L. (2001). A computational approach to first-passage-time problems for Gauss-Markov processes. *Advances in Applied Probability*, 33:453–482.

Dombeck, D. A., Khabbaz, A. N., Collman, F., Adelman, T. L., and Tank, D. W. (2007). Imaging large-scale neural activity with cellular resolution in awake, mobile mice. *Neuron*, 56(1):43–57.

Ermentrout, G. and Kopell, N. (1986). Parabolic bursting in an excitable system coupled with a slow oscillation. *SIAM Journal on Applied Math*, 2:233–253.

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

Frank, L., Eden, U., Solo, V., Wilson, M., and Brown, E. (2002). Contrasting patterns of receptive field plasticity in the hippocampus and the entorhinal cortex: An adaptive filtering approach. *J. Neurosci.*, 22(9):3817–3830.

Gerstein, G. and Mandelbrot, B. (1964). Random walk models for the spike activity of a single neuron. *Biophysical Journal*, 4:41–68.

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

Haskell, E., Nykamp, D., and Tranchina, D. (2001). Population density methods for large-scale modelling of neuronal networks with realistic synaptic kinetics. *Network: Computation in Neural Systems*, 12:141–174.

Huys, Q., Ahrens, M., and Paninski, L. (2006). Efficient estimation of detailed single-neuron models. *Journal of Neurophysiology*, 96:872–890.

Huys, Q. and Paninski, L. (2007). Model-based smoothing of, and parameter estimation from, noisy biophysical recordings. *PLOS Computational Biology*, Under review.

Iyengar, S. (1985). Hitting lines with two-dimensional Brownian motion. *SIAM Journal on Applied Mathematics*, 45:983–989.

Iyengar, S. and Liao, Q. (1997). Modeling neural activity using the generalized inverse Gaussian distribution. *Biological Cybernetics*, 77:289–295.

Izhikevich, E. (2007). *Dynamical Systems in Neuroscience*. MIT Press.

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

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

Kass, R. and Ventura, V. (2001). A spike-train probability model. *Neural Comp.*, 13:1713–1720.

Kass, R. E., Ventura, V., and Brown, E. N. (2005). Statistical issues in the analysis of neuronal data. *J Neurophysiol*, 94:8–25.

Keat, J., Reinagel, P., Reid, R., and Meister, M. (2001). Predicting every spike: a model for the responses of visual neurons. *Neuron*, 30:803–817.

Knight, B. (1972). Dynamics of encoding in a population of neurons. *J. Gen. Physiol.*, 59:734–766.

Knight, B., Omurtag, A., and Sirovich, L. (2000). The approach of a neuron population firing rate to a new equilibrium: an exact theoretical result. *Neural Computation*, 12:1045–1055.

Koyama, S. and Kass, R. (2008). Spike train probability models for stimulus-driven leaky integrate-and-fire neurons. *Neural Computation*, In press.

Koyama, S., Kass, R., and Paninski, L. (2008). Efficient computation of the most likely path in integrate-andfire and more general state-space models. *COSYNE*.

Kulkarni, J. and Paninski, L. (2007). Common-input models for multiple neural spike-train data. *Network: Computation in Neural Systems*, 18:375–407.

Lansky, P., Sanda, P., and He, J. (2006). The parameters of the stochastic leaky integrate-and-fire neuronal model. *Journal of Computational Neuroscience*, 21:211–223.

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

Moreno, R., de la Rocha, J., Renart, A., and Parga, N. (2002). Response of spiking neurons to correlated inputs. *Physical Review Letters*, 89:288101.

Mullowney, P. and Iyengar, S. (2007). Parameter estimation for a leaky integrate-and-fire neuronal model from interspike interval data. *Journal of Computational Neuroscience*, doi:10.1007/s10827-007-0047-5.

Nikitchenko, M. and Paninski, L. (2007). An expectation-maximization Fokker-Planck algorithm for the noisy integrate-and-fire model. *COSYNE*.

Nykamp, D. and Tranchina, D. (2000). A population density approach that facilitates large-scale modeling of neural networks: analysis and an application to orientation tuning. *Journal of Computational Neuroscience*, 8:19–50.

Okatan, M., Wilson, M., and Brown, E. (2005). Analyzing functional connectivity using a network likelihood model of ensemble neural spiking activity. *Neural Computation*, 17:1927–1961.

Paninski, L. (2004). Maximum likelihood estimation of cascade point-process neural encoding models. *Network: Computation in Neural Systems*, 15:243–262.

Paninski, L. (2005). Log-concavity results on Gaussian process methods for supervised and unsupervised learning. *Advances in Neural Information Processing Systems*, 17.

Paninski, L. (2006a). The most likely voltage path and large deviations approximations for integrate-and-fire neurons. *Journal of Computational Neuroscience*, 21:71–87.

Paninski, L. (2006b). The spike-triggered average of the integrate-and-fire cell driven by Gaussian white noise. *Neural Computation*, 18:2592–2616.

Paninski, L. (2007). Inferring synaptic inputs given a noisy voltage trace via sequential Monte Carlo methods. *Journal of Computational Neuroscience*, Under review.

Paninski, L., Fellows, M., Shoham, S., Hatsopoulos, N., and Donoghue, J. (2004a). Superlinear population encoding of dynamic hand trajectory in primary motor cortex. *J. Neurosci.*, 24:8551–8561.

Paninski, L., Haith, A., and Szirtes, G. (2007a). Integral equation methods for computing likelihoods and their derivatives in the stochastic integrate-and-fire model. *Journal of Computational Neuroscience*, 24:69–79.

Paninski, L., Pillow, J., and Lewi, J. (2007b). Statistical models for neural encoding, decoding, and optimal stimulus design. In Cisek, P., Drew, T., and Kalaska, J., editors, *Computational Neuroscience: Progress in Brain Research*. Elsevier.

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

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

Pillow, J. and Paninski, L. (2007). Model-based decoding, information estimation, and change-point detection in multi-neuron spike trains. *Under review, Neural Computation*.

Pillow, J., Paninski, L., Uzzell, V., Simoncelli, E., and Chichilnisky, E. (2005). Prediction and decoding of retinal ganglion cell responses with a probabilistic spiking model. *Journal of Neuroscience*, 25:11003–11013.

Pillow, J., Shlens, J., Paninski, L., Simoncelli, E., and Chichilnisky, E. (2008). Visual information coding in multi-neuronal spike trains. *Nature*, In press.

Plesser, H. and Tanaka, S. (1997). Stochastic resonance in a model neuron with reset. *Physics Letters A*, 225:228–234.

Press, W., Teukolsky, S., Vetterling, W., and 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. *Proceedings of the IEEE*, 77:257–286.

Reich, D., Victor, J., and Knight, B. (1998). The power ratio and the interval map: Spiking models and extracellular recordings. *The Journal of Neuroscience*, 18:10090–10104.

Ricciardi, L. (1977). *Diffusion processes and related topics in biology*. Springer.

Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialek, W. (1997). *Spikes: Exploring the neural code*. MIT Press, Cambridge.

Risken, H. (1996). *The Fokker-Planck Equation*. Springer.

Schnitzer, M. and Meister, M. (2003). Multineuronal firing patterns in the signal from eye to brain. *Neuron*, 37:499–511.

Seshadri, V. (1993). *The inverse Gaussian distribution*. Clarendon, Oxford.

Siegert, A. (1951). On the first passage time probability problem. *Physical Review*, 81:617–623.

Simoncelli, E., Paninski, L., Pillow, J., and Schwartz, O. (2004). Characterization of neural responses with stochastic stimuli. In *The Cognitive Neurosciences*. MIT Press, 3rd edition.

Smith, A. and Brown, E. (2003). Estimating a state-space model from point process observations. *Neural Computation*, 15:965–991.

Snyder, D. and Miller, M. (1991). *Random Point Processes in Time and Space*. Springer-Verlag.

Srinivasan, L., Eden, U., Willsky, A., and Brown, E. (2006). A state-space analysis for reconstruction of goal-directed movements using neural signals. *Neural Computation*, 18:2465–2494.

Stein, R. (1965). A theoretical analysis of neuronal variability. *Biophysical Journal*, 5:173–194.

Stevens, C. and Zador, A. (1998). Novel integrate-and-fire-like model of repetitive firing in cortical neurons. *Proc. 5th joint symp. neural computation, UCSD*.

Truccolo, W., Eden, U., Fellows, M., Donoghue, J., and Brown, E. (2005). A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *Journal of Neurophysiology*, 93:1074–1089.

Tuckwell, H. (1989). *Stochastic Processes in the Neurosciences*. SIAM.

Ventura, V., Cai, C., and Kass, R. (2005). Trial-to-trial variability and its effect on time-varying dependence between two neurons. *Journal of Neurophysiology*, 94:2928–2939.

Vogelstein, J. and Paninski, L. (2007). Model-based optimal inference of spike times and calcium dynamics given noisy and intermittent calcium-fluorescence imaging. *Biophysical Journal*, Under review.

Yu, B., Afshar, A., Santhanam, G., Ryu, S., Shenoy, K., and Sahani, M. (2006). Extracting dynamical structure embedded in neural activity. *NIPS*.