

Spike Train Probability Models for Stimulus-Driven Leaky Integrate-and-Fire Neurons

Shinsuke Koyama and Robert E. Kass

Department of Statistics and Center for the Neural Basis of Cognition
Carnegie Mellon University
Pittsburgh, PA 15213, U.S.A.

October 23, 2007

Abstract

Mathematical models of neurons are widely used to improve understanding of neuronal spiking behavior. These models can produce artificial spike trains that resemble actual spike-train data in important ways, but they are not very easy to apply to the analysis of spike-train data. Instead, statistical methods based on point process models of spike trains provide a wide range of data-analytical techniques. Two simplified point process models have been introduced in the literature: the time-rescaled renewal process (TRRP) and the multiplicative inhomogeneous Markov interval (m-IMI) model. In this article we investigate the extent to which the TRRP and m-IMI models are able to fit spike trains produced by stimulus-driven leaky integrate-and-fire (LIF) neurons.

With a constant stimulus the LIF spike train is a renewal process, and the m-IMI and TRRP models will describe accurately the LIF spike train variability. With a time-varying stimulus, the probability of spiking under all three of these models depends on both the experimental clock time relative to the stimulus and the time since the previous spike, but it does so differently for the LIF, m-IMI, and TRRP models. We assessed the distance between the LIF model and each of the two empirical models in the presence of a time-varying

stimulus. We found that while lack of fit of a Poisson model to LIF spike-train data can be evident even in small samples, the m-IMI and TRRP models tend to fit well, and much larger samples are required before there is statistical evidence of lack of fit of either the m-IMI or TRRP models. We also found that when the mean of the stimulus varies across time the m-IMI model provides a better fit to the LIF data than the TRRP, while when the variance of the stimulus varies across time the TRRP provides the better fit.

1 Introduction

The leaky integrate-and-fire (LIF) model is one of the fundamental building blocks of theoretical neuroscience (Dayan & Abbott, 2001; Gerstner & Kistler, 2002; Koch, 1999; Tuckwell, 1988). Its use in examining spike-train data is limited, however, because its full parameter vector can not be estimated uniquely from spike trains in the absence of sub-threshold measurements and estimation of a reduced set of parameters is somewhat subtle (Iyengar & Mullen, 2007; Paninski, Pillow & Simoncelli, 2004). An alternative is to use likelihood methods based on the conditional intensity function $\lambda(t|H_t)$, where H_t is the complete history of spiking preceding time t (Kass, Ventura & Brown, 2005 and references therein). The stimulus-driven LIF model depends only on the experimental clock time t , relative to stimulus onset, and the elapsed time $t - s_*(t)$ since the preceding spike $s_*(t)$, i.e., it satisfies

$$\lambda(t|H_t) = \lambda(t, t - s_*(t)). \quad (1.1)$$

Models of the general form of Equation 1 have been called “inhomogeneous Markov interval (IMI) models” by Kass & Ventura, 2001, following Cox & Lewis, 1972, and “1-memory point processes” by Snyder & Miller, 1991. Two special cases of Equation 1 have been considered in the literature: multiplicative IMI models (Kass & Ventura, 2001, and references therein) and time-rescaled renewal-process models (Barbieri, Quirk, Frank, Wilson & Brown, 2001; Koyama & Shinomoto, 2005; Reich, Victor & Knight, 1998). We label these, respectively, m-IMI and TRRP models. Both the m-IMI and TRRP models separate the dependence on t from the dependence on $t - s_*(t)$, but they do so differently. The purpose of this paper is to investigate their relationship to the stimulus-driven LIF model. Specifically, given a joint probability distribution of spike trains generated by an LIF model, we ask

how close this distribution is to each of the best-fitting m-IMI and TRRP models, where closeness is measured using Kullback-Leibler divergence. This quantifies the extent to which the m-IMI model and the TRRP model can capture the dynamics of a stimulus-driven LIF neuron. It also provides an interpretative distinction between the m-IMI and TRRP models themselves.

2 Method

2.1 Leaky integrate-and-fire model

The LIF model is the simplest model that retains the minimal ingredients of membrane dynamics (Dayan & Abbott, 2001; Gerstner & Kistler, 2002; Koch, 1999; Tuckwell, 1988). The dynamics of the model are represented by the equation,

$$\tau \frac{dV(t)}{dt} = -V(t) + I(t), \quad (2.1)$$

where $V(t)$ is the membrane potential of the cell body measured from its resting level, τ is the membrane decay time constant, and $I(t)$ represents an input current. When the membrane potential reaches the threshold, v_{th} , a spike is evoked and the membrane potential is reset to v_0 immediately.

By suitable scale transformation, the original model can be reduced to a normalized one,

$$\frac{dX(t)}{dt} = -X(t) + I(t). \quad (2.2)$$

The threshold value and reset potential are given by x_{th} and x_0 , respectively. While $I(t)$ represents an external input, x_{th} and x_0 could be interpreted as “intrinsic” parameters of the neuron model and are directly related to biophysical properties (Lansky, Sanda & He, 2006). In this article, we consider stimuli which have the form

$$I(t) = \mu(t) + \sigma(t)\xi(t), \quad (2.3)$$

where $\xi(t)$ is Gaussian white noise with $E[\xi(t)] = 0$, $V(\xi(t)) = 1$, and $Cov(\xi(t), \xi(t')) = 0$, for $t \neq t'$.

2.2 Probability models of spike trains

A point process can be fully characterized by a conditional intensity function (Daley & Vere-Jones, 2003; Snyder & Miller, 1991). We consider the following two classes of models.

- m-IMI model

The conditional intensity function of the m-IMI model has the form

$$\lambda(t, s_*(t)) = \lambda_1(t)g_1(t - s_*(t)). \quad (2.4)$$

Here, $\lambda_1(t)$ modulates the firing rate only as a function of experimental clock while $g_1(t - s_*(t))$ represents non-Poisson spiking behavior.

- TRRP model

This model has the form

$$\lambda(t, s_*(t)) = \lambda_0(t)g_0(\Lambda_0(t) - \Lambda_0(s_*(t))), \quad (2.5)$$

where g_0 is the hazard function of a renewal process, and $\Lambda_0(t)$ is defined as

$$\Lambda_0(t) = \int_0^t \lambda_0(u)du. \quad (2.6)$$

In this article, we call λ_0 and λ_1 *excitability functions* to indicate that they modulate the amplitude of the firing rate, and we call g_0 and g_1 *recovery functions* to indicate that they affect the way the neuron recovers its ability to fire after generating an action potential. The fundamental difference between the two models is the way the excitability interacts with the recovery function. In the m-IMI model the refractory period represented in the recovery function is not affected by excitability or firing rate variations. In the TRRP model, however, the refractory period is no longer fixed but is scaled by the firing rate (Reich, Victor & Knight, 1998).

Note that in both the m-IMI model and the TRRP, the excitability and recovery functions are defined only up to a multiplicative constant: replacing $\lambda_1(t)$ and $g_1(t - s_*(t))$ by $c\lambda_1(t)$ and $c g_1(t - s_*(t))$, for any positive c , leaves the model unchanged, and similarly for λ_0 and g_0 . This arbitrary constant must be fixed by some convention in implementation.

2.3 Kullback-Leibler divergence

We use the Kullback-Leibler (KL) divergence to evaluate closeness between the LIF model and the two empirical models. The KL divergence is a coefficient measuring a non-negative asymmetrical “distance” from one probability distribution to another, and the model distribution with a lower value of the KL divergence approximates the original probability distribution better. For two probability densities p_f and p_q of a spike train $\{t_1, \dots, t_n\}$ in the interval $[0, T']$, the KL divergence per spike between the two densities is given by

$$D(p_f||p_q) = \lim_{T' \rightarrow \infty} \frac{1}{E[n]} \sum_{n=1}^{\infty} \int_0^{T'} \int_{t_1}^{T'} \cdots \int_{t_{n-1}}^{T'} p_f(t_1, \dots, t_n) \log \frac{p_f(t_1, \dots, t_n)}{p_q(t_1, \dots, t_n)} dt_1 dt_2 \cdots dt_n, \quad (2.7)$$

where $E[n]$ is the mean spike count in the interval $[0, T']$, the expectation being taken over replications (trials). For simplicity, in considering first $\mu(t)$ and then $\sigma(t)$ to be time-varying functions, we will assume in each case that they are periodic with period T , so that spike train segments across time intervals of the form $[kT, (k+1)T)$ for nonnegative integers k may be considered replications (in other words, the spike train generated from the LIF neuron becomes periodically stationary). Let $\theta(t) = t \bmod T$ be the phase of the periodic stimulus where we take the phase of the stimulus to be zero at $t = 0$, and $\lambda(t, s_*)$ be a conditional intensity of an IMI model. The conditional inter-spike interval density given the previous spike phase θ , $q(u|\theta)$, is obtained as

$$q(t - s_*|\theta(s_*)) = \lambda(t, s_*) \exp \left(- \int_{s_*}^t \lambda(x, s_*) dx \right). \quad (2.8)$$

Since the LIF model belongs to the class of IMI model, it is completely characterized by a conditional inter-spike interval density, $f(u|\theta)$. Then, as shown in appendix A, under the condition of $E[n] = \int_0^{T'} \lambda_0(t) dt \rightarrow \infty$ as $T' \rightarrow \infty$, where $\lambda_0(t)$ is the instantaneous firing rate given by $\lambda_0(t) = E[\lambda(t, s_*(t))]$, equation 2.7 is reduced to

$$D(p_f||p_q) = \int_0^T \left(\int_0^\infty f(u|\theta) \log \frac{f(u|\theta)}{q(u|\theta)} du \right) \chi(\theta) d\theta, \quad (2.9)$$

where $\chi(\theta) \equiv p\{\text{spike at } \theta | \text{one spike in } [0, T)\}$ is a stationary spike phase density. Under the periodic stationary condition the KL divergence between

two probability densities of spike trains of IMI models (equation 2.7) is reduced to the KL divergence between the conditional inter-spike interval densities (equation 2.9). Note that the KL divergence of the conditional inter-spike interval density is averaged over the spike phase distribution, $\chi(\theta)$, since spikes are distributed by $\chi(\theta)$ over time.

For calculating the KL divergence given by equation 2.9, we need to calculate $f(u|\theta)$ and $\chi(\theta)$ which can be obtained as follows. Since the membrane potential of the LIF model is a Markov diffusion process, $f(u|\theta)$ satisfies the renewal equation (van Kampen, 1992)

$$p(x, t|x_0, 0; \theta) = \int_0^t f(t'|\theta)p(x, t|x_{th}, t'; \theta)dt', \quad x \geq x_{th}, \quad (2.10)$$

where $p(x, t|x_1, t_1; \theta)$ is the conditional probability density that the voltage is x at time t if it is x_1 at time $t_1 < t$. This conditional probability density can be obtained by solving the stochastic differential equation 2.2 as

$$p(x, t|x_1, t_1; \theta) = \frac{1}{\sqrt{2\pi\eta(t)}} \exp \left[-\frac{(x - \gamma(t))^2}{2\eta(t)} \right], \quad (2.11)$$

where

$$\gamma(t) = x_1 e^{t_1-t} + \int_{t_1}^t e^{t'-t} \mu(t' + \theta) dt', \quad (2.12)$$

and

$$\eta(t) = \int_{t_1}^t e^{2(t'-t)} \sigma^2(t' + \theta) dt', \quad (2.13)$$

see van Kampen, 1992. Inserting equations 2.11-2.13 into equation 2.10 we can solve equation 2.10 numerically to obtain $f(u|\theta)$. (See Burkitt & Clark, 2001 for numerical solution of the integral equation.)

The stationary spike phase distribution $\chi(\theta)$ can be obtained as follows. The phase transition density, which is the probability density that a spike occurs at phase θ' given the previous spike at phase θ , is given by

$$g(\theta'|\theta) = \int_0^\infty f(u|\theta) \delta([u + \theta]_{\text{mod } T} - \theta') du. \quad (2.14)$$

Following the standard theory of Markov processes (see van Kampen, 1992), the stationary spike phase distribution $\chi(\theta)$ is obtained as a solution of

$$\chi(\theta') = \int_0^T g(\theta'|\theta) \chi(\theta) d\theta. \quad (2.15)$$

The stationary phase distribution $\chi(\theta)$ is the eigenfunction of $g(\theta'|\theta)$ corresponding to the unique eigenvalue 1. In practice, the spike phase distribution $\chi(\theta)$ can be calculated by discretizing the phase and calculating the eigenvector corresponding to eigenvalue 1 of the transition probability matrix using standard eigenvector routines (Plesser & Geisel, 1999).

2.4 Fitting IMI model

We fit the m-IMI model and TRRP model to spike trains derived from the LIF model via maximum likelihood.

2.4.1 Fitting the m-IMI model

We follow Kass & Ventura, 2001 to fit the m-IMI model to data. For fitting m-IMI model in equation 2.4, we use the additive form

$$\log \lambda(t, s_*(t)) = \log \lambda_1(t) + \log g_1(t - s_*(t)). \quad (2.16)$$

We first represent a spike train as a binary sequence of 0s and 1s by discretizing time into small intervals of length Δ , letting 1 indicate that a spike occurred within the corresponding time interval. We represent $\log \lambda_1(t)$ and $\log g_1(t - s_*(t))$ with cubic splines. Given suitable knots for both terms, cubic splines may be described by linear combinations of B-spline basis functions (de Boor, 2001),

$$\log \lambda_1(i\Delta) = \sum_{k=1}^M \alpha_k A_k(i\Delta), \quad (2.17)$$

$$\log g_1(i\Delta - s_*(i\Delta)) = \sum_{k=1}^L \beta_k B_k(i\Delta - s_*(i\Delta)), \quad (2.18)$$

where M and L are the numbers of basis functions that are determined by the order of splines and the number of knots. Note that the shapes of B-spline basis functions $\{A_k\}$ and $\{B_k\}$ also depend on the location of knots. Fitting of the model is accomplished easily via maximum likelihood: for fixed knots the model is binary generalized linear model (McCullagh & Nelder, 1989) with

$$\log \lambda(i\Delta, s_*(i\Delta)) = \sum_{k=1}^M \alpha_k A_k(i\Delta) + \sum_{k=1}^L \beta_k B_k(i\Delta - s_*(i\Delta)), \quad (2.19)$$

where $\{A_k(i\Delta)\}$ and $\{B_k(i\Delta)\}$ play the role of explanatory variables. The coefficients of the spline basis elements, $\{\alpha_k\}$ and $\{\beta_k\}$, are determined via maximum likelihood. This can be performed by using a standard software such as R and Matlab Statistics Toolbox.

In the following simulations we chose the knots by preliminary examination of data. We conducted the fitting procedure for several candidates of knots, and then chose the one which minimizes the KL divergence between the estimate and the LIF model.

2.4.2 Fitting the TRRP model

We begin by noting that $\lambda_0(t)$ is a constant multiple of the trial-averaged conditional intensity (see Appendix B). To fix the arbitrary constant we normalize so that the constant multiple becomes 1. We may then first estimate $\lambda_0(t)$ from data pooled across trials (effectively smoothing the PSTH), which we do by representing it with a cubic spline and using a binary regression generalized linear model. Then we apply the time-rescaled transformation given by equation 2.6 to spike train $\{t_i\}$ to obtain a rescaled spike train, $\{\Lambda_0(t_i)\}$. Finally we determine an inter-spike interval distribution of $\{\Lambda_0(t_i)\}$ by representing the log density with a cubic spline and again using binary regression.

3 Results

3.1 Time-varying mean input

We first considered the case that the input to the LIF neuron was

$$I(t) = a \sin \frac{t}{\tau_s} + \xi(t), \quad (3.1)$$

where a and τ_s are amplitude and time scale of the mean of the input, respectively. We took $(x_{th}, x_0) = (0.5, 0)$ and $\tau_s = 5/\pi$. In the following simulation we used the Euler integration with a time step $\Delta = 10^{-3}$. We first simulated the LIF model over the time interval $10^7\Delta$ to generate a spike train, and fitted the m-IMI model and TRRP model to data. For each model we used 4 knots for fitting the excitability function, and 3 knots each for fitting the recovery function and the inter-spike interval distribution. Then we calculated the KL divergence between the LIF model and those fitted probability

models as described in Section 2.3. We repeated this procedure ten times to calculate the mean and the standard error of the value of the KL divergence.

Figure 1(a) depicts the KL divergences as a function of the amplitude of the mean of the input. We also show in this figure the KL divergence between the LIF model and the inhomogeneous Poisson process for comparison. As shown in this figure, the KL divergence of the m-IMI model is the smallest among three models. When the amplitude of the mean input is increased, the KL divergences of these models get increased. It is remarkable that even when the firing rate is highly modulated as seen in Figure 1(b), the KL divergence of the m-IMI model remains small.

Figure 1(c) and (d) display rescaled interval distributions which are calculated from recovery functions of the m-IMI model and the TRRP model, respectively, as

$$p(t) = \hat{g}(t) \exp \left(- \int_0^t \hat{g}(u) du \right), \quad (3.2)$$

where $\hat{g}(t)$ denotes the fitted recovery function. Note that this is not the actual interval distribution of spike trains, but the one which is extracted from a spike train after removing the effects of the stimulus. (Here we show $p(t)$ but $\hat{g}(t)$ since the shape of $p(t)$ is more stable in the tail of the distribution: the tail of $p(t)$ converges to zero as $t \rightarrow \infty$, while the estimation of the tail of $\hat{g}(t)$ is rather variable because there are few spikes at large t in data.) The gray dashed line in (c) and (d) represents the inter-spike interval distribution of the LIF model with $a = 0$. The interval distribution plot shows less variation for the m-IMI model than for the TRRP model as the amplitude of the stimuli is increased, especially in the range of short time scale of the refractory period. This indicates that the dynamics of LIF model in the range of short time scale is not affected by the stimuli, and the recovery function of the m-IMI model can capture the stimulus-independent spiking characteristics of the LIF model.

Figure 2 displays the results for various values of τ_s and fixed value of a . Parameter values of the LIF model were taken as $(x_{th}, x_0) = (0.5, 0)$ and $a = 1$. The results are qualitatively the same as Figure 1, i.e., the m-IMI model fits better than the TRRP model, and the interval distribution of the m-IMI model shows less variation than that of the TRRP model. These results depicted in Figures 1 and 2 suggest that for various amplitudes and time scales of a stimulus the fitted m-IMI model is closer than the TRRP model to the true LIF model when a stimulus is applied to mean of the LIF model.

3.2 Statistical interpretation of the KL divergence

The scale of KL divergence is, by itself, somewhat difficult to interpret. However, in general terms, the value of the KL divergence is related to the probability that a test will reject a false null hypothesis: as the KL divergence increases, the probability of rejecting the hypothesis increases. We used the Kolmogorov-Smirnov test to provide a statistical interpretation of the KL divergence, and we obtained corresponding K-S plots (Brown, Barbieri, Ventura & Kass, 2002) to provide visual display of fit in particular cases.

The K-S plot begins with rescaled inter-spike intervals

$$y_i = \Lambda(t_i) - \Lambda(t_{i-1}), \quad (3.3)$$

where $\Lambda(t)$ is the time-rescaling transformation obtained from the estimated conditional intensity

$$\Lambda(t) = \int_0^t \hat{\lambda}(u, s_*(u)) du. \quad (3.4)$$

If the conditional intensity were correct, then according to the time-rescaling theorem the y_i s would be independent exponential random variables with mean 1. Using the further transformation

$$z_i = 1 - \exp(-y_i), \quad (3.5)$$

z_i s would then be independent uniform random variables on the interval $(0, 1)$. In the K-S plot we order the z_i s from smallest to largest and, denoting the ordered values as $z_{(i)}$, plot the values of the cumulative distribution function of the uniform density, i.e., $b_i = \frac{i-1}{n}$ for $i = 1, \dots, n$, against the $z_{(i)}$ s. If the model were correct, then the points would lie close to a 45-degree line. For moderate to large sample sizes the 95% probability bands are well approximated as $b_i \pm 1.36/n^{1/2}$ (Johnson & Kotz, 1970). The Kolmogorov-Smirnov test rejects the null-hypothetical model if any of the plotted points lie outside these bands.

Using this procedure, we first generated spike trains from the LIF model with parameter values $(x_{th}, x_0) = (0.5, 1)$ and $(a, \tau_s) = (1.4, 5/\pi)$ (the same as in Figure 1), and fitted the empirical models to the data. The integral in the time-rescaling transformation (equation 3.4) was computed with discrete time step, $\Delta = 10^{-3}$. We also confirmed that the results did not change with the temporal precision $\Delta = 10^{-4}$. In order to calculate the rate at which the empirical models are rejected by the Kolmogorov-Smirnov test,

we generated one hundred sets of repeated spike train trials, with varying numbers of trials—and, thus, varying total numbers of spikes (ISIs).

Figure 3 displays the rate of rejecting the models as a function of the number of spikes. The probability of rejecting the inhomogeneous Poisson process rises very fast, while it takes very much larger data sets to reject the m-IMI model and the TRRP. Figures 4(a) and (b) depict examples of Kolmogorov-Smirnov plots. With only 200 ISIs (part (a)), there is clear lack of fit of the Poisson model. With 7000 ISIs (part (b)) lack of fit of the TRRP becomes apparent, though the m-IMI model continues to fit the mean-modulated LIF data. From Figure 3, in this example, roughly 10^4 ISIs are necessary to reject the TRRP with probability 0.8.

3.3 Time-varying input variance

Next we considered the input whose variance varies across time,

$$I(t) = \sigma(t)\xi(t), \quad (3.6)$$

where

$$\sigma^2(t) = 1 + a \sin \frac{t}{\tau_s}. \quad (3.7)$$

This form could be interpreted as a signal in the variance of the input, of which Silberberg et al considered the possibility in the cortex (Silberberg, Bethge, Markram, Pawelzik & Tsodyks, 2004). We set the parameter values of the LIF neuron to $(x_{th}, x_0) = (0, -4)$ and $\tau_s = 10/\pi$. We simulated the LIF model over the time interval $10^7\Delta$ to generate a spike sequence, fit the three probability models to data, and calculated the KL divergence.

The result of the KL divergence is displayed in Figure 5(a). In contrast to the case of time-varying mean input, the TRRP model fits the best to the LIF neuron for this case. The value of the KL divergence of the TRRP model remains small even when the firing rate varies largely (Figure 5(b)).

Figures 5(c),(d) depict the rescaled interval distributions of the m-IMI model and the TRRP model, respectively. The interval distributions of the TRRP model are almost identical even when the amplitude of the variance modulation is changed whereas interval-distributions of the m-IMI model show variation.

Figure 6 displays the results for various values of τ_s and fixed value of a . The parameter values of the LIF model in this figure are $(x_{th}, x_0) = (0, -4)$

and $a = 0.8$. It is confirmed from this figure that the results are qualitatively the same as Figure 5. That is, the time-rescaled renewal process shows the best agreement with the LIF neuron, and the best fitted interval distribution of the TRRP model is invariant for different values of τ_s .

3.4 Another example

So far we have examined the ability of the statistical models to accommodate data from the LIF model driven by sinusoidal stimuli. We performed another simulation to confirm that the results are robust against different stimuli. In this example, we took $f(t)$ to be a stimulus which has the form represented in Figure 7(a). Two cases were considered: the mean of the input to the LIF neuron varies across time, $I(t) = f(t) + \xi(t)$, and the variance of the input varies across time, $I(t) = f(t)\xi(t)$. The parameter values of the LIF neuron were taken to be $(x_{th}, x_0) = (0.5, 0)$. We simulated the LIF model to generate sequences of spikes, fitted the m-IMI, TRRP, and the Poisson model to data, and conducted the KS test. Figures 7(b) and (c) depict examples of Kolmogorov-Smirnov plots. For the mean-modulated case (part (a)), the lack of fit of the TRRP model becomes apparent, though the m-IMI model continues to fit the LIF model with 6000 ISIs. For the case of variance-modulated version (part (b)), however, the m-IMI model shows the lack of fit with 1000 ISIs while the TRRP model continues to fit the LIF model.

4 Discussion

Our results examined the extent to which regularity and variability in spike trains generated by a stimulus-driven LIF neuron could be captured by two empirical models, the m-IMI model and the TRRP. Although the LIF model involves a gross simplification of neuronal biophysics, it remains very widely applied in theoretical studies. This is one reason we investigated the performance of statistical models for stimulus-driven LIF spike trains. Our more fundamental motivation, however, was that LIF spike trains serve as a vehicle for quantifying the similarity and differences between the m-IMI and TRRP specifications of history effects. For a constant stimulus the LIF model becomes stationary, generating a renewal process of spike trains, and the m-IMI and TRRP models become identical. The qualitative distinction between the m-IMI and TRRP models in the non-stationary case may be understood by

considering the way the refractory period is treated: if, during an interval of stationarity, for which the neuron has a constant firing rate, there is a refractory period of length δ , then when the firing rate changes the TRPP model will vary the refractory period away from δ while the m-IMI model will leave it fixed. More complicated effects, described under stationarity by a renewal process will, similarly, according to the two empirical models, either vary in time with the firing rate or remain time-invariant. Thus, it is perhaps not surprising that when a stimulus was applied to the mean of the LIF model, the fitted m-IMI model was closer than the TRRP model to the true LIF model that generated the data. The interval distribution plots in Figures 1 and 2 are consistent with this interpretation.

On the other hand, when the variance in the LIF model is temporally modulated the renewal effects, such as the refractory period, become distorted in time. In this case we observed that the TRRP model fits better than the m-IMI model, and the interval distribution plot in Figure 5 shows less variation for the TRRP model than for the m-IMI model.

We also computed the power of the Kolmogorov-Smirnov (K-S) test as a function of the number of spikes. We found that with relatively small sample sizes it can be easy to reject the hypothesis of Poisson firing, but very hard to distinguish m-IMI from TRRP models. This is demonstrated in the K-S plot of Figure 3(b). For the very large data set used in the K-S plot of Figure 3(c), the m-IMI model provided a satisfactory fit while the TRRP did not. Overall, we would conclude first, that in most practical circumstances it is unlikely to matter much whether one uses the m-IMI model or the TRRP to produce empirical fits to spike-train data but, second, m-IMI models would be preferred when a mean-modulated LIF conception might be thought to represent reality better than the variance-modulated version, and vice-versa.

Statistical models used to characterize such things as a receptive field, or the effect of an oscillatory local field potential, must account for spike history effects. A successful approach has been to include m-IMI or TRRP terms in a loglinear model for the conditional intensity (Brown, Frank, Tang, Quirk & Wilson, 1998; Okatan, Wilson & Brown, 2005; Paninski, 2004; Truccolo, Eden, Fellows, Donoghue & Brown, 2005). As special cases of Equation 1, the m-IMI and TRRP models studied here may be considered simplified versions of the more realistic models used in the literature. Because we are focusing on history effects, we would anticipate that results for more complicated settings would be similar to those reported here.

A Derivation of the KL divergence

In this appendix we derive equation 2.9 from equation 2.7 under the conditions that p_f and p_q are periodically stationary with period T and $E[n] \rightarrow \infty$ as $T' \rightarrow \infty$. Let $\{t_1, \dots, t_n\}$ a sequence of spikes, $\theta_i = \theta(t_i)$ be the phase of spike time t_i and $u_i = t_i - t_{i-1}$ be i -th inter-spike interval. The probability density of a spike sequence $\{t_1, \dots, t_n\}$ in the time interval $[0, T')$ for $T' \rightarrow \infty$ whose conditional ISI density is $f(u|\theta)$ is given by

$$\begin{aligned} p_f(t_1, \dots, t_n) &= Q(n) \chi(\theta(t_1)) f(t_2 - t_1 | \theta(t_1)) \cdots f(t_n - t_{n-1} | \theta(t_{n-1})) \\ &= Q(n) \chi(\theta_1) f(u_2 | \theta_1) \cdots f(u_n | \theta_{n-1}), \end{aligned} \quad (\text{A.1})$$

where $Q(n)$ is a probability distribution of spike count in the interval $[0, \infty)$. Equation (A.1) satisfies the normalization condition:

$$\begin{aligned} Q(0) + \lim_{T' \rightarrow \infty} \sum_{n=1}^{\infty} \int_0^{T'} \int_{t_1}^{T'} \cdots \int_{t_{n-1}}^{T'} p_f(t_1, \dots, t_n) dt_1 dt_2 \cdots dt_n \\ = Q(0) + \sum_{n=1}^{\infty} Q(n) \int_0^T \chi(\theta_1) d\theta_1 \int_0^\infty f(u_2 | \theta_1) du_2 \cdots \int_0^\infty f(u_n | \theta_{n-1}) du_n \\ = \sum_{n=0}^{\infty} Q(n) = 1. \end{aligned} \quad (\text{A.2})$$

In equation 2.7, taking the limit of $T' \rightarrow \infty$,

$$\begin{aligned} &\int_0^{T'} \int_{t_1}^{T'} \cdots \int_{t_{n-1}}^{T'} p_f(t_1, \dots, t_n) \log \frac{p_f(t_1, \dots, t_n)}{p_q(t_1, \dots, t_n)} dt_1 dt_2 \cdots dt_n \\ &\rightarrow Q(n) \int_0^T \int_0^\infty \cdots \int_0^\infty \chi(\theta_1) \left(\prod_{i=2}^n f(u_i | \theta_{i-1}) \right) \sum_{j=2}^n \log \frac{f(u_j | \theta_{j-1})}{q(u_j | \theta_{j-1})} d\theta_1 du_2 \cdots du_n \\ &= Q(n) \sum_{i=2}^n \int_0^T \chi(\theta_1) \left(\prod_{j=2}^{i-1} \int_0^\infty f(u_j | \theta_{j-1}) du_j \right) d\theta_1 \\ &\quad \times \left(\int_0^\infty f(u_i | \theta_{i-1}) \log \frac{f(u_i | \theta_{i-1})}{q(u_i | \theta_{i-1})} du_i \right) \left(\prod_{k=i+1}^n \int_0^\infty f(u_k | \theta_{k-1}) du_k \right) \\ &= Q(n) \sum_{i=2}^n \int_0^T \int_0^\infty \chi(\theta_{i-1}) f(u_i | \theta_{i-1}) \log \frac{f(u_i | \theta_{i-1})}{q(u_i | \theta_{i-1})} d\theta_{i-1} du_i \\ &= Q(n)(n-1) \int_0^T \int_0^\infty \chi(\theta) f(u | \theta) \log \frac{f(u | \theta)}{q(u | \theta)} d\theta du, \end{aligned} \quad (\text{A.3})$$

where we use

$$\int_0^T \chi(\theta_1) \left(\prod_{j=2}^{i-1} \int_0^\infty f(u_j | \theta_{j-1}) du_j \right) d\theta_1 = \int_0^T \chi(\theta_{i-1}) d\theta_{i-1}. \quad (\text{A.4})$$

Substituting equation A.3 into equation 2.7 leads to

$$\begin{aligned} D(p_f || p_q) &= \lim_{T' \rightarrow \infty} \frac{1}{E[n]} \sum_{n=1}^{\infty} Q(n)(n-1) \int_0^T \int_0^\infty \chi(\theta) f(u|\theta) \log \frac{f(u|\theta)}{q(u|\theta)} d\theta du \\ &= \lim_{T' \rightarrow \infty} \left(1 - \frac{1 - Q(0)}{E[n]} \right) \int_0^T \int_0^\infty \chi(\theta) f(u|\theta) \log \frac{f(u|\theta)}{q(u|\theta)} d\theta du \\ &= \int_0^T \int_0^\infty \chi(\theta) f(u|\theta) \log \frac{f(u|\theta)}{q(u|\theta)} d\theta du. \end{aligned} \quad (\text{A.5})$$

B TRRP model

In this appendix we show that the rescaling function of the TRRP model, $\lambda_0(t)$, corresponds to the trial-averaged conditional intensity function, up to an arbitrary multiplicative constant. Let $g_0(u)$ be the hazard function of a renewal process and, to avoid transient start-time effects, suppose that the renewal point process starts from $u = -\infty$. Let $e_*(u)$ be the event time prior to u . By the renewal theorem (Theorem 4.4.1 in Daley & Vere-Jones, 2003), the expectation of $g_0(u - e_*(u))$ is constant:

$$Eg_0(u - e_*(u)) = c \quad (\text{B.1})$$

for some positive c . Now transform time from u to t with a monotonic time-rescaling function $t = \Lambda_0^{-1}(u)$, where $\Lambda_0(t)$ is given by Equation 2.6. By the change of variables formula, the conditional intensity as a function of time t is given by Equation 2.5. Taking expectations of both sides and applying Equation B.1 we get

$$E[\lambda(t, s_*(t))] = \lambda_0(t)c. \quad (\text{B.2})$$

If we replace $g_0(t)$ with $h(t) = g_0(t)/c$ then $\lambda_0(t)$ becomes the expected (trial-averaged) conditional intensity function. This is convenient because $\lambda_0(t)$ may then be estimated by pooling data across trials, i.e., by smoothing the PSTH.

Acknowledgment

We thank Taro Toyozumi for helpful discussions.

References

- Barbieri, R., Quirk, M. C., Frank, L. M., Wilson, M. A., & Brown, E. N. (2001). Construction and analysis on non-Poisson stimulus-response models of neural spiking activity. *Journal of Neuroscience Methods*, 105, 25-37.
- Brown, E. N., Barbieri, R., Ventura, V., Kass, R. E., & Frank, L. M. (2002). The time-rescaling theorem and its application to neural spike train data analysis. *Neural Computation*, 14, 325-346.
- Brown, E. N., Frank, L. M., Tang, D., Quirk, M. C., & Wilson, M. A. (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.
- Burkitt, A. N., & Clark, G. M. (2001). Synchronization of the neural response to noisy periodic synaptic input. *Neural Computation*, 13, 2639-2672.
- Cox, D., R., & Lewis, P. A. W. (1972). Multivariate point processes. In *Proc. 6th Berkeley Symp. Math. Statist. Prob. 3*, (pp.401-448). Berkeley: University of California Press.
- Daley, D., & Vere-Jones, D. (2003). *An introduction to the theory of point processes*. New York: Springer-Verlag.
- Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience*. Cambridge, AM: MIT Press.
- de Boor, C. (2001). *A practical guide to splines* (Revised Ed.). New York: Springer.
- Gerstner, W., & Kistler, W. M. (2002). *Spiking neuron models*. Cambridge: Cambridge University Press.
- Iyengar, S., & Mullen, P. (2007). Inference for the Ornstein-Uhlenbeck model for neural activity. preprint.
- Johnson, A., & Kotz, S. (1970). *Distributions in statistics: Continuous univariate distributions*. Vol. 2. New York: Wiley.
- Kass, R. E., & Ventura, V. (2001). A spike-train probability model. *Neural Computation*, 13, 1713-1720.

- Kass, R. E., Ventura, V., & Brown, E. N. (2005). Statistical issues in the analysis of neuronal data. *Journal of Neurophysiology*, 94, 8-25.
- Koch, C. (1999). *Biophysics of computation: Information processing in single neurons*. New York: Oxford University Press.
- Koyama, S., & Shinomoto, S. (2005). Empirical Bayes interpretations of random point events. *Journal of Physics A: Mathematical and General*, 38, L531-L537.
- Lansky P., Sanda, P., & He J. (2006). The parameters of the stochastic leaky integrate-and-fire neuronal model. *Journal of Computational Neuroscience*, 21, 211-223.
- McCullagh, P., & Nelder, J. A. (1989). *Generalized linear models* (2nd ed.). New York: Chapman and Hall.
- Okatan, M., Wilson, M. A., & Brown, E. N. (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., Pillow, J. W., & Simoncelli, E. P. (2004). Maximum likelihood estimation of a stochastic integrate-and-fire neural encoding model. *Neural Computation*, 16, 2533-2561.
- Plesser, H. E., & Geisel, T. (1999). Markov analysis of stochastic resonance in a periodically driven integrate-and-fire neuron. *Physical Review E*, 59, 7008-7017.
- Reich, D. S., Victor, J. D., & Knight, B. W. (1998). The power ratio and interval map: Spiking models and extracellular recordings. *Journal of neuroscience*, 18, 10090-10104.
- Silberberg, G., Bethge, M., Markram, H., Pawelzik, K., & Tsodyks, M. (2004). Dynamics of population rate codes in ensembles of neocortical neurons. *Journal of Neurophysiology*, 91, 704-709.
- Snyder, D. L., & Miller, M. I. (1991). *Random point processes in time and space* (2nd Ed.) New York: Springer-Verlag.
- Truccolo, W., Eden, U. T., Fellows, M. R., Donoghue, J. P., & Brown, E. N. (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. C. (1988). *Introduction to theoretical Neurobiology*. Cambridge: Cambridge University Press.

van Kampen, N. G. (1992). *Stochastic processes in physics and chemistry* (2nd ed.). Amsterdam: North-Holland.

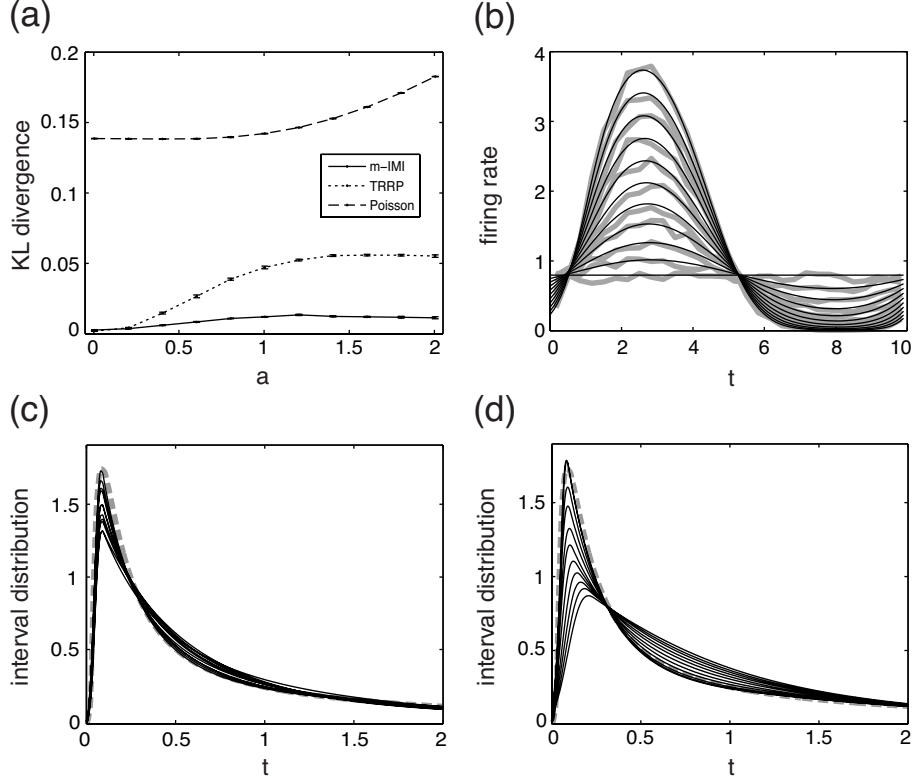


Figure 1: Results for the case that the mean input is varying in time. (a) The KL divergence as a function of the amplitude of the mean input, a . Solid line, dotted line and dashed lines represent the KL divergence of the m-IMI model, the TRRP model and the inhomogeneous Poisson process, respectively. The mean and the standard error at each point were calculated with ten repetitions. The KL divergence of the inhomogeneous Poisson process is much larger than that of the m-IMI model and the TRRP model. (b) The instantaneous firing rates of the LIF neuron for various values of a . Thin solid lines represent the instantaneous firing rates of the LIF model, and the gray thick lines are raw traces. The amplitude of the instantaneous firing rate is increasing as a is increasing. (c) Solid lines represent rescaled interval distributions of the m-IMI model for various values of a , which are obtained from the recovery function by equation 3.2. The gray dashed line is the interval distribution of the LIF model for $a = 0$. (d) Same as (c) for the TRRP model. The rescaled interval distributions both in (c) and (d) are departing from the interval distribution of the LIF model as a is increasing, but it shows less variation for the m-IMI model than for the TRRP model.

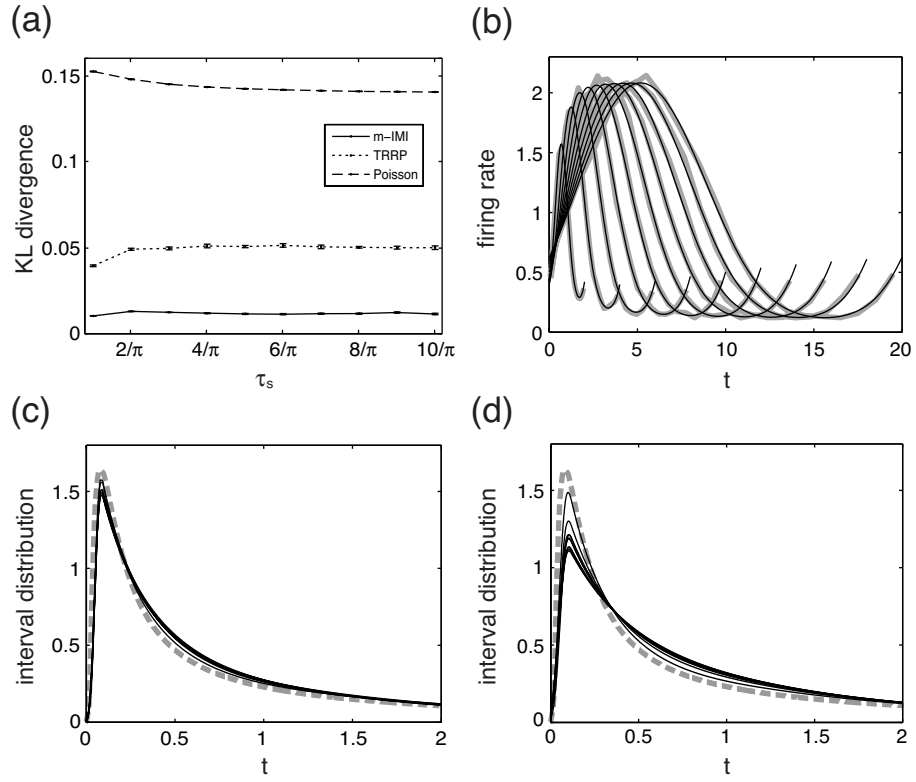


Figure 2: Results for various values of the time constant of the time-varying mean input, τ_s . (a)-(d) are the same as Figure 1.

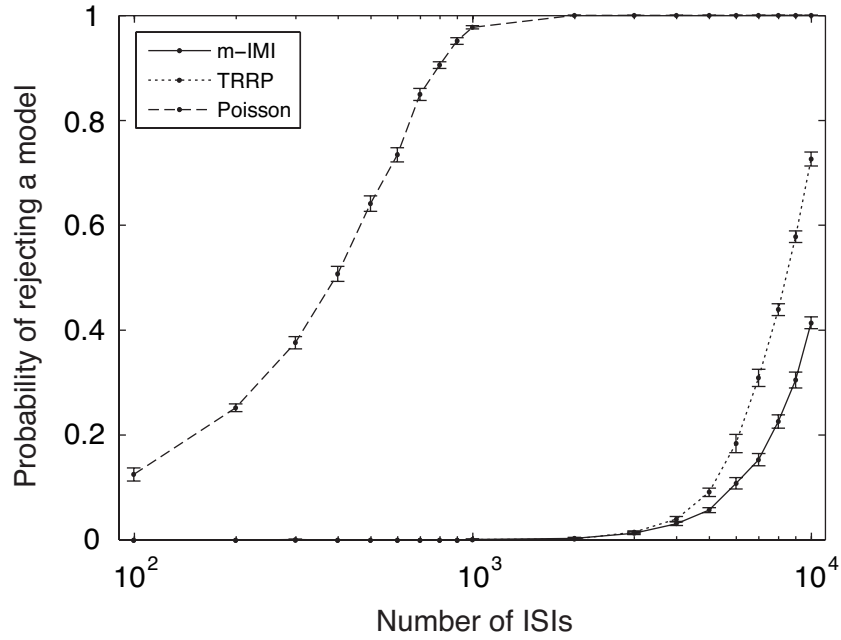


Figure 3: The probability of rejecting probability models as a function of the number of ISIs. Solid line, dotted line and dashed represent the m-IMI model, the TRRP model and the inhomogeneous Poisson process. The mean and the standard error at each point were calculated with 10 repetitions. With relatively small sample sizes it can be easy to reject the hypothesis of Poisson firing, but very hard to distinguish m-IMI model from TRRP models.




Figure 4: (a) The Kolmogorov-Smirnov plot using 200 ISIs showing the Poisson model does not fit, but the two other models do. (b) Same as (a) using 7000 ISIs showing that the m-IMI model fits, but the two other models do not.

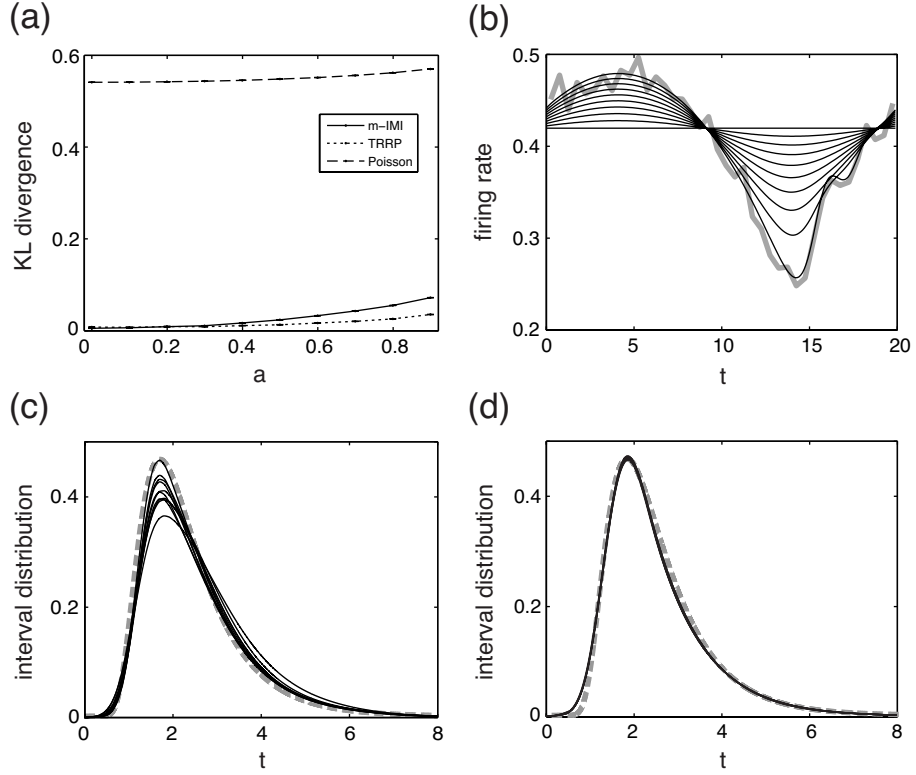


Figure 5: Results for the case that the input variance is varying in time. (a)-(d) are the same as Figure 1. Only the raw trace of the firing rate for $a = 0.9$ is shown in (b). In this case the KL divergence of the TRRP model is the smallest among three models, and the interval distribution for the TRRP model shows less variation than for the m-IMI model.

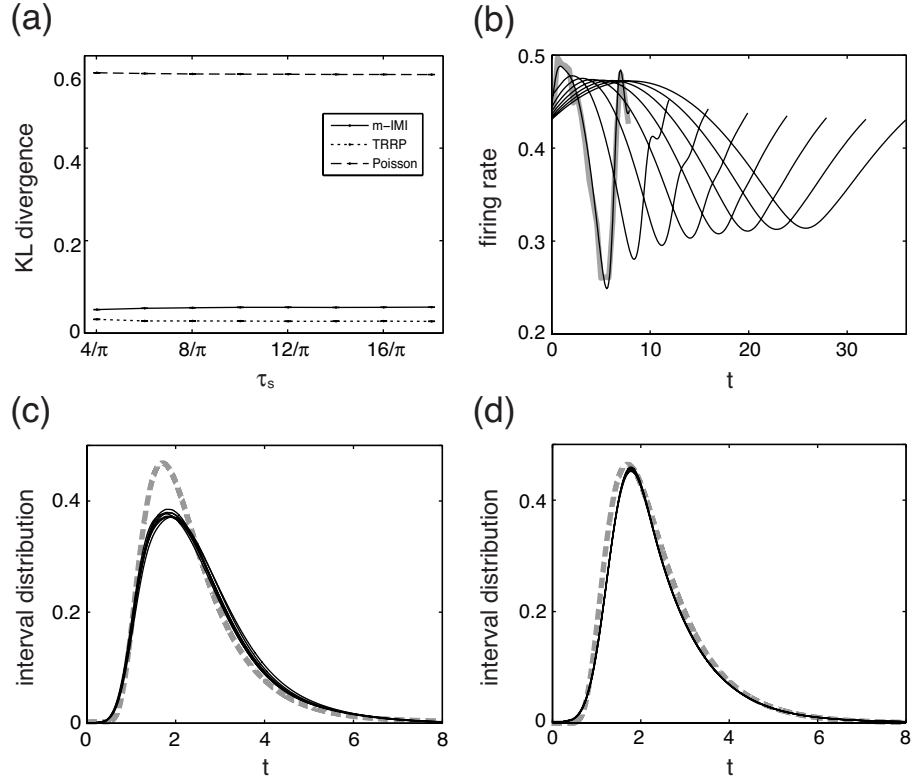


Figure 6: Results for the various values of the time constant of the time-varying variance of the input, τ_s . (a)-(d) are the same as Figure 5. Only the raw trace of the firing rate for $\tau_s = 4/\pi$ is shown in (b).

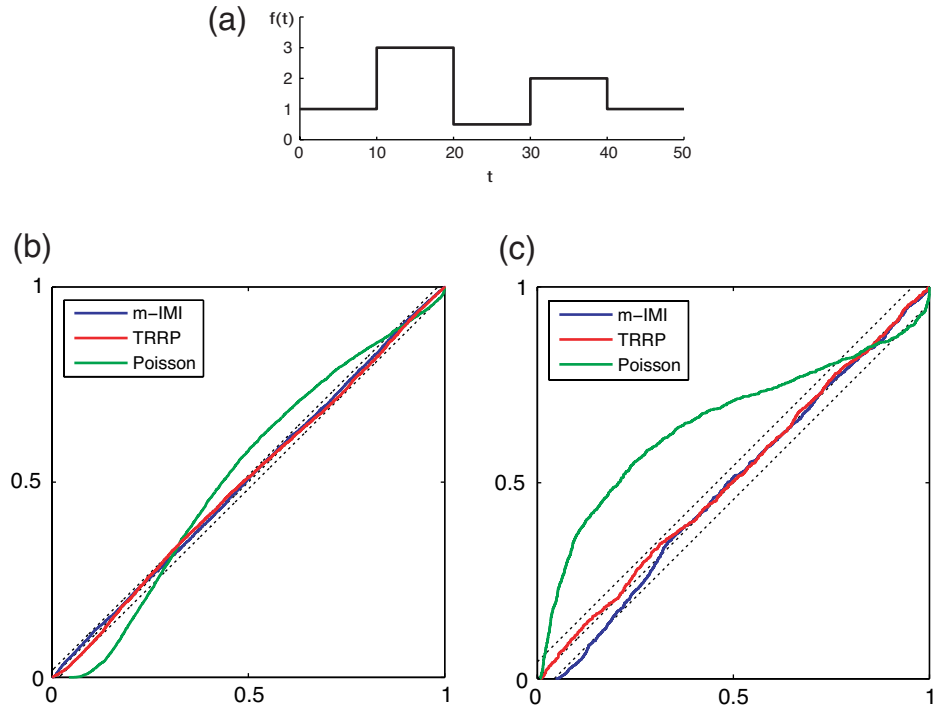


Figure 7: (a) The shape of the stimulus, $f(t)$. (b) Kolmogorov-Smirnov plot using 6000 ISIs showing that the m-IMI model fits, but the two other models do not for the mean-modulated stimulus. (c) Same as (b) using 1000 ISIs showing that the TRRP model fits, but the two other models do not for the variance-modulated stimulus.